www.forgetit-project.eu

ForgetIT

Concise Preservation by Combining Managed Forgetting and Contextualization Remembering

Grant Agreement No. 600826

Deliverable D2. 3

Work-package	WP2: Foundations of Forgetting and Remembering
Deliverable	D2.3: Foundations of Forgetting and Remembering: Intermediary Report
Deliverable Leader	Robert Logie
Quality Assessor	Vasileios Mezaris
Dissemination level	PU
Delivery date in Annex I	2015-01-31
Actual delivery date	2015-02-10
Revisions	10
Status	Final
Keywords:	Human Memory and Forgetting, Personal Preservation, Organizational Preservation, Digital Memory Support for Human Memory

Revision History

Version	Major changes	Authors
0.1	Draft of Text on Human Autobiographical Memory	Elaine Niven
0.2	Draft of formal modelling to complement human memory	Walter Allasia
0.3	Draft of text on organizational preservation and forgetting	Viktor Mayer-Schönberger,
0.4	First overall draft collating and editing text from above versions plus defining overall structure and outline of content	Robert Logie
0.5	Addition of text on experimental studies	Elaine Niven
0.6	Addition of text on surveys	Maria Wolters
0.7	Matching text structure to standard template	Elaine Niven
0.8	Editing and finalizing complete draft for internal project review	Robert Logie
0.9	Internal quality control/internal project review	Vasileios Mezaris,
1.0	Implement suggested edits and submit	Robert Logie, Maria Wolters, Elaine Niven

List of Authors

Partner Acronym	Authors
UEDIN	Robert Logie, Elaine Niven, Maria Wolters
UOXF	Viktor Mayer-Schönberger
EURIX	Walter Allasia

Table of Contents

Executive Summary

This deliverable from Workpackage 2 is a report on:

(a) The intermediate results of a review of research literature on human memory and forgetting (section 2).

(b) The intermediate results of two phases of an experimental study that collected objective data on how samples of human users recall details of personally experienced events and how they use and organise digital photographs of those events. Photographs were grouped and contextualised by the human volunteers themselves and, in a collaboration between UEDIN (WP2) and CERTH (WP4), were automatically grouped and labelled by the pictorial categorisation algorithms developed by CERTH. In a further collaboration between UEDIN (WP2) and DFKI (WP9), individual users also sorted and managed photographs using the Personal Information Management systems developed by DFKI, in an initial evaluation of its usability (section 3.3).

(c) The intermediate results of a public, multi-country and multi language survey collecting self report data on how human users organise, and make keep/delete decisions on their collections of personal digital photographs. The survey also gathers opinions on what users might expect of a system that might support their selective preservation of their digital photographs. A complementary survey of information storage and management within organisations has also been available online, in collaboration between UEDIN (WP2) and DKD (WP10), (section 4).

(d) The intermediate report of work on formal modelling of human memory that might be implemented in a digital preservation and managed forgetting system (section 5).

(e) The intermediate results of a review of research literature of organisational memory, and the development of conceptual modelling of the interaction between human and digital information storage and management within organisations (section 6).

1 Introduction

A general review of empirical evidence and theoretical approaches to understanding the organisation, conceptualization and functioning of human memory was presented in D2.2. In summary, human memory relies on using a lifetime accumulation of knowledge about the world and about the self, collectively referred to as ‘semantic memory’ to set the context for both preserving and retrieving details of events, referred to as ‘episodic memory’. The latter refers specifically to a combination in memory recall of what event happened, where it happened, when it happened, and how the event is related to memories of other public and personal events. Episodic memory is characterized by rapid loss of detail shortly after an event. Semantic memory is more robust over time and is much less affected by forgetting. Evidence for this distinction has come from a wide range of studies with healthy adults (e.g. see review in Baddeley, Anderson & Eysenck, 2015), and from studies of individuals who have suffered memory impairments as a result of brain damage. For example, there are numerous studies of brain damaged individuals who have a serious problem with recall of experienced events, but tend to have intact access to knowledge about themselves, their family history, their occupation, and unimpaired knowledge of their native language (reviews in Baddeley, Kopelman & Wilson, 2002). Moreover, retrieval of experienced events involves a process of context based reconstruction, drawing context from semantic memory, not the recall of a veridical record (originally established by Bartlett, 1932). The interaction between episodic memory and semantic memory is the core of this reconstruction process. The process is very efficient for storage and retrieval and works extremely well for most scenarios, but it results in rapid and substantial forgetting of detail of specific events. The reconstruction process can lead to errors and is subject to interference from subsequent and previous experiences, with the individual being unaware that their recall is inaccurate. Key to the success of a future ForgetIT system will be the identification of how it might support the weaknesses of human memory and avoid the common error in the design of digital systems of attempting to replace functions for which humans are highly efficient. In D2.2, we considered the relationship between semantic and episodic memory in general terms. Here we build on that previous discussion. In this first main section of the deliverable, we consider the nature of this relationship for personal memories, or autobiographical memory, in greater detail.

A further important issue is generalisability of findings from studies of human memory. It is clear that each individual has different memories of experiences, and different contents of their semantic and episodic memory systems. Also, people vary in the efficiency with which they can learn and retrieve new knowledge, and which they can encode, forget, preserve or retrieve details of events that they have experienced. However, the study of human memory is based on the assumption that the general principles that govern human memory organisation and function is the same across all healthy human adults, and are the result of the effect of evolution on the human brain. So, there are common principles across all humans of learning, forgetting, encoding, preserving, retrieval and so on. Many of these principles also apply to many animals. An analogy would be that the principles by which other aspects of human physiology function, such as the heart, the liver, the kidneys, the lungs, the immune system or the endocrine system, are the same across all healthy human adults, even if these systems differ in their efficiency between individuals. On this principle, if we are collecting objective memory performance data, we need only study the principles of functioning of human memory in a single healthy adult and these principles should generalise to all healthy human memory. However, experimental data are inevitably noisy and an individual who is studied might have some underlying anomaly such as subtle undetected brain damage, or be simply unco-operative. Therefore, in practice multiple healthy participants are typically recruited and allocated randomly to different experimental groups for the purposes of comparison. Using this approach, the numbers in each group need not be large, although there is usually a matching in the age range and educational level to mitigate the variability from differences in memory efficiency. Data are then averaged across participants within each group to reduce the impact of possible idiosyncrasies of any one participant. Therefore, unlike studies that involve subjective opinions, self-report or survey data, there is no need in experimental studies of this kind to be concerned about sampling from cross sections of the population.

2 Literature Review: Personal preservation

2.1 Studies of Autobiographical Memory

Autobiographical memory refers primarily to memory for real-life personal events in contrast to episodic memory for more artificial events in a cognitive psychology laboratory (e.g. remembering words and pictures presented in an experiment). It involves the interaction between episodic details of a specific event (what-where-when) and accumulated previous life experiences and accumulated knowledge (semantic memory) (e.g. Conway & Pleydell-Pierce, 2000; Conway & Loveday, 2010). Research on autobiographical memory variously focuses on more semantic or more episodic elements of an individual's personal history. The most influential factors in determining the episodic versus semantic composition of the reconstruction of a memory for an event have been identified as *age of memories, event frequency, rehearsal* and *age of participants* (Cabeza & St Jacques, 2007). The real world context of a person's life history and the factors relating to time since the event, frequency, rehearsal, and age do not operate in isolation, rather they interact. Detailed memory for events throughout the lifespan also requires a complex interplay of a number of aspects of the memory system (see e.g. Conway, 2009).

The study of episodic memory and specifically of autobiographical memory, is as varied as the approaches to defining these forms of memory: from studying episodic-like memory based only on the requirement of representing 'what-where-when' information (e.g. Clayton & Russell, 2009), through incorporating temporal-spatial based autobiographical-event memory (e.g. Maguire, 2001), to others who rigorously delineate episodic memories as a specific subset or type of memory. Conway (2009) sets out nine defining features of what makes a memory an episodic memory, all of which should be present, namely:

- Contain summary records of sensory-perceptual-conceptual-affective processing
- Retain patterns of activation/inhibition over long periods
- Often represented in the form of (visual) images
- They always have a perspective (field or observer)
- Represent short time slices of experience
- They are represented on a temporal dimension roughly in order of occurrence
- They are subject to rapid forgetting
- They make autobiographical remembering specific
- They are recollectively experienced when accessed.

2.2 Methodologies for the study of autobiographical memory

The systematic study of autobiographical memory is complicated by the large number of factors that have to be considered, and the lack of experimental control over spontaneous or even planned events that occur in everyday life. Individuals also differ in the extent to which they retain details, even of the same event, and which details are retained. Finally, rarely is there a full objective record of an event that can be used to assess the accuracy of the individual's memory for that event. This last feature is particularly problematic given the reconstructive nature of memory.

An early approach to the study of autobiographical memory involved participants keeping a diary or log of events. After delays of days, weeks, months or years, one detail from a recorded event can be used as a cue for participants to recall other details of those events. The original record can be used to check for recall accuracy (e.g. Linton, 1975). A related approach has been to prompt participants at random intervals (e.g. via a pager; Brewer, 1988) to record details of the most recent events in their daily life, and again use these records as sources to cue specific events for later recall and for checking accuracy. However, these methods suffer from two major limitations. First, the fact that participants choose the events and the details of those events, and take the time to record them acts as a form of ‘mental rehearsal’ of the events that would make it more likely that those events will be recalled when tested at a later date. Second, given the difficulties of running studies over more than a few weeks or months, the method does not allow testing of memory for events from the more distant past. The Linton (1975) study took place over six years, but that involved a single participant (Linton herself) and few research projects have the resources for this length of study, or for the numbers of participants that would be required for a comprehensive study.

When there has been no specific documented record of an event, other sources such as family and friends can be used to check recall accuracy, but they too will be subject to memory recall errors. If it was a public event then there may be publically available documentation, video or audio records to verify memory accuracy. This method has been used very successfully to demonstrate errors in recall of major public events, such as errors in the recall by President George Bush junior of the tragic events of September 11th, 2001 (Greenberg, 2004). Therefore many researchers have focused their investigations on memory for, or surrounding public events (see D2.2 for a review of the relevant literature).

The review of memory for public events in D2.2 indicated that the importance of an event for the individual, or their emotional response to an event is a major factor in the level of detail that can be recalled. Kopple and colleagues (2013) found that after delays of around four months after a major public event, the ratings of 'emotional intensity’ and ‘personal importance' regarding the events were the best predictors of the level of consistency in their recall of the event. However, the fact that an event resulted in an emotional response or was of personal importance would also lead to the individual repeatedly thinking about or talking about the event, thereby having repeated recalls and rehearsal of the event before the formal test of their memory within the study. The extent of media coverage of the event may also be important in driving this form of repeated recall.

A further approach has been to simulate natural autobiographical events by staging them within a controlled experimental setting. This has the advantage that the same or very similar events can be experienced by a large number of participants, thereby removing the confounding factor that some events are more memorable than others or include different amounts of detail (e.g. Marsh & Roediger, 2013). Such experiences and associated memory will not be ‘pure’ or free from influence of a participant’s personal, autobiographical knowledge. For example, certain elements of a constructed experience may resonate strongly with some participants more than others because of their individual past experiences. However, if a sufficiently large number of participants is included then averaging the results across participants can help mitigate the influence of these individual influences on recall. Moreover, constructed situations allow researchers to manipulate components of experiences, enabling generalisable insight into the features of interest in any given experiment such as delay until recall, nature of the event, number of times recall is repeated, effects of different memory cues and so on.

Difficulties inherent in trying to characterize, measure, quantify, and understand the information that remains available and accessible from an experienced event still persist even when many of the external characteristics of the experience are under the control of an experimenter. Williams, Conway and Baddeley (2008) made an attempt to address some of these factors by tackling the question of event boundaries in memory: when does an event begin and end, and does this change over time? The approach builds on the work of Zacks and colleagues (e.g. Zacks et al. 2007; Zacks & Swallow, 2007; Kurby & Zacks, 2008), regarding boundaries of events when looking at and perceiving the world, and conceptual features (for example, hierarchically nested goals, which can determine ‘fine or coarser grained’ events). In the context of Conway’s general framework (e.g. 2005, 2009) episodic memory is considered a record, or the product of, carrying out short term goals in the service of increasingly greater over-arching goals; while moment-to-moment goals are necessarily dynamic and short-lived, goals further up the hierarchy are more stable and enduring. Williams et al. (2008) asked participants to freely recall their morning commute on the day of testing, to identify how they would segment this into ‘discrete memories’ and then to recall the same commute one week later. After analyzing the memory elements that began and ended a segment, Williams et al. (2008) determined that actions were more likely (but not exclusively) to characterize a beginning, with a fact (more likely but not exclusively) ending the segment. Other features changed with time of recall. This so-called ‘action-fact’ structure to the memories was replicated for three memories from a more remote time of a recent holiday, and was taken as support of memories defined around goal structure, with actions indicating an initiation of the current short-term goal. Notably, the authors observed that despite the general pattern of facts ending a segment, the properties with which a segment could end were more variable than those with which they would start, and this indicated that goals may be terminated in various ways.

Hohman, Peynircioglu and Beason-Held (2013) tested event memory for events from different time points in the lives of individual participants. They found that college age and middle age adults were more likely to be flexible in their attribution of event boundaries with increasing time – that is, once a memory was recalled and probed for further information, the more time that had passed since the event, the more likely the adults were to accept information outside of the original boundaries for the event as belonging to the original event (therefore moving event boundaries). Overall, older adults were more likely than college age adults to move event boundaries, indicating an age related component to event structure. Such results are important in demonstrating that not only does the nature of creation of event boundaries at encoding determine how we should approach our memories for events, but so too does the nature of event boundaries at retrieval/reconstruction.

2.3 Photograph use and event memory

Within the ForgetIT project, our focus is on the interaction between human memory and digital information storage, with digital photographs as the use-case scenario. In D2.2, we reviewed the key previous research studies on the use of photographs in the study of human autobiographical memory. As that research provides a context for the experimental work reported in the next section, we first summarise that previous research here.

Some of the previous studies have investigated whether viewing photographs after an event helps or hinders subsequent recall of that event. Across several studies, Koutsaal and colleagues (1998; 1999) and Schacter and colleagues (1997) among others tested participants’ memory for a previously shown video or set of events. If participants were shown photographs taken from the video, this helped later recall of the details from the video that also appeared in photographs. However, in those same studies, showing new photographs of scenes that were not in the video led participants later to falsely recall that details in the new photographs had actually appeared in the original video. That is, people remembered the details they had seen but could not remember the original source of those details, leading to false memories. A related study by St Jacques and Schacter (2013) showed very similar results following a tour of a museum, wearing a camera that automatically took pictures. Again, showing photographs of events from the tour helped subsequent recall of those events, but photographs of objects that had not been seen on the tour led to subsequent false memories that the tour had included those objects. These results clearly illustrate that memories are reconstructed and subject to change, and this process of reconstruction can lead to false memory for the original source.

A further study by Henkel (2013) showed that the act of taking photographs can impair memory. Following a tour of an art museum, participants showed that they had rather poor memory for objects that they had been asked to photograph compared with objects that they were simply asked to look at. Asking people to zoom in to objects when taking the photograph counteracted the negative effect on memory. Henkel suggested that reliance on digital memory storage has negative effects for our own memory, and that taking a photo serves as a cue to forget because there is a digital record, unless some additional cognitive effort is involved, such as zooming in. The cognitive effort and attention expended on a zooming task appears to counteract this effect. However, as acknowledged by Henkel, this apparent negative impact on memory may not extend to situations in which people make a choice about which objects they photograph. She therefore called for further research to disentangle the factors that may contribute to the contribution of photographs to memory in real-life situations. The experimental studies described in the next section comprise, in part, a response to that call, in addition to providing insight into digital photograph use for the ForgetIT project. However, the results do suggest that the act of storing information digitally might, under some circumstances, make it more likely that the individual will forget what has been stored.

3 Experimental Studies on Personal Digital Photograph Use

In order to design a future ForgetIT system that will be adopted by potential users, it is crucial to identify how such potential users currently use digital storage for personal information, and how that use impacts on their biological memory for the original source information, such as an experienced event. As for other project partners (e.g. DKFI, CERTH) we have focused on personal digital photograph use. In this section we report on experimental studies to collect objective data on this use-case scenario.

As reported in D2.1, in August 2013 the UEDIN team conducted a study centred around a section of the annual Edinburgh Fringe Festival, a large street fair (musicians, acrobats, street theatre, stand-up comedy) that takes place in the centre of Edinburgh, UK over three weeks in August each year. In this study participants spent an hour at the street fair and were tasked with taking snapshots every three minutes to document this experience. Through using the Edinburgh Festival Fringe as a setting for studies centred around photograph use to support memory for an event, we were able to take advantage of a great deal of activity happening in a fairly well defined area: approximately 400 metres along one street. This provides a substantial amount of potential information that participants may want to remember, and also provides a rich opportunity and variety of contexts with which to link their photographs upon subsequent reviewing. More specifically, this enables investigation of a series of research questions about photograph use when experiencing and mentally revisiting an event. D2.1 was completed very shortly after completion of our data collection in August 2013, and so provided a general description of our methodology, but with limited analysis. Here we report a more detailed analysis of data collected in 2013, provide a preliminary analysis of a follow up in which participants in that study returned to the laboratory approximately 12 months after the event, and the methodology and preliminary analysis of a second study carried out during the Edinburgh Festival Fringe in 2014.

3.1 August 2013 Festival Fringe experiment

In summary, we examined three factors regarding participants' memory and photographs of the events they experienced during the one hour visit to the street fair: the contribution of episodic detail to their recall, which photographs they chose to keep, and their methods for organising photographs into meaningful groups. These elements were investigated for changes over time, and episodic recall was investigated for potential changes over time when having viewed or not viewed photographs as part of a post-event review.

For testing, participants were split into four pseudo-randomly selected groups (random allocation constrained by dates upon which participants were able to commit to returning; see table 3.1 for demographic details per group). Immediately after their hour of festival experience finished (hereafter referred to as Time 1), all participants returned to be interviewed: they orally recalled their experience in as much detail as possible. After this, Groups 1-3 were required to view and make a series of decisions about their photographs. Groups 1-3 then repeated the process (verbal recall, followed by photograph review) after a day (Group 1), a week (Group 2), or a month (Group 3). Group 4 followed the same schedule as Group 3, but did not review their photographs at Time 1. Second interviews are hereafter collectively referred to as Time 2.

Table 3.1: Demographic details of participant groups

Group	1	2	3	4
Number of participants	20	18	18	18
Age (years)	31.65 ±5.52 (19-68)	38.61 ±20.41 (19-71)	33.94 ±17.54 (20-72)	35.17 ±15.02 (19-59)
Males	3	8	5	10

3.1.1 Recall data

As noted in section 1, rapid loss of access to episodic detail in memory immediately following an event is often emphasised as an inherent feature of this type of memory. However, while the large volume of loss is well established, the nature of what is preserved longer term is not entirely clear. Moreover, reviewing an event – either through verbal recall, or through use of photographs - can benefit memory for that event, including the amount of episodic detail that can be remembered. Additionally, previous studies have demonstrated that key details of an event can be maintained consistently over periods spanning a number of months after an initial recall has taken place (e.g. Wynn & Logie, 1998). The outcome of this present study will therefore contribute new evidence regarding preservation of detail by addressing the research questions of whether episodic details change over time and with photograph viewing at post-event review.

For the 2013 Festival Fringe study, the Autobiographical Interview (Levine et al., 2002) scoring technique was applied to recall data generated at Time 1 and Time 2 from each of the four groups. This method is very time consuming in that it requires transcription of audio recording of the oral recall of each participant, and detailed content analysis of the transcriptions by a trained researcher. However, it is a well established and rigorous method for assessing the contribution of episodic information to recall. Comparing the proportion of recall that is episodic in nature at Time 1 and Time 2 across the four groups allows investigation of the effect of time, and the effect of presence of photographs at post-event review, on episodic recall. Additionally, content of recall and how this changes over time in relation to whether or not photographs have been viewed at initial interview also were addressed. Within each interview session, participants were asked to choose from all of their photographs those that they would like to keep and those that they would like to delete.

3.1.2 Photograph data

As reported in D2.1, thematic analysis was used to assess the types of categories that people used to organise their photographs, and this revealed that, with the present data set, categories that were based upon properties of external entities, such as buildings or characteristics of the location and the event, appeared to remain relatively stable over time, while groupings of photographs that were made according to personal criteria – such as particularly salient experiences (which may lead to initial but perhaps not lasting fine-grained categorization) - or other idiosyncratic categories showed more fluctuation with time.

Since the report in D2.1, more detailed analyses of the data have been undertaken, and some of these analyses have been published (Wolters, Niven & Logie, 2014). The decisions of participants to keep or delete each individual photograph that they took throughout the hour spent in the festival (*mean* = 19, *range* = 12-20 photographs) were of particular interest in order to infer preservation value changes over time (from Time 1 to Time 2). As represented in Figure 3.1, those participants who made decisions about their photographs at Time 1 (groups 1–3) selected to delete, on average, 28% of their photographs (median: 26%, range 0%-58%, SD 15%). At Time 2 (a day, week, or month later) these participants deleted on average 33% of their photographs (median: 30%, range: 0-70%, SD 17%). This difference of 5 percentage points (SD: 12) corresponds to one photograph and is statistically significant ($t(55)=3.3352$, $p < 0.002$, 95% CI for median [2%,9%]).

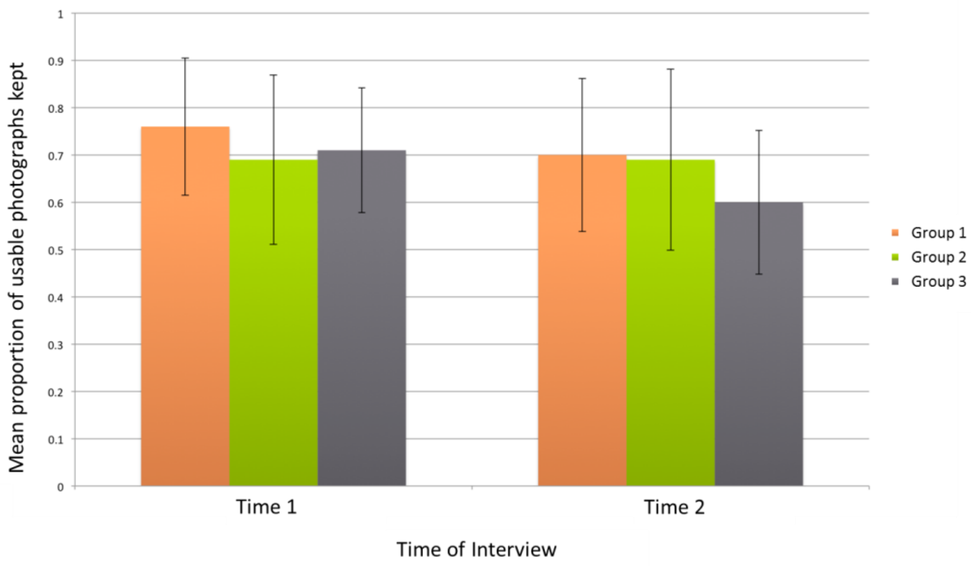

Figure 3.1. Bar chart of proportion of usable photos kept, at Time 1 and Time 2, as broken down by group. Standard error bars are shown.

In further probing of the keep or delete decisions, it was observed that 7 participants (13%) decided to keep exactly the same photographs when interviewed at both Time 1 and Time 2, while 7 other participants (13%) deleted four or more photographs at Time 2 that they had previously kept at Time 1. Overall, 32 participants (59%) did not choose to keep any photographs at Time 2 that they had not already selected as being kept at Time 1, while 37% decided to keep 1–3 additional photographs at Time 2 compared to Time 1, and two outliers added 7 and 9 photographs at Time 2 compared to those they had opted to keep at Time 1. Of note, the percentage of photographs deleted at Time 2 did not differ between the different groups of participants (one-way ANOVA, $F(1)=0.94$, $p < 0.4$), indicating (within the time frame tested) no effect of length of time since initial interview on deletion decisions.

Participants in Group 4 (that is, those who first saw and interacted with their photographs a month after the experience) deleted on average 33% of their photographs (median: 31%, range: 0%–78%, SD 19%) at Time 2. This number is comparable to that observed in Groups 1-3 at Time 2, indicating that viewing photographs immediately after an experience does not produce a strong influence on deletion decisions made at a later date.

In order to better understand the decision process of participants, following their keep/delete decisions participants were probed with further questions. Of those photographs that they opted to keep, participants were asked for five pictures to provide an explanation for why they wanted to keep that particular photograph, and of those that they chose to delete participants were likewise asked for the motivation behind deleting that particular picture. Answers were noted, and, using thematic analysis, six categories of reasons for keeping photographs were derived from all provided explanations. Likewise, six reasons for deletion were derived. Results are summarized in Table 3.2. This method enabled participants to provide more than one 'reason' (category) per explanation or per photograph, for example "Enjoyed their music. Not a good photo composition but it's a reminder of the atmosphere" (which includes the categories of reasons 'Reaction', and 'Personal').

As noted in the published paper reporting this material (Wolters et al., 2014), in sum the 'reasons' that people provided as motivation behind keep and delete decisions for their small photograph collection of this specific event were in part highly predictable, while in part also very subjective. Therefore, it appears that criteria measures that mimic or simulate these judgments could potentially be computed (semi-)automatically (see Section 5). Others require views that help users make the comparisons that are at the heart of aspects of the decision making process. These are highly subjective value judgments and would at least require an automated digital system to learn the preferences of a given user over time, therefore possibly needing substantial user input, in order to either support or simulate user decision making.

Table 3.2: Definitions and relevant examples for each category of reason produced by thematic analysis of participants' keep and delete decision explanations.

Decision	Category of reason	Definition of category	Example quote*
Keep	Typical	Photo as something typical of Festival/ Edinburgh/ Scotland etc.	"Shows Scotland"
	Aesthetics	The composition, quality or staging (or aesthetics otherwise) of the photograph.	"Came out well"
	Reaction	Something the person found funny/ interesting/ scary etc.	"It shocked me the way his eyes are painted – different and unusual"
	Sharing	Would be used to tell stories about the person's experience of street fair	"The man whose mask kept falling off! ...would tell people the story when showing them the photo"
	Subject matter	Specific act, item or view etc.	"It's a Storm Trooper!"
	Personal	Personal relevance to person, their life, or time at street fair.	"Start of Journey. Not successful photo, but start of journey"
Delete	Aesthetics	Bad quality (such as, blurry, fingers blocking view)	"It's fuzzy"
	Surplus	Content was already covered (better) in another picture	"I've already got a picture of them"
	Random	Pictures were just taken because the task alarm went off, indicating a picture should be taken	"buzzer shot"
	No shows	Photograph did not capture content participant had intended to photograph	"Wanted a picture of the street but here (there is a) bus and no view"
	Untypical	Photograph included intended subject matter, but on viewing did not represent the city/ country/ festival	"The photo could also be Christmas, not necessarily Fringe"
	Reaction	Photograph evoked negative feelings or connections, or were uninteresting	"Didn't like act"

* note that example quotes may demonstrate more than the example category they are listed alongside.

Figure 3.2 shows the percentage of deletion explanations containing a justification/reason per category, collapsed across participants in Groups 1-3. For example, from all of Group 1-3's delete explanations at Time 1, 39.5% referenced poor photo quality/aesthetics. Additionally, not necessarily exclusively (that is, a photo could have been both poor quality and something else), 42% mentioned that they evoked a negative reaction. Across Time 1 (that is, on day of testing) and Time 2 (here, collapsed across day, week and month later) it is evident that there are slight but not drastic changes. The largest differences emerged in 'aesthetic' and 'random' deletion reasons; specifically, participants' aesthetic judgments appeared to feature more in decision making at Time 2 than they did at Time 1, while they were less likely to suggest 'random' content of a picture led to it being deleted at Time 2 than they were at Time 1.

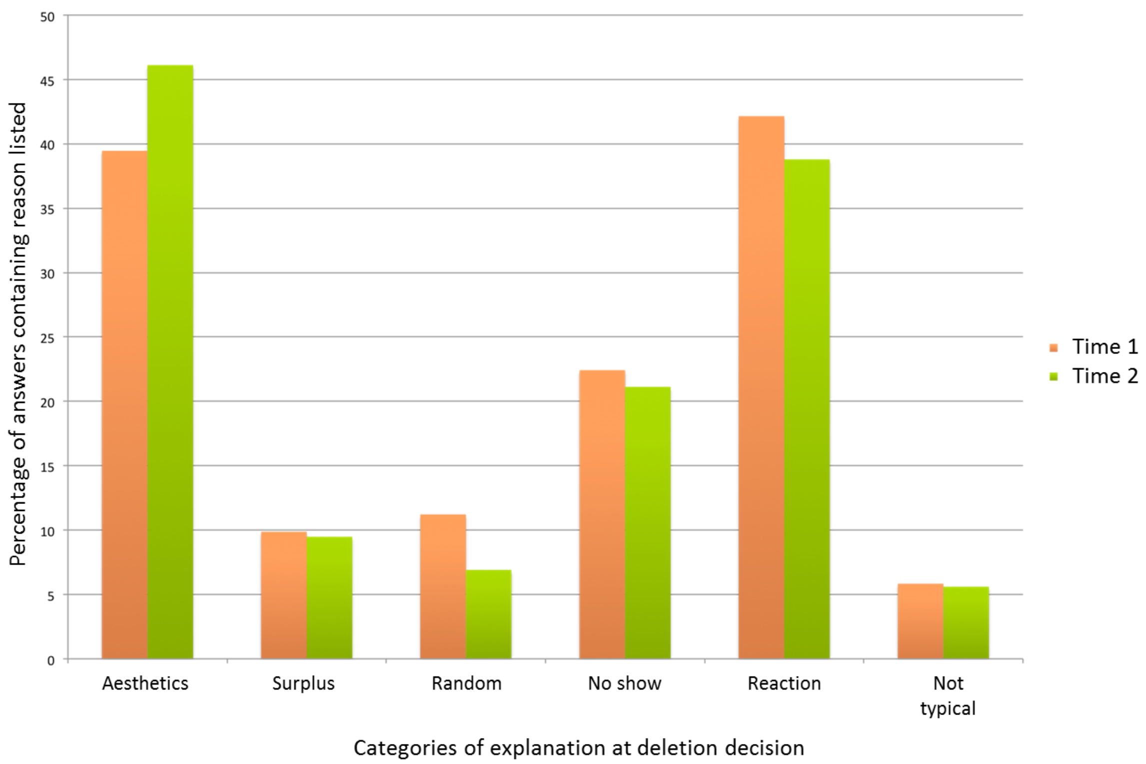

Figure 3.2: Percentage of explanations for delete decisions containing thematic analysis categories of reasons, observed at Time 1 (day of testing) and Time 2 (day, week, month later), collapsed across Groups 1-3.

It is important to note that our participants' decisions to keep or delete photographs were often explicitly made in relation to other photos – for example, "(I have) Another one that's better as mid performance" or "Was walking out of frame; have other better shots of her". Moreover, this factor is also evident when looking more closely at descriptions that did not explicitly state a reason for deletion that could be identified as surplus – for example, the description "Not a good angle, things included are okay but e.g. too much of chairs – would rather reshoot" does not preclude the possibility that this non-optimal photograph would be kept if other photographs of the event did not exist. That is, sometimes a non-optimal photograph is better than no photograph. Therefore, these data suggest that to best support keep/delete decisions, interfaces should make it easy for users to browse sets of related pictures. 'Related' pictures may refer to, among other properties: pictures taken at a certain time, of a certain location, of particular buildings, or event pictures that represent visually similar views, despite originating in different places. As noted in Wolters et al. (2014), it is important to offer users flexibility in the views they can define. Relations between photos can be highly subjective, in particular when photos remind people of previous experiences that might not be documented in the photo management system. While automatic image annotation can support the process, flexible tagging facilities are also required. These latter aspects (relating photographs, automatic image annotation, and tagging facilities) are of particular interest to the ForgetIT project and the beginning of analyses to address these processes feature in sections 3.2.2.1 and 3.3.1.

Consistency of groupings

In the study, participants were required to sort their photos into meaningful groups several times. Participants were asked to produce two sorts at Time 2, and those who saw their photos at Time 1 (Participant Groups 1-3) additionally sorted them once on the day of the study.

In collaboration with student Zeyu Wang at UEDIN, we are currently investigating to what extent the size, composition, and meaning of groups changed between sorts when participants repeatedly sort the same material into groups. Quantifying the variability in human judgments is important when evaluating automatic methods for grouping complex digital artifacts, such as photos, when the resulting groups are intended to be browsed, searched, and manipulated by people. Instead of right or wrong groups, we have more or less acceptable, useful, and meaningful groups.

Similarity between two groupings G_1, G_2 is quantified as follows.

First, for each pair of groups g_i, g_j, with $g_i \in G_1$ and $g_j \in G_2$, we compute the similarity $S(g_i,g_j)$ using different indices proposed in the literature, such as the Jaccard index, the Sørensen-Dice index (Sørensen, 1948), and the symmetrical version of the Tversky index (Tversky, 1977)

We then map G_1 onto G_2 using the Kuhn-Munkres matching algorithm for bidirectional weighted graphs (Kuhn, 1955/2010) and use this mapping to compute the overall similarity between G_1 and G_2.

With these data, we will investigate three research questions:

- How similar are the initial photo groupings that were created at times 1 (same day), 2 (recall day/ week / month later), and 3 (year later)?
- How dissimilar are the two groupings that were created at time 2?
- To what extent is the quantitative degree of similarity reflected in the verbal descriptions that participants gave each group, as judged by human analysts?

We also plan to compare human-generated groupings to the outcome of two cluster analyses generated by the CERTH team, one based on the output of the concept detection system (Deliverable D4.2), and one based on visual information, time, and geolocation,, as described in D4.3.Results of the consistency analysis will be reported in D2.4.

3.2 One year follow-up of 2013 Festival Fringe study

One year after the experienced event and initial interview, a third stage of interview was completed with a subset of participants. This stage will hereafter be referred to as Time 3. All participants who took part in the 2013 August festival study were contacted to see if they would return for a ‘related study’; of those who replied we randomly chose 6 from each of the 4 experimental groups that existed (those who were re-interviewed 1 day after initial interview, 1 week after, 1 month after, or 1 month after initial interview yet having never seen or made decisions about their photographs at initial interview). We thus had 24 participants who agreed to take part; this study was designed to enable further exploration of the way in which memory for an event and perceived context and preservation value changes over time (here, one year). The original grouping construction allows the added consideration of long term (here, one year) influence of time since an event at re-viewing (comparison of Groups 1-3) on memory and photograph judgments for an event, as well as long term influence on these same measures of having viewed photographs immediately after an event or at a delayed time point (comparison of Groups 1 and 4).

3.2.1 Recall data

The recall data generated from the Time 3 interviews have been transcribed, and are currently undergoing the detailed content analysis to assess – via proportion of recall that is deemed episodic in nature - the long-term (one year) consequences of time since an event (day/week/month) on reviewing an event, and of reviewing an event with or without photographs. As before, the Autobiographical Interview (Levine et al., 2002) scoring technique is being used to enable this analysis. Results of this analysis will be reported in D2.4

3.2.2 Photograph data

3.2.2.1 Consistency of groups

Participants were asked, following their recall of events at Time 3, to *organise their photographs of the experience into meaningful groups*. This replicates the section of the August 2013 interviews detailed in D2.1.

3.2.2.2 Keep and delete decsions

The decision processes behind keeping or deleting photographs after a long (one year) period of time were addressed at Time 3 by requiring participants to indicate on a 10 point scale how much they would want to keep or delete a photograph, rather than make absolute judgments about keeping or deleting as was required the previous year. These decisions – which will be compared with Time 1 and Time 2 decisions – were further probed through the use of three additional 10 point scales per photograph, designed around ascertaining *meaningfulness* of the photograph to the individual, and how concerned the individual would be about *findability* and *speed of retrieval* of the item. The answers to these three questions will be analysed in relation to the keep/delete scale answers to determine the degree of overlap of these concepts with the observed *preservation value.* Additionally, participants were asked to provide notes per picture in order to suggest information that they wish they had been able to record with the photograph at Time 1, and information that they would want to record with the picture at Time 3 should they revisit the photograph collection again. The content and volume of these notes will not only be used to infer a measure of preservation value, they will also be used to suggest the types of context that participants wish to use to better document an experience for revisiting. Results will be reported in D2.4.

3.3 August 2014 Festival study

In August of 2014, the UEDIN team conducted a second, independent study during the Edinburgh Festival Fringe. This involved 22 participants who were given a camera phone for 24 hours, and were asked to experience and to take digital photographs to document a day at the Festival. With a constraint to produce between 40 and 80 photographs, this study sought to develop the August 2013 festival study into an even more naturalistic setting. The study was over a longer time per participant, specifically 24 hours rather than an hour as in 2013, and participants *chose* when and how to take pictures to document their experience, rather than as in 2013, being required to take a photograph every three minutes when prompted by a beep from the phone. As in August 2013, the August 2014 study group were interviewed immediately following the 24 hours during which they experienced the festival (referred to as Time 1), and then when they returned a month later, for a second interview (Time 2). The primary research questions were the same as those listed in D2.1, namely: how does episodic recall change over time, how do people categorise the photos they take, how generalisable are categorisations, and do these categories change over time.

As with the previous studies, interviews with participants at Time 1 and Time 2 assessed their biological memory for their (24 hour) experience of the festival. These recall data will provide the basis for investigations of changes in episodic recall over time from the analyses currently under way.

At Time 1 and 2 Participants were again asked about the keep/delete, *meaningfulness, findability* and *speed of retrieval* values they would ascribe to each photograph of a pseudo-randomly selected subset of their photographs. These data will be used to determine overlap or independence of these listed concepts and the effect on these value judgments of time passed (here, one month). Numerous additional data measures have been collected in response to photographs at Time 1 and Time 2; participants were presented with suggested tags for their pictures (for example, as derived through previously used groupings or labels, through image analysis generated by the CERTH team, and through data collected from Twitter – as described in more detail below) as well as presented with the opportunity to provide notes with their pictures. At Time 2 participants utilized the Mobile PIMO application, provided by the DFKI team, to write their own notes and to make preservation value judgments about their photographs. These data will be analysed in the coming months to inform the final summative evaluation and the implementation of memory buoyancy in the PIMO; findings will be reported in D9.4 and D2.4.

3.3.1 Automatic contextualization using Twitter

As we report in Section 4 regarding the Survey study, most people who take digital photos do not add manual context information, and most digital photos are taken on smartphones, which can automatically store rich information about time and geolocation. Because social media updates are often tagged with time and location, we expected that time and place matched Twitter updates might provide a useful source for automatic contextualization.

Similar ideas have already been explored by Shamma, Kennedy, and Churchill (2010) for enriching records of broadcast TV shows, and by McParlane and Jose (2014) for contextualizing photo sets that were taken at large events. There is also a substantial literature on event detection using social media (e.g. Marcus et al, 2011; Chakrabarti & Punera, 2011; De Choudhury, 2012).

For 18 of the 22 participants, we collected all tweets that were posted to Twitter on the same day and from a location that was within a bounding circle calculated from the location metadata associated with the photo collection. When no such metadata were available, we used an area that comprised Central Edinburgh and the Meadows area (a large city centre park) of Edinburgh, where most Festival events take place. Data were gathered using the Twitter search API (https://dev.twitter.com/rest/public/search). Due to technical problems with setting up the Twitter search, we have no data for the first 4 participants.

Table 3.3 shows the frequency of the 20 most frequent hash tags in the complete data set. There were five types, Festival (generic hash tags that relate to the Edinburgh Festival and the Edinburgh Festival Fringe), Location (hash tags mentioning places and countries), Politics (hash tags that refer to the Scottish Independence Referendum Campaign), Promotion (tags that refer to small business promotions or Twitter marketing events), and Event (tags that refer to events taking place during the Edinburgh Festivals).

Table 3.3. Frequency of Hash Tags in our Twitter Sample

Type	Total	Hashtags
Festival	4197	edfringe (2911); unbored (389); edfringe2014 (248); fringe (244); edinburghfringe (142); edinburghfestival (121); fringefest2014 (78); comedy (64)
Location	1063	edinburgh (758); scotland (205) \\
Politics	1013	indyref (389); bbcindyref (232); voteyes (157); patronisingbtlady (156); yes (79) \\
Promotion	279	sbutd (114); bamforxmas (101); udobiz (64) \\
Event	172	dginalba (102); travestiplay (70) \\

In our ongoing data analysis, we are investigating the following research questions:

Q1 Would time and location matched Twitter hash tags be useful for semi-automatic photo annotation?
Q2 Should hash tags be supplemented with a list of named entities (i.e., people, places, shows, events, etc.) that have been frequently mentioned in a day's tweets?

For Q1, we have explicit judgments from human participants, which were gathered as part of the Time 2 interview of the 2014 Festival study group. Q2 will be tested using transcribed recall interviews and group descriptions. If frequent mentions appear in the interview and group data, then they would be good candidates for photo set annotation.

4 Survey

The aim of the ForgetIT survey is to establish an overview of existing personal preservation practices that individuals report, using the example of photo collections. This complements the objective, experimental data, and allows collection of responses from large numbers of participants over the internet. A copy of the survey is available in the supplementary web documentation.

A full cross sectional sample would target many people who would be very unlikely to adopt a ForgetIT system, and for the same reasons would be very difficult to persuade to complete the survey. Since we did not recruit participants through a market research company, and given that the survey was distributed online, most of our respondents will have an interest in the topic of preserving their own personal photo collection. This means that the survey is likely to reflect the views of potential early adopters, i.e., people who are interested in preserving their own photos and who are interested in considering IT solutions. In D2.2, we outlined the development process of the survey. The main part of the survey questions focused on finding out what respondents do (or rather, what they think they do) when taking, managing, and preserving photos. In the following discussion of the survey, we will refer to these three sets of activities as photowork.

We did not aim to categorise respondents into pre-existing types, or measure properties or personality traits. Rather, it was more important to ensure that answer options covered a range of common practices, hardware, and software options around photo work.

Items for the survey were sourced from the wealth of existing literature on personal photo taking and management (see D2.2 for detailed references), the free text options of the pilot survey, and several semi-structured interviews with people interested in the subject.

After reducing the survey in size so that it could be answered within 15 minutes, it was tested face to face with 74 participants in the Festival study described in D2.2 in order to ensure items were easy to understand.

The deployed version represents a minor revision of the face-to-face survey.

In the demographic part of the survey, we use two instruments to assess relevant participant characteristics. The first instrument is an inventory of privacy concerns designed by Buchanan et al. and is well validated (Buchanan et al, 2007). The second instrument is a brief assessment of attitudes to technology which is based on Schulenberg and Melton's Computer Attitudes, Aversion, and Familarity Index (CAAFI, Schulenberg and Melton, 2008). The original 30 questions of the CAAFI were reduced to 8, and these were updated to reflect changes in technology since 2008. Preliminary data on this assessment are presented in Section 4.2.4. Scores on each item range from -3 to 3.

4.1 Planned analyses

Recruitment will finish at the end of March. We have planned three main analyses of the survey data.

a) Preserving Digital Photos – Comparing Two Generations
Based on our sample of younger and older people from Sweden and the UK, we will seek to determine whether there are any generational differences in the way potential early adopters manage and preserve their photos.

b) Preserving Digital Photos Across Cultures
We are actively recruiting students from four different European countries, the UK, Sweden, Italy, and Germany and students from five non-European Countries, China, Japan, Qatar, Egypt, and

Taiwan. Respondents will be matched for age group and education. In this study, we aim to investigate potential cross cultural differences in photo preservation.

c) Personas for Preservation

Based on an analysis of the complete data set, we will define an initial, small set of prototypical users ("personas"[1]) of ForgetIT personal preservation solutions. Personas are a popular tool for designers and developers that make it easy to describe use cases and user stories (Cooper, 2004; Pruitt & Adlin, 2005). They are also a good way of presenting the findings of a large survey to a wider audience.

These personas will be derived from a cluster analysis of responses to survey questions that cover photo management and preservation. Unlike the personas that were developed in D10.1 for the Organisational Preservation use case, these personas will be data driven. For each persona, we will flesh out the initial sketches using semi-structured interviews with 10 respondents whose responses fit the profile and who agreed to be re-contacted for interviews and studies.

4.2 Survey to date

4.2.1 Recruitment

We conducted separate recruitment drives for each of the three sub studies mentioned in the previous section. In order to accommodate the different sources of participants, we created several versions of the main survey that differed in translations available, compensation options, and additional test items.

In order to maximize recruitment across partner sites and collaborators of partners, the survey was translated into 10 languages. We only covered languages that were either represented within the consortium or for which we were able to recruit a translator among partners and colleagues. The number of respondents per language option can be found in Table 4.1.

Table 4.1: Recruitment by survey language (as of January 5th, 2015)

Language	N	Percent
English	619	52.2%
Swedish	229	19.3%
Turkish	74	6.2%
Czech	57	4.8%
German	55	4.6%
Greek	42	3.5%
Italian	42	3.5%
Chinese (Modern Spelling)	39	3.3%
Chinese (Traditional Spelling)	13	1.1%
Japanese	13	1.1%
Arabic	2	0.2%

The survey has been distributed through five channels.

- Snowball sampling: Project partners forward the survey URL to collaborators, friends, and family, post about it on blogs and social media
- Participant panels: The survey URL has been distributed to two participant panels, Botnia Living Lab, Sweden, and the Psychology Participant Panel, University of Edinburgh

[1] In this deliverable, we use the plural "personas" instead of the correct Latin "personae", because this is the term that is commonly used in the HCI literature.

comprising volunteer members of the general public. Edinburgh participants were not compensated; the Botnia Living Lab participants were offered the opportunity to compete for a prize.

- Targeted emails to students: The survey is offered at the University of Edinburgh for participation credit and is also distributed to students in Sweden, Italy, Germany, Japan, China, Taiwan, Egypt, and Qatar. Students in Sweden, Italy, and Germany can choose to participate in a book voucher draw or competition, while students from Japan, China, Taiwan, Egypt, and Qatar can choose to receive a signed and personalized postcard from the investigator team.
- Web sites: The survey can be found on the project web site (https://www.forgetit-project.eu/survey/index.php/759861/lang-en); it has also been linked from a Cloud platform service provided by Turk Telekom. There is no compensation for participation.
- Crowdsourcing: Participants were recruited through Amazon Mechanical Turk (US) and Crowdflower (US, Canada, Australia, New Zealand, UK, Ireland). Compensation was set at $1 per completed survey.
-

As of January 5, the total number of complete responses was 1185. A response was judged to be complete if the participant had completed all survey items that related to photowork.

Table 4.2 lists the number of responses by compensation option.

Table 4.2: Recruitment by compensation options

None	512	43.2%
Book Voucher Draw	158	13.3%
Post card offer	71	6.0%
Crowdsourcing	444	37.5%

4.2.2 Participants

Due to our sampling strategy and the strong participant panels in Sweden, most of our participants chose the English or the Swedish version of the survey (c.f. Table 4.1), with the other languages trailing behind. In our final push, we hope to further increase the numbers of German and Arabic participants.

The demographics of the full sample are given in Table 4.3. Gender distribution is balanced. 52.7% are employed or self-employed, while 31.9% are students. Overall, respondents tend to be aged between 18 and 34, which is partly due to the large number of students who were recruited. In order to estimate education levels, we asked respondents to estimate the number of years they spent in full-time education, a standard practice in psychology. The response categories offered corresponded to primary education only / short secondary education (up to 8 years), secondary education including apprenticeship (9-13 years), university and further education (13-18 years), and postgraduate degrees (19 years and more).[2] Our sample is disproportionately highly educated, with 23.9% stating that they spent 19+ years in full time education.

Almost all respondents (1136, 95.9%) reported taking digital photos themselves. In addition, 588 (49.6%) get digital photos by email, 428 (36.1%) download them from the web, and 304 (25.7%) receive collections of digital photos from others.

[2] The overlapping categories 9-13 and 13-18 reflect a compromise between different education systems.

Table 4.3: Demographics of the full sample

Age	18-24	344	29.0%
	25-34	359	30.3%
	35-44	236	19.9%
	45-54	138	11.6%
	55-64	67	5.7%
	65+	25	2.1%
	Not stated	8	6.8%
Gender	Female	591	49.9%
	Male	578	48.8%
	Not stated	7	0.6%
Occupation	Employed F/T	487	41.1%
	Employed P/T	63	5.3%
	Self-employed	74	6.2%
	Student	377	31.9%
	Homemaker	69	5.8%
	Retired	34	2.9%
	Unemployed	40	3.4%
	Not stated	41	3.5%
Education	Up to 8 years	23	1.9%
	9-13 years	157	13.2%
	13-18 years	697	58.8%
	19+ years	283	23.9%
	Not stated	25	2.1%

Table 4.4 shows how often respondents take digital photos with each of four main devices, smartphones, non-smart camera phones, tablets, and digital cameras. While 90.3% take pictures on their smartphone, and 88.6% on their digital cameras, the cameras tend to be reserved for special occasions, while 62.2% take smartphone photos at least every week.

Table 4.4: Usage patterns for digital photography devices

	Daily	Weekly	Monthly	Rarely	Never / Don't own
Smartphone	323 (27.3%)	413 (34.9%)	180 (15.2%)	105 (8.9%)	115 (9.7%)
Cameraphone	15 (1.3%)	35 (3.0%)	27 (2.3%)	159 (13.4%)	900 (75.9%)
Digital Camera	62 (5.2%)	197 (16.6%)	380 (32.1%)	362 (30.5%)	135 (11.4%)
Tablet	26 (2.2%)	64 (5.4%)	131 (11.1%)	324 (27.3%)	591 (49.9%)

4.2.3 Quality control

When administering surveys, it is important to ensure that participants pay attention to the wording of questions and answer diligently.

One method involves Instruction Manipulation Checks (Oppenheimer, Meyvis, & Davidenko, 2009), where participants are given instructions that are counterintuitive or where giving the correct answer requires a careful reading of the original question. We used a variant of Attention Checking Questions in the crowdsourcing variant of the survey. At the end, participants were presented with a 4-6 digit number that had been generated randomly. They had to note down this number and enter it into a form on the survey administration web site (Amazon Mechanical Turk or Crowdflower). They were only paid if the number they had entered matched the number stored in the survey logs.

While Amazon Mechanical Turk has its own participant panel, Crowdflower distributes surveys to crowdsourcing websites around the globe. For each of these websites, Crowdflower displays the percentage of trusted judgements obtained. The ForgetIT survey was only distributed to web sites with rates of 80% or better, which is likely to lead to a higher response quality (Peer, Vosgerau, & Acquisti, 2014).

For the full sample, we use post-hoc checks of answer distributions to identify participants who may not have answered the survey diligently. This approach is similar to the data quality checks that we have used in earlier work on speech intelligibility data collected on Amazon Mechanical Turk (Wolters, Isaac, & Renals, 2010).

The main survey contained three mandatory questions that consisted of several statements – a question about frustrations when managing digital photos, a question about digital photo management practices, and a question about digital photo preservation practices. For each statement, the user had to answer on a four-item scale that deliberately excluded a central neutral option. In each case, the statements were presented in an array format, which makes it tempting for respondents to tick the same column for each statement and move on.

For each of the three indicator questions, we identify respondents who ticked the same option every time (one-column answers). If respondents answer truthfully and thoughtfully, one-column answers should be rare, and several one-column answers in one survey return would be cause for concern.

This is borne out by the data. In the survey data collected so far, 135 people (11.4%) provided one-column answers to at least one question, but only 25 people (2.1%) gave one-column answers to two or more. Of these 25, 13 come from the Crowdflower data set (3% of the total number of responses collected through that source), 2 from the Chinese student data (5%), 3 from the British student data (2%), and 7 from the web data (2%).

4.2.4 Preliminary analyses

In order to assess the Attitude to Technology items, which were not mandatory and are listed in Table 4.5, we used data from 1177 respondents who had answered all eight items. The automated item selection procedure for Mokken Scale Analysis (van der Ark, 2007) suggests two underlying constructs, which we will call technology aversion (Items AVOID and LEARN), and technology literacy (Items ENJOY, IMPT, SITES, GADGET, and READ). Item EMAIL appears to reflect a third construct.

Table 4.5: Attitude to technology items

Item name	Item text	Scale
ENJOY:	I enjoy using technology	Technology Literacy
IMPT:	Being able to use technology is important to me	Technology Literacy
GADGET:	I keep up with the latest technology	Technology Literacy
READ:	I enjoy reading about technology.	Technology Literacy
AVOID:	I avoid using technology whenever possible	Technology Aversion
EMAIL:	Email is an easy way to communicate.	Not Assigned to Scale
SITES:	I often visit web sites.	Technology Literacy
LEARN:	I hate learning to use new software and	Technology Aversion

	applications.	

The technology literacy sub-scale has good internal consistency (Cronbach's alpha 0.78), but the consistency of the technology aversion sub scale is moderate to poor (Cronbach's alpha 0.59). The mean value of the technology literacy scale for our sample is 2 (SD: 1.0) with a minimum of -3 and a maximum of 3. Therefore, most of our respondents are highly technology literate and represent a sample of potential early adopters.

We piloted the methodology to be used for determining personas on a subsample of 236 students that had been recruited from European Universities and completed the survey in English, Swedish, Italian, or German. Here, we briefly describe our method and the main findings; results on the complete survey will be presented in detail in Deliverable D2.4.[3]

In order to find clusters of related response patterns, using the k-means algorithm as implemented in the R package cluster. Respondents were clustered according to their responses to two sets of items, one covering photo management practices (Table 4.6) and the other covering photo preservation practices (Table 4.7). Each practice was described by a brief statement, and for each statement, respondents were asked to say whether it described their practices very accurately, somewhat accurately, somewhat inaccurately, or very inaccurately.

We explored cluster solutions ranging from 2 to 10 clusters and chose the solution that proved most stable, judging by the four stability criteria provided by the R clValid package (Brock, Pihur, Datta, and Datta, 2008). The stability measures examine the extent to which the result of a cluster analysis changes when one variable at a time is removed. Those stability criteria were average proportion of non-overlap, average distance between means, average distance, and figure of merit.

The coherence of each of the final clusters was then assessed using average silhouette width (Rousseeuw, 1987), where 1 indicates high coherence and -1 high dissociation.

Table 4.6: Photo management items

[Auto Labels] I add information automatically.
[Deletion] I delete most digital photos.
[Filenames] I use file and folder names.
[Keywords] I add keywords and titles.
[Manual Labels] I label people, places, and objects.
[Stay Organised] I keep photos organised.

For both photo management and photo preservation, the stable solutions consisted of two clusters, one consisting of respondents for whom the statements tended to be accurate, and one with respondents for whom the statements tended to be inaccurate. Although the silhouette width for each cluster is relatively low (ranging from 0.17 to 0.33), the clusters are interpretable and provide a meaningful partition of the full data set.

[3] At the time of writing, the results of this pilot analysis had been submitted to a conference for publication in the works in progress section.

Table 4.7: Photo preservation items

[Autom. Backup] I use automated backups.
[Careful Filing] I file photos carefully.
[Copy New Media] I move photos to new storage media.
[Man. Backup] I make manual backups.
[Mult. Copies] I keep multiple copies.
[Mult. Locations] I keep copies in multiple places.
[Print Digital] I print photos on high-quality paper.
[Read Checks] I check whether photos are still readable.
[Safekeeping] I give copies to others for safekeeping.
[Scan Photos] I archive printed photos by scanning.
[Secure Store] I store photos privately and securely.

Figure 4.1 shows the results for photo preservation. The middle line gives statistics for the complete data set, the upper and lower line represent the clusters.

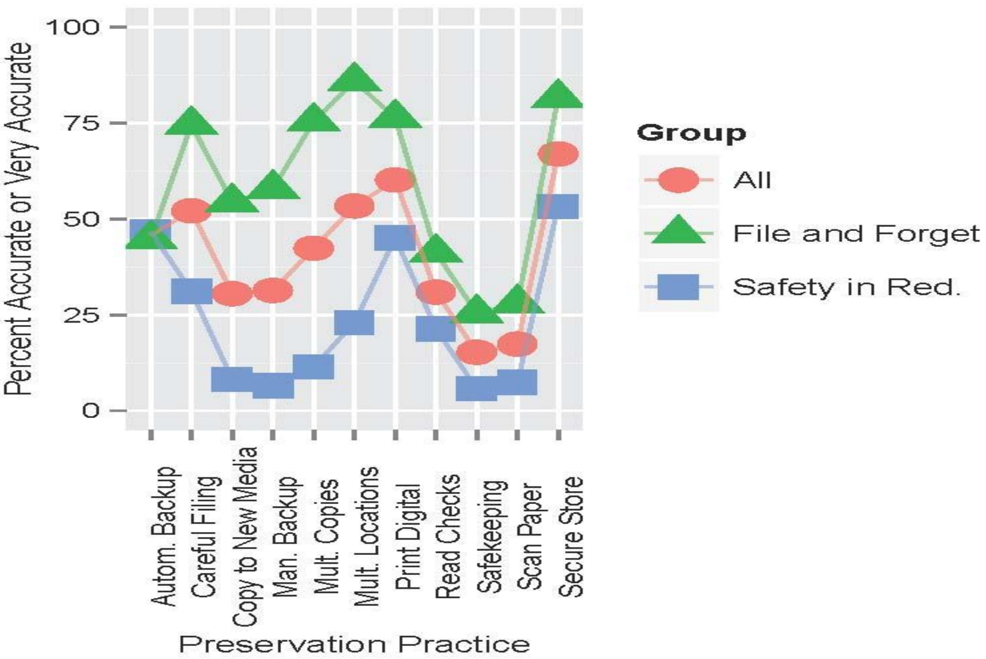

Figure 4.1: Frequency of Practices Related to Photo Preservation

While Group A (48%, n=114, avg. silhouette 0.17) is very diligent in transferring photos to new media, printing them off, and storing them safely, Group B (52%, n=122, avg. silhouette 0.21) is far more sanguine. The core strategies of Group A (Safety in Redundancy) all build on redundancy, guarding against the failure of a storage unit. They were far less aware of the need to regularly check whether photos are still readable. The main preservation strategies for Group B (File and Forget) are secure storage, automated backups, and printing on paper.

The results for photo management can be found in Figure 4.2

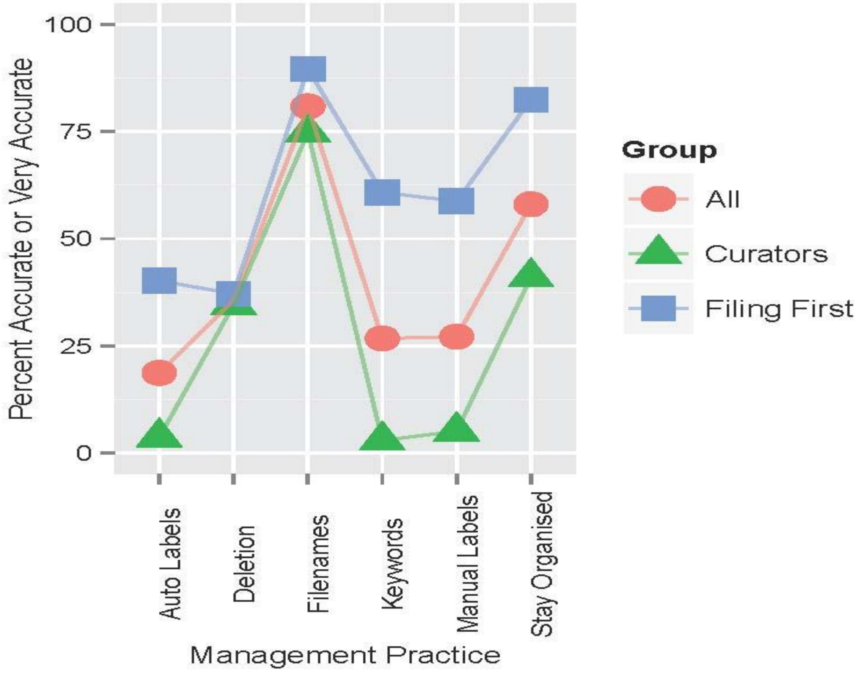

Figure 4.2: Frequency of Practices Related to Photo Management

Group A (Filing First, n=139, 59%, avg. silhouette 0.3) relies almost exclusively on files and folders, and Group B (Curators, 41%, avg. silhouette 0.22) is more likely to curate their photos using both manual (keywords) and automatic (face recognition) options. In both groups, around a third manage the overload of digital photos through regular deletion of most of the photos they take.

We conclude that cluster analysis is a viable means for inferring potentially interesting user groups from the data, but it needs to be supplemented with additional analytical tools that we will investigate in Year 3 of the project. Also, even though k-means has numerous advantages (easy to interpret, easy to find prototypical users for each cluster in the data), it might not be the most appropriate solution for uncovering the internal structure in the data.

5 Towards human and machine information preservation: Archival and retrieval

In this section, we report preliminary work on a collaboration between UEDIN and EURIX (Allasia) to develop a formal model that is inspired by the understanding of human memory and could be implemented in a ForgetIT system. This work investigates a mathematical conceptual model for representing digital objects according to our human memory function. Intentional retrieval of human memories is a process that obeys a set of rules investigated in cognitive sciences, and described in D2.2 and in section 1 of the current deliverable. At the same time digital information retrieval techniques are enabling complex searches in the “digital world”, including natural language processing, concept detection from images and query by content. The current work aims to propose a candidate link connecting the models conceptualized in cognitive sciences to the information technology methodologies adopted for archiving and retrieving digital information.

5.1 Background

The study of human memory has a very long history, with scientific approaches starting in the late 19th century, and there is substantial conceptual understanding of many of the basic functional principles both at a conceptual level, and over the last few decades, at the level of the neurobiology of the brain. Encoding and storage at the conceptual level involves extracting key features of personal experiences throughout the lifetime of the individual to develop a database of knowledge about the world. This lifetime database is used to set a context for encoding and retrieving details of individual items or events. However, although there is agreement among researchers regarding the broad principles of human memory function, there are many different conceptual models, some more formal than others. The less formal and less clearly specified conceptual models tend to be the least contentious and Figure 5.1 illustrates the main features of such a conceptual model which serves as the basis for ongoing research by the UEDIN team. The rationale and evidence for this model were reviewed in D2.2. The concepts of Episodic and Semantic Memory are described in Section 1 of the current deliverable. In sum, the information from auditory, visual and other forms of human perception (e.g. tactile) activates stored information (Knowledge Base) accumulated over a lifetime regarding knowledge about the world and about the self (Semantic Memory), and preserved information about individual events (Episodic Memory) related to the perceived stimuli. Episodic Memory is subject to rapid loss of detail immediately following an event. Semantic Memory is much less affected by forgetting. Working memory refers to temporary buffer storage (verbal/acoustic and visual/actions) and on-line processing with second to second updating when interacting with the current environment. The “Executive Functions” perform content processing, as well as providing access to currently activated knowledge from the Knowledge Base. Outputs from the executive processing can update the contents of Semantic and Episodic memory. All the associations linking concepts are stored in Semantic Memory.

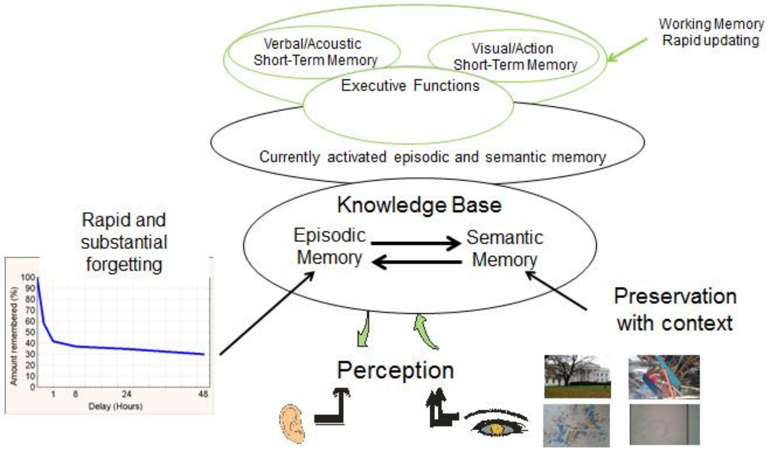

Figure 5.1. A conceptual model of human memory, adapted from Logie (2003; 2011)

Information Retrieval (IR) is a well-established branch of research on digital information technologies, primarily focused on low level information (features) that can be extracted from digital contents. Unlike the human case, there are common approaches to IR with retrieval based on the indexing of the low level content features in an appropriate digital structure and recall based on similarity, i.e. distance functions applied to multidimensional spaces. Figure 5.2 illustrates how an image is represented in IR.

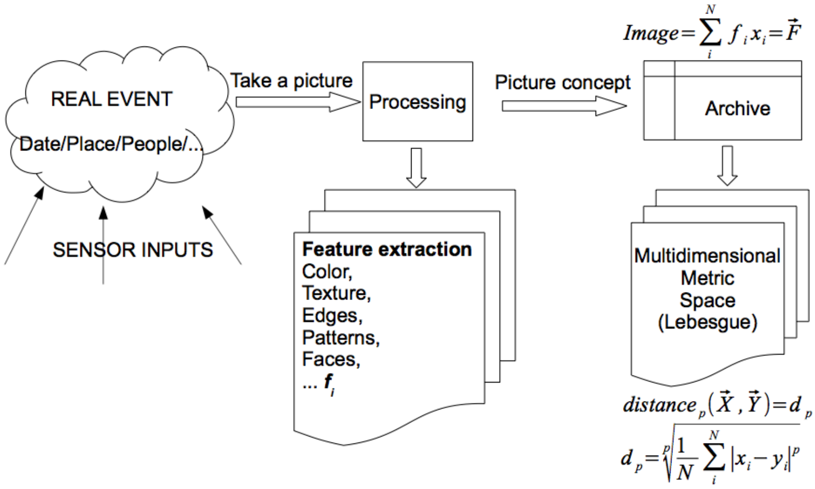

Figure 5.2 The classic approach of information processing, storage and retrieval for Digital Object Features

As described by Conway & Loveday (2010), we can summarize a general event, an episode, as made up of several Episodic Elements (e.g. what-where-when) bound together with the respective Frame, which is relevant information extracted from the repository of the related contextual knowledge about the world and about the self.

5.2 The processing flow

Figure 5.3 illustrates the information flow. The human processes are represented on the left of the Figure 5.3, starting from the Sensory input and Perception. Initial data are captured as sensory inputs (e.g. our eyes and ears). The functions of perception identify basic features of the event (what), followed by activation of stored relevant previous experiences and knowledge. This activated knowledge sets a context for the current sensory input and generates cues that trigger retrieval of further relevant information. Information relevant to the current task (what-when-where) is held in the temporary memory buffers, while executive processes undertake reasoning, problem solving, language comprehension, planning and other processes relevant for observing, interacting with the current environment. The result of the processing may be encoded in, and added to Episodic and Semantic memory.

The remainder of Figure 5.3 illustrates the digital counterpart involving a digital device (such as a camera) taking a picture. Here, the processing flow is focused on digital images, but this could apply equally to other digital contents such as text or numerical data. A digital camera can provide information about current date (when), geographical position from GPS (where). Some information about ‘what’ can be gleaned from technical details such as the field of view (FoV) and focal length that provide information about which aspects of the scene, the picture was intended to capture (something close to the user or to the infinity). Further indications of ‘what’ that can be extracted automatically from the camera data could be ‘New Year’s Day’ from the date, or “Edinburgh Castle” from GPS. Other data can be recorded such as audio at the time, plus human entries of annotations and tags or web posts. Together, the automatic extraction, the human entries, and the additional automatic processing form an episodic record or cue for preservation, referred to here as a hook.

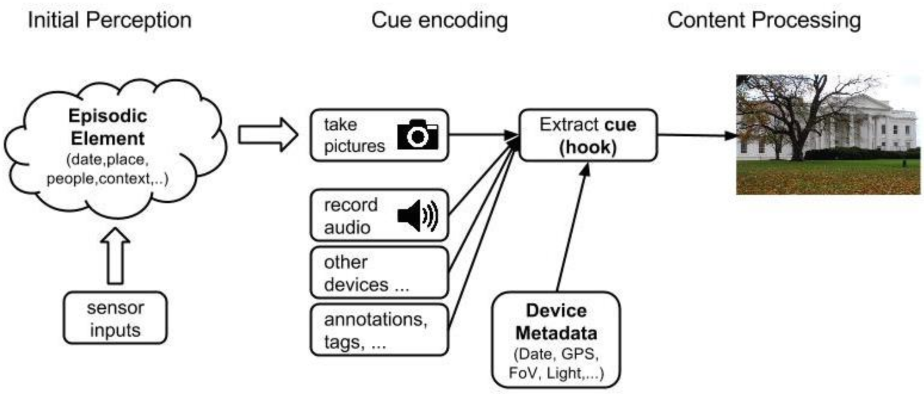

Figure 5.3. Information Extraction for digital images. Automatic extraction of date, field of view, focal length, and GPS location can extract ‘when-where-what’ for preservation along with the image, complementing human processing and memory for scenes

The extracted cue represents the hook that Greenwood et al. (2014, D6.2) describe as “hooks” processed in the “contextualization method”, while the initial data provided by the digital devices can be considered the conceptual units of “contextualization sources” (Figure 1 of D6.2).

The next step of processing flow is the content processing. In our example the content is a simple image. Figure 5.4 shows a classic IR image processing flow, based on extracting low level features and apply some reasoning on them in order to extract some latent semantics such as concept detection or quality assessment. Figure 5.5 shows how detection of individual concepts can lead to a higher level concept. Preliminary results from an example processed by partners CERTH (Markatopoulou et al., 2014; Mezaris et al., 2014, D4.2) are depicted.

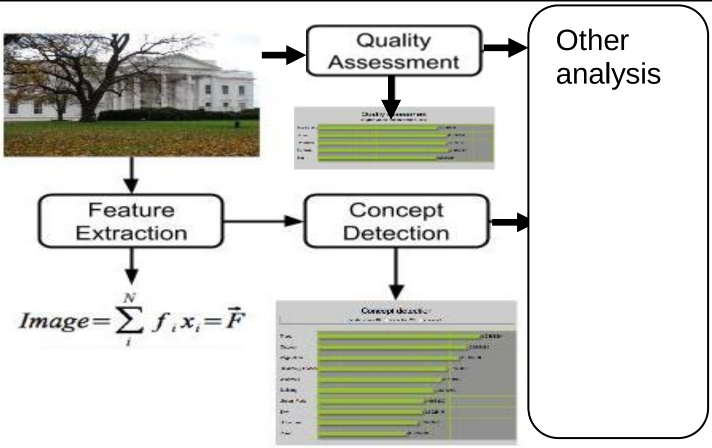

Figure 5.4. The information extraction (IE) from the image involves several processes such as low level feature extraction, concept detection, quality assessment and other analysis such as automatic contextualization using external databases (Markatopoulou et al., 2014; Mezaris et al., 2014, D4.2).

Other analysis in Figure 5.4 is a placeholder for further reasoning and elaborations made on summary representation of the image as well as information collected from referenced sources. The contextualization process (Greenwood et al. (2014, D6.2) is a specific analysis performed on images that make use of a model vector representation, and attempt to guess a candidate context for the processed image according to similarity measures evaluated in pre-archived and contextualized items (see section 4 of D6.2). For example additional automatic processing of textual annotations can allow identification of key words and phrases that may be linked with entries in external databases (partner Sheffield) such as Wikipedia. Automatic processing of audio recording might allow further identification of cues available at the time the picture was taken, such as 'dog barking', 'child crying', 'female singing' etc. Any and all cues encoded with the image can later be used individually or in different combinations to retrieve the entire episodic record. This is analogous to the principle of encoding specificity (Tulving, 1972) that cues present at the time of experiencing an event may be encoded with other details of the event, and these cues can later be used to help retrieve details of the original event. When a cue is presented subsequently it activates relevant knowledge, and that activation spreads to related knowledge, allowing a broad spectrum of associations to provide a context for retrieval, and to allow reconstruction of elements of the episode that had been forgotten. For example, the word 'picnic' would act as a cue for episodes of specific picnics, but the activation would spread to general knowledge about picnics. Therefore, if many details of a specific picnic have been forgotten, the general knowledge allows plausible guesses that the event took place outside, likely in a field or park, involved food that can be eaten without cutlery, and involved friends or family.

Once we have all the extracted and processed information, Figure 5.5 describes the archival phase where the digital content enters the digital repository

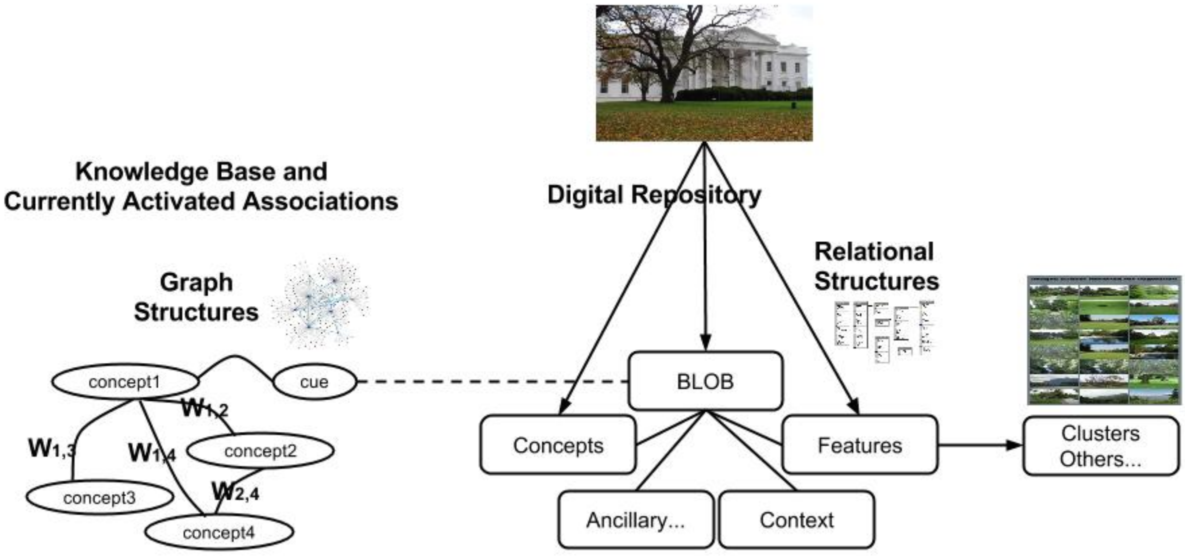

Figure 5.5. The processed image with extracted features and added information is archived in a digital repository where more elaboration can be applied in order to evaluate similar contents based on Lebesgue spaces (distance at the top right). At the same time, according to the encoded cue (hook) the content (described as Binary Large OBject) can be associated to its spread of activated associations, based on the knowledge base of the user. W*ij* are the weights for each association connecting the concept *i* with concept *j*

In the classic approach the digital repository can be considered as a content management system with a relational structure. The digital image of our example can be represented as a Binary Large OBject (BLOB in Figure 5.5) associated with extracted features, concepts and also the context, plus other potential relational data (ancillary placeholder in the picture). The features can be used again in a model vector representation in order to perform a refinement of clustering or extract other features based on similarities making use of a distance function in a Lebesgue space *Lp* such as:

$$d_p = \sqrt[p]{\frac{\sum_{i}^{N} \left| x_i - y_i \right|^p}{N}} \quad (1)$$

where x and y are the contents between which we wish to evaluate the distance, having features x_i and y_i respectively with dimension *N*. When *p* = 2 equation (1) gives the euclidean distance, widely adopted in IR.

Beside the digital repository, Figure 5.5 shows the Knowledge Base that is more suitably represented in a graph structure. More precisely Figure 5.5 shows an example of its Currently Activated Associations, where the concepts related to the digital image are stored according to their links and encoded with the initial "cue" (hook), representing the connection between the BLOB and the Knowledge Base.

5.3 The model

What is currently missing is the connection between Working Memory behaviours and IR in order to be able to represent the episodic elements and associated knowledge base for current and ongoing tasks with suitable mathematics in a unified representation of digital objects in memory.

5.3.1 First attempt

There are multiple formal theoretical approaches available for describing processing frameworks in cognitive science, perhaps the best known of which being Adaptive Control of Thought (ACT-R; ACTR (e.g. Taatgen & Anderson, 2008), but these functions at an higher level than required for our current purposes, because we are looking for a representation allowing the feature descriptions of digital objects. Other approaches (e.g. Du & Swamy, 2014) describe the Associative Memory Networks by mean of Neural Networks, more precisely with a multi layer perceptron approach, making use of the Hopfield network based on the classic Hebbian rule. Even if close to the cellular behaviour, the neuronal cells with axons connected to dendrites, the mathematical approach resides in the perceptron or artificial neural network, i.e. matrices and vectors.

Starting from the low level features that we can compute from digital objects and the results provided by information extraction techniques, we can create multidimensional spaces where we can perform similarity searches, clustering and other latent semantic analysis. As described above, from the human perspective, images can be considered as part of an episodic element and must be archived together with their frame or context: we need to represent the digital image together with the associated knowledge base, the related function for the spread of activation through the network.

Hence we are looking for a formalism able to:

1.represent the image with information available from computer vision techniques,
2.represent the knowledge base,
3.represent the currently activated knowledge base spreading (according to the spreading activation model)

Candidates for representing these features and related analysis results are “tensors” (http://en.wikipedia.org/wiki/Tensors; Weinberg, 1998), allowing a more flexible mathematical representation compared to simple vector spaces. We can move from the image representation in classic IR:

$$Image\ X = \sum_{i}^{N} f_i x_i \qquad (2)$$

to a tensor representation such as:

$$Image\ x^{a} = F_{b}^{a} x^{b} \qquad (3)$$

where *F* is the tensor representing the features that enables the computation of the associated knowledge base, regardless of dimensions and lengths, following the Einstein summation convention:

$$Image\ x^{a} = K_{c}^{a} F_{b}^{c} x^{b} \qquad (4)$$

where *K* is the knowledge base tensor.

Even if (4) provides a quite simple and efficient representation, tensors for the knowledge base do not provide the best mathematical solution: the knowledge base describes the association between concepts in our semantic memory and graphs with triple notation are more suitable (e.g. Berners-Lee & Connolly, 2011; Owl Working Group, 2012; see also W3C standards)

As shown in Figure 5.5, the knowledge base is better represented in a graph structure rather than in a relational structure.

5.3.2 The selected approach

In order to be able to best represent the knowledge base and related computations, we introduce the concept of graph, more precisely a direct graph, because the connections in our spreading activation path cannot (in principle) be considered bidirectional: concept A connects concept B and not necessarily vice versa.

Concepts (objects, words and so on) in our knowledge base have all of their features archived in our semantic memory, representing the “semantic nodes”. Semantic nodes are connected by associations that have a specific direction. Moreover associations can be reinforced (e.g. Bugelski, 1956; Sutton & Barto, 1998). Hence we need to specify different weights on associations. More generally, the semantic memory can be represented as a weighted network of directed graphs. A graph in our context can be defined as any set (C,A)

$$\text{semantic graph} \equiv (C,A) \quad (5)$$

where C is the set of Concepts, our Semantic Nodes, and A is the set of Associations between concepts. They are directed and have weights. Moreover, they are time dependent because according to the spreading activation model, an association can be reinforced at any time as well as overwritten, and subject to decay over time.

The association

$$a_{i,j}(t) \in A \qquad (6)$$

represents the link between concept “i” c_i and concept “j” c_j where c_i , $c_j \in C$, i.e. every concept is taken from the set of concepts C. In (6) we have pointed out the time dependency of association.

The total number of concepts we have in the semantic memory can be considered N at a specific time and is usually referred as the size of the network, i.e. the size of our semantic memory. Making use of combinatorics algebra, we can calculate the maximum possible number of associations S we can have in a semantic memory of size N. The associations are connecting two concepts, hence the number of two-concepts element subsets taken within a concept is the binomial N coefficient

$$S = 2\binom{N}{2} \qquad (7)$$

where we have multiplied by 2 because every pair of two concepts can be navigated in two different ways since our associations are directed. The binomial coefficient can be computed as:

$$S = 2\binom{N}{2} = 2\frac{N!}{2!\,(N-2)!} = N(N-1) \qquad (8)$$

giving us the important result that if a semantic memory has size N (i.e. made up of N concepts), the maximum number of ordered, directed and weighted associations linking them is $S = N(N-1)$. An example of the network graph of concepts that can either be extracted automatically or part of the stored knowledge base is shown in Figure 5.6.

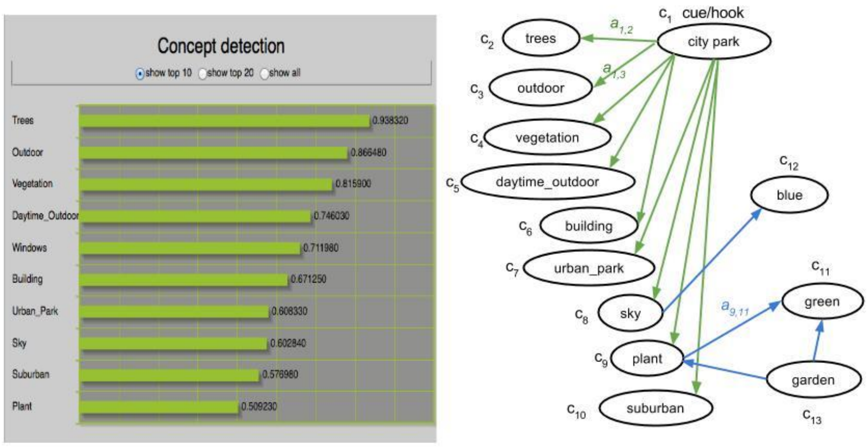

Figure 5.6. Concepts automatically extracted from the image analysis are illustrated on the left for the picture depicted in Figures 3 and 4 together with strength values. On the right is shown the activation graph connecting the cue (hook), for example 'citypark' to the concepts c_j. Examples of some associations, the $a_{i,j}$ arches are reported. In green are the associations extracted automatically from the picture, and in blue are the associations already connecting the concepts stored in the knowledge base.

The image analysis can provide the concepts available in the image with a specific strength, the probability of having the specific concept. We can consider the strength as the weight of association. As example the "city park" concept ($c1$) is connected to the "trees" concept ($c2$) with the association at the time we have $a_{1,2}$ = 0.938320 from analysis of the image. If we had already the association of these two concepts, the analysis result has to reinforce $a_{1,2}$.

More generally and according to the Ebbinghaus (1913) curve illustrated for Episodic Memory of the lower left of Figure 5.1, we can express the association as:

$$a_{i,j}(t) = w_{i,j}\, x_{i,j}\, e^{-\frac{t}{\beta}} \qquad (9)$$

where $w_{i,j}$ is the association weight, the current amount, decreasing with a retention factor β over time and $x_{i,j}$ is the adjacency matrix X (Barrat et al., 2013).

$$X = \{x_{i,j}\} \quad \text{where} \quad x_{i,j} = 1 \; if \; a_{i,j} \in A \; and \; x_{i,j} = 0 \; if \; a_{i,j} \notin A \qquad (10)$$

The adjacency matrix for Figure 5.6 is provided in the following Table 5.1 where the semantic nodes, the concepts, are placed in rows and columns and for every directed association connecting concept "i" to concept "j" there will be a "1" in the corresponding row "i" and column "j". The matrix has *NxN* elements where *N* is the size of our concept network.

It is worth noting that directed associations require an asymmetric adjacency matrix, i.e.

$$x_{i,j} \neq x_{j,i} \qquad (11)$$

Table 5.1: Adjacency matrix referred to the example reported in Figure 5.6.

C	1	2	3	4	5	6	7	8	9	10	11	12	13
1	0	**1**	**1**	**1**	**1**	**1**	**1**	**1**	**1**	**1**	0	0	0
2	0	0	0	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0	**1**	0
9	0	0	0	0	0	0	0	0	0	0	**1**	0	0
10	0	0	0	0	0	0	0	0	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0	0	0	0	0
12	0	0	0	0	0	0	0	0	0	0	0	0	0
13	0	0	0	0	0	0	0	0	**1**	0	**1**	0	0

In complex networks the degree of a concept c_i is defined as the number k_i of associations incident on the concept “i”, either arriving or leaving the semantic node.

(Barrat et al., 2013):

$$k_i = \sum_{j}^{N} (x_{j,i} + x_{i,j}) \qquad (12)$$

Considering Figure 5.6 as an example, concept *c*11 has degree 2 (2 arriving associations) as well as *c*13 (2 leaving associations), while concepts 2, 3, 4, 5, 6, 7, 10 have degree 1.

The first concept, the cue, has degree 9 (9 leaving associations) and the second highest is *c*9 with degree 3 (2 arriving and 1 leaving associations).

Table 6.2 reports the degree of each concept (semantic node) shown in Figure 6.6, evaluated by (12).

Table 5.2: The degree k of every concept c_i shown in Figure 6.6.

C_i	1	2	3	4	5	6	7	8	9	10	11	12	13
k	9	1	1	1	1	1	1	2	3	1	2	1	2

Having introduced the degree (12) of our semantic network, we can define the associated ***degree distribution as** $P(k)$ **, the probability that a randomly chosen concept has degree k*** (Barrat et al., 2013). The degree distribution enables statistical evaluation of the network behaviour.

We can compute the *n*-th moment of the degree distribution such as (Barrat et al., 2013):

$$< k^n > \;=\; \sum_k k^n P(k) \;\;\Rightarrow\;\; \int_d^{\infty} k\, P(k)\, dk \qquad (13)$$

Where we have moved the degree from discrete to continuum and $d \geq 1$ is the lowest degree in the network. With n = 1 we have the average value of k, which is also

$$< k > \;=\; \frac{1}{N} \sum_i k_i \;\leq\; \frac{S}{N} \qquad (14)$$

From (14) and (8) we can say that the average degree cannot be bigger than (N-1). Hence for semantic networks with $< k > \ll N$ we have “sparse graphs” (Barrat et al., 2013) while for $< k > \approx N$ we have very “dense graphs”.

Applying (14), the example reported in Figure 5.6 and Tables 5.1, 5.2 provides $< k > = 2$ with $N = 13$ and $S = 156$.

Since $< k > \ll N$, i.e. $2 \ll 13$, the example of Figure 5.6 can be considered a sparse graph.

With (13) we can now start to investigate the behaviour of sparse and dense graphs in a semantic network and evaluate what can happen statistically when we remove some associations randomly from a sparse graph and from a dense graph. This work, which is also sumarised in D8.2, is planned for year 3 of the project as a means to implement some of the known characteristics of human memory into the development of the ForgetIT preservation and managed forgetting system.

6 Organisational Memory and Forgetting

This report covers work undertaken by the University of Oxford together with other ForgetIT partners to draft a conceptual framework for organisational remembering and forgetting, and take steps to have this conceptual framework operationalized and implemented into digital tools developed as part of the ForgetIT project. It will first frame the conceptual challenge in developing such a framework, offer a recap of extant literature, put forward the core elements of the conceptual framework, and how this framework is being operationalized; it will conclude with a short outlook on planned future work.

> Organisational memory is sometimes considered as analogous to human memory in that there is accumulated knowledge about the organisation within the individual staff. However, this analogy offers only partial insight, given that organisational documentation, policies and practice can be used to compensate for corporate forgetting, such as when an employee leaves, or there is updating and restructuring. Additional crucial aspects of individual and organisational managed preservation and forgetting of information are the processes of interaction between human users and digital information storage and management systems.

6.1 Framing the Research Challenge

While human remembering and forgetting is a cognitive process inside an individual's brain and thus not directly visible and recordable, indirect, objective experimental techiques and modeling have identified the basic principles of human memory and forgetting. Over more than thirteen decades, researchers have added much detail to the understanding of human remembering function, so that today for instance we know that forgetting is steepest at the beginning, repeated rehearsal and recall can improve memory retention, but that the recall process is reconstructive. A review of the relevant literature is given in D2.2. and section 1 of the current deliverable.

Overall, individual remembering and forgetting follows a human remembering function, a complex mechanism that relies on numerous inputs to make decisions about what to retain and what to get rid of. This function is the "baseline", the standard of how individuals remember and forget. And thus, if one wanted to build a digital tool that incorporates the same mechanism of human remembering and forgetting, one "only" has to rebuild in software this human remembering function. The result will be a system that remembers and forgets like an average human being.

While certainly not perfect, this human mechanism of remembering and forgetting is the result of tens of thousands of years of a Darwinian selection process entailing improvements and refinements. Thus, incorporating core elements of this human memory function in the ForgetIT system would not only enable ForgetIT to mimic aspects of human remembering and forgetting, it would also infuse into ForgetIT a mechanism of remembering and forgetting that incorporates tens of thousands of years of Darwinian "learning". In this important sense, therefore, ForgetIT with elements of the human memory mechanism built in will not just be anthropomorphic; it will arguably also be effective and efficient. Moreover, a ForgetIT system may be designed to support the weaknesses of human memory that evolution did not provide for functioning in a complex digital age.

The situation, however, is very different when it comes to organisational remembering and forgetting. Here, it is not sufficient to understand how organisations remember and forget (much like it may be sufficient to understand how individual human beings remember and forget). Current organisations are not the product of tens of thousands of years of a natural selection process that led to frequent improvements of the processes of organisational remembering and forgetting, or at least to a convergence of a dominant remembering and forgetting mechanism. On the contrary, currently there is no dominant mechanism of organisational remembering and forgetting, but rather a tremendous heterogeneity among organisations' remembering and forgetting.

Therefore, the first research challenge is to overcome the lack of a credible and widely accepted "baseline" of organisational memory. The easiest way to do this, of course, would be to identify, much like with individual remembering and forgetting, a single function. So far, however, a quest to find such a universal organisational memory function has been futile, largely because the universe of human organisations covers so many different needs and preferences of organisational memory that one standard function does not fit all circumstances, and perhaps because organisations have not had enough time to converge on a standard function. This does not suggest that a highly adaptive standardized organisational memory function could not be created, but rather only that to date we have been unsuccessful in doing so. The fact that some organisations survive whereas others do not suggest that there might be factors in organisations that contribute to their natural selection, although the factors related to organisation memory might have little to do with the overall success or failure of an organisation even if they may affect efficiency, scope for change and innovation, and responses to challenges that are internal or external to the organisation.

6.2 Existing research in organisational memory

One of the first major strands of research into organisational memory took as its departure point the aircraft industry after World War II. As demand for military aircraft plummeted, large aircraft producers such as Boeing drastically reduced their workforce, including design engineers and experienced workers on the assembly line. The result was a steep decline both in innovation in aircraft design as well as in efficient aircraft production, as relevant expertise within the organisations had been lost. (Benkard, 2000)

The same phenomenon was later found at NASA after the end of the Apollo program and the ramp down of the Space Shuttle design. Thus when the shuttle accidents happened, learning from them was stunted. And later, when the new Orion crew system and the new heavy launch vehicle were commissioned, NASA had few design engineers still within the organisation that would have the expertise from the Apollo days to know how to build such a system, leading to a number of false starts, the need to call back engineers from retirement, and decisions to replicate trusted and tried (and well documented!) main Apollo design features rather than to innovate. (Mahler 2009)

The role of organisational memory (or lack thereof) for aircraft and spacecraft designs are just one of many areas that have since been studied. More recent studies include, for instance, the impact of employment churn in large IT companies on their ability to innovate.

In this strand of research, organisational memory is regularly (but not exclusively) seen contained in the individuals that are working within the organisation. The obvious strategy to retain the organisation's memory therefore is to retain the individuals within the organisation that are identified as the key memory repositories.

For some types of organisations, however, improving employee retention is not a viable strategy, for instance when churn is unavoidable given the nature of the organisation or the nature of the jobs. The former is particularly true in the case of volunteer organisations, and the latter is a frequent feature of for instance fast food organisations. In these situations, research has pointed towards incorporating key pieces of organisational memory into the structures of and processes within the organisation itself. For instance, in a global fast food organisation, the average tenure of employees in outlets in measured in a handful of months, and thus the organisation aims to embed as much as possible of its expertise in the processes rather than the employees.

Putting these and similar approaches into a wider and more generalizable context, we suggest that three distinct, but often intertwined strategies are at play in organisations that emphasize memory external to the individuals working for them:

- documentality (loosely following Ferraris 2010)
- procedurality
- physicality

Documentality denotes the strategy to externalize memory and expertise in (often written) form that it can be used by others. A classic case of documentality is the requirement in many organisations to keep extensive (and sometimes standardized) written logs of work done. A more recent application of documentality were the so-called knowledge management systems of the 1980s and 1990s. Employed in numerous organisational contexts, the aim was to have knowledge that traditionally rested in individuals externalized and placed in these systems. Their success was largely limited, mainly because of misaligned incentives and a lack of user experience for both creating and retrieving such digital organisational memory.

Procedurality implies that the processes put in place with organisations encapsulate relevant expertise, so that individuals within the organisation only need to follow the procedure in order to make the right decisions. Checklists are a well-known example of procedurality, and have been employed in a wide variety of contexts, from flying airplanes to treating patients, and have shown – for instance in the medical context – to significantly improve outcomes over methods relying on the memory and expertise of individuals within an organisation (Gawande, 2011).

Physicality finally signifies that the physical structures utilized by an organisation enable some and constrain other behavior thereby (ideally at least) representing the memory and expertise of the organisation. Fast food outlets use physicality, too, by designing food preparation devices that only permit to be used a certain (correct) way. Similarly, some commercial aircraft (notably those by Airbus, but deliberately not those by Boeing) filter all inputs from the pilots and only “permit” those that the flight management system assumes will not endanger the safety of the aircraft.

Documentality, procedurality, and physicality represent organisational memory within the organisation but external to individuals. They are thus more immune to churn among members of the organisation. At the same token, their relative immutability also comes at a cost. In particular, procedurality and physicality also greatly limit the ability of the organisation to change and to innovate as any such variation requires the modification of the procedures or the physical environment in place. This has been recognized as the importance for organisations to unlearn existing practices (i.e. to forget) and has become the focus in contemporary organisational memory research. (Esterby-Smith and Lyles 2011)

And because so much hinges on the organisational processes and structures if a good bit of organisational memory is incorporated in them rather than in individuals within the organisation, creating and maintaining these processes and structure will also be costly, both in terms of time and effort as well as financial resources, again, too, militating against changeability.

Organisational memory rooted in individuals’ memory and organisational memory represented external to individuals in documents, processes, and (physical) structures represent two distinct approaches to organisational memory, each with its own set of advantages and drawbacks. Organisations will have to choose between them, or – perhaps more precisely – choose the mix between them that they think is “right”. This of course prompts the question of what is “right”.

6.3 Core elements of a conceptual framework of organisational remembering and forgetting

Any conceptual framework of organisational memory will have as one of its core elements the actual organisational remembering / forgetting function. As organisations are human constructs, consisting of human beings, this function in principle resembles the human forgetting function from individual memory, reflecting different levels of recall over time.

A major insight, however, is that this function alone is insufficient to create a comprehensive conceptual framework for organisational memory, as it fails to account for the lack of a “baseline” in organisational memory, i.e. a standard of remembering and forgetting that is both established and at least claimed to be effective (if not optimal).

The “baseline challenge” thus has two parts to it:

(a) non-standard: It suggests that different organisations may have different organisational memory functions, and neither of them is *per se* better (or worse) than the others. Consequently, there is no standard organisational memory function that applies to all organisations, which could be recreated easily in a digital tool such as those developed by the ForgetIT project.

(b) dual phase: As each organisation has its own organisational memory function, it is unclear whether this function is actually effective for this particular organisation. Organisational memory functions are regularly not the result of a long Darwinian process of selection, but rather (among others) the product of a combination of a creative act *ab initio*, accumulated and reactive actual hacks/practices, and the reflection of past (or present) power structures within the organisation. If one were to explore and capture a particular organisation’s memory function and represent it in digital tools, these tools would at best mimic organisational reality at the time of tool creation, but with little ability to evolve, and even more importantly with no reference to any goals or aspirations the organisation may have with respect to its memory.

In the conceptual framework of organisational memory, we thus have to take both challenges into account. We do this by suggesting that conceptual framework has two distinct qualities that also differentiate it from the individual memory framework.

(a) Customization of the organisational memory function:

We cannot simply build into the digital tools of the ForgetIT project a standard organisational memory function. Rather, to recreate an organisation’s memory function, as every organisation differs, we first have to study how the organisation deals with memories, especially with knowledge and expertise.

The ultimate goal here is to develop an organisational memory function that has a standard form, but can (and must!) easily be customized by changing various factors and constants to reflect case specificities. This would permit the conceptual framework to scale without complex and comprehensive case studies of each organisation. It would not be perfect but perhaps good enough. In the absence of such a possibility, organisations would have to be studied individually, in order to recreate digitally their specific memory function.

Important elements to examine and measure would be:

- the initial decay of memory obtained by the organisation,

- the mechanisms and specificities (such as frequency and depths) of memory being recalled (either in its original or in an “embedded” form),

- and the balance between an organisation’s memory resting external versus internal in its members / employees (for which a more traditional human memory function combined with the probability of churn could be used).

It is important here to emphasize that the aim is to capture the status quo of organisational memory, the actual treatment of memory within the organisation rather than reflecting normative rules of dealing with memory prescribed for instance by compliance norms. There are two reasons for this. The first, and obvious is that without knowing how a specific organisation actually remembers or forgets, no baseline of organisational reality can be captured. Without knowing where an organisation stands, it cannot know in which direction it needs to move even if the ultimate goal it wants to reach is clear. Second, and perhaps more subtle is the reason that representing compliance rules within ForgetIT tools would fail to provide these tools with a competitive advantage in the market. There exists already quite a developed market for preservation tools that encapsulate various sets of standard compliance norms, and by just incorporating them ForgetIt would fail to stand out.

Ultimately, of course, the idea would be to capture as much as possible information flows and related behaviors within an organisation digitally, so that just by observing these over time, the appropriate features and contours of that organisation's memory function could be gleaned (as much as possible) automatically. This may work better in some, and worse in other organisations – and in general worse in organisations that on balance have more of its memory retained in individuals. But it would also have the very significant advantage of adaptability – as the tools not only shape what is being done, but also how the members of the organisations interact with them, the digitally embedded organisational memory function could lose its immutability, thus incorporating a truly "learning" organisational memory model.

(b) The Normative Nature of the Conceptual Framework:

Capturing organisational memory one organisation at a time, as outlined above is a first step to capture organisational memory. But, unlike most human cognitive processes organisations are not relatively fixed; they are much more plastic and thus malleable. Especially in organisations with significant parts of its memory contained external from its individual members, the effectiveness that this organisational memory is being created, retained, and recalled can be shaped by organisational design.

So, with a different organisational design a particular organisation's memory can be enhanced or constrained and thus exogenous goals (such as expertise retention or innovation capabilities) improved. Therefore one can foresee a second step in recreating organisational memory, by highlighting the normative dimension of the undertaking: organisational memory can be shaped (at least to an extent) through how memory mechanisms are implemented.

This suggests that in a second stage, organisations, after reviewing the actual state of their organisational memory as reflected in their distinct memory function, could also articulate a normative aspiration of how they want their organisational memory to be.

This could give digital tools, such as those developed through the ForgetIT project, also a unique competitive advantage in the market space, as organisations with ForgetIT tools would now be able to improve their organisational memory towards their aspired preferences and goals. This would elevate ForgetIT tools from enabling to replicate memory reality in organisations to a mechanism to shape the organisation according to its stated goals. It could thus contribute to increased revenues and/or improved profits, going far beyond purely replicating the status quo of remembering and forgetting within organisations.

The previous paragraphs use the subjunctive deliberately. There is no immediate need for organisations to move to the second stage of prescriptively shaping their organisational memory if they do not see a need. But the possibility exists, and building the capability into the conceptual framework, and perhaps into the digital tools themselves offers an enticing new use dimension.

In fact, ForgetIT partner DKD is actively pursuing the design of a business strategy to employ the ForgetIT tools they are developing into something that could be used prescriptively. They see a chance to move their business model of providing digital content management systems based on TYPO3 to include consulting services on how to improve organisational memory by implementing a prescriptive organisational memory strategy.

6.4 Operationalizing organisational memory – taking the framework to practice

In accordance with the DoW we have worked with other ForgetIT partners, in particular DKD, to operationalize the core elements of the conceptual framework and incorporate it into digital tools that are being developed as part of the project.

(a) Operationalizing the “Is”

We defined capturing organisational memory realities as the first important step above. In principle, this can be done through traditional methods such as in-depth interviews and comprehensive case studies. Partner UEDIN is using this methodological approach going forward to test actual cases for the remembering and forgetting framework developed in the ForgetIT project. Two such case studies have commenced with commerical organisations in Germany, and two public organisations have been identified for case studies in year 3 of the project.

Such an approach is suitable for gathering insight within a research project on organisational memory, but taking advantage of those insights and applying the same in depth case study approach would not be feasible in actual commercial practice, as it would require a hefty investment in resources that organisations might not be willing to make. In addition, in order to adapt an organisation’s memory function to changing realities, such case studies would have to be repeated at regular intervals, constraining even further this usability of this approach.

As a consequence, partner UOXF has been collaborating with partner DKD to suggest that organisational practices (and thus organisational preferences reflected in them) necessary to model the specific contours of an organisation’s memory function ought to be captured indirectly, through the recording and analysis of subtle signals of preservation, recall and neglect. At the same token, existing explicit rules (such as stemming from compliance norms) should be disregarded in favour of capturing the actual organisational memory practices. (If a compliance norm is successfully internalized in the organisation’s practices, it will be reflected anyway in behaviour that is being captured the way we suggest.)

As part of ongoing collaboration with DKD, including a special Organisational Memory Workshop, UOXF developed a concept of organisational preservation based on machine learning concepts, and have mapped out a preliminary structure on how to capture organisational preservations (or neglect) intentions derived from data use preferences in the context of the use of TYPO3, the content management system that is DKD’s mainstay.

This structure identifies measurable proxies of preservation preferences that map onto the three organisational memory perspectives we identified in an earlier report. For instance, DKD is capturing behavioural processes in the use of TYPO3, reflecting a behavioral perspective of organisational memory. At the same time, DKD is also experimenting with a memory decay function derived from individual remembering, to reflect the cognitive perspective view that organisational memory is similar to human memory. And finally, DKD is also capturing relevant inputs aligned with the social perspective of organisational memory, such as the quality and type of social network ties among TYPO3 users within an organisation.

Specifically, DKD has suggested four main elements of capturing organisational preservation intentions digitally.

- Quotas: These include observing the setting and adjustment of storage quotas to reframe digital remembering and forgetting within organisations as intra-organisational trade-offs.
- Preservation metrics: DKD is also implementing measuring organisational preservation/ forgetting proxies such as information production cost estimates, content popularity metrics, and information distance from core business metrics (comparing ex ante, ex post, and dynamic as part of their machine learning approach).

- Social graph: DKD is working on utilizing social graph theory, especially as put forward by Burt (2007), to capture salient information about what content an organisation implicitly suggests it wants preserved or forgotten.
- “information gravity”: Finally, DKD is experimenting with proxies that could capture the “gravity” of information, that is reflect its implied salience.

It is hoped that DKD’s implementation work continues apace, so that in the spring 2015 testing could commence for a system that captures the memory practices of an organisation through the use of suitable digital proxies (and of course in the context of a TYPO3 environment).

(b) Operationalizing the “Ought”

We suggested that an important second step is to be able to capture the memory aspirations (the “ought”) of an organisation so as to compare the realities and the aspirations and through the use of digital preservation tools help an organisation to develop towards its goal. This in turn, as we have noted, could improve ForgetIT’s market appeal.

To that end, DKD has initiated an initiative that aims to present the “ought” of organisational memory is a very simplified model of three distinct dimensions of how organisational memory is currently being used within an organisation. Organisational leadership can then identify their goals for organisational memory on these three dimensions, and organisational memory tools can shape organisational memory practices in this direction. The ‘”ought” of organisations involved in the detailed case studies will also be identified to allow for an assessment of any consistency across methods and types of organisations.

Thus, the DKD initiative aims to deliver a diagnosis tool, as well as a tool to remedy / address organisational memory desiderata of the organisation’s leadership. Both, the diagnosis and the tool to shape practices DKD believes would offer revenue potential for itself (and similarly situated information consultants), as well as revenue / profit opportunities for the actual organisations themselves (to the extent that they are commercial). While some of the details have been fleshed out already, DKD wants to develop this further before presenting it in its comprehensive form, which will likely take place in the next deliverable from WP10. There will be cross reference to reports on the organisational in depth case studies in D2.4.

6.5 Future work

Future work for UOXF focuses on fine-tuning the conceptual model based on feedback received from the case studies undertaken by UEDIN, as well as from the partners operationalizing the framework, especially DKD. This will provide us with plenty of feedback to revise the conceptual framework.

Conceptually, the focus is on dealing with heterogeneous goals and preferences within organisations expressing a particular “ought”, an organisational memory aspiration. The specific question here is who to ask for an organisation’s memory goals. Even if one agrees that goals for an organisation are regularly set by its leadership, there often is a heterogeneity of goals among organisational leaders, and the issue is how to deal with that.

Another, and large component of future work is to incorporate social remembering and forgetting into the overall framework. This was intended to take place in this reporting period, but project partners, especially DKD, have voiced very strong preferences to focus on working at the details of the organisational memory framework first as they see a more pressing use case here. Thus the incorporation of societal memory is commencing now.

7 References

Baddeley, A.D., Anderson, M. & Eysenck, M. (2015). *Memory*. Hove, UK: Psychology Press.

Baddeley, A.D., Kopelman, M.D. & Wilson, B.A. (2002). *The Handbook of Memory Disorders.* Chichester, UK: Wiley.

Barrat, A., Barthélemy, M. & Vespignani, A. (2013). Dynamical Processes on Complex Networks. Cambridge, UK: Cambridge University Press.

Bartlett, F.C. (1932). *Remembering*. Cambridge, UK: Cambridge University Press.

Benkard, C. Lanier. (2000) Learning and Forgetting: The Dynamics of Aircraft Production. American Economic Review 90:1034.

Berners-Lee, T. & Connolly, D. (2011). http://www.w3.org/TeamSubmission/n3/. Accessed on 10-February-2015.

Brewer, W.F. (1988) Memory for randomly sampled autobiographical events. In U. Neisser & E. Winograd (Eds.), Remembering reconsidered: Ecological and traditional approaches to the study of memory (pp.177-221). New York: Cambridge University Press.

Brock, G., Pihur, V., Datta, S., & Datta, S. (2008). clValid : An R Package for Cluster Validation. *Journal Of Statistical Software*, *25*, 1–28. doi:citeulike-article-id:2574494

Buchanan, T., Paine, C., Joinson, A. N., & Reips, U.-D. (2007). Development of measures of online privacy concern and protection for use on the Internet. *Journal of the American Society for Information Science and Technology*, *58*(2), 157–165. doi:10.1002/asi.20459

Bugelski, B.R. (1956). The psychology of learning. Oxford, UK: Henry Holt.

Burt. Ron. (2007) Brokerage and Closure: An Introduction to Social Capital. Oxford University Press.

Cabeza, R., & St Jacques, P. (2007). Functional neuroimaging of autobiographical memory. *Trends in Cognitive Sciences*, 11(5), 219–227.

Chakrabarti, D., & Punera, K. (2011). Event Summarization Using Tweets. In Proc. *ICWSM* (pp. 66–73).

Clayton, N. S., & Russell, J. (2009). Looking for episodic memory in animals and young children: Prospects for a new minimalism. *Neuropsychologia*, 47(11), 2330–2340.

Conway, M. A. (2005). Memory and the self. *Journal of Memory and Language*, 53(4), 594–628. doi:10.1016/j.jml.2005.08.005

Conway, M. A. (2009). Episodic memories. Neuropsychologia, 47(11), 2305–2313. doi:10.1016/j.neuropsychologia.2009.02.003

Conway, M. A., & Loveday, C. (2010). Accessing Autobiographical Memories. In J. H. Mace (Ed.), The Act of Remembering: Toward an Understanding of How We Recall the Past (pp. 56–70). Chichester, West Sussex: Wiley-Blackwell.

Conway, M. A., & Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. Psychological Review, 107(2), 261–288.

Cooper, A. (2004). *The Inmates Are Running the Asylum*. *The Inmates Are Running the Asylum*. Indianapolis, IA: SAMS. doi:10.1007/978-3-322-99786-9_1

De Choudhury, M., Diakopoulos, N., & Naaman, M. (2012). Unfolding the event landscape on twitter. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work - CSCW '12* (p. 241). New York, New York, USA: ACM Press. doi:10.1145/2145204.2145242

Du, K.L. & Swamy, M.N.S. (2014). Neural Networks and Statistical Learning. London: Springer Verlag.

Ebbinghaus, H. (Ruger, H. A. & Bussenius, C. E., Trans. 1913). Memory: A contribution to experimental psychology. New York, NY, US: Teachers College Press.

Esterby-Smith, Mark and Marjorie Lyles (2011). In Praise of Organisational Forgetting. Journal of Management Inquiry 20:311.

Ferraris, Maurizio (2010). Documentality - Why It Is Necessary to Leave Traces. Fordham University Press.

Gawande, Atul (2011). The Checklist Manifesto. Picador.

Greenberg, D.L. (2004). President Bush's false [flashbulb] memory of 9/11/01. *Applied Cognitive Psychology,* 18, 363-370.

Greenwood, M.A., Mezaris, V., Papadopoulou, O., Solachidis, V. Niederee, C. & Ceroni, A. (2014). ForgetIT Deliverable D6.2 First Release of Tools for Contextualization, Feb 2014, ForgetIT project, http://www.forgetitproject. eu/fileadmin/fmdam/deliverables/ForgetIT_WP6_D6.2.pdf

Henkel, L. A. (2013). Point-and-Shoot Memories: The Influence of Taking Photos on Memory for a Museum Tour. Psychological Science, 25, 396-402.

Hohman, T. J., Peynircioğlu, Z. F., & Beason-Held, L. L. (2013). Flexibility of event boundaries in autobiographical memory. Memory, 21(2), 249–260.

Koppel, J., Brown, A. D., Stone, C. B., Coman, A., & Hirst, W. (2013). Remembering President Barack Obama's inauguration and the landing of US Airways Flight 1549: A comparison of the predictors of autobiographical and event memory. *Memory*.

Koutstaal, W., Schacter, D. L., Johnson, M. K., & Galluccio, L. (1999). Facilitation and impairment of event memory produced by photograph review. Memory & Cognition, 27(3), 478–493.

Koutstaal, W., Schacter, D. L., Johnson, M. K., Angell, K. E., & Gross, M. S. (1998). Post-event review in older and younger adults: improving memory accessibility of complex everyday events. Psychology and Aging, 13(2), 277–296.

Kuhn, H. W. (2010). The Hungarian method for the assignment problem. In *50 Years of Integer Programming 1958-2008: From the Early Years to the State-of-the-Art* (pp. 29–47). doi:10.1007/978-3-540-68279-0_2

Kurby, C. A., & Zacks, J. M. (2008). Segmentation in the perception and memory of events. *Trends in Cognitive Sciences*, *12*(2), 72–79.

Levine, B., Svoboda, E., Hay, J.F., Winocur, G., & Moscovitch, M. (2002) Aging and autobiographical memory: Dissociating episodic from semantic retrieval. *Psychology and Aging, 17,* 677-689.

Linton, M. (1975). Memory for real-world events. In D.A. Norman & D.E. Rumelhart (Eds.), Explorations in Cognition (pp. 376-404). San Francisco: Freeman

Logie, R.H. (2003). Spatial and Visual Working Memory: A Mental Workspace. In D. Irwin and B Ross (Eds.) *Cognitive Vision: The Psychology of Learning and Motivation*, Vol 42, pp 37-78. Elsevier Science (USA).

Logie, R.H. (2011). The functional organisation and the capacity limits of working memory. *Current Directions in Psychological Science*, 20(4), 240-245.

Maguire, E. A. (2001). Neuroimaging studies of autobiographical event memory. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 356(1413), 1441–1451.

Mahler, Julianne with Maureen Casamayou (2009). Organisational Learning at NASA: The Challenger and Columbia Accidents. Georgetown University Press.

Marcus, A., Bernstein, M. S., Badar, O., Karger, D. R., Madden, S., & Miller, R. C. (2011). Twitinfo: aggregating and visualizing microblogs for event exploration. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11* (pp. 227–236). New York, New York, USA: ACM Press. doi:10.1145/1978942.1978975

Markatopoulou, F., Mezaris, V. & Kompatsiaris, I. "A Comparative Study on the Use of MultiLabel Classification Techniques for ConceptBased Video Indexing and Annotation", Proc. 20th Int. Conf. on MultiMedia Modeling (MMM'14), Dublin, Ireland, Springer LNCS vol. 8325, pp.112. January 2014.

Marsh, E. J., & Roediger, H. L., III. (2013). Episodic and autobiographical memory. In A. F. Healy (Ed.), Handbook of psychology: Volume 4, Experimental Psychology (2nd ed., Vol. 4, pp. 472–494). New York: John Wiley & Sons.

McParlane, P. J., & Jose, J. (2014). Exploiting Twitter and Wikipedia for the annotation of event images. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval - SIGIR '14* (pp. 1175–1178). New York, New York, USA: ACM Press. doi:10.1145/2600428.2609538

Mezaris, V., Solachidis, V. & Papadopoulou, O. (2014). ForgetIT Deliverable D4.2: Information analysis, consolidation and concentration techniques, and evaluation First release, Feb 2014, ForgetIT project, http://www.forgetitproject.eu/fileadmin/fmdam/deliverables/ForgetIT_WP4_D4.2.pdf

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, *45*(4), 867–872. doi:10.1016/j.jesp.2009.03.009

Owl Working Group (2012). http://www.w3.org/2001/sw/wiki/OWL. Accessed on 10-February-2015.

Peer, E., Vosgerau, J., & Acquisti, A. (2014). Reputation as a sufficient condition for data quality on Amazon Mechanical Turk. *Behavior Research Methods*, *46*(4), 1023–31. doi:10.3758/s13428-013-0434-y

Pruitt, J., & Adlin, T. (2005). *The Persona Lifecycle: Keeping People in Mind Throughout Product Design*. *Group*. Morgan Kaufmann. doi:10.1016/B978-012566251-2/50005-8

Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, *20*, 53–65. doi:10.1016/0377-0427(87)90125-7

Schacter, D. L., Koutstaal, W., Johnson, M. K., Gross, M. S., & Angell, K. E. (1997). False recollection induced by photographs: a comparison of older and younger adults. *Psychology and Aging*, *12*(2), 203–215.

Schulenberg, S. E., & Melton, A. M. A. (2008). The Computer Aversion, Attitudes, and Familiarity Index (CAAFI): A validity study. *Computers in Human Behavior*, *24*(6), 2620–2638. doi:DOI: 10.1016/j.chb.2008.03.002

Shamma, D. A., Kennedy, L., & Churchill, E. F. (2010). Tweetgeist: Can the twitter timeline reveal the structure of broadcast events. *CSCW Horizons*, 589–593.

Sørensen, T. (1948). A method of establishing groups of equal amplitude in plant sociology based on similarity of species and its application to analyses of the vegetation on Danish commons. *Kongelige Danske Videnskabernes Selskab*, *5*(4), 1–34.

Sutton, R.S. & Barto, A.G. (1998). *Reinforcement Learning: An Introduction*. Camrdige, MA: MIT Press

Taatgen, N.A. & Anderson, J.R. (2008). ACT-R. In R. Sun (ed.), *Constraints in Cognitive Architectures.* Cambridge University Press, pp 170-185.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (eds.) *Organisation of Memory*. London: Academic Press.

Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327–352. doi:10.1037/0033-295X.84.4.327

Weinberg, S. (1998). *Gravitation and cosmology: principles and applications of the general theory of relativity*. Chichester: John Wiley & Sons

Williams, H. L., Conway, M. A., & Baddeley, A. D. (2008). The Boundaries of Episodic Memories. In T. F. Shipley & J. M. Zacks (Eds.), Understanding Events: From Perception to Action (pp. 39–52). New York: Oxford University Press.

Van der Ark, L. A. (2007). Mokken Scale Analysis in R. *Journal of Statistical Software*, *20*, 1–19. doi:10.1007/s11336-007-9034-z

Wolters, M. K., Isaac, K. B., & Renals, S. (2010). Evaluating Speech Synthesis Intelligibility Using Amazon Mechanical Turk. In *Proc. SSW-7*.

Wolters, M., Niven, E. & Logie, R.H. (2014): The Art of Deleting Snapshots. In: *Proceedings of the ACM SIGCHI Conference on Computer-Human Interaction 2014,* Toronto. http://dx.doi.org/10.1145/2559206.2581287

Wynn, V.E. & Logie, R.H. (1998). The veracity of long-term memories - Did Bartlett get it right? *Applied Cognitive Psychology*, 12, 1-20.

W3C Standards. http://www.w3.org/standards/techs/rdf. Accessed on 10-February-2015.

Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: A mind-brain perspective. Psychological Bulletin, 133(2), 273–293.

Zacks, J. M., & Swallow, K. M. (2007). Event segmentation. Current Directions in Psychological Science, 16(2), 80–84.