

ForgetIT

Concise Preservation by Combining Managed Forgetting and Contextualized Remembering

Grant Agreement No. 600826

Deliverable D2.4

Work-package	WP2: Foundations of Forgetting and Remembering
Deliverable	D2.4: Foundations of Forgetting and Remembering - Final Report
Deliverable Leader	Robert Logie
Quality Assessor	Vasilis Solachidis (CERTH)
Dissemination level	PU
Delivery date in Annex I	31.1.2016
Actual delivery date	11.3.2016
Revisions	7
Status	Final
Keywords	organisational memory, societal remembering, human memory, preservation, forgetting

Disclaimer

This document contains material, which is under copyright of individual or several ForgetIT consortium parties, and no copying or distributing, in any form or by any means, is allowed without the prior written agreement of the owner of the property rights.

The commercial use of any information contained in this document may require a license from the proprietor of that information.

Neither the ForgetIT consortium as a whole, nor individual parties of the ForgetIT consortium warrant that the information contained in this document is suitable for use, nor that the use of the information is free from risk, and accepts no liability for loss or damage suffered by any person using this information.

This document reflects only the authors' view. The European Community is not liable for any use that may be made of the information contained herein.

© 2015 Participants in the ForgetIT Project

Revision History

Date	Version	Major changes	Authors
3.12.2015	v0.1	Evaluation Method	Cristina Alexandru, Stephen Rhodes
4.12.2015	v0.2	Conceptual Framework Organisa- tional and Societal Memory	Viktor Meyer- Schönberger
7.12.2015	v0.3	Structure	Maria Wolters
14.01.2016	v0.4	Festival Study Text	Elaine Niven
25.01.2016	v0.5	Case Studies of Organisational Mem- ory	Jason Rutter
10.02.2016	v0.6	Revision of v0.5	Robert Logie
16.02.2016	v0.7	Introduction and Conclusion	Robert Logie

List of Authors

Partner Acronym	Authors
UEDIN	Robert Logie, Cristina Alexandru, Elaine Niven, Stephen Rhodes, Jason Rutter, Maria Wolters
UOXF	Viktor Mayer-Schönberger

Table of Contents

Table of Contents	4
Executive Summary	7
1 Introduction	8
1.1 Target Audience	8
1.2 Overview and Structure of Deliverable	8
2 Experimental Studies on Personal Digital Photograph Use	10
2.1 The August 2013 Festival Fringe Study	10
2.1.1 Recall data	12
2.1.2 Consistency of Groupings	19
2.2 August 2014 Festival Study	21
2.2.1 Demographics	23
2.2.2 Consistency of Events and Groups	23
2.2.3 Automatic Contextualisation	25
2.2.4 Keep and Delete Decisions	26
2.3 Conclusion	30
3 Survey	31
3.1 The Final Sample	31
3.2 Personas for Preservation	34
3.3 Outlook: Preserving Digital Photos Across Cultures and Generations	38
4 A Conceptual Framework for Organisational as well as Societal Remembering and Forgetting	40
4.1 Development of Conceptual Frameworks	40
4.1.1 Existing research in organisational memory	41
4.1.2 Existing research in societal remembering and forgetting	43

4.1.3	Assessing Similarities - From Frameworks to Framework	45
4.1.4	The Conceptual Framework – A Process-Oriented Approach	48
4.2	Operationalizing Social Memory – Taking the Framework to Practice	49
4.3	What the Future may hold	51
5	Organisational Memory and Forgetting: Review and Case Studies	52
5.1	Literature Review	52
5.1.1	Organisational Memory	52
5.1.2	Organisational Forgetting	54
5.2	Methods	57
5.2.1	Aims of the Research	57
5.2.2	Data Collection	58
5.3	Findings	59
5.4	Considerations for Managing Digitised Organisational Memory	62
5.4.1	Motivation for Managing Memory	62
5.4.2	Form of Organisational Memory	62
5.4.3	Access to Organisational Memory	63
5.4.4	Managing Memory Context	64
5.4.5	Practice	64
6	Evaluation Methodology	66
6.1	Personal Preservation Use Case: Design and Methodology	66
6.1.1	Design and Motivation	66
6.1.2	Method	66
6.2	Organisational Preservation Use Case: Design and Methodology	68
6.2.1	Design	68
6.2.2	Method	69
7	Conclusion	72
7.1	Assessment of Success Indicators	72

7.2	Lessons Learned	72
7.3	Vision for the Future	74
	References	75
	Acronyms	79
A	Sample Callback Sheet From 2014 Festival Study	80
B	Schedule for ForgetIT Semi-Structured Interviews	82
B.1	Preamble	82
B.2	Interview Topics and Questions	82
B.3	Afterward	84
C	Material for the Organisational Evaluation	85
C.1	Participant Information Sheet	85
C.2	Intake Interview	88
C.3	Debriefing Interview	90
C.4	Interview Schedule for Odd Participant Numbers	91
C.5	Interview Schedule for Even Participant Numbers	94
C.6	List of Themes and Subthemes	96
C.6.1	Themes and subthemes for the TYPO3 evaluation	96
C.6.2	Themes and subthemes for the evaluation of the annotation tool	97

Executive summary

This deliverable from Workpackage 2 is a report on:

1. The final results of two phases of an experimental study that collected objective data on how human users recall details of personally experienced events and how they use and organise digital photographs of those events. The study involved collaborations between UEDIN (WP2) and CERTH (WP4), USFD (WP6), and DFKI (WP9). Technologies developed by CERTH and USFD were used to contextualise the photographs, and photographs were managed using using the Personal Information Management systems developed by DFKI (Section 2).
2. The final results from 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. (Section 3)
3. A summary of the conceptual model of organisational and societal remembering and forgetting that was developed for the project. (Section 4)
4. A report of four case studies examining the feasibility of the conceptual model in practice, involving a learned society, a journal, a museum, and two companies. (Section 5)
5. The design and methodology used for the formative evaluation of the personal and organisational preservation methods developed in WP9 (DFKI) and WP10 (dkd). The results of the evaluations, which UEDIN carried out for both partners, are reported in D9.5 (WP9) and D10.4 (WP10). (Section 6)

1 Introduction

1.1 Target Audience

This deliverable is intended for computer scientists and psychologists interested in how an understanding of human memory and forgetting can be used to inform the design of information preservation and managed forgetting systems.

1.2 Overview and Structure of Deliverable

WP2 has been concerned with the fundamental principles and functioning of human remembering and forgetting, and how those principles might inform the design of managed remembering and forgetting in digital preservation systems. The overall motivation has been to help ensure that digital preservation services fulfil the requirements of human users of those services both over the long term, preferably across generations, as well as over short periods and for current active use. The approach within the overall ForgetIT project has been to use the understanding of human remembering and forgetting from WP2 to guide the design of the ForgetIT Framework and demonstrators. Human remembering and forgetting have been considered at individual, group, and societal levels, with the major focus on the individual use case. WP2 has been led by partner UEDIN, with contributions from partner UOXF. Over the course of the ForgetIT project, there have been collaborations with WP3, WP4, WP6, WP8, WP9 and WP10, as well as contributions to WP1 and WP11.

As stated in the Description of Work, the primary activity within WP2 has been at the level of the individual as a potential user of digital preservation. This has driven our review and evaluation of previous research, the design and conduct of our objective empirical studies, as well as subjective interviews and surveys, the development of a conceptual model, and the user evaluations of the personal preservation system (PIMO) developed by partner DFKI in WP9. The majority of that work has already been reported in D2.1, D2.2 and D2.3. Additional results from the empirical studies and survey are reported in Section 2 and 3 of the current document. The methodology for evaluations of the PIMO is described in Section 6.1. The results of those evaluations are reported in D9.5.

A second focus has been at the level of organisations. In the Description of Work, this activity was intended to be a focus for partner UOXF in collaboration with partner DKD, but with UEDIN undertaking an evaluation of an initial digital preservation system for organisations to be developed by partner DKD in WP10. Towards the end of the second year of the project, the decision was taken to move the investigations and in-depth interviews of staff within organisations to the team at UEDIN. This work is described in Section 5 of the current document. The methodology for evaluation of the initial DKD system for web page preservation by organisations is described in Section 6.2 and the results are given in D10.5. The development of a conceptual model of remembering and forgetting

in organisations remained the responsibility of UOXF, and the outcome of this work is reported in Section 4. The development of a ForgetIT preservation system to function at the societal level was outside the scope of the Description of Work for the project, but UOXF undertook planned work on the development of a conceptual model of remembering and forgetting in societies, reported in Section 4.

2 Experimental Studies on Personal Digital Photograph Use

The UEDIN team conducted two experimental studies with the aim of collecting objective data that would inform the design of a ForgetIT system for personal preservation. The studies focused on how potential users curate and categorise their digital memories, and how use of digital memories affects their biological memory for the original source information (here: an experienced event). For these studies, we worked closely with project partners DFKI, CERTH, and USFD.

The experienced events we chose for both studies were linked to the Edinburgh Fringe Festival, a large annual arts festival that takes place for three weeks in August each year in Edinburgh, UK. Participants were supplied with a simple smartphone, and were asked to take regular photographs of the events they experienced. They then proceeded to revisit these photographs as part of a protocol to let them mentally revisit the original events at various time periods ranging from a few hours through to 11 months after the events were originally experienced.

The first study, which started in August 2013 and was completed with the 11 month follow up in July 2014, focused on a large street fair that is associated with the Festival. 74 participants spent an hour at the street fair and were prompted by the smartphone to use it to take snapshots every three minutes to document this experience. The design of this study is described in detail in D2.1; initial results are reported in D2.3.

For the second study, which was carried out in August–September 2014, 22 participants were asked to experience and document a day at the Festival by taking photos using the supplied smartphone. Participants were asked to produce between 40 and 80 photographs, taken whenever they felt there was an appropriate event to photograph. The design of this study is reported in D2.3.

In Section 2.1, we summarise the main findings of the 2013 Festival study, with an emphasis on the analysis of autobiographical recall. Results on grouping patterns will only be discussed briefly, because preliminary findings have been reported in Deliverables D2.1, D2.2, and D2.3. We have omitted a detailed report on additional findings regarding preservation decisions for the July 2014 follow up, as the results are very similar to the findings for the 2014 study.

Section 2.2 focuses on the main findings of the Festival 2014 study that are relevant for the design of ForgetIT personal preservation solutions, the usefulness of automatic contextualisation, the factors that affect preservation decisions, and the grouping of photos according to events and categories.

2.1 The August 2013 Festival Fringe Study

We report results for four aspects of our data set:

- the contribution of episodic (i.e., event-related) detail to participants' recall;
- the extent to which participants' recall centred around information present in their photographs
- which photographs participants chose to keep or delete at different time intervals;
- the criteria participants used for organising photographs into meaningful groups.

We focused on evidence for changes over time, and the effect of having viewed or not viewed photographs as part of a post-event review.

As noted in D2.3, participants were split into four pseudo-randomly selected groups. Their allocation was constrained by dates upon which participants were able to commit to returning (for demographic details, see Table 1). 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.

Eleven months 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. Participants who took part in the 2013 August festival study were asked in their original session whether they would be willing to participate in related experiments. Of those who agreed, 6 per experimental group were randomly selected and contacted to see if they would return for a 'related study'. Where one of these participants was unable to return, a replacement from the same group was subsequently randomly chosen. This follow-up study allowed us to investigate:

- the way in which memory for an event and perceived context and preservation value changes over a longer time frame
- the long-term influence of having reviewed photographs immediately after the event or later

Group	1	2	3	4
Number of participants	20	18	18	18
Age (years)	32 ± 5.5	39 ± 20	34 ± 17.5	35 ± 15
Age range	19–68	19–71	20–72	19–59
Males	3	8	5	10

Table 1: Demographic details of participant groups

2.1.1 Recall data

As noted in D2.1 and D2.2, our autobiographical knowledge—our knowledge of our personal history and experiences—is attributed to both initial and re-experience of events. Memory for events is often, however, characterised by loss of access to details immediately following the experience—indeed, rapid loss of episodic detail is often emphasised as an inherent feature of this type of memory. This therefore leaves open the question of how we can build up our autobiographical knowledge over time, and how we can re-visit experiences (for example, to recall what happened, to imagine alternative outcomes if changes in behaviour were made, to consider whether there were clues present during the event about something which we have only recently learned, and so on).

While the rapid large volume of loss of details is well established, the nature of what is preserved longer term is not entirely clear (e.g., [Conway, 2009]). Moreover, reviewing an event—either through oral recall, or through use of photographs—can benefit memory for that event, including the amount of episodic detail that can be remembered [Koutstaal et al., 1998]. Additionally, previous studies have demonstrated that key details of an event can be maintained consistently over periods spanning a number of months, when an initial recall has taken place some time after the event (e.g., [Wynn and 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, and whether the details that are maintained are affected by photograph viewing (for example, whether information present in reviewed photographs may come to be the dominant content in recall of the event).

For the 2013 Festival Fringe study, the Autobiographical Interview [Levine et al., 2002] scoring technique was applied to recall data generated at Times 1, 2 and 3 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 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 was also addressed. Within each interview session (excluding Time 1 interview session for group 4), 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.

Results of recall for Groups 1, 2 and 3 are shown in Figure 1. This representation of recall shows the amount of information produced by participants that was episodic that is, internal or key features of events (for example, details of event, perceptual, emotional, location and temporal information), as a proportion of the overall amount of information the participant generates in their interview (for example, non-episodic information provided could include semantic/general knowledge statements, non content-full utterances, and

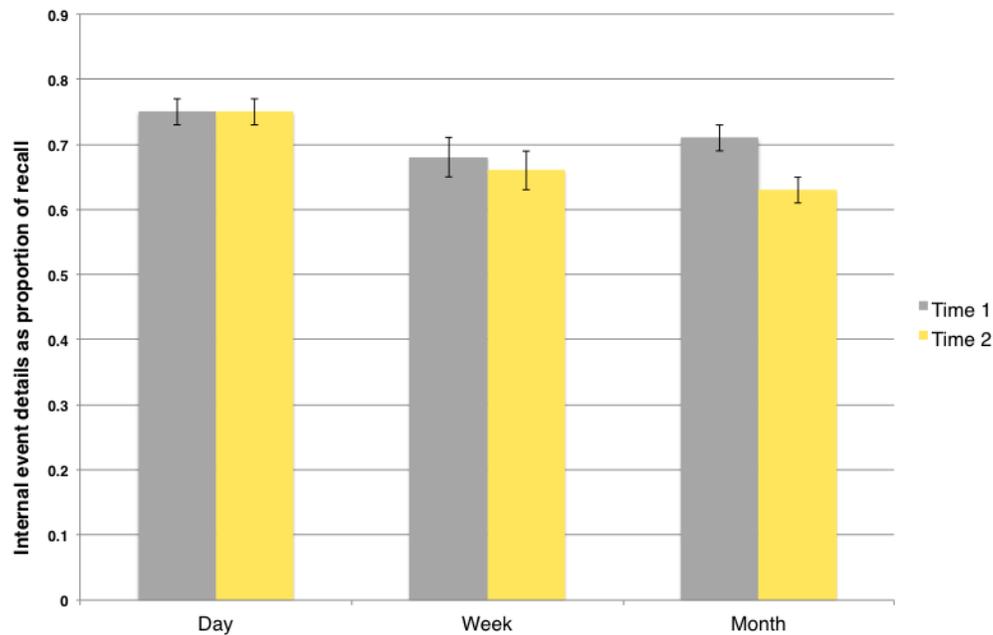


Figure 1: Proportion of details recalled from original event, at interview Time 1 and interview Time 2, compared across different delays.

memories of other related events). Statistical analysis showed that the amount of recall at Time 2 that was episodic in nature was lower than that produced at Time 1 overall ($F(1,53) = 4.42, p = .04$)¹. While the Groups 1-3 differed overall in the proportion of recall generated that was episodic in nature ($F(2,53) = 5.70, p < .01$), when broken down the three groups did not significantly differ from each other individually and, importantly, the amount of change in recall from Time 1 to Time 2 did not differ between the groups. This therefore indicated that while recall became less episodic after immediate recall, this change did not appear to vary across further time frames (of a day, a week and a month).

While the above noted results from the Festival Study appear to demonstrate that recall becomes less episodic after immediate interview, this impression could be the result of two different changes recall after a delay may include less episodic information or it could include more non-episodic information while the same amount of episodic information is maintained. Therefore, the number of episodic details generated at Time 1 and Time 2 were considered. Analyses of this information, as presented in Figure 2, reveals that the number of episodic details produced during recall did not significantly differ between Time 1 and Time 2, nor did it differ across the three different delay groups. Moreover, the lack of change between Time 1 and Time 2 was consistent across the delay periods.

In sum, the comparison of the above two sets of findings indicates that participants are providing the same amount of episodic details after a delay, but it is proportionately less of

¹The F value obtained from analysis of variance (ANOVA) summarises the ratio of systematic variance (i.e. that explained by group difference) to unsystematic (error) variance. The corresponding p value gives the probability of an F -ratio the size observed or larger under the null hypothesis of no group differences. By convention a p value smaller than 0.05 is taken as significant evidence against the null hypothesis.

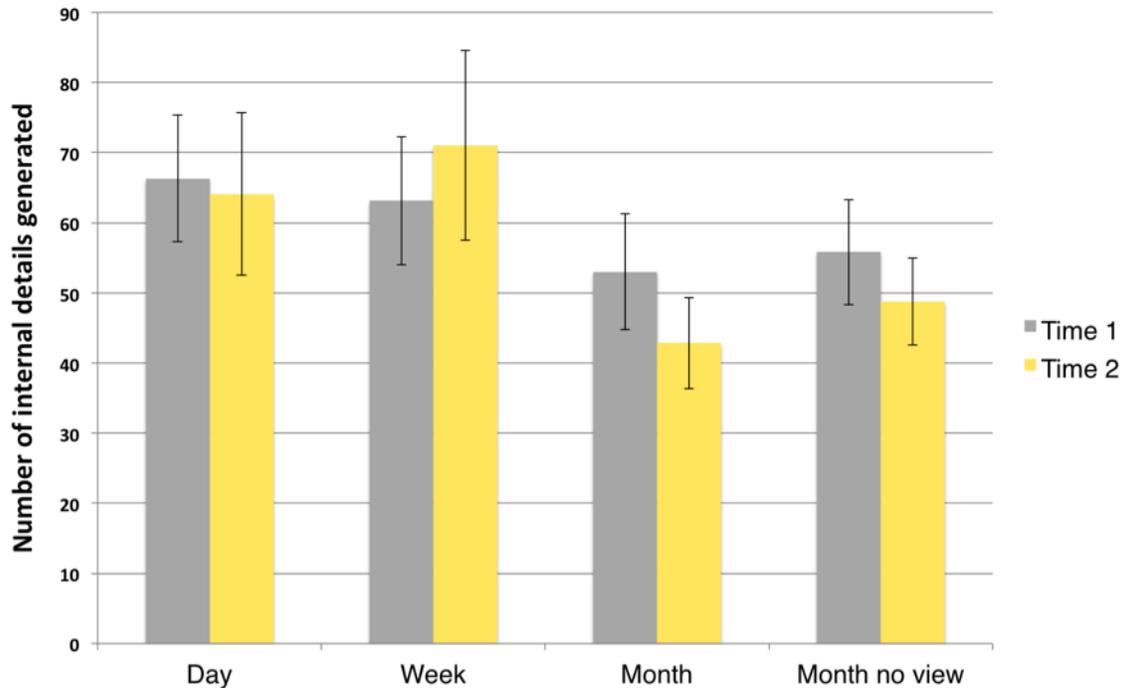


Figure 2: Number of episodic details recalled in interview at Time 1 and Time 2, compared across different delays.

their recall at second interview. This production of more non-episodic details after a delay period could be for a number of reasons, such as because participants feel they should be able to recall more information and therefore volunteer contextual details (for example, providing a knowledge of a history of the Festival) to enhance their stories, or because they are adding in lifetime context (such as Festival-related things that have happened to them since their first interview), or because they were initially too tired after experiencing the event to be able to provide all details that they could remember at interview. Further analyses of the data (such as applying rating scales of episodic richness of descriptions; [Levine et al., 2002]) will be able to investigate this finding further. However, currently these results indicate that memories for events may still prove to be rich in episodic detail over time periods of up to a month, following an initial interview.

As noted above, review of photographs has previously proven a means of support for reminiscence, with increase in the amount of information that can be recalled while viewing the photographs (e.g. [Koutstaal et al., 1998]). In the Festival 2013 study, the influence on recall of having previously viewed photographs was investigated; comparison of recall between Groups 3 and 4 allowed comparison of recall over the same time frame (one month; see Figure 3) with one group having viewed their photographs following recall at Time 1 (Group 3) while the other group did not view their photographs between recall at Time 1 and Time 2 (Group 4). Analysis of proportion of recall that was episodic in nature showed a significant reduction between first and second interview ($F(1,34) = 6.30, p < .05$), this reduction did not differ between the two groups: viewing or not viewing photographs following initial recall did not change proportion of episodic information

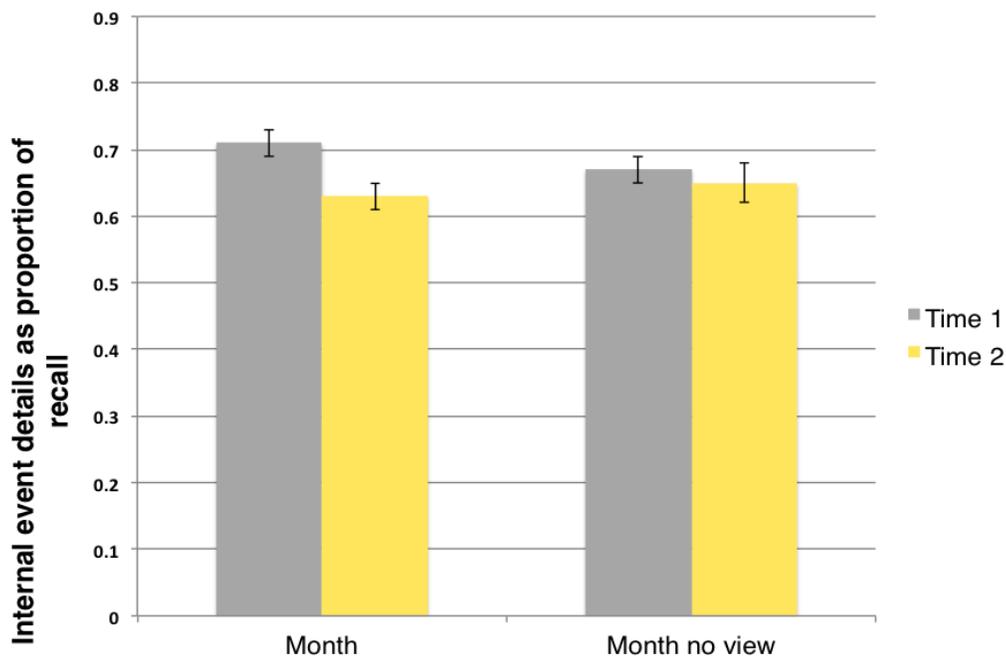


Figure 3: Proportion of details recalled from original event, at interview Time 1 and after one month (interview Time 2), for participants who did and did not view their photographs at immediate interview.

maintained over a one month period ($F(1,34) = 2.20, p = ns$). As in the comparison of time frames carried out above, the data was then subject to comparison of number of episodic details generated at initial and second interview (see Figure 2 above). These comparisons revealed that there was no difference in the number of event details generated on the basis of time of interview (Time 1 and Time 2) nor was there a difference in the number generated by Group 3 or Group 4. Moreover, viewing (Group 3) or not viewing (Group 4) photographs after initial recall (Time 1) had no significant effect on the consistency of amount of episodic details produced at Time 1 and Time 2.

While these analyses demonstrate that the number of internal or episodic details did not significantly change as a consequence of time, nor of photograph review, they do not establish that the same details are recalled at Time 2 as at Time 1. That is, details which were remembered at first recall may be lost, but new information—that present in photographs reviewed—could be incorporated into recall and thus be present at second interview. This could—for example—indicate selective rehearsal effects due to ‘post event reviewing’.

To investigate the potentially biasing effect of photograph review to recall, a second and entirely separate scoring procedure was applied to the interview data; this new scoring procedure sought to look at content inclusion and changes over time—specifically, is recall of photograph content more evident in participants who did see their photographs at Time 1 (Group 3) compared to those who did not see their photographs at Time 1 (Group 4) when comparing recall over the same time scale (a month). Independent coders were



Figure 4: Image captured by a participant during their experience in the Festival study and used throughout the interview and decision making requirements for that participant.

trained on a procedure to detail the content of every individual photograph each participant had generated. For example, the coder would state the main subject of the photograph in Figure 4 as a performer outside a known landmark (St Giles), dressed in a red and black jumpsuit, wearing a wig, with secondary content (used to aid later decisions made by subsequent coders) noted as that the performer is being observed by a large group of people, including a substantial number of children. A separate coder then addressed all recall content produced by participants and coded (in binary form) whether the main content of each photograph was present, at each interview time point, in the recall descriptions of the participant who took (and thus later reviewed) the photograph.

The proportion of photographs which were spoken about—per person—at initial interview (Time 1) and second interview (Time 2) was then considered (see Figure 5). Analyses revealed that the proportion of an individual's photographs that were noted in recall at Time 1 did not significantly differ from the proportion noted as present in recall at Time 2, and did not differ overall between those who did or did not see their photographs after initial interview. Moreover, there was no significant indication that viewing of photographs immediately after first recalling the experienced event had any effect on inclusion of their content in recall one month later.

Recall after one year (Time 3) from those participants who returned for interview (24 participants) was subject to the same three analyses detailed above, in order to determine any long term effects that viewing or not viewing pictures immediately after recall may have, when participants subsequently recall and review their pictures one month after the event (Groups 3 and 4 respectively). Similarly, recall after one year was considered on the basis of the delay period between initial event and recall (Time 1) and subsequent recall/review (Time 2).

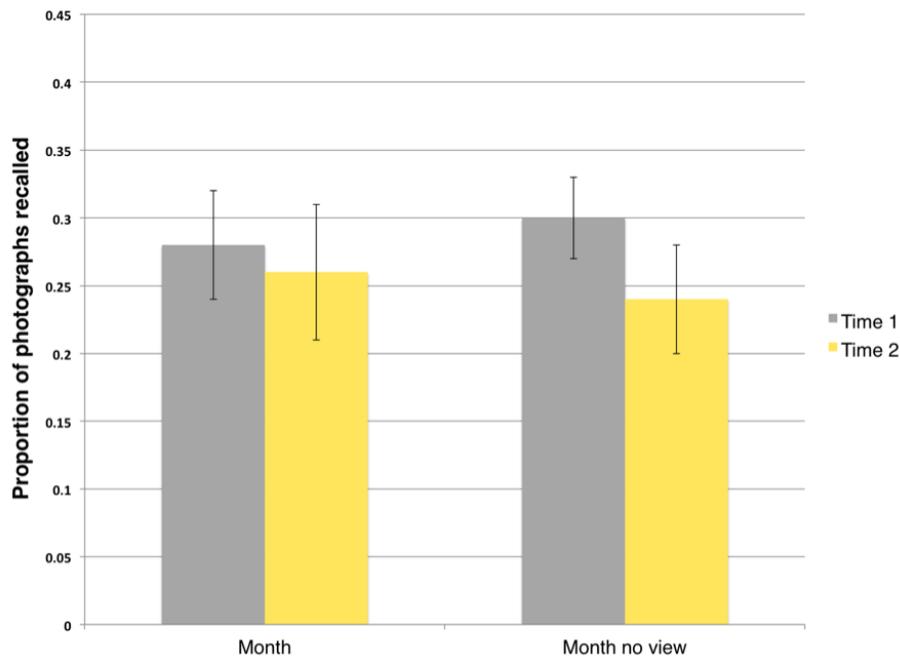


Figure 5: Proportion of an individual's photographs that were spoken about in recall, at Time 1 and Time 2, of those participants who did and did not see their photographs following initial recall.

As demonstrated in Figure 6 (upper panel) proportion of recall that was episodic in nature was found to be lower at Time 3 than at Time 1, ($F(1,15) = 56.20, p < .001$). However, there was no overall difference between the three review groups $F(2, 15) = 2.52, ns$, and this spacing of delay of second review (that is, the time period associated with Time 2 for each of the three groups) had no significant influence on proportion of recall that remained episodic in nature after one year ($F(2,15) = 1.03, ns$).

This pattern effects was also true when the number of unique episodic details provided at recall was considered; there was a loss of detail after one year, but the timing of second review did not significantly influence this loss.

A comparison of those who did and did not see their photographs at immediate review, and came back for a second interview after a month (see Figure 6, lower panel) shows that, one year later, proportion of recall that is episodic in nature is lower than that at Time 1. This drop was found to differ on the basis of having seen photographs ($F(1,10) = 6.57, p = .03$); people who initially viewed their photographs showed a larger drop in recall after one year than people who did not view their photographs immediately after the event.

However, when the absolute number of unique episodic details produced was considered, the difference between Time 1 and Time 3 remained, yet there was no influence on this change of having viewed or not viewed photographs at Time 1 ($F(1,10) = 0.10, ns$).

Interpretation of these results could suggest that those participants who saw their photographs after initial recall produce more non-episodic details after one year, perhaps

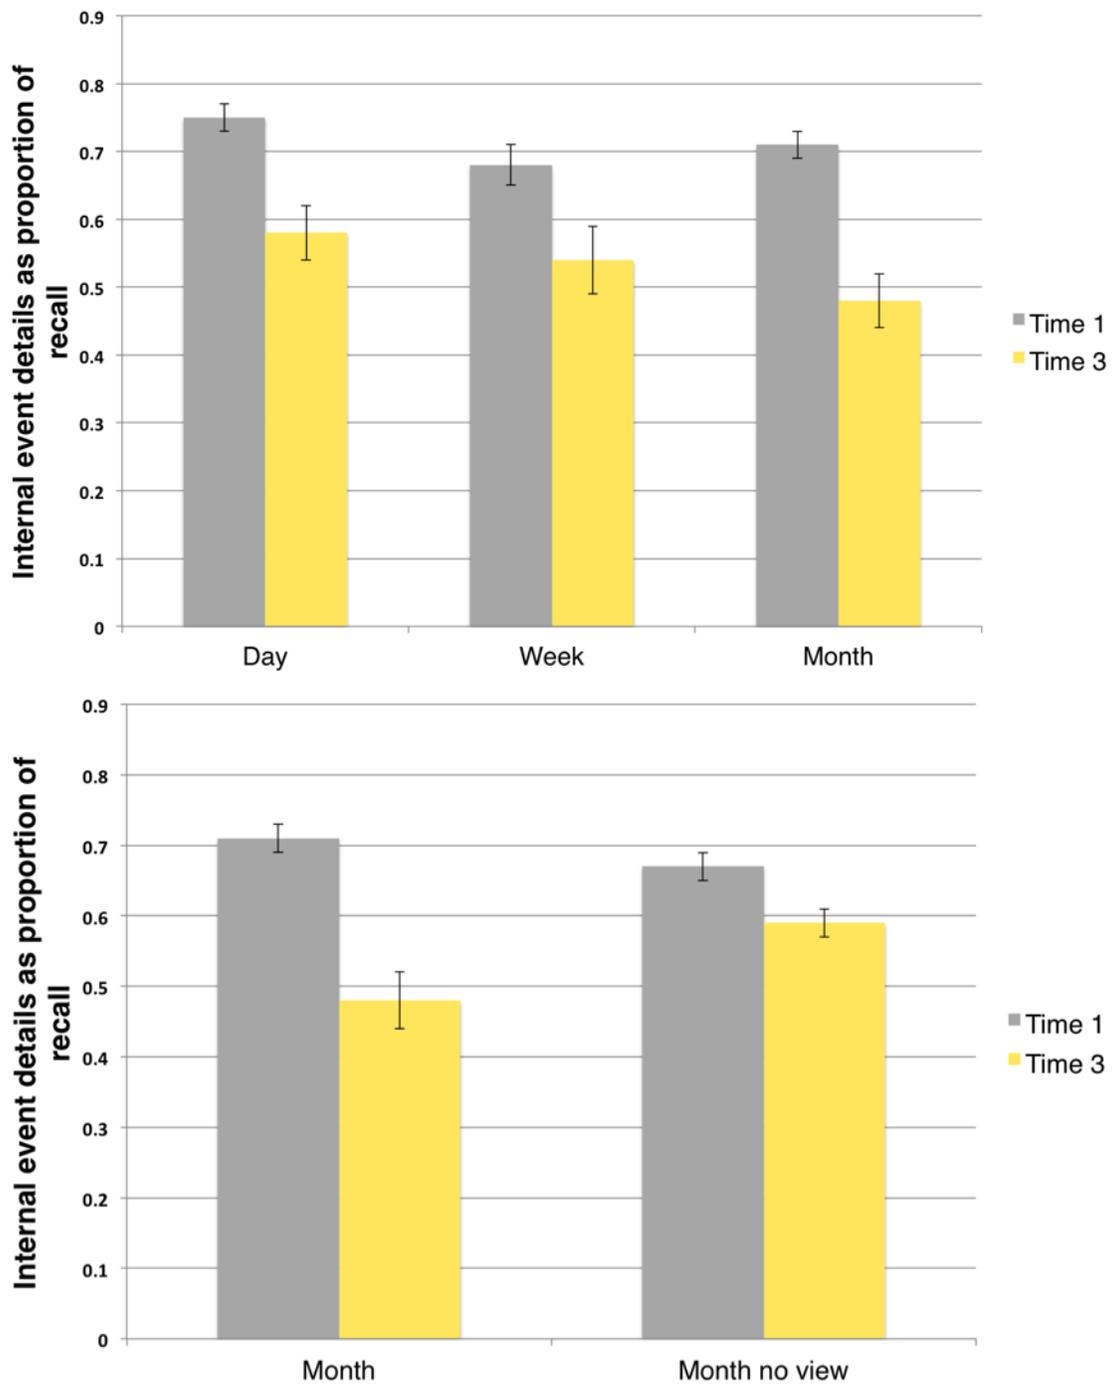


Figure 6: Proportion of details recalled from events at Time 1 and Time 3, compared across different delays (Groups 1, 2 and 3; top panel), and between participants who did and participants who did not view their photographs at immediate interview (bottom panel).

relating to the time of the second interview—for example, reflecting on how they felt they ‘performed’ at the month interview and/or their surprise at, on seeing their photographs at Time 2, remembering having already seen them at Time 1.

While the above comparisons demonstrate that participants—across all groups—remember less after one year, and proportionally and absolutely less of it is episodic the question remains of what information does persist in memory? Content analysis—assessing the prominence of participants’ photographs in their recall—found that, for Groups 3 and 4, the proportion of participants’ photographs that are mentioned within their recall did not appear to clearly differ between interviews produced at Time 1 and interviews one year later. Moreover, this pattern held for both those participants who had viewed their photographs immediately after initial recall and those who only saw their photographs one month after the initial event and recall.

In summary, the Festival study data indicated reduction in proportion of episodic detail available about an event beyond the immediate past; comparison over a one-month period offered no indication that this reduction was affected by length of delay, nor by viewing or not viewing photographs. There was, however, no significant change in number of episodic details generated over this time period. While lack of influence of length of time since an event is not without precedent [Wynn and Logie, 1998], the current work adds to the literature in demonstrating that this consistency is found for a recently experienced event (with initial recall taking place immediately after the event). Additionally, these findings expand on existing work by considering change across three different frames (a day, a week and a month).

After one year, drop in episodic detail (either by proportion or by sheer number of details) produced between immediate and delayed conditions in the one year follow-up study suggests a vulnerability in event memories not fully protected by (repeated) verbal or photographic review. Of note, however, and as visible in Figure 6, approximately 50% of recall is still episodic in nature after one year. Such a figure is substantial, and this ability to recall detail is not often acknowledged in the existing literature on event memories, particularly in comparison to emphasis placed on loss of detail [Conway, 2009]. Of particular consideration for the current project, the study does not provide any evidence that repeated viewing of photographs—when carried out following unsupported recall—would bias the content of an individual’s recall of an experience.

2.1.2 Consistency of Groupings

In the 2013 study, participants were required to sort their photos into meaningful groups several times. Participants who saw their photos at Time 1 (Participant Groups 1-3) sorted them once on the day of the study (Sort 1). All participants then produced two sorts at Time 2 (sort 2.1 and 2.2). Finally, those participants who returned after a year also sorted their photos into groups (Sort 3).

The instructions for Sort 1, 2.1, and 3 were always the same—participants were asked to assign photos to meaningful groups in such a way that each photo belonged to exactly

Table 2: Group Similarity for Participants Who Reviewed Their Photos at Time 1, Sorts at Time 1 and Time 2, Jaccard Index. Similarities range between 0 and 1.

	Similarity between			Significant Differences
	Sort 1 / 2.1 (S1S21)	Sort 1 / 2.2 (S1S22)	Sort 2.1 / 2.2 (S21S22)	
All	0.75 ± 0.2	0.49 ± 0.2	0.47 ± 0.1	S1S21 > (S1S22, S21S22)
Day	0.80 ± 0.2	0.53 ± 0.25	0.49 ± 0.2	S1S21 > (S1S22, S21S22)
Week	0.77 ± 0.2	0.50 ± 0.2	0.49 ± 0.1	S1S21 > (S1S22, S21S22)
Month	0.68 ± 0.2	0.43 ± 0.2	0.40 ± 0.1	S1S21 > (S1S22, S21S22)

one group. For Sort 2.2, the second sort at Time 2, participants were asked to produce a meaningful sort that was different from Sort 2.1.

Here, we report results for a research question which is particularly relevant to the ForgetIT context: How similar are the initial photo groupings that were created at Time 1 (same day), 2 (recall day/ week / month later), and 3 (year later)? If groupings are relatively similar, this indicates that individuals show consistent patterns when grouping photos which can be learned from data.

As the three similarity measures proposed in D2.3 roughly yielded the same results, we will focus on the Jaccard Index here, as this is the most straightforward of the similarity measures. Our statistical analyses focused on differences in similarity between sorts. Since participants had been instructed to make the two sorts at Time 2 different from each other, the similarity between Sort 2.1 and Sort 2.2 provides a baseline, i.e., the values we obtain when two sets of groupings are intended to be dissimilar.

Table 2 shows our findings. Perhaps the most surprising finding is that S1 is similar to both of the sorts at Time 2, S21 and S22. If people had one consistent way of categorising their photos, we would expect this default to emerge in the first sort at each time. We would also expect this default to be substantially different from S22, which is designed to be dissimilar from the default. This pattern is exactly what we are seeing in the data. The first sorts at each time (S1 and S21) are relatively similar to each other, and both are equally dissimilar from S22 (paired Mann-Whitney test, $p < 0.001$).

While Table 2 suggests that the similarity between S1 and S21 decreases as the interval between sorts increases, a Kruskal-Wallis test suggests that this difference is not particularly meaningful ($p < 0.3$).

Figure 7 gives an overview of group similarities between Time 1, Time 2, and Time 3 after a delay of 11 months. Due to the small number of participants per cell, we did not compute significance tests. The two main results are that S3 is only similar to S1 and S21 if participants returned after a day. For longer delays (week, month), the similarity between S1 and S3 is close to our baseline for dissimilar sorts.

The findings for those participants who returned after a month, and did not view their photos at Time 1 (MonNo or Month No Sort) are different yet again. The sort they produce

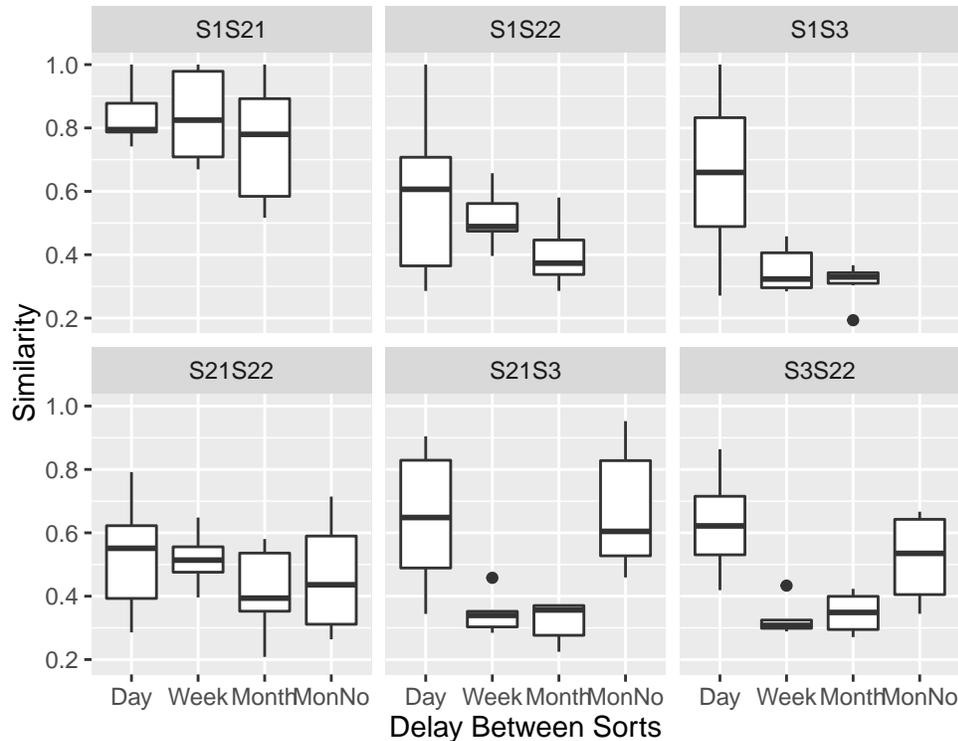


Figure 7: Group Similarities between Time 1, Time 2 Sort 1 (S21), Time 2 Sort 2 (S22), and Time 3 for all four participant groups

at Time 3 is similar to both of the sorts at Time 2.

We conclude that while there is evidence that people have a preferred way of sorting their digital photos into groups, these preferences can change over time. The stability of the preferences might also potentially be affected by rehearsal patterns, such as the last time a person viewed their photos, and how deeply they engaged with their photos at that time.

2.2 August 2014 Festival Study

In August of 2014, the UEDIN team conducted a second, more naturalistic study during the Edinburgh Festival Fringe. This time, participants were asked to document a day at the Festival, including an hour spent on the Royal Mile Street Fair.² Participants chose when and how to take pictures to document their experience, with the only constraint that they should take between 40 and 80 photos. While this is far more than one would take in a normal situation, it allowed us to compare the photographic evidence with participants' accounts for the whole 24 hours.

Participants were interviewed twice, once immediately after their time at the Festival (Time

²In future analyses, this will allow us to compare the data collected in the 2013 Festival study, where participants were asked to take a photo every three minutes, with a situation where participants had more freedom.

1) and once a month later (Time 2). Ratings were collected using pen and paper at Time 1, and at Time 2, participants used both pen and paper and a preliminary version of the DFKI PIMO5 system to annotate, categorise, and organise their photos.

The primary research questions were as follows:

- How does episodic recall change over time? This time, we investigated a longer stretch of time, with several distinct events.
- How do people categorise the photos they take, both in terms of groups and in terms of attributes? To what extent can automatic categorisation be used?
- To what extent are people's decisions to keep or delete the photos they take consistent over time, and what other attributes of photos (such as landmark status or meaningfulness) affect those decisions?

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 provide the basis for investigations of changes in episodic recall over time.

After the interviews, participants were asked to assign keep/delete values to their photographs, and to sort them twice, once according to events, and once according to non-event categories. For example, at Time 1, Participant 20 created five event-related groups, and six category-related groups. The events were "Grassmarket (first show she went to see)", "Princes Street" (a large shopping street in Edinburgh), "Royal Mile" (the hour on the Royal Mile), "inside the Tron Kirk" (former church on the Royal Mile), and "on her way home". The categories were "murals", "street performers", "landscape", "posters and advertising", "the two shows she went to", and "stained glass in the Tron Kirk".

For each event, each group, and the whole day itself, participants chose a landmark photo. For those participants who took more than 80, we chose a random subset of 80 for grouping and preservation value decisions. We then sampled another subset of 30 photos that was designed to be representative of the day.

For these 30 photos, participants were asked about meaningfulness, desired ease of retrieval (findability) and desired speed of retrieval. At Time 1, participants were asked to make handwritten notes, add their own tags, and select appropriate tags from a list of 20 frequent CERTH picture descriptors, which had been derived from automatic analysis of all 2013 photos. Mark Greenwood from partner USFD generated an additional list of 20 frequently mentioned concepts in the 2013 Festival Study interviews, such as Tron, Castle, Statue, Performer, Funny, and Comedy.

For Time 2, a month later, each participant received a custom account on a preliminary version of DFKI's PIMO5 system, which was pre-populated with concepts from participants' group and event descriptions, and the 30 photos that had been the focus of more detailed annotations at Time 1.

Using PIMO5, participants were asked to make keep/delete decisions and to write annotations in the text fields. In order to motivate their annotations and keep/delete decisions,

we said that we would use the data to provide them with a PDF scrapbook of their photos, complete with annotations.

Meaningfulness and ease and speed of retrieval judgements were collected using pen and paper. For each photo, participants were also provided with the top 12 Twitter hashtags from their day at the Festival and the top 6 CERTH categories, and asked to tick those hashtags and categories they would want to use to describe their own photo in their own notes. Appendix A provides a sample feedback form for the first two photos of Participant 1.

The following sections provide preliminary results from this study. Due to the volume of data, we have focused on those analyses that are most relevant for ForgetIT applications, namely the value of automatic categorisations (Section 2.2.3), and preservation decisions (Section 2.2.4). The analysis of the consistency with which photos were assigned to event and non-event related groups is presented in Section 2.2.2.

2.2.1 Demographics

A total of 22 participants took part at Time 1, and all but one ($n = 21$) returned at Time 2. 15 (68%) of the original 22 participants were female, 7 (32%) were male. The mean age was 25 years (SD: 5, range: 19–41); there were no significant age differences between male and female participants (Asymptotic Wilcoxon Mann-Whitney Rank Sum test, $Z = -0.3575, p < 0.73$). 7 (32%) of all participants were not native speakers of English; all of those non-native speakers were female.

Participants returned with an average of 50 photos (SD: 15, range: 28–80; since three participants took over 80 photos, which meant that they were presented with a sample of 80 photos from their total).

We have grouping and keep/delete data for a total of 1097 photos from 22 participants at Time 1, directly after their day at the Festival, and grouping data from 1052 photos from Time 2, a month later (excluding 45 photos contributed by the participant who did not return), and keep/delete data from 628 Time 2 photos, as the keep/delete decisions were only made for the subset of 30 photos that were reviewed using PIMO5.

For 658 photos, we collected more detailed judgements, tags, and annotations at Time 1, and these data were collected again for 628 photos at Time 2.³

2.2.2 Consistency of Events and Groups

Overall, participants created a similar number of event and non-event groupings both immediately after the day itself and after a month. At time 1, our 22 participants divided their photos into a total of 118 distinct events; at time 2, 21 participants created 109

³One participant contributed only a total of 28 photos, which means we did not have sufficient data to draw a sample of 30.

events. This translates into a median of reported 5.5 events per participant at Time 1, (IQR: 3.25-7.75, range: 2–8), and a median of 5 events at Time 2 (IQR: 3–6, range: 2–11). The difference in number of events between Time 1 and Time 2 is not significant (Rank Sum test, $Z = 0.527$, $p < 0.6$).⁴The total number of non-event categories was the same each time, namely 82. At time 1, participants created a median of 4 (IQR: 2-5, range: 2-8) groups, at time 2, the median number of groups is 3 (IQR: 2-5, range: 2–7). Again, the difference in medians is not significant (Rank Sum test, $Z = -0.2484$, $p < 0.81$).

When looking at the number of photos assigned to each category-related photo group and each event-based photo group, we also find very little change. For Time 1, the average number of photos per event for each participant is 11.4 (SD: 5.75, range: 4–23.5), and for Time 2, it is 11.9 (SD: 7.5, range: 3.2–23.5). The mean size of category-related photo groups is 16 at Time 1 (SD: 7.5, range: 4.4–27), and 15 at Time 2 (SD: 6.5, range: 5.8–26.7). Neither of the differences is significant (events: $p < 0.75$, categories: $p < 0.68$, Rank Sum test as before).

Photos that were used as landmarks for the whole day, an event, or a non-event category at Time 1 are far more likely to be chosen again at Time 2 (day: odds ratio [Odds Ratio (OR)]=65.8, category: OR=7.6, event: OR 17.6), but the actual levels of agreement are low: 44% for day landmarks (whole collection), 34% for event landmarks, and 19% for category landmarks.

Implications for Practice:

Participants were consistent with regard to the number of events and categories they used to sort their photos, and with regard to the number of photos per event / category.

For event-related photo groups, this is not surprising, because participants were encouraged to attend shows and visit locations at the Edinburgh Festival. Therefore, the day will naturally fall into a sequence of events structured around major activities and shows. If a ForgetIT system has access to information about the event (time and location) when a photo was taken, this could be a reliable basis for summarisation.

At first glance, non-event categories also appear relatively stable, but for both event and category-related photo groups, we need to conduct further formal analyses into the degree of overlap before we can confirm this assumption.

When it comes to selecting landmarks to represent an event or a category, there are no right and wrong answers. Rather, landmarks tend to be chosen from meaningful photos that are likely to be kept.

⁴This particular test is one of a set of non-parametric tests implemented in the R package `coin`.

Table 3: Most frequently used tags and descriptors

Rank	Text / Tweets				CERTH Picture Descriptors				
	Time 1		Time 2		Time 1		Time 2		
	Tag	Freq.	Tag	Freq.	Tag	Freq.	Tag	Freq.	Freq.
1	Edinburgh	75.5%	edinburgh	53.5%	Building	22.5%	Outdoor	18.9%	
2	Festival	66.6%	edfringe	43.5%	Crowd	18.2%	Person	13.7%	
3	Fringe	64.3%	edfringe2014	43.3%	Group of people	18.1%	Daytime	8.9%	
4	Atmosphere	30.7%	scotland	26.6%	Person	16.7%	Outdoor	6.5%	
5	Royal Mile	29.6%	fringe	24.0%	Cityscape	15.2%	Building	4.6%	
6	Performer	22.9%	edinburghfringe	21.5%	Actor	10.5%	Sky	3.7%	
7	Building / Architecture	20.2%	edinburghfestival	12.4%	Road	10.5%	Urban Scenes	2.9%	
8	Crowd	14.9%	unbored	9.7%	Musician	9.7%	Text	2.4%	
9	Weather	13.5%	fringefest2014	7.3%	Doorway	9.4%	Indoor	2.1%	
10	Show	11.6%	comedy	0.80%	Sky	7.9%	Trees	1.6%	
							Apartments		

2.2.3 Automatic Contextualisation

Here, we focus on the suitability of using automatically assigned context descriptions for photos. At both times, participants were presented with tags that were based on the concepts yielded by the CERTH image analysis (Picture Descriptors), tags derived from common concepts in the 2013 episodic recall interviews, and tags derived from Twitter hash tags (Content Tags). As described above, at Time 1, these tags were generic (i.e., derived over a whole set of photos and the whole of the 2013 Festival), and at Time 2, these tags were specific (i.e., based on the concepts used to describe the specific photo and on the most frequently used hashtags in Edinburgh that day).

Participants were asked to imagine that the tags had been generated by photo classification software, and asked which tags they wanted to keep as useful descriptors.

At Time 1, the overall median number of tags (both content tags and picture descriptors) that participants assigned per photo was 6 (Inter-Quartile Range (IQR): 4–8, range: 0–16); at Time 2, this number halved to 3 (IQR: 2-4, range: 0–11).

Most of this difference is due to the content tags. At Time 1, participants selected a median of 4 tags (IQR: 3-5, range: 0-12), at Time 2, this was reduced to a median of 2 (IQR: 1-4, range: 0-8). The 10 most popular tags at each time are listed in Table 3. All of the top 10 tags that our participants used at Time 2 referred to the Festival itself; the highest-placed political tag came in at number 11 (indyref, 0.6%). Table 4 reproduces the most frequent Twitter hashtags used during the time of the 2014 Festival study. In contrast, the content tags derived from previous interviews, which were used at Time 1, also covered locations, photo content, and event indicators.

Participants used fewer of the picture descriptors at Time 2, even though this time, the

Table 4: Top Twitter Tags During 2014 Festival Study (Table 3.3, D2.3)

Type	Total	Hashtags
Festival	4197	edfringe (2911); unbored (389); edfringe2014 (248); fringe (244); edinburghfringe (142); edinburghfestival (121); fringe2014 (78); comedy (64)
Location	1063	edinburgh (758); scotland (205)
Politics	1013	indyref (389); bbcindyref (232); voteeyes (157); patronisingbtlady (156); yes (79)
Promotion	279	sbutd (114); bamforxmas (101); udobiz (64)
Event	172	dginalba (102); travestiplay (70)

descriptors were photo-specific. However, we only presented participants with the 6 concepts that had been assigned the highest scores by the CERTH image analysis software (as documented in D4.2) for that specific photo. At Time 1, participants used a median of 1 descriptor (IQR: 1-3, range: 0-8). At Time 2, a month later, the median number of descriptors was again 1, but participants rarely selected more than that (IQR: 0-1, range: 0-5). The top 5 descriptors chosen at both times are listed in Table 3.

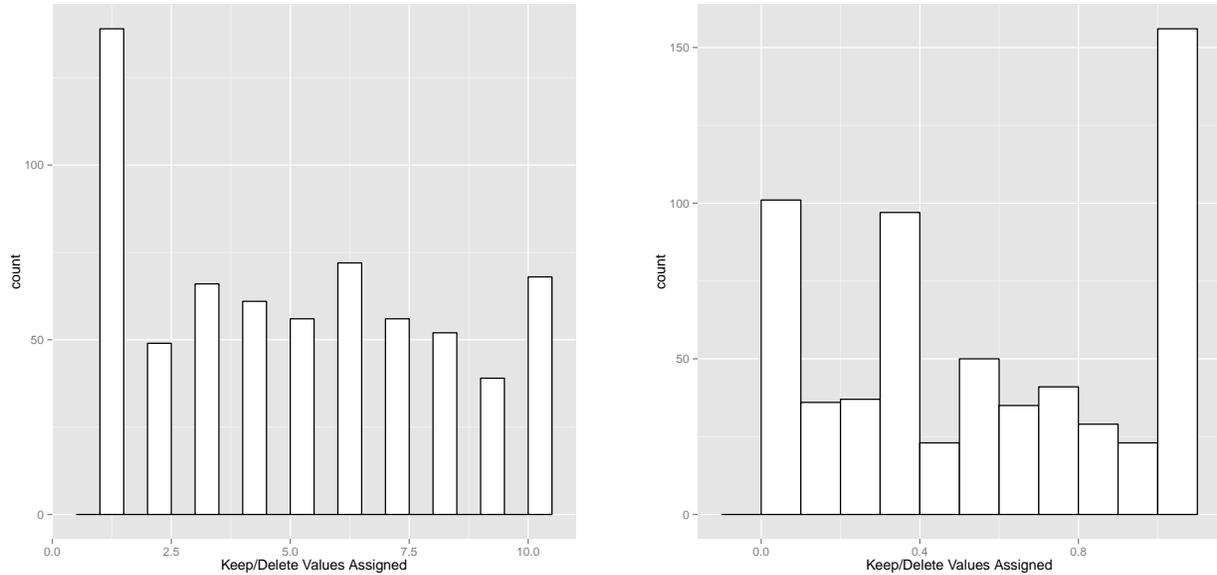
Implications for Practice:

The most striking initial result appears to be that participants used far fewer tags at Time 2 than at Time 1. However, this might well be an artefact of the study design, as participants added comments and descriptions to their photos using PIMO5 before they chose the tags.

Looking at Time 2, concepts that describe the content of a photo tend to be more useful than concepts which are mined from the tweets of that particular day. The concepts are specific to what participants chose to depict, while Twitter hashtags reflect what happened on Twitter that day, which is often multifaceted and encompasses many aspects, such as politics, Twitter-centric conversations, or promotions, that are not relevant to participants' experiences. What dominates Twitter will not necessarily dominate personal experience. This could also be the reason while participants used concepts derived from actual event descriptions (presented at Time 1) far more frequently than Twitter hashtags (presented at Time 2).

2.2.4 Keep and Delete Decisions

The patterns of keep/delete judgements at Time 1 and Time 2 were quite different. At Time 1, participants provided judgements for all photos (up to 80) using pen and paper, and they saw their photos in a folder on a computer screen. At Time 2, on the other hand, participants saw each photo in the PIMO5 interface, together with a box for annotations, an option for marking the photo as a favourite, and a scroll bar for making the keep/delete



(a) Time 1 (pen and paper; 1 = definitely delete; 10 = definitely keep).

(b) Time 2 (PIMO5; 0 = definitely keep; 1 = definitely delete).

Figure 8: Histogram of Keep/Delete Decisions for the 30 Core Photos at Time 1 and Time 2.

decisions.

For the rest of this section, we will focus on the 30 photos that participants used at both Time 1 and Time 2, in order to highlight differences and commonalities between both time points.

As the histograms in Figure 8 show, at Time 1, the only clear peak is around 1, complete deletion, and there is no corresponding peak that could be associated with a clear “keep” decision, except maybe at 10. At Time 2, on the other hand, three clear categories emerge from the data: deletion (1), keeping (1), and a third, intermediate category (around 0.7), which inspired the “hide” option in the final version of PIMO5 documented in D9.5.

At both times, most participants used the full scale for their ratings. At Time 1, 4 participants (3, 10, 12, and 16) only used 1-9, participant 2 restricted themselves to 2-9, and participant 7, who had a strong deletion bias (median rating: 1, definitely delete), used 1-7. At Time 2, there were only three exceptions, namely participants 3, 12, and 15, whose lowest ratings were 0.3, 0.1, and 0.25 respectively.

These data suggest an important difference between Time 1 and Time 2. When participants made their keep/delete decisions while engaging more deeply with a smaller number of photos (30), they were less likely to delete, whereas previously, when they were asked to judge 40–80 photos, they were less likely to keep.

Despite this change, there is a strong correlation between judgements at Time 1 and Time

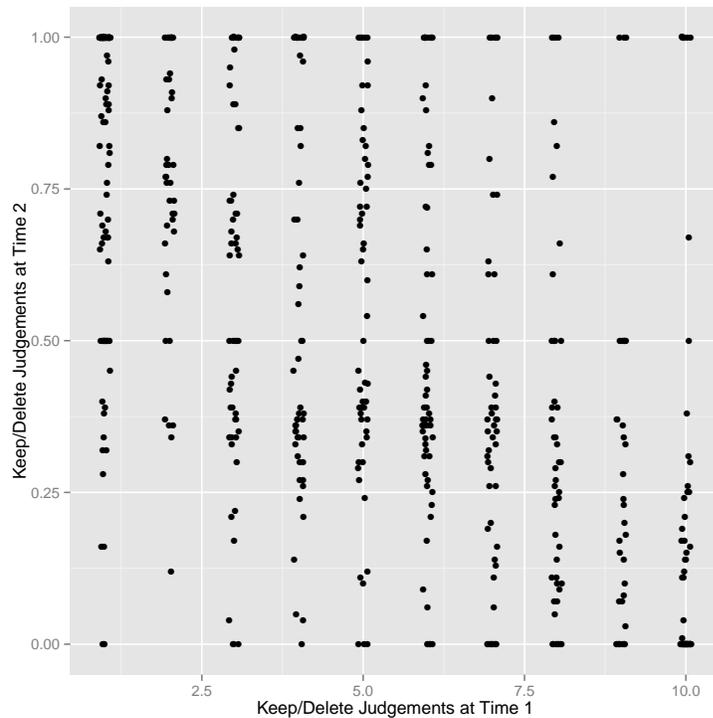


Figure 9: Correlation Between Keep/Delete Decisions for 30 Core Photos at Time 1 and 2. Scatterplot uses jitter to ensure separation of data points.

2 (Spearman's $\rho = -0.584$, $p < 0.0001$, Asymptotic Spearman Correlation Test stratified by participant), as illustrated by the scatter plot in Figure 9.

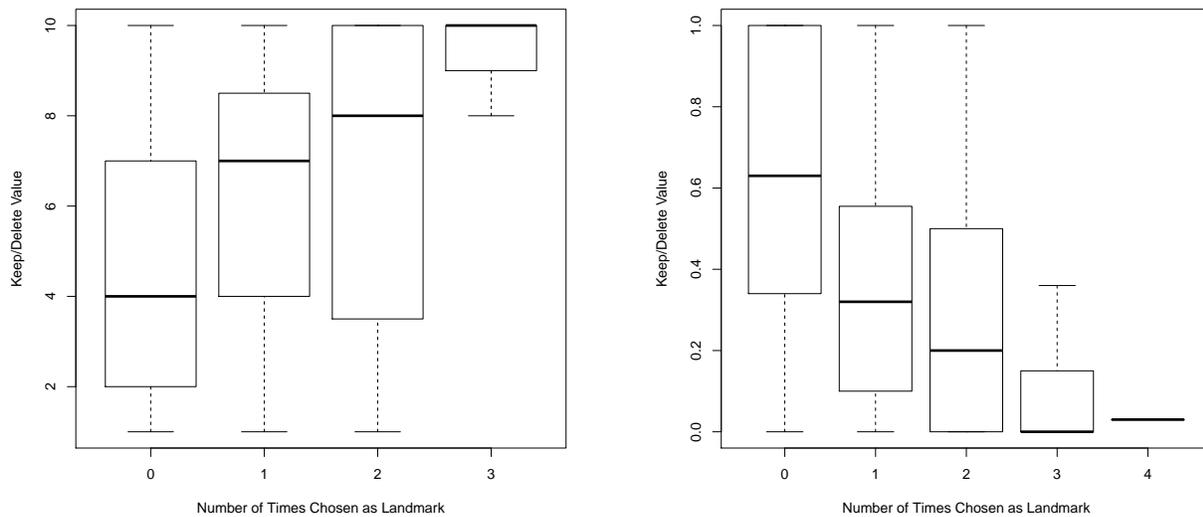
Both at Time 1 and Time 2, photos that were selected as landmarks were more likely to be kept, as Figure 10 shows. There appear to be two important thresholds. Choosing a photo as a landmark once (for a collection, event, or a non-event category) makes people more likely to keep it than to delete it, while photos that are landmarks for an event, a non-event category, and the entire collection ($n=3$) will definitely be kept.⁵

Creating a statistical model of the keep/delete decisions is quite complex, as the target distributions are bimodal (Time 1) and trimodal (Time 2). Therefore, the factors that affect keep/delete ratings will not be integrated into a full ANOVA (which assumes a normally distributed outcome variable) or a generalised linear model. Instead, we report relevant correlations. For significance testing, the Asymptotic Spearman Correlation Test stratified by participant was used.⁶

We examined six potential predictors, meaningfulness, ease of finding, speed of retrieval, the total number of tags assigned to a photo, the number of CERTH concept tags assigned, and the number of Twitter-derived tags assigned. While the first three predictors

⁵The photo at Time 2 that was chosen four times as a landmark was a landmark for an event, an event subcategory, a non-event category, and the entire collection.

⁶Confidence intervals are not reported for Spearman rank correlation coefficients, as these cannot be estimated easily.



(a) Time 1 (pen and paper; 1 = definitely delete; 10 = definitely keep).

(b) Time 2 (PIMO5; 0 = definitely keep; 1 = definitely delete).

Figure 10: Keep/Delete Values for Photos Selected as Landmarks at Time 1 and Time 2.

reflect subjective judgements, the last three predictors can be observed objectively. The corresponding correlations are reported in Table 5. All of these correlations are significant at $p < 0.0001$.

Of the subjective predictors, meaningfulness shows the strongest correlation with keep/delete decisions, closely followed by findability and retrievability. Participants also assigned more tags to photos that they were likely to keep. The number of CERTH concepts chosen correlates more highly with preservation decisions than the number of Twitter tags.

Table 5: Correlations between Keep/Delete Judgements and Photo properties. Correlation coefficient: Absolute values of Spearman’s ρ rank correlation

Criterion	Time 1	Time 2
Meaningful	0.7278	0.6937
Should be		
fast to retrieve	0.6187	0.6293
easy to find	0.6167	0.6512
Number of tags		
all	0.4990	0.5779
Twitter	0.2244	0.1806
CERTH Concepts	0.5461	0.6171

Implications for Practice:

The data indicate that strict binary keep/delete judgements might not be realistic. Instead, participants appeared to use three categories, “definitely keep”, “potentially keep”, and “definitely delete”. Photos that should definitely be kept are also far more likely to be chosen as landmarks. This is reminiscent of the way in which Google creates stories from sets of photos, where highlights of a photo collection associated with a certain time and location are assembled into a slide show.

The bias towards keeping a photo at Time 2 suggests that when participants had the time to engage more deeply with a smaller number of photos, they were less likely to delete them. If this finding is confirmed in more formal studies, it suggests innovative ways for eliciting data on personal preservation preferences, such as editing automatically assembled slide shows, or creating scrap books that combine photos with videos, comments, and other documents.

2.3 Conclusion

In the 2014 Festival study, we moved from a more controlled to a less controlled study design. Our design is unrealistic in that it forces participants to engage far more deeply and thoroughly with their digital photographs than they normally would. As our results on keep/delete decisions show, this may affect results. In particular, we do not have a good model for people who tend to point, shoot, and then forget about their photos.

Nevertheless, some of our findings will hold more generally. For example, preservation decisions tend to be probabilistic rather than absolute, and cluster around three intuitive categories, which can be mapped to more fine-grained labelling systems. The favourite/show/hide/delete categories developed by DFKI split the “definitely keep” category into favourite and show, and the “maybe keep” category into show and hide. The levels of preservation proposed by partner L3S provide a more fine-grained view of the “definitely keep” category.

3 Survey

The aim of the ForgetIT survey was to establish an overview of existing personal preservation practices using the example of photo collections. The survey complements the objective, experimental data collected in the two Festival studies (Sections 2.1 and 2.2). It was conducted online in order to allow the recruitment of large numbers of participants around the world. A PDF copy of the survey is available in the supplementary web documentation.

Instead of reaching out to a representative cross sectional sample through a market research organisation, we focused on recruiting people with an interest in photo preservation and ForgetIT technology, who would be potential early adopters.

The structure and development process of the survey was covered in D2.2 and D2.3. Survey questions focused on finding out what respondents do (or rather, what they think they do) when taking, managing, and preserving photos. Questions were designed to cover a range of common practices, hardware, and software options. In the following discussion of the survey, we will refer to these three sets of activities as *photowork*.

We also assessed two potentially relevant user characteristics, attitude to technology and privacy concerns. Privacy concerns were assessed using a well-validated questionnaire designed by Buchanan et al. [Buchanan et al., 2007]. Attitude to technology was assessed using a revised version of Schulenberg and Meltons Computer Attitudes, Aversion, and Familiarity Index (CAAFI, [Schulenberg and Melton, 2008] 2008). The original 30 questions of the CAAFI were reduced to 8, and these were updated to reflect changes in technology since 2008.

Demographic data included country of origin, native language, age group, occupation (employed full-time, employed part-time, retired, homemaker, student, other), and years of full-time education. The demographic items and the privacy and attitude to technology scales were not mandatory.

3.1 The Final Sample

We collected a total of 1383 complete responses from over 20 countries in Europe, North America, and Asia. Participants were recruited through crowdsourcing, targeted calls for participants in the UK, China, and other European countries, and the ForgetIT web site. The number of returned surveys is given in Table 6.

These responses were then filtered using a simple heuristic to exclude people who did not appear to answer honestly. The heuristic relies on 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, participants had to answer on a four-item scale that deliberately excluded a central neutral option. For each question, the state-

Table 6: Survey Returns by Source

Source	Total	Included	
		n	Percent
<i>Link From ForgetIT Web Site</i>			
Direct Web	515	454	88.2
<i>Crowdsourcing</i>			
Amazon Mechanical Turk	48	44	91.7
Crowdfunder	396	336	84.8
<i>Country-Specific Recruitment</i>			
Europe	198	183	92.4
UK	150	138	92.0
China	42	36	85.7
Other	34	30	88.2
Total	1383	1221	88.3

ments were presented in an array format, which makes it tempting for respondents to tick the same column for each statement and move on. Next, we identified respondents who ticked the same option every time (one-column answers). This led us to exclude data from 162 participants (11.7% of the total). The remaining number of participants in the final sample is 1221.

The demographics of the final sample are given in Table 7. Most of our participants are aged between 18 and 34, are working (employed full- or part-time), and have 13-18 years of education, which is equivalent to an undergraduate degree or a vocational qualification. We have an equal number of male and female respondents, but they differ in terms of education, occupation, and age. Male participants are more likely to have completed 19+ years of education, to be older, and to be working ($p < 0.0001$, χ^2 test).

The five most common native languages were English ($n=484$, 39.6%), Swedish ($n=177$, 14.5%), German ($n=80$, 6.6%), Turkish ($n=50$, 4.1%), and Italian ($n=48$, 3.9%). 112 participants (9.2%) spoke other European languages, 109 (8.9%) spoke other non-European languages, and 161 participants (13.2%) did not provide any data.

The five best represented countries were the UK ($n=199$, 16.3%), Sweden ($n=183$, 15.0%), the United States ($n=153$, 12.53%), Germany ($n=78$, 6.4%), and Canada ($n=60$, 4.9%). 210 respondents (17.2%) provided no or incorrect responses, 171 (14.0%) came from other European countries, including Italy, the Czech Republic, and Greece, and 167 (13.7%) were from other non-European countries, including China, Japan, and Taiwan.

Table 7 shows the demographic distribution for the UK, Sweden, and the US. It is clear that the demographics were heavily skewed by recruitment method. There were two big recruitment drives among UK and Swedish students, which leads to the recruitment of more younger people, while most of the US sample comes from Crowdsourcing platforms (Crowdfunder and Amazon Mechanical Turk), which skew towards a middle-aged population. Therefore, for the comparative analyses reported below, we will use matched

Table 7: Demographics of the Final Survey Sample

Category	All		UK		Sweden		US		
	n	Percent	n	Percent	n	Percent	n	Percent	
<i>Age</i>									
18-24	364	29.8	80	40.2	43	23.5	17	11.11	
25-34	350	28.7	17	8.5	56	30.6	41	26.80	
35-44	224	18.4	18	9.0	33	18.0	43	28.10	
45-54	140	11.5	24	12.1	30	16.4	30	19.61	
55-64	89	7.3	31	15.6	16	8.7	13	8.50	
65+	46	3.8	28	14.1	5	2.73	9	5.88	
not stated	8	0.66	1	0.5	0	0.00	0	0.00	
<i>Gender</i>									
Female	607	49.71	132	66.33	83	45.4	93	60.8	
Male	602	49.30	66	33.17	100	54.6	60	39.2	
not stated	10	0.82	0	0.0	0	0.0	0	0.0	
<i>Occupation</i>									
working	617	50.5	73	36.7	88	48.1	95	62.09	
student	410	33.6	72	36.2	79	43.2	42	27.5	
other	159	13.0	51	25.6	13	7.1	14	9.15	
not stated	14	1.1	1	0.5	0	0.0	2	1.31	
<i>Years of Education</i>									
up to 13	165	13.5	28	14.1	28	15.3	25	16.3	
13-18	742	60.8	132	66.3	121	66.1	85	55.6	
19+	299	24.5	38	19.1	33	18.033	43	28.1	
not stated	14	1.1	1	0.5	0	0.0	0	0.0	

samples of the complete survey data set.

Overall, our participants were highly privacy conscious (cf. Table 8). When comparing our youngest group (18-24, digital natives) to participants aged 55 and older, we see that both are equally concerned about privacy in general, sharing too much personal information, online identity theft, and data mining. The older group is generally less concerned about theft, sharing, or being tagged in photos.

Mokken analysis revealed three subscales in the attitudes to technology items discussed in D2.3, a technology acceptance scale (Cronbach's $\alpha=0.78$, good reliability), a technology avoidance scale (Cronbach's $\alpha=0.58$, moderate reliability), and a technology use scale (Cronbach's $\alpha=0.30$, poor reliability). Table 9 shows the association between items and scales as well as basic statistics about the scores.

Overall, our sample has an overwhelmingly positive attitude towards technology, which fits in well with our intention to reach potential early adopters. For both the technology acceptance and the technology avoidance scale, we also computed mean average scores for each participant. The mean technology avoidance score of our participants is 2.5 (SD:

Table 8: Concern with privacy threats: All respondents, young respondents, older respondents.

Text	All	18-24	55+
Your privacy while using the Internet.	891 (73.2%)	269 (74.1%)	99 (73.9%)
Being asked for too much personal information when you use an online service.	960 (78.8%)	291 (80.2%)	110 (82.1%)
Online identity theft	958 (78.7%)	286 (78.8%)	109 (81.3%)
Other people stealing your digital photos	687 (56.4%)	224 (61.7%)	53 (39.6%)
People you do not know obtaining information about you from your online activities	906 (74.4%)	287 (79.1%)	99 (73.9%)
Being tagged or marked in a digital photo without your consent	672 (55.2%)	206 (56.7%)	63 (47.0%)
People you do not know seeing digital photos that you have taken without your consent	684 (56.2%)	219 (60.3%)	54 (40.3%)
Other people sharing digital photos of you without your consent	751 (61.6%)	243 (66.9%)	66 (49.3%)
Total	1218	363	134

1.5, median: 2, IQR: 1–3.5, range: 1–7), while the mean technology acceptance score is 5.6 (SD: 1.1, median: 5.75, IQR: 5–6.5, range: 1–7).

Almost all respondents (1173, 96.1%) reported taking digital photos themselves. In addition, 640 (52.4%) receive photos by email, 463 (37.9%) download them from the web, and 337 (27.6%) receive collections of digital photos from others. 41 (3.4%) of respondents mentioned additional methods of obtaining photos such as via social media, file sharing through Dropbox or Google Drive, or WhatsApp.

Most of our respondents took digital photos using a smartphone (89.7%) or a digital camera (89.3%). Just under half used a tablet, and one in four took photos with a simple mobile phone with a camera. As Table 10 shows, smartphones are the device of choice for frequent photo taking. All other devices tend to be used a couple of times a month or less.

925 (75.8%) of participants found it important or very important to preserve photos for coming generations, while 789 (64.6%) were worried or very worried about losing photos. This further confirms that our survey has reached many potential early adopters of ForgetIT technology.

3.2 Personas for Preservation

As reported in [Wolters et al., 2015] for a subset of the full data set, we examined whether there were characteristic patterns in the way respondents organised their photos and

Table 9: Descriptive Statistics of the Technology Avoidance, Acceptance, and Use Scales. 1 = absolutely false for me now; 7 = absolutely true for me now

Text	Scale	Median	IQR	Range
I enjoy using technology	Acceptance	6	6–7	1–7
Being able to use technology is important to me.	Acceptance	7	6–7	1–7
I keep up with the latest technology.	Acceptance	6	5–6	1–7
I avoid using technology whenever possible.	Avoidance	1	1–3	1–7
I enjoy reading about technology.	Acceptance	5	4–6	1–7
Email is an easy way to communicate.	Use	7	6–7	1–7
I often visit web sites.	Use	7	6–7	1–7
I hate learning to use new software and applications.	Avoidance	2	2–5	1–7

Table 10: Devices Used for Taking Digital Photos. Percentage Owned/Used refers to full cohort; frequency of taking digital photos based on number of people who own/use a device.

Device	Owned / Used		Frequency of Taking Digital Photos					
	N	%	Daily		Weekly		Monthly or Less	
			N	%	N	%	N	%
Smartphone	1095	89.7%	314	28.7%	424	38.7%	309	28.2%
Camera Phone	282	23.1%	16	5.7%	27	9.6%	191	67.7%
Digital Camera	1090	89.3%	53	4.9%	190	17.4%	799	73.3%
Tablet	603	49.4%	20	3.3%	64	10.6%	471	78.1%

Table 11: Photo Management Practices. The answer options were very accurate, accurate, somewhat inaccurate, and very inaccurate.. Difference between Clusters tested using χ^2 test. *:p<0.01, **:p<0.001, *:p<0.0001**

Item	Text	Very or Somewhat Accurate					
		All N	All Percent	File N	File Percent	Curate N	Curate Percent
Delete	I delete most of the digital photos I take.*	367	30.06%	207	31.55%	160	28.32%
Org	I keep all my digital photos meticulously organised.***	641	52.50%	293	44.66%	348	61.59%
Auto	I add as much information as possible to my photos automatically.***	344	28.17%	113	17.23%	231	40.88%
Keywords	I add keywords and titles to my photos.***	374	30.63%	75	11.43%	299	52.92%
ManLabel	I label people, locations, and objects in photos.***	329	26.95%	80	12.20%	249	44.07%
Filenames	I use file and folder names to find digital photos.***	949	77.72%	475	72.41%	474	83.89%

preserved them. In order to determine patterns of **photo organisation**, we clustered responses to all items included in the question “How Do You Manage Your Own Photos”, listed in Table 11. For patterns of photo preservation, we used all items included in the question “How Do You Archive Your Own Digital Photo Collection” (c.f. Table 12).

The cluster analysis algorithm used was agglomerative hierarchical clustering (agnes, [Kaufman and Rousseeuw, 1990]). Starting with single data points, this method merges the two most similar data points or clusters until all data points are subsumed within a single main cluster. There are several methods for determining the distance between clusters. Here, we assume that the distance between two clusters is the maximum distance between any item i_1 in Cluster 1 and any item i_2 in Cluster 2. This favours relatively compact clusters. Other metrics produce unbalanced solutions, where one of the final 2 clusters only contains less than 10 data points.

Table 11 shows the clusters that were found in the data on photo management practices. Curators are far more likely to organise their photos carefully, use keywords, and assign manual labels than Filers, who tend to store their photos in folders without additional information.

The two clusters that emerged for the photo preservation items are summarised in Table 12. Those who focus on Safe Storage are more likely to follow archiving recommendations such as checking that old photos are still readable and storing photos in multiple places. They are also more likely to make manual backups. Members of the Basic Storage group, on the other hand, are far less concerned with preserving their data. Instead, just under half state that they file their photos carefully and securely, and use multiple locations. These two clusters have implications for potential preservation strategies, as

Table 12: Photo Preservation Practices. The answer options were very accurate, accurate, somewhat inaccurate, and very inaccurate. All differences significant at $p < 0.0001$

Item	Text		Very or Somewhat Accurate				
			All	Safe Storage		Basic Storage	
Location	I keep copies of my photos in multiple places	739	60.52%	451	71.70%	288	48.65%
File	I file my digital photos carefully	695	56.92%	412	65.50%	283	47.80%
Copy	I keep multiple copies of my photos	588	48.16%	382	60.73%	206	34.80%
MoveNew	I regularly move archived photos from old to new storage media	430	35.22%	287	45.63%	143	24.16%
Manual	I regularly make manual backups of my digital photos	482	39.48%	326	51.83%	156	26.35%
Auto	I rely on an automated backup system	459	37.59%	265	42.13%	194	32.77%
Sec	I store my digital photos privately and securely	697	57.08%	399	63.43%	298	50.34%
Read	I regularly check whether my old digital photos are still readable.	346	28.34%	244	38.79%	102	17.23%
Scan	I archive printed photos by scanning them in.	307	25.14%	213	33.86%	94	15.88%
Paper	When I print photos, I choose high-quality paper.	673	55.12%	391	62.16%	282	47.64%
People	I give copies of important photos to other people for safekeeping.	256	20.97%	201	31.96%	55	9.29%

Table 13: Preservation Priorities of Four Personas Defined By Clusters.

Group	Worried About Photo Loss	Preservation Important
Basic Curator	65.1% (138)	76.4% (162)
File and Forget	58.2% (221)	65.3% (248)
Safe Curator	68.8% (243)	83.9% (296)
Safe Filer	67.8% (187)	79.3% (219)
All	64.6% (789)	75.8% (925)

discussed in WP9, and they were the basis for modifying preservation value calculations within the final implementation of PIMO5.

A person's photo preservation strategy does not predict how they manage their photos. While 353 respondents (28.9%) both curate and value safe storage (**safe curators**), 212 are curators who follow a basic storage strategy (17.4%, **basic curators**). 276 (22.6%) are **safe filers**, while 380 (31.1%) "**file and forget**", opting for a basic storage strategy and not curating their photos carefully. As Table 13 shows, those who file and forget tend to be less concerned about losing digital photographs and are less likely to find photo preservation important, while safe curators tend to value the preservation of digital photos for future generations.

Likewise, as Table 14 shows, both curators and filers tend to notice the same problems when managing their photos. Curators are more likely to complain about lack of information about the people in their photos, but otherwise, patterns are fairly similar. The main problems are around filing. Photos are unfiled, filed in the wrong place, or filed in a place that the person can no longer remember.

This suggests that easy filing and retrieval might be an important side-benefit that could induce people to use a ForgetIT-based preservation solution. Since the ForgetIT system already contextualises all archive items, this context information can be used to facilitate information retrieval. In fact, ease of filing and retrieval of information was one of the main benefits of the DFKI system as highlighted by the PANIC students (cf D9.5).

3.3 Outlook: Preserving Digital Photos Across Cultures and Generations

Our data set provides rich insights into the photo management and preservation practices from over 1200 potential early adopters. The core insight to be derived from these data are the four distinct personas, who will need to be fleshed out in further work. As we saw in D10.1, data-driven personas are an important foundation for further design work. These personas can also serve as a baseline for the design of semi-automated preservation strategies, as demonstrated in D9.5.

Fully describing and analysing the data set would be beyond the scope of this deliverable.

Table 14: Problems with Managing Photos for All Participants and for Personas Derived From Photo Management and Preservation Practices (Difference between Clusters tested using 2 test. *:p<0.01, **:p<0.001, *:p<0.0001)**

Text	Very or Somewhat Frustrating					
	All		Curator		Filer	
	N	%	N	%	N	%
Too many						
... similar photos	542	44.39%	261	46.19%	281	42.84%
... unfiled photos*	735	60.20%	323	57.17%	412	62.80%
It takes a lot of time to find a photo	748	61.26%	336	59.47%	412	62.80%
Lack of Information about						
... people in the photo***	411	33.66%	198	35.04%	213	32.47%
... location or place of photo	451	36.94%	208	36.81%	243	37.04%
... objects in photo	306	25.06%	153	27.08%	153	23.32%
... time photo was taken	450	36.86%	207	36.64%	243	37.04%
I don't know what search criteria to use*	560	45.86%	249	44.07%	311	47.41%
Photos are filed in the wrong place*	642	52.58%	314	55.58%	328	50.00%
I cannot						
... remember where I filed a photo	767	62.82%	365	64.60%	402	61.28%
... search photos efficiently	718	58.80%	320	56.64%	398	60.67%
... filter photos by quality***	411	33.66%	208	36.81%	203	30.95%

In ongoing analysis for publication, we focus on two aspects, a more detailed analysis of the digital photography habits of each of our four persona prototypes, and an investigation into potential cultural and generational factors. Delving into cultural and generational differences will require a more complex statistical analysis in order to take into account the biases introduced through our sampling procedure.

Table 15: The general human remembering function adapted from Ebbinghaus 1913. P.E.m stands for probable error of the mean and summarises the expected deviation of means from repeated observations.

No.	I After X Hours	II So much of the series learned was retained that in relearning a saving of $Q\%$ of the time of the original learning was made	III P.E.m	IV The amount forgotten was thus equivalent to $v\%$ of the original in terms of time of learning
	$X=$	$Q=$		$v=$
1	0.33	58.2	1	41.8
2	1.	44.2	1	55.8
3	8.8	35.8	1	64.2
4	24.	33.7	1.2	66.3
5	48.	27.8	1.4	72.2
6	6×24	25.4	1.3	74.6
7	31×24	21.1	0.8	78.9

4 A Conceptual Framework for Organisational as well as Societal Remembering and Forgetting

This report covers work undertaken by the University of Oxford together with other ForgetIT partners to develop conceptual frameworks for organisational as well as societal remembering and forgetting, and take steps to have these conceptual frameworks operationalized and implemented into digital tools developed as part of the ForgetIT project. The primary aim here was to focus on how organisations might record, monitor and review ‘corporate’ remembering and forgetting on a routine basis rather than as part of the detailed research focus on evaluating the initial conceptual framework for organisational memory within the ForgetIT project. This latter is described in Section 5. Consequently, the report has two main parts: (a) the development of conceptual frameworks, and (b) suggestions and strategies of how these frameworks can be operationalized within digital tools.

4.1 Development of Conceptual Frameworks

While human remembering and forgetting is a cognitive process inside an individual’s brain and thus not directly visible and recordable, the basic principle of human memory that forgetting most generally correlates with the passing of time has long been identified. Over the last five decades, researchers have added much detail to the so-called human remembering function, so that today for instance we know that forgetting is steepest at the beginning, while accurate recall improves memory retention.

Overall, individual remembering and forgetting is following a general human remembering

function (see Table 15), a complex mechanism that relies on numerous inputs to make decisions about what to retain and what to get rid of. Detailed reviews of different aspects of the relevant literature on human memory are provided in D2.1, D2.2 and D2.3. This function is the 'baseline', the standard of how individuals remember and forget. And thus, if one wants 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 is a system that remembers and forgets like an average human being.

Moreover, the subject of the study of human memory is relatively clearly defined: it is the human brain, and how it functions when committing information to memory, and when querying and recalling memory.

While certainly not perfect, this human mechanism of remembering and forgetting arguably is the result of tens of thousands of years of evolution entailing improvements and refinements. Thus, incorporating core elements of how this human memory functions into the ForgetIT system would not only enable ForgetIT to mimic 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.

The situation, however, is very different when it comes to organisational as well as societal remembering and forgetting, not the least because the study of memory beyond the confines of the human brain is a comparatively recent phenomenon. A detailed review of the research literature on organisational memory is provided in Section 5. Here, there is an overview of the key issues that have arisen from that research and that are relevant for the UOXF contribution.

4.1.1 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 its workforce, including design engineers and experienced workers on the assembly line. The result was a steep decline both of innovation in aircraft design as well as 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 is 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 keep within the organisation the individuals 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 low-paid, low-skill jobs, such as those for instance offered in the fast food industry.

In these and many other situations, research has pointed towards capturing and infusing 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 is measured in a handful of months, and thus the organisation aims to embed as much as possible of its expertise in processes rather than employees.

Putting these and similar approaches into a wider and more generalizable context, we suggest – loosely following Ferraris [Ferraris, 2010] – that three distinct, but often intertwined strategies are at play in organisations that emphasize the creation and maintenance of memory external to the individuals working for them:

- Documentality
- 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 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 individuals' memory [Gawande, 2010]. So-called standard operating procedures (SOPs) similarly aim to encapsulate an organisation's memories of best practices in mandatory processes, laid out in handbooks and manuals (and crossing over, to an extent, into documentality).

Physicality finally signifies that the physical structures utilized by an organisation enable some and constrain other behaviour, thereby (ideally at least) representing and reflecting the memory and expertise of the organisation. Fast food outlets use physicality, too, by designing food preparation devices that only permit use in 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 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 of organisations to unlearn existing practices (i.e. to forget) has become a focus in contemporary organisational memory research [Easterby-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. This creates a variety of localities of organisational memory, much beyond the memories shared by the individuals within an organisation. To understand an organisation's memory, it is not sufficient to ask individuals to recollect their relevant memories and look for common threads. Rather, researchers also have to examine the written artefacts, processes and even (physical) structures within an organisation in order to capture that organisation's memory. Just such a detailed examination is reported in Section 5 for a small number of targeted organisations.

4.1.2 Existing research in societal remembering and forgetting

At first glance, the very notion of collective memory may strike some as odd. How can human memory, so linked with a person's brain, be anything but individual? And yet, remembering has always had a hugely important social dimension. Thousands of years ago, at a time when external tools to preserve memory (such as language and script) were hardly developed, let alone broadly available, remembering was a deeply social process

of telling and listening to shared and common tales, creating and maintaining something that people had in common - glue binding communities together. Only as remembering turned into a more individual act, something done by one person herself or two people outside of a deeply social and communal process, did the concept of human memory become dominated by the concept of individual memory, competing with, and at times superseding more traditional social memory [Douglas, 1986].

The point here is not to argue that in human history we always moved from social to individual memory as we advanced towards modernity (or, as some seem to prefer, towards 'logocentrism' [Hacking, 1998], but that the contours of human memory are significantly more complex and multi-faceted than the view of equating human memory with individual memory would posit.

Similarly, as the role of time in human culture changed drastically over the centuries, moving from something that is preordained and immutable to what Anderson famously coined 'homogenous, empty time', stressing simultaneity of time across societies and cultures, memory as time past could be redefined and rewritten to fit a particular (political) agenda in the present. It is no coincidence therefore, as Anderson has pointed out, that new regimes are often tempted to put in place new calendars, new structures to organise and view time present as well as time past [Anderson, 2004].

Rewriting a society's memory is an important element in the rise of nation-states. By changing how the past is being documented (through not only new calendars, but also dictionaries purporting to record the roots of a language while actually largely inventing it), by putting in place processes, such as celebrations or remembrance days, of specific past events, and by creating memorials, authorities successfully developed a rich toolkit not only to shape a society's memory, but to bring about collective memory. Much like the communal memory of the old, ritually remembered through singing, telling, practicing in a communal context, this collective memory is anchored (at least partially) outside the brain of any one individual. And it can be created and shaped, written and rewritten, interpreted and reinterpreted by those in power, or to be more precise by those that hold the interpretative authority over such externalized memory.

Consequently, 'collective memory', a term made popular by Maurice Halbwachs in the early 20th century, is at the very least highlighting the fact that all human memory is socially framed. Even as individuals we remember not in a social vacuum; what we remember for instance depends on the social groups we belong to, the social values we subscribe to, and the social practices we adhere to [Halbwachs and Coser, 1992]. Who thus shapes the social context influences human memory.

This may sound obvious, but it does offer the powerful within society a process, if not also a reason to deliberately shape (and reshape) human memory. World War I was initially remembered through solemn memorials highlighting the tragic experience of human suffering. But these very memorials were later reinterpreted as symbols of national heroism, facilitating and supporting the rise of authoritarian regimes.

At first one may think of social memory as something durable aimed at transcending time

– especially relative to individual memory. But in fact, perhaps one of the most remarkable features of social memory is its malleability.

Some have argued for an even stronger collective approach to social memory, arguing not only for individual memory to be socially framed, but to be inextricably connected with, perhaps even inseparably founded on collective memory.

The history of collective memory studies is deeply linked with our attempts to come to grips with the horror of the Shoah. After that important phase collective memory as a field it has been suggested got out of fashion, only to be revived by a combination of tragic world events in the 1990s (Rwanda, Srebrenica et al) and the rise of digital tools facilitating the massive creation and sharing of external memory. But this narrative of the field may itself be a good illustration of the social dimension of memory, and its reconstruction from the vantage point of the present.

Irrespective of the quibbles over the master narrative, collective memory studies have led to countless useful and valuable contributions in the literature examining the conditions of successful creation or reinterpretation of collective memory, and what processes and structures this necessitated and/or gave rise to. Expectedly given the variety of social contexts, this has not led to a simple ‘how to’ guidebook for how to create or shape collective memory. But it has led to at least the beginning of an understanding of what social elements are at play in these dynamics. And arguably most crucially, collective memory studies have highlighted the importance of the social dimension for human memory.

4.1.3 Assessing Similarities - From Frameworks to Framework

Other research that has been undertaken into organisational and societal remembering is described in Section 5. But the elements highlighted above already identify a number of key themes and threads that point towards similar contours in both organisational and societal memory studies.

These are:

- the (likely) variability of the memory function
- the importance of remembering processes and structures
- the malleability of organisational and societal memory
- the normative quality of non-individual memory

Humans may remember different things, and with varying duration and accuracy. But these differences are understood as the results of variations in the inputs during the process of storage (both committing to and keeping in storage) and remembering. The underlying fundamental memory function of how human brains remember and forget is likely universal, and itself the result of thousands of generations of evolutionary adaptation. As noted in D2.1, D2.2 and D2.3, human memory function has been studied by researchers for over 150 years. Although many questions remain, this extensive research has allowed

sufficient understanding as to what researchers have to build when they want to emulate individual remembering and forgetting through digital tools.

As organisations and societies 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. But organisations and societies are not the product of tens of thousands of years of a natural selection process that led to a convergence of a dominant remembering and forgetting mechanism. In fact, so far no dominant mechanism of organisational or societal remembering and forgetting has been identified. Instead, one can find a tremendous heterogeneity among how organisations and societies remember. Neither of them is per se better (or worse) than the others.

Therefore, the first communality between organisational and societal memory is their conceptual difference from individual memory as they lack a credible and widely accepted 'baseline' memory function. This poses a unique challenge – especially when compared with individual memory – for those aiming to build organisational and/or societal remembering into digital systems: if no common 'baseline' function has yet been identified, what exactly should then be emulated?

To be clear, we are not suggesting that in fact no 'baseline' memory function common to all organisations or all societies exists, but that so far we have not been able to identify it. The detailed research in Section 5 comprises an initial attempt to do so. Remembering and forgetting may differ between different organisations and different societies not only in the content of what is remembered and forgotten, but also in the principles that govern such remembering and forgetting. However, this does not preclude the use of a conceptual framework within which to test hypotheses about those principles of remembering and forgetting that could be universal rather than specific to particular organisations.

Organisational and societal memory are not anchored in an individual's brain, but are influenced by the memories of individuals even if they are external to those memories. But what encapsulates external, non-individual memory? With organisational memory, we propose that different ways of remembering may be identified and encapsulated in the categories of documentality, procedurality, and physicality.

Intriguingly perhaps, very similar strategies of remembering can be found in the context of societal remembering. Calendars and dictionaries, well-documented in societal memory studies, aid and facilitate documentality. Communal and societal rituals, from coming-of-age rites to processions and parades on remembrance days are procedures of creating and recreating shared, collective memory about things past. Similarly, monuments and museums are examples for the physicality of where and how societal remembering may take place.

To be sure, there are differences in how these different types of remembering work in the contexts of organisations and societies. As a crude rule of thumb, for instance, it is easier for the leadership of a smaller organisation to influence remembering compared with doing the same for an entire society. But, arguably, there are no fundamental differences between the two. If anything, differences in scale, control and authority exist not

just between organisations and societies but among them.

As a consequence, therefore, studies of different types of such external, non-individual memory may hold valuable insights for understanding organisational as well as societal memory, and as demonstrated in Section 5, salient elements of the remembering processes identified in these studies may be applicable to both contexts.

Equally important is the insight that the memory function of organisations and societies is deeply plastic. We are not suggesting that organisational and societal memory is more (or less) malleable than individual memory, only that it is flexible and open to change. At first, this may sound counter-intuitive; after all organisational and societal memory is specifically designed to be stable and persistent, and thus to overcome some of the fluidity of individual memory. But such a view overlooks the very nature of memory outside the human brain: if anything such external memory is constructed, framed and then enlivened through actual organisational or societal practice. Such memory only attains its place as it is interpreted by humans, and thus reconstructed and reshaped. It is not only that processes, structures, and physicality of organisational and societal memory can be altered, it is that these types of external non-individual memory only work when they are used, and thus appropriated.

This plasticity can be seen as a weakness, as remembering in organisations and societies may be less durable, and more vulnerable to shaping especially by those with authority or in control. On the other hand, such malleability can also be seen as an advantage, as it enables organisations and societies to change how and what they remember over a relatively short period of time.

Capturing how an organisation or a society remembers is an important first step. But one can also foresee a possible second step, highlighting the normative dimension of the undertaking: the deliberate shaping (at least to an extent) of organisational and societal memory through the memory mechanisms that are implemented.

This suggests that after reviewing their distinct memory function, organisations perhaps even societies could articulate a normative aspiration of how they want their external, non-individual remembering to be. This could give digital tools, such as those developed through the ForgetIT project, also a unique competitive advantage in the market space and elevate ForgetIT tools from enabling to replicate memory reality in organisations and societies to a mechanism to shape remembering in such contexts.

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 memory strategy.

These multiple similarities in the underpinnings of a conceptual framework for organisational as well as societal remembering point towards an important supposition: that rather than having a remembering framework each for organisations and societies, it is possible to describe a common framework for (at least partially) external, non-individual

remembering, with key elements being memory function variability, malleability and normativity, as well as the central importance of capturing external remembering processes and structures.

4.1.4 The Conceptual Framework – A Process-Oriented Approach

One cannot simply build into the digital tools of the ForgetIT project a standard organisational or societal memory function. Rather, to recreate an organisation's or society's memory function, one first has to study the organisation's or society's remembering. It is important here to emphasize that the aim is to capture the status quo, the actual treatment of memory within the organisation or society rather than reflecting normative rules. Without knowing how a specific organisation or society actually remembers or forgets, no 'baseline' can be discerned.

Research into individual remembering is often done through experiments: for instance, people are exposed to certain stimuli, later asked what exactly they remember. Recently, brain imaging technologies have augmented this approach. But at the core, research is routinely indirect, capturing not the process of remembering but its results.

In contrast, gathering insights into how an organisation or a society remembers may work differently. The process of remembering is not confined to an individual's brain, but takes place in a social context. It is one in which people interact with reality, and in which external memory is being appropriated by humans. As a result, it is easier to seize the act of remembering rather than just its outcome. This is an important difference to individual memory and its study.

Moreover, because of the deeply social nature of non-individual memory, it is also insufficient to look solely at the external memory artefacts, for instance to study an organisation's standard operating procedures and guidelines, or to describe and study a society's war memorials. The social dimension of external, non-individual memory implies that its essence can only be gleaned through the act of memory being utilized. Standard operating procedures may not be followed the way an objective outsider would understand them, having been reinterpreted and thus 're-remembered'. Similarly, a war memorial's meaning may change even though physically it remains the same (for an even stronger emphasis on the performative dimension see [Connerton, 1989]). This is why it is crucial for the understanding of the remembering of an organisation or a society to capture its actual processes of remembering, and for the conceptual framework to be fundamentally process-oriented.

Some of the aspects to examine and evaluate could be:

- the initial decay of memory obtained by the organisation or society,
- the mechanisms and specificities (such as frequency and depths) of memory being recalled (either in its original or in an 'embedded' form),
- and the tension between an organisation's or society's memory versus the memory

of its individual members.

Ultimately, of course, the goal is to capture digitally as much as possible the interactive, social processes of remembering through information flows and behaviours, so that by observing these over time, the appropriate features and contours of that organisation's or society's memory function could be gleaned (as much as possible). This may work better in some, and worse in other contexts. But it would offer the significant advantage of adaptability – as the tools not only shape what is being done, but also how the members interact, the digitally embedded memory function could lose its technical immutability, thus incorporating a truly 'learning' memory model.

4.2 Operationalizing Social Memory – Taking the Framework to Practice

In accordance with the Description of Work we have worked with other ForgetIT partners, in particular DKD, to operationalize core elements of this conceptual framework and incorporate it into digital tools that are being developed as part of the project. As DKD strives to address organisational remembering needs, the focus of these operationalization efforts has been on organisational rather than also societal remembering.

We defined capturing organisational memory realities as the first important step. In principle, this can be done, as specified in the DOW, through observation, in-depth interviews and comprehensive case studies. This approach has the potential to offer both considerable insight into the organisations that are the targets for detailed study, and a detailed evaluation of the conceptual framework for remembering and forgetting within organisations. The University of Edinburgh is using this approach to help develop the remembering and forgetting framework developed in the ForgetIT project.

This approach has been important and appropriate for investigating and attempting to understand the characteristics of target organisations within the context of the ForgetIT project, in particular to assess whether the framework of Documentality, Procedurality and Physicality fully encompasses corporate memory within the organisations studied. However, such an approach would not be feasible in actual commercial practice for the purposes of regular monitoring and review, as it would require hefty investments 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, absorbing even more resources.

In cooperation with DKD we therefore suggest that for the purposes of regular monitoring and review, organisational practices (and thus organisational preferences reflected in them) that are 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. The aim is for software to 'observe' an organisation's practice of remembering. This obviously works best when an organisation relies on digital tools in daily work.

We developed a concept of organisational preservation based on machine learning concepts, and have mapped out a 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. For instance, DKD is planning to capture behaviours in the use of TYPO3, reflecting a process perspective of organisational memory. At the same time, 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.

An important further step is to understand an organisation's memory aspirations (the 'ought'). To that end, DKD is starting to pursue an initiative that aims to present the 'ought' of organisational memory in a very simplified model of three distinct dimensions of how external, non-individual memory is currently being utilized within an organisation in order to identify goals for organisational memory, and to shape organisational memory practices in the appropriate direction. This approach will not capture the ideal level of granularity or a comprehensive set of relevant data; but what it lacks in sophistication it attains in ease of use and resource-efficiency. It may not be a perfect diagnostic, but if the data collected routinely by these means are found to generate even modest insight that organisations find useful for enhancing their business, then it would be better than having no effective process of monitoring and review regarding how organisational memory is, can be, or ought to be preserved.

The operationalization of the conceptual framework of remembering does not need to be limited to organisations, even if in the context of the ForgetIT software tools developed it is the main focus. Rather, one could envision the conceptual framework being operationalized for societal remembering as well. One could, for instance foresee software tools that by 'observing' the use of digital memory tools in a society capture how certain types of external memory are being appropriated and (re)interpreted. This would offer tantalizing opportunities for collective memory studies: examining at scale a society's collective memory dynamic.

To achieve this ambitious goal, a number of key challenges would have to be met. First, one would have to take stock of the multitude of external memories of different types (denoted above with the labels documentality, procedurality, and physicality). Then, digital systems would have to observe based on the specific qualities of such a memory type, how a particular piece of collective memory is being utilized and appropriated by the people, and most importantly how this appropriation changes across time (and perhaps within societal subgroups).

Initially, it is perhaps likely that rather than analysing collective memory comprehensively, collective memory research may look at certain memory types, or even just certain memory pieces, not only to gain first glimpses into collective memory dynamics at scale, but also to further develop the methodological and technical toolkit necessary.

Longer-term, though, the creation of processes that capture collective memory dynamics, and the resulting repository of societal memory practices may prompt us to rethink the

task and role of societal archives. If in the past they collected artefacts of external memory – documents, images, recordings – in the future, we may see archives gathering and preserving data on the actual use of such external memory. As a result, we may no longer see archives as sources for societal history, but also places of societal memory.

4.3 What the Future may hold

Partner UOXF set out to develop conceptual frameworks for organisational and societal remembering and forgetting. The hope was that this could complement what we know of individual memory, to facilitate comprehensive operationalization in the ForgetIT tools of remembering, from the individual to societies. The quest for these frameworks led us to discover important similarities in the contours of organisational and societal remembering, letting us propose a common framework founded on the key elements of memory function variability, malleability, normativity, and the importance of external processes and structures. Despite the early suggestion that there may be no commonality across organisations, emerging from this foundation is a common conceptual framework that is process-oriented, reflecting the social dimension of remembering.

Its future operationalization through software tools may offer a much desired ‘baseline’ for remembering in organisations, and perhaps even societies. It may also offer a potentially valuable diagnostic, especially when complemented with tools to capture the ‘Ought’, the aspirations for remembering. And as a by-product it may also breathe new life into archives and similar institutions of remembering.

At the same token, we must not forget that much remains to be done. Many questions, especially those centred on questions of power, authority and control over the appropriation of memory in a social context remain to be studied. So do the qualities and interactions of the types of memory we identified. But it was not the aim of the ForgetIT project to solve or to tackle them directly, although ForgetIT’s tools may turn out to be the instruments with which further research into this important domain will be enabled.

At the end of the project, one very large question continues to loom. It is whether there is scope to think of a remembering function covering both individual and social (that is organisational and societal) remembering. There are many fundamental differences, as we have pointed out. And yet, most recent discoveries of how human memory actually works seem to point towards complexity of structures, importance of processes, and the need for pragmatic steps resonating with what we know from social memory studies. Perhaps there is more common ground than we thought.

5 Organisational Memory and Forgetting: Review and Case Studies

5.1 Literature Review

5.1.1 Organisational Memory

In the previous section, there was a focus on the means by which an organisation or society may retain a corporate memory for information and procedures that are essential for its business: documentation, procedures, and physical machinery. In the current section, we focus on understanding what factors may determine whether information is retained or forgotten in organisations. In principle, these factors may apply to documentation, procedures or the design of machinery. Research on organisational knowledge, learning and memory is recognised as both disciplinary and methodologically fragmented [Berends and Antonacopoulou, 2014, Corbett, 1997, Rowlinson et al., 2009, Spender, 1996, Walsh and Ungson, 1991]. As a result the terminology and definitions used when approaching organisational memory vary significantly between authors and fields. In many cases, knowledge, information and memory are used across the literature to mean similar things, and differences are often missed or ignored. Here we consider organisational memory as more an organisational process than the recording, retention and recall of data.

With a basis in wartime production work and standardised, routinized labour, research on organisational memory has often focused on organisational memory as the successful storage and retrieval of facts and reproducible behaviour (such as remembering how to fulfil a function on a production line). Developments, especially in Europe and North America, towards a knowledge-based and service-based economy has prompted a change in approaches to organisational memory as successful firms rely increasingly on innovation and problem solving in times of change rather than standardisation of knowledge and workplace routines.

Most famously, organisational memory has been defined as, ‘stored information from an organisation’s history that can be brought to bear on present decisions’ [Walsh and Ungson, 1991, pp. 61]. The metaphors associated with organisational knowledge and memory draw repeatedly on the notion of collection whether it be in storage bins [Olivera, 2000, Spender, 1996, Walsh and Ungson, 1991], bathtubs [Dierickx and Cool, 1989, Fernández-Molina and Peis, 2001] or the accumulation of knowledge stock [Black et al., 2004, Deeds and Decarolis, 1999, Mills, 1940, Weick and Quinn, 1999]. The assumption is that the power of organisational memory is the knowledge that comes through its accumulation [Jackson et al., 2006] and that more knowledge and quicker access to that knowledge is de facto a good thing [Mills, 1940].

The ‘systematic collection of administrative, financial, diplomatic and family documents’ as part of routinized administration can be seen as a bureaucratization of memory [Jedlowski, 2001,

pp. 37]. The use of new technologies from writing to photography, audio-video recording and currently electronic data capture with sensors have contributed to this exteriorisation of memory. The tools allow a recording of the past with which is associated accuracy and objectivity.

Research on, and theoretical approaches, to memory in organisations have largely taken a functional, even mechanical [Rowlinson et al., 2009], approach. Memory is seen as a resource in need of management in order to support efficiency, enhance problem solving and streamline routine business activities. Feldman & Feldman point out the 'Firms are increasingly viewed as knowledge-based, hence the belief that 'knowledge management' brings a competitive advantage to a company,' [Feldman and Feldman, 2006, pp.862]. As such, approaches to understanding organisational memory have been strongly aligned with management objectives 'the approach taken by most has been problem-focused and strategy-driven' [Feldman and Feldman, 2006, pp.864]. For example, Spender focuses primarily on the firm as a system for generating revenue when he asserts that, 'organizations must surely be defined as systems of purposive activity,' [Spender, 1996, pp. 64].

That is, organisational memory is seen as a tool used to engage with organisational problems whether they are associated with manufacturing efficiency or service innovation. It has prioritised organisational memory not as an entity per se but as sets of knowledge which can be accessed and applied to a given decision or process. This functional view of memory has had a tendency to see memory as information which can be drawn out of storage silos and the value of researching and theorising this drawing out of memory is to make the system more streamlined, efficient and measurable. As such, although individual human memory is often used as a metaphor for organisational memory the metaphor is often strained as individualist approaches are developed which exclude a more social and collected view of organisational memory. Nevertheless, as discussed in Section 4, collective memory is heavily influenced by the memories of individuals within an organisation. Moreover, as was demonstrated in D2.3 Section 5 for individual memory, a detailed understanding of human memory may operate as more than a metaphor by offering principles of function for a conceptual model at the organisational as well as the individual level.

In such a functional model, a key issue has been attempting to establish efficient ways in which information can be collected and codified in a manner considered to be veridical and the expertise needed (usually for managers) to access that information in a timely and appropriate manner. A weakness of the silo model is that, in the first instance, it creates an understanding of information types as being essentially discrete and as such benefit from being categorised and stored in separated containers. It segregates information and knowledge by virtue of criteria such as physical form, level of encoding, discipline or specialisation or access level within a company. It places an emphasis on information management rather than practice and knowledge. Organisational knowledge is seen as a matter of creating perfect and complete records existing outside any social or political context, it sees information as objective and the memory/recording of this as neutral. However, as pointed out Casey and Olivera, memory is 'continually constructed and re-

constructed by humans interacting with each other and their organizational environment,' [Casey and Olivera, 2011, pp. 306].

Further, the systems-like approach tends to emphasise the role of individuals as parts of the organisational memory with little consideration of social, interactive or cultural elements of organisations. Drawing on Kansteiner [Kansteiner, 2002], Rowlinson et al. [Rowlinson et al., 2009] (have argued that the emphases in organisational memory largely remain "within a methodologically individualist framework of 'collected memory', rather than 'collective memory'" (p.78).

However, research on human memory has demonstrated that memory is not simply a process of storage and retrieval of veridical data but rather the synthesis of different sources of information [Loftus and Palmer, 1974]. Arguing more recently that it is not sufficient to regard organisational memory and knowledge as relatively homogenous and individualistic, Heckler has argued that organisational knowledge is 'irreducibly embedded in a collective practice that underlies even individual knowledge and action' [Hecker, 2012, pp. 425] and Adorasio [Adorasio, 2014] has suggested that organisational memory (in her case banks) is neither homogenous nor based entirely upon the recall of facts, being instead made up from information selected to fit within stories or narratives. Walsh & Ungson [Walsh and Ungson, 1991] have had their storage bin approach to memory criticised precisely because it does not seem able to respond to the possibility that memory's relationship to information and knowledge is contextual [Spender, 1996].

As such, we are currently at a stage where work is responding to Spender's warning against seeing organisational memory as a simplistic accumulation of facts, and the presumption that it is 'made up of discrete and transferable granules of understanding about reality which can be added to an extant heap of knowledge' [Spender, 1996, pp. 64]. Sociological and anthropological approaches have suggested a focus on the processes of organisational memory rather than focus primarily on shaping outputs. This, less functional approach to organisational memory, suggests a less linear, production line approach which looks at memory and knowledge work as a social and network practice. Rather than a flow, memory of organisations 'engage in acts and interactions involving language and objects' [Feldman and Feldman, 2006, pp. 863].

5.1.2 Organisational Forgetting

While the majority of work done on the ways in which organisations manage their memory has focused on storing, accessing and using memory to improve organisational efficiency, a smaller body of work has looked at the opposite side of the memory process namely forgetting. Framed as the opposite of managing an efficient organisational memory, forgetting can be seen as an undesirable phenomenon with distinct and measurable negative impacts on productivity and efficiency [Benkard, 2000].

However, as in the case of human forgetting, organisational forgetting is not always seen as negative, detracting from efficiencies attributed to managing storage and recall. There is a growing body of work which regards forgetting not only as a prob-

lem and a failure to successfully encode or recall knowledge but as something that, when managed appropriately, is a necessity which provides value to an organisation especially during times of change or innovation [Bettis and Prahalad, 1995, Carlile, 2004, Easterby-Smith and Lyles, 2011, Lyles and Schwenk, 1992].

As indicated in Section 4, in times of uncertainty, innovation, or technological disruption it is often not beneficial for organisations to routinely apply knowledge of past events and processed to current contexts [Easterby-Smith and Lyles, 2011]. Lessons learned from an organisation's past, and the knowledge resources that agents share may fail to competently describe current developing situations and as such developing new practices based upon this outdated knowledge can lead to costs in development of new products and services and failures in maintaining the relevance of existing products within the market [Carlile, 2004].

In such contexts organisational forgetting is no longer framed in its usual manner in management and organisational literature as a 'bad thing' [Easterby-Smith and Lyles, 2011, pp. 311] but an essential criterion for managing change and supporting innovation.

Of course, not all forgetting is positive as companies build experience, capabilities and established practices many of which are crucial to successful operation. The key issue is how forgetting is managed. This is one reason why it is important to understand different types of forgetting within an organisational framework and enable a framework onto which to build systems for making decisions related to organisational memory and managed forgetting. There are undoubtedly risks to managing organisational forgetting with fears that crucial knowledge and expertise will be lost, established networks decimated, identity compromised or even that individuals and organisations will fail to be accountable for previous actions.

Easterby-Smith and Lyles offer a distinction between two types of loss of knowledge by organisations: 'unlearning' and 'forgetting'. For them, 'unlearning refers to deliberate attempts to dispose of unwanted knowledge, whereas forgetting refers to a loss of knowledge that is not necessarily planned or intended' (p.311). Inspired by this characterisation, we argue that it becomes vital to recognise organisational memory and forgetting as practices rather than attributes [de Holan, 2011].

Knowledge of the past can become problematic from an organisational perspective when it:

- presents a barrier to adopting new knowledge [de Holan, 2011]
- supports outdated practices [Easterby-Smith and Lyles, 2011]
- is used to develop future strategies based on information which has been superseded
- reinforces unproductive boundaries across teams and within an organisation [Carlile, 2004]
- encourages the rejection of innovation and innovative practices as illegitimate [Easterby-Smith and de Holan, 2011]

While types of organisational memory are well detailed and debated within organisational and management literatures, there remains a lack of systematic unpacking of what type of forgetting can be observed within organisations. From a management perspective, managing forgetting involves a management of knowledge practices as well as knowledge resources. We can begin to engage with managing memory and forgetting in organisations as part of a set of processes instigated by factors such as innovation, out-dated knowledge [Fernandez and Sune, 2009] rather than isolated actions or causally linked stimuli and responses and that these processes are interlinked.

If we are to view the management of forgetting as a valuable resource within organisations (and not merely the consequence of failures to manage organisational memory) then it is important to begin to develop a taxonomy of forgetting. There is value in beginning to theoretical distinguish and empirically explore different type of forgetting. Our approach has been to use causes of forgetting in human memory to inspire categories of forgetting in organisational memory. A similar approach was used successfully in D2.3 Section 5, to generate a formal model of individual memory. The reader is referred to that previous document, and to reviews of human memory research in D2.1, D2.2 and D2.3 for an explanation of the terms used below to refer to individual memory.

- Failure to encode: Knowledge is not retained beyond its immediate use. Organisations may acquire information but fail to retain it [Fernandez and Sune, 2009].
- Failure to retrieve: Information is stored in analogue or electronic archives to which key agents do not have access [Easterby-Smith and Lyles, 2011] or fail to access. Retrieval failure is common in individual human memory, typically because the context used for retrieval fails to match the context used for initial encoding.
- Deterioration: Old knowledge becomes impossible to retrieve in an acceptably reliable fashion or the veridical qualities of the information is questionable [Easterby-Smith and Lyles, 2011]. This has been demonstrated repeatedly for human memory when memories fragment and become increasingly inaccurate over time.
- Overwriting: New knowledge and information replaces previous memories. Information is updated, replaced [de Holan, 2011] and new practices become institutionalised as common practice. This has been demonstrated repeatedly for human memory when subsequent experiences or information changes the memory for the original event.
- Removal: A key member of a team leaves or a resource is destroyed. This can mean not only that the knowledge they have embodied is lost but their role in maintaining collective knowledge is lost through the removal of their interconnectedness within the organisation. There is a disruption both by the loss of a key member of a team and by the arrival of any new members [Easterby-Smith and Lyles, 2011]. Here, the equivalent in human memory would be the loss of a memory for past experiences because of brain damage.
- Rewriting: Attempts can be made to revise the collective view of an organisation's past in order to rationalise decisions for current strategies or changes in approach.

This may be associated with changes in key agents within an organisation and put in place for political means. [Easterby-Smith and Lyles, 2011]. In human memory, much of the process of retrieval involves reconstruction and this can often include rationalisation or reinterpretation of a previous event.

5.2 Methods

Building on the research on organisational memory, the research undertaken within the ForgetIT project focused on the practices of managing memory and forgetting in organisations. As stated in the previous section, the concept of organisational memory is a difficult one to observe directly in an empirical manner. In part, this is due to organisational memory being created, retrieved and acted upon by a range of agents within the workplace. Further, organisational memory as a singular is a misnomer with there being little evidence to support the idea that organisations produce and maintain a homogenous set of memories. As Kleiner, Nickelsburg, & Pilarski suggest, 'organizational knowledge is idiosyncratic' [Kleiner et al., 2012, pp. 69].

What we have explored is the processes of how memory is achieved within organisations so as to illustrate the organisational practice and routine management of organisational memory and managed forgetting. We have done so through investigating in detail how people manage their work in organisations.

5.2.1 Aims of the Research

The research on organisational memory has explored empirically the processes used by individuals in organisations to manage the encoding, retrieval and forgetting of information as well as of contextualized remembering. It has aimed at detailing specific organisational policies and practices that shape the use of information, encoded and tacit knowledge, and collective memories within the organisations used as case studies.

The challenge has been to observe and chart the memory and forgetting practices within case study organisations in order to a) develop a deep understanding of practice, behaviour, routines and systems, and b) move towards generating a systematic abstraction of these observations in order to inform our general understanding of organisational memory and work towards producing a general and applicable schema.

With special reference to the development of a digital system for organisational memory, this research has addressed the following:

1. What is the theoretical contribution of "forgetting" to the knowledge-based theory of 'the firm' and to the practice of organisational learning?
2. What mechanisms and practice are in place that support the retention and communication of knowledge within organisations?

3. How can we describe the different methods for forgetting and also different measures?
4. Do managers ignore forgotten past experiences or routines when evaluating strategic issues? Is the past irrelevant? What does this imply about the importance of knowledge retrieval within firms?
5. The strategic management literature demonstrates the importance of organisational learning on organisational performance. How does forgetting benefit or hinder organisations?
6. What is the interrelationship between forgetting and job roles? Could these be applied to the wider organisation? What is preventing their implementation?

5.2.2 Data Collection

The responsibility for empirical elements of the ForgetIT Project's work on organisational memory were transferred from UOXF to UEDIN with agreement of the ForgetIT partners and the EU Scientific officer during the latter part of the second year of the project. A member of staff was recruited at UEDIN to take on this work beginning month 25. Although this process involved delays in the initial scheduling of the data collection, interviews and analysis were completed on time using the reallocated budget.

This change in responsibility did necessarily involve a deviation from the initial Description of Work concerning the case studies developed. As such the study remained committed to completing three case studies in organisations heavily involved in managing organisational, societal and administrative memory.

The three organisations for case studies were:

- A British national museum responsible for managing a large and mixed collection of items including natural sciences, science & technology, art and design, national history and archaeology, and world cultures.
- An international academic journal which has its current editorial base in the UK but draws upon an international network of assistant editors and reviewers as well as authors to produce paper and online versions of the journal.
- An international (but USA-weighted) scientific organisation which had within the last few years engaged in a shift in administrative structure as part of a process to refocus the organisation on broadening membership and the management of academic outputs.

Semi-structured interviews were undertaken with individuals working in the three organisations with the aim of getting a deep view of how memory and forgetting are managed by individuals working within these organisations and the view they had of institutional practices. Working with gatekeeper/s in the organisations we identified individuals within each

organisation who played different and significant roles in engaging with and managing the organisation's memory. The initial aim was to complete 15 interviews, however, this was exceeded with 19 completed across the final 12 months of the project.

The interviews were divided as follows:

- 11 at the national museum – including curatorial staff, librarians, and legal department
- 3 with the international journal – including editor, editorial assistant, assistant editor
- 5 with the learned society – including chair, chair elect, secretary-treasurer, executive director

Interviewees were briefed on the project and given a printed overview of the project and key contact details. Issues with anonymity were discussed and participants were informed that while they would not be named in any forthcoming outputs from the research, the specialised nature of their jobs meant that identification might be possible. Following this, interviewees were given a consent form to sign which also detailed the withdrawal procedure to sign before agreeing to participate in the interview.

Interviews were based upon an agenda of questions and themes that were agreed in advance and which had received ethical approval through UEDIN's Psychology Ethical Review Board. (A copy of this agenda is attached in Appendix B). Given the nature of the research, the interview-style was managed to be flexible enough to be responsive in engaging with interviewees' expertise of their role within the organisation. The intention with this approach was to explore possible areas of relevance not previously identified.

Interviews taking place in the UK were undertaken face-to-face usually in the interviewees' workplace. International interviews took place online via Skype. In both cases the interviews were audio recorded and transcribed. The interviews took about 90 minutes each.

Interviews were then transcribed and coded using the framework of documentality, procedurally and physicality to identify key practice and perspective offered by the interviewees. Further open coding was done to capture issues which were not satisfactorily captured by our initial framework.

5.3 Findings

What is clear is that the concepts of Documentality, Procedurality and Physicality do apply to all the organisations explored. While we might consider the sample of organisations limited in its focus on service-orientated, knowledge-based institutions, because of the framework's basis in existing literature, we hypothesise that although organisations would vary in the relative importance of each of these three factors, they would be apparent across institutions. This, we suggest, is not necessarily a weakness of the framework but rather a recognition that different organisations will have varying requirements regarding

the management and application of organisational memory. Industries focused on producing standardised and quantified products will require different operational procedures from those who services are more responsive and customised. Such standardisation is also important in organisations where high staff turnover necessitates externalising procedures to support quick replacement of members of the workforce and have them fit in to an existing organisational structure.

However, in each of our organisations, the level of experience of staff was notably long-term with staff often having been in role for decades. In such a situation, prescriptive encoding of organisational memory and its management can create tensions as the control of organisational memory, its encoding and retrieval is strongly linked to the power asymmetries and organisational hierarchies within an institution. The non-reflexive capturing of organisational memory can cause problems as staff seek to control what they see as 'their' data and the manner in which they see the process as one which seeks to deprofessionalise their role and downplay the value of their accumulated tacit knowledge. As one respondent said of this in a museum context:

Computers allow the recording of the collection. So all that apprenticeship that people used to go through. Having curators who learned their collections by being immersed in the collection were in post for a long time. And it's through that time they spent working with the stuff that they've developed a level of tacit knowledge and tacit expertise. I think there is a perception that that is no longer necessary because all the information on the collection is in the computer. And so you could take a subject specialist straight off the street straight out of a university and they can more or less hit the ground running because the data are there. [. . .] It undervalues the tacit knowledge which can only be professionalised by working deeply with the material.

Although research intensive, it would be possible to identify and map the key aspects of an organisation's memory encoding to develop an overview of their memory encoding and management procedures.

However, there are some areas in our initial framework that the empirical work has suggested as potential limitations if the framework is to be used to its maximum impact as an analytical management tool. By adopting an interview methodology rather than one which looks primarily at management optimisation, we have noted how access and use of organisational memory is varied and limited. Some of these have been technical or procedural issues – such as staff not being able to access organisational memory systems because of the way IT systems are standardised through the organisation regardless of staff requirements – or information missing from systems because resources are not available to ensure information integrity. No data here are considered better than poor data for a standardised system.

What is clear is that the value (and sunk costs) of documentality is only clear through its contextual and situated retrieval. The organisational memory contained within an institution's documents only realises its potential when it is retrieved in order to be applied in a

new context or to a new problem. By seeing the contents of memory systems as social objects rather than data it has become apparent that staff use of systems is not about rote recall but about engaging with information using existing internal knowledge resources to gain relevant knowledge which can be applied to specific situational contexts.

This suggests that our model may be lacking in application to certain organisational contexts in that:

- The framework tends to emphasise encoding rather than retrieval of data.
- The model lacks an approach to the context of retrieval and use of organisational memory.
- Its application encourages a view of memory encoding as a single event.
- The focus of physical/technical solutions encourages assumptions about the fixing of data being of ultimate importance.
- The emphasis on collecting and fixing organisational memory can devalue the importance of immediate working memory.
- The emphasis is on a standardisation of memory encoding and retrieval which is not responsive to change, innovation and problem solving.
- It is poorly suited to understanding the relative importance of different memories to users.
- The focus on memory collection marginalises the importance of forgetting.
- The focus on collection of organisational memory can under-value the importance of the tacit knowledge contained within the heads of an organisation's staff.

What this suggests is that there is value to be gained from viewing organisational memory more like human memory. This enables us to explore how the social as well as technological context of memory is important when looking at how strategies are built up to manage technologically-mediated memory and forgetting.

What has been observed is that most of the memory systems used by the organisations in the case studies were effectively designed for the capturing of specific and specialised type of information. However, what is less clear is that these systems were designed with equal consideration of the retrieval of information by a range of users with different intentions and expertise. Building on our initial framework and considering the issue listed above it is possible to present a series of issues that can be considered for both understanding similarities and differences across organisations' management of memory

What this knowledge allows us to do, however, is build upon our previous framework by suggesting key aspects of the management of organisational memory and forgetting which are worth considering when developing new systems or analysing existing organisational practices. This we do in the next section.

5.4 Considerations for Managing Digitised Organisational Memory

It is clear from the above that both the encoding of organisational memory and its retrieval is evidently fractured and varied. This presents significant challenges for a system which seeks to capture digital data in order to manage organisational memory and managed forgetting. Below we present a series of considerations which can be used to understand the practice of organisational memory and developing strategy from managing these.

5.4.1 Motivation for Managing Memory

Benefit can be seen in developing a strategy for managing organisational memory which is based on predicting re-use value of information. The recording, archiving and maintaining organisational memory has costs associated with it as does the accessing of these resources. Consideration can be given to issues including:

- **Regulatory requirements:** Certain data, such as financial reporting, have minimum periods for records being kept only after which they may be consolidated or disposed of.
- **Institutional requirements:** Data key to the successful running and survival of the organisation such as stock/archival holdings, customer/supplier details, and current work processes (e.g. the stage at which individual projects are).
- **Risk management:** Is data likely to be needed to be called upon – for example during legal action, insurance or to respond to a Freedom of Information request? Is so is the value in keeping or deleting the data? What are the consequences for providing or failing to provide the information?
- **Knowledge Management:** What plans, resources or strategies are dedicated to the upkeep and analysis of this codified organisational memory in order to maximise its use, recognise potential for exploitation and identify areas for development.

5.4.2 Form of Organisational Memory

Codified and digitised organisational memory should ideally be in a form which can be integrated, translated or interrogated by other systems for example during software changes, organisational mergers or as part of open knowledge projects.

- **Interoperability:** Does the archive need to be recalled in part or in whole by other organisational memory systems?
- **Diversity:** Recognition needs to be shown to the diversity and fractured systems used within an organisation to capture and manage organisational memory given that memory systems tend to have limited scope and specific focus.

- **Integration:** Linked to the above, users may have to navigate several systems regarding a single operation such as email accounts⁷, diary, logs, and non-digital sources as well as recognised systems for managing organisational memory. Technical systems (such as the solution offered in the ForgetIT PIMO) may be able to integrate such digital data sources. Additionally, awareness may be shown in recognising practice which encourages the fragmentation of organisational memory (such as curators' personal object registers in museums which are developed as a personal worktool when the central register is not readily available or data added do not fit in with the standard structure – e.g. photographs which do not conform to the organisational standard).

5.4.3 Access to Organisational Memory

Linked to the technical form of the documenting and archiving of organisational memory is consideration that needs to be shown for the requirements and practices for accessing and retrieval of the data. Organisational memory which is not accessible by appropriate staff especially in a timely and verifiable manner is a wasted resource which essentially presents a cost to un-managed forgetting.

- **Delivering Access:** Digitising organisational memory offers massive potential for maximising impact across an organisation. There is a not inconsiderable challenge in delivering 'community memory applications' [Ackerman and McDonald, 1996] to enable user access to collective knowledge. The issue is of building systems which allow users to retrieve data in a manner appropriate to their own context.
- **Access Level Control:** Conversely, while the ability to retrieve information digitised in organisational memory systems by those that require it is generally considered a positive process, completely free-access to all data may not be considered advantageous. While the siloing of data – e.g. separating work done by design and engineering departments – may negatively impact innovation and productivity, systems need to be responsive to issues such as security, intellectual property protection, and restriction of personal data. While access to organisational memory systems is often organised following established organisational structures, such administrative units may not be ideally suited for supporting the flow of information.
- **Contextual Integrity:** Informational recall should respect privacy norms such that data encoded can only be retrieved by those people who might generally be expected to have been given access to that information [Barth et al., 2006, Nissenbaum, 2004]. For example, while a journal editor might reasonably expect access to information regarding the submission and review of a paper, they would not normally be given access to any personal information contained in emails between an author and a sub-editor even where this information is contained in the same correspondence.

⁷Email remains a central memory technology in the organisations studied despite sitting separate from what were regarded as the key organisation memory management systems.

5.4.4 Managing Memory Context

A key difficulty in archiving data is the attempt to foresee for what use and interrogation a digital archive may be required. This is not a trivial endeavour not only given the problems of crystal ball gazing when it comes to forecasting archival and research demands of the future but also given that the foreseen needs of archivists and other user group and stakeholders may be different. This means that the structure of digital archives must necessarily be stable and organised to support future extraction of information from the archived data and future 'functions' [Kim and Ross, 2008] that may be required.

- **Encoding Context:** The use of metadata offers potential for adding contextual information to digitised organisational memory. However, this information is generally semantic in nature. This means the 'what' of memory is captured but there is lack of the 'why' and the contextual narrative of episodic memory. For example, while museum registers have rich and (semi-)standardised metadata associated with object records, this is limited to object descriptors. While a record of items displayed in exhibitions can record what is included, it poorly captures why the objects were chosen and the narrative behind the exhibition.
- **Continuing Context:** The notion of documentarily may encourage us to think of organisational memory archives as fixed. However, archives, like human memories, change over time and meaning of contained information. They are added to, changed and restructured but also may suffer from data loss or other types of bit rot. Version control and the ability to retrieve different versions of organisational memory is important in order to verify data and understand how and why the data and its associated meaning may have changed.
- **Recall Context:** Although some respondents were reluctant to suggest their work involved a 'normal day' evidence exists to suggest the importance of time of routine being used to manage work. Interviewees would allocate certain time of the day to particular activities (for example dealing with particular emails first thing in the morning; working at specific locations at certain times). Systems that respond to use context for example temporal or spatial can build upon human contextual memory.

5.4.5 Practice

Procedural organisational memory can often be codified and standardised both through operational procedure and technological affordance. This memory involves the routine and repetitive 'doing' of a role which is subject to little change over time. While organisational loss of this memory can have short-term economic impact, the documenting of procedures and their standardisation often means capability can be established relatively quickly. However, while many aspect of operational behaviour can be codified and documented in training manuals and performance handbooks, other aspects of organisational memory need to be enacted for optimal performance. For example in fast food restaurants while food contents and preparation may be standardised, aspects such as staff

management, kitchen cleaning or customer relations rely on other forms of practice.

The silo model of organisational memory focuses on memory encoding and retrieval as discrete events. However, both human and organisational memory are processes which are often performed over periods of time. Memories are rehearsed, re-contextualised and changed. The time and context of retrieval as well as the background knowledge of experience of the individual can also change what information retrieved can mean. As such, memory management systems may benefit from considering the following:

- **Professional knowledge:** Memory systems would benefit from a user-orientated design where the knowledge resources that users will bring to understanding and applying information retrieved from digital memory archives are considered. In a museum context, administrative staff and curators will have different competences to understand and apply knowledge retrieved and have different requirements for the ways they want to interrogate the data.
- **Trust:** The indicators necessary for users to be able to trust the information retrieved from organisational memory systems need to be identified as well as the data entry/encoding practices that can be considered reliable? What processes are in place to trust users to encode data into the system and what is the acceptable level of user error?
- **Training:** Are users expected to learn-by-doing when using organisational memory systems or is training provided? What level/s and frequency are these training courses offered? Are they tailored to specific users groups? What costs are associated with the training? Does it have to be outsourced to system developers who may not be familiar with the organisation's requirements?
- **Engagement:** Use value needs to be clear to users (not just management). Are there identifiable benefits for the user committing to using the organisational memory digital system rather than relying on personal memory or records?

6 Evaluation Methodology

6.1 Personal Preservation Use Case: Design and Methodology

6.1.1 Design and Motivation

This evaluation was the third time that feedback from end users had been gathered. The first study was the PANIC study reported in D9.5, which informed development of the DFKI system as part of an agile process, and the second occasion was the Festival2014 study, when people used the DFKI system for very limited tasks under the supervision of an experimenter. Both of these studies informed the final design that was evaluated in late 2015.

The 2015 evaluation was designed to be formative, not summative. Whilst the DFKI system was sufficiently stable for participants to use it on their own phone, it was clearly not a product that could be moved to an open beta. For this, as we argued in D9.5, substantial additional programming would be required.

Therefore, the evaluation had three aims:

1. establish overall usability judgements for the use case
2. systematically collect suggestions for further usability improvements
3. provide feedback on preservation strategy

6.1.2 Method

The evaluation study was split into three sessions. For **Session 1** participants were instructed to bring 40–80 photographs on a USB stick from a holiday or trip that were not organised into specific events (i.e. all in one folder) and, where possible, included blurry images and duplicates. The first part of Session 1 was devoted to organising the participants' photographs into specific events; participants were instructed to create new folders, each given a brief descriptive title, and collect photographs capturing a particular aspect of the trip (e.g. visits to specific locations, certain memorable events). Participants then discussed each of these folders in turn with the experimenter who used this time to extract key concepts. Following this participants were given the *ForgetIT survey* to complete while the experimenter uploaded the participants' event folders to the PIMO5 system (<https://pimofringe.opendfki.de/pimo5>).

Session 1 also included a detailed tutorial on the PIMO. Participants were guided verbally by the experimenter through a series of actions showcasing the core functions of the photo collection manager. For the tutorial 22 photographs from the 2014 fringe festival study were selected. Participants added the following text as a note describing the collection:

Here are some photos from my trip to the Fringe festival in Edinburgh. I went with my friends Dave and Sarah. We walked along the Royal mile and saw comedy shows in the Gilded Balloon.

This text was designed to contain concepts requiring disambiguation (e.g., “Edinburgh”, “Fringe festival”), those that could be *manually annotated* (e.g., “Royal Mile”, “Gilded Balloon”), and those that required the creation of a new PIMO concept (implemented as a *thing*, e.g. “Dave”, “Sarah”)) thus demonstrating the annotation functionality of the PIMO. The tutorial then introduced preservation values, including the initial show/ hide decisions made by the system at upload and the effect of choosing favourite photographs. Participants also annotated a single image with the following text:

The Ferris wheel in Princes Street Gardens

which also required participants to disambiguate and manually annotate concepts. In the final part of the tutorial participants were shown how to apply PIMO ‘things’ as filters and search for specific photographs; this was done with both the CERTH automatically generated concepts (e.g. “person”) and with user-defined concepts (e.g. “Dave”).

The final part of the first session involved participants reviewing the initial show/ hide decisions for their own collections. When a decision was changed participants were invited to give a brief explanation of their decision. Session 1 took approximately 1 hour to complete.

Session 2 took place, on average, 3.5 days after Session 1 (range 2–4). Participants were asked about their usage of the system between sessions before being asked to complete a series of tasks with their two largest photo collections. The 5 tasks were as follows:

1. To review and change preservation values.
2. For the first favourite photograph apply one of the CERTH concepts as a filter.
3. Add a note describing the collection as a whole complete with key concepts.
4. Annotate individual photographs using concepts to make them more searchable.
5. Search the PIMO for a user-defined concept.

Participants were instructed to think aloud while performing these tasks—that is, describe their current goals and how they were aiming to achieve them, with particular focus on any difficulty experienced—and this was recorded along with a screen cast using the BB TestAssistant (<http://www.bbtestassistant.com/>). Following completion of the tasks participants completed the System Usability Scale [Brooke, 1996] and NASA Task Load Index [Hart and Staveland, 1988] prior to an approximately 8 minute (range 5.5–11) long semi-structured interview. This interview covered the users’ overall impression of

the system with a particular focus on the annotation and search functionality. Participants were asked to comment in detail on the main issues or difficulties they had when using the software and any improvements they would make. Session 2 lasted approximately 1 hour.

Session 3 took place approximately 5.6 days following the second session (range 4–7) and began with a brief interview probing usage, if any, between sessions. The remainder of Session 3 was devoted to the PIMO's time capsule. The results of the ForgetIT survey collected in Session 1 were used to assign participants to one of four categories described in [Wolters et al., 2015] which in turn determined the thresholds used to create the time capsule. Participants then reviewed the decisions for each photo collection either moving items to or from the time capsule; when an image was moved participants were asked to give a brief rationale for this change. The final session lasted approximately 30 minutes.

In return for completing the three sessions each participant was compensated with £30 worth of One4All gift card vouchers (<http://www.one4allgiftcard.co.uk/>).

6.2 Organisational Preservation Use Case: Design and Methodology

6.2.1 Design

While there have been several previous evaluations of the DFKI system, this was the first time that a working version of the DKD system was tested with users. For organisational preservation, previous user feedback sessions had consisted of talks to the community and video demonstrations of system mock ups to the two companies involved in user requirements gathering, Spielwarenmesse and ako.

The video demonstrations represented design sketches of a more or less fully functional system. Therefore, the user interface presented there differed substantially from the interface that was assessed in the 2015 user test.

As the 2015 evaluation of the organisational preservation system was the first formal feedback on a working implementation, it was designed to provide mainly qualitative, formative feedback to guide further development of the system and the user interface.

The final study had three main aims:

1. Provide internal and external evidence that can help our partner L3S create an algorithm to compute memory buoyancy and preservation value. As the system had not been deployed in a company yet, we designed the study to provide some of the data that had been identified as sources of information for preservation value, such as number of likes, views, and edits.
2. Provide systematic qualitative feedback on the annotation tool developed by DKD, which was the main part of the final user interface already visible to users.

Table 16: Allocation of editing tasks per participant

Participant	Web site	Fish (main area) / accessories (minor area)
1	web2	freshwater; aquariums
2	web3	saltwater; decoration
3	web4	freshwater; aquariums
4	web5	saltwater; decoration
5	web6	freshwater; aquariums
6	web2	saltwater; decoration
7	web3	freshwater; aquariums
8	web4	saltwater; decoration
9	web5	freshwater; aquariums
10	web6	saltwater; decoration

Table 17: Plan of web site visits for each participant

Task	P1, P6	P2, P7	P3, P8	P4, P9	P5, P10
freshwater like	web3	web4	web5	web6	web2
seawater like	web4	web5	web6	web2	web3
news like	web5	web6	web2	web3	web4
web site like	web6	web2	web3	web4	web5

3. Collect information on the usability of the TYPO3 CMS itself, because we needed to understand what usability issues were due to the annotation tool itself, and what issues might stem from the user interface of the TYPO3 CMS.

6.2.2 Method

DKD created 5 identical instances of a fake business website, “Fake’s Famous Fish Shop”, for the purposes of the study: web2.fish-shop.net, web3.fish-shop.net, web4.fish-shop.net, web5.fish-shop.net, and web6.fish-shop.net (in the following, shortened as web2 – web6). The web site was professionally designed; more information can be found in D10.4. The websites consisted of areas for freshwater fish, seawater fish, aquariums, decorations, news, and information about the fake shop. We considered the two fish areas as main areas, and the aquariums/decorations as minor areas.

During the study, each participant was asked to edit a main and minor area (freshwater-aquariums or seawater-decoration) from one of the websites, and perform some browsing tasks on a different area in each of the remaining 4 websites. The task distribution was designed such that, by the end of the study, each website had been edited by 2 participants, each of whom focused on a different main-minor area pair, and visited by the other 8. Tables 16 and 17 present an outline of the editing and browsing task allocation per participant.

DKD also prepared 2 sets of training YouTube videos: one about using TYPO3 for rele-

Table 18: Allocation of participants per tasks and website areas

Participant	Main area	Minor area	Editing		Browsing		
1	freshwater	aquariums	web2	web3	web4	web5	web6
2	seawater	decorations	web3	web4	web5	web6	web2
3	freshwater	aquariums	web4	web5	web6	web2	web3
4	seawater	decorations	web5	web6	web2	web3	web4
5	freshwater	aquariums	web6	web2	web3	web4	web5
6	seawater	decorations	web2	web3	web4	web5	web6
7	freshwater	aquariums	web3	web4	web5	web6	web2
8	seawater	decorations	web4	web5	web6	web2	web3
9	freshwater	aquariums	web5	web6	web2	web3	web4
10	seawater	decorations	web6	web2	web3	web4	web5

vant basic website editing functionality, and one about using their annotation extension to TYPO3 for adding metadata to webpage contents. The videos use a slideshow format, without narration.

Participant recruitment

We recruited 10 participants with experience in editing, maintaining or creating web sites from within the Edinburgh University School of Informatics. Commercial web experience or familiarity with TYPO3 were not a requirement. We identified willing participants by using the Inf-general Informatics mailing list, which is frequently used for these purposes. We also distributed flyers within EUSA (Edinburgh University Students Association) buildings, but no participants replied through this channel.

Data Collection

Participants were asked to attend a 2-hour session split into 4 parts. In Part 1, the Intake Interview, they described their experience with web development, editing, maintenance; level of expertise in different platforms and programming languages; experience of creating the information architecture of a web site, and experience with metadata, tags, annotations, and Search Engine Optimisation.

Part 2, participants edited and annotated their assigned web site using the TYPO3 6.1 CMS (site editing) and the ForgetIT annotation interface (annotation). This generated backend data relevant to preservation values.

First, participants edited the web site. This was intended to familiarise them with TYPO3 itself. Before editing, they watched the training videos provided, and then changed text and added images. After completing their edits, they filled in the SUS questionnaire (as already used for the DFKI evaluation) based on their experience of TYPO3. While the tasks and the number of web pages to be edited was fixed, participants were free to be

creative with their actual edits.

Next, participants moved on to annotation. After watching the training videos, participants annotated a number of pages from their assigned areas. The tasks were designed to cover all relevant functions of the annotation interface including searching the annotations.

During both editing and annotation, they were encouraged to “think aloud”, verbalising their actions, observations, and any problems that they encountered. In addition, the experimenter observed, took notes, and was available to step in at any time. Participants were free to refer back to the training videos at any time, and to take notes during the videos. As for the DFKI evaluation, the sessions were recorded using BB Test Assistant.

In Part 3, participants browsed and ‘fishliked’ pages on the remaining 4 websites, focusing on a different area in each, according to Table 2. The ‘fishlikes’ were a version of the usual social media liking mechanism that were not linked to a social media account.

Finally, in part 4, the debriefing interview, we asked participants for their feedback about TYPO3 and the annotation process. In particular, we were interested in their general impression of the software, what the main issues were, what the main benefits were, and what could be improved. The intake and debriefing interview schedules, and the different sets of tasks are available in the Appendices.

The cumulated traces of edits from change logs constitute internal evidence, and the numbers of visits and ‘fishlikes’ external evidence, to be used by L3S for the computation of preservation value and memory buoyancy (aim 1).

The Edinburgh team analysed the qualitative data for aims 2 and 3. In particular, the data generated from the study interviews and tasks was first transcribed from handwritten notes and audio recordings, and enriched with detailed observations from the video recordings. The transcripts were then imported in NVivo (<http://www.qsrinternational.com/>) and analysed thematically. For the intake and debriefing interviews, we started from a pre-defined list of themes representing the main questions which had been asked, but we were open to adding new themes as emerging from the text. For the tasks, we derived the themes and subthemes from the text, the focus being on the problems that had been identified by the participants, and their suggestions. The generated list of themes and subthemes is available in Appendix C.6. A more detailed discussion of the findings can be found in Deliverable D10.4.

Table 19: Success Indicators from Revised D2.1 (Appendix)

Outcome	Indicator	Del.
1) Principles of human and organisational preservation and forgetting that can be implemented in digital systems	conceptual principles	D2.2
2) User evaluation framework	number of users; analysis of results	D2.3/D2.4
3) Human information preservation and forgetting when interacting with digital systems	revised conceptual models	D2.3
4) Conceptual framework for organisational/societal forgetting	in-depth case studies and validated conceptual framework	D2.4

7 Conclusion

7.1 Assessment of Success Indicators

Table 19 provides an overview of the success indicators listed in D2.1, and the deliverables where the corresponding frameworks and models were described. All of tasks assigned to the success indicators have been completed successfully. The only major change concerns the deliverable where the revised conceptual model of human forgetting is presented; this work, a collaboration with EURIX, was summarised in Deliverable D2.3. Documentation of the second indicator, the user evaluation framework, is spread across several deliverables. The overall approach is documented here, in D2.4, whilst the actual results are described and discussed in detail in Deliverables D9.5 (personal preservation) and D10.4 (organisational preservation).

7.2 Lessons Learned

A major outcome from WP2 has been a much greater understanding of the individual use case for both human and digital preservation and managed forgetting of real life events, and of digital photographs as records of those events. That understanding has been incorporated in the overall ForgetIT Framework, and a specific conceptual and computational model of a system for supporting preservation and managed forgetting of a personal digital photograph collection was reported in D2.3.

Additional results reported in this document have revealed more precisely what individuals

recall from real life events and demonstrate that reviewing digital photographs appears to have little influence on the episodic details that are recalled up to 11 months after an event. Notable too was that while the automatic analysis, classification, and contextualizing of digital photographs can be very helpful for managing a photograph collection, human input (e.g. annotations) is still required for personal contextual reasons, such as an emotional response to an event even if the photograph of that event is of poor quality.

The survey demonstrated the range of personas that there may be among potential early adopters of a ForgetIT system, requiring some tailoring to user preferences for any such system. The user evaluations provided useful feedback to help further improve the usability and utility of the DFKI Personal Information Manager (PIMO). They also demonstrated that the PIMO system has significant potential for supporting human users in their management of digital photograph collections, and to complement the automated systems for classifying and contextualizing digital material within the ForgetIT Framework.

As shown in the review of previous research in Section 5.1, the current understanding in the research literature of organisational remembering and forgetting is at a very early stage, relative to the understanding of human memory in individuals. The discussion of conceptual modelling of remembering in organisations, discussed in Section 4 has offered a framework for understanding the means by which organisations may retain corporate memory that is necessary for their business, through documentation, procedures and design of the physical workplace.

The in-depth interviews conducted with three different kinds of organisations, and reported in Section 5 have added to knowledge as to what might determine whether or not information is remembered or forgotten within an organisation. One common cause of forgetting (e.g. in a national museum) appears to be a failure to retrieve because individuals who require to access certain classes of digitized (documented) information have either no authority to do so, or do not know how to do so. Other constraints on organisations arise from the legal or statutory requirements for preservation. Finally, organisations remain vulnerable to loss of staff who are the only repositories of skills and knowledge that have not been preserved by other means. This is especially true of scientific societies and other organisations in which the key personnel are in elected positions for fixed periods and then are replaced, but is also true of commercial organisations when making policy decisions to reduce staffing levels in parts or all areas of activity, or when staff themselves decide to leave. The design of a digital system to support remembering and managed forgetting within organisations could incorporate documentation and procedures, but for almost any organisation there would be a requirement to incorporate some means to capture the knowledge and skills held within the biological memory systems of individuals.

We note in Section 5 that the in-depth interviews conducted within the ForgetIT project have offered some insight into features that determine success or failure of organisational remembering and forgetting, and indicate that these features may be common across different types of organisations. This offers the basis on which to identify which of the possible causes of weakness or failure of digital preservation systems in organisations.

That is, we now have an idea of what to look for. The challenge is how to identify those weaknesses within organisations on a more routine and non-invasive basis.

One approach now being developed by partner dkd is to capture interactions with digital systems automatically. This will reveal search patterns and possible search failures when using digital repositories and web pages that have already been set up and are in use. It will of course not reveal what key information is held only in the biological repositories (the brains) of the individuals who make up the organisation.

Our approach in D2.3 to generating a conceptual model of the individual interacting with a digital preservation system emphasized a synergetic relationship with the human and the machine each having a role. This followed the overall approach within the project of attempting to design digital preservation systems that complement the human user, not attempt to replace the human. Although organisations and societies are more than the sum of the individual members, the results from WP2 point to the very substantial additional gains in effective digital preservation that may be achieved by taking a similar synergetic approach to the design of digital preservation that fully incorporates an understanding of, and an interaction with, not a replacement for, the human users in their respective roles within organisations and within society.

7.3 Vision for the Future

As outlined in D11.5, both UEDIN and UOXF already have funded follow-on projects that follow up lines of research that began within ForgetIT. At UEDIN, we are also preparing additional grant proposals that continue the ForgetIT theme of harnessing results on human memory and forgetting to improve the usability and usefulness of IT systems. Members of the UEDIN team are also actively researching the collection, management, and analysis of health and wellbeing data, which lies at the intersection of personal preservation (the individual who is collecting health data) and organisational preservation (the organisations involved in treating the individual or providing health insurance).

References

- [Ackerman and McDonald, 1996] Ackerman, M. S. and McDonald, D. W. (1996). Answer garden 2: merging organizational memory with collaborative help. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work*, pages 97–105. ACM.
- [Adorisio, 2014] Adorisio, A. L. M. (2014). Organizational remembering as narrative: storying the past in banking. *Organization*, 21(4):463–476.
- [Anderson, 2004] Anderson, B. (2004). *Imagined communities: Reflections on the origin and spread of nationalism*. Verso, London, UK.
- [Barth et al., 2006] Barth, A., Datta, A., Mitchell, J. C., and Nissenbaum, H. (2006). Privacy and contextual integrity: Framework and applications. In *2006 IEEE Symposium on Security and Privacy*, pages 184–198. IEEE.
- [Benkard, 2000] Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *American Economic Review*, 90:1034–1054.
- [Berends and Antonacopoulou, 2014] Berends, H. and Antonacopoulou, E. (2014). Time and organizational learning: A review and agenda for future research. *International Journal of Management Reviews*, 16(4):437–453.
- [Bettis and Prahalad, 1995] Bettis, R. A. and Prahalad, C. K. (1995). The dominant logic: Retrospective and extension. *Strategic management journal*, 16(1):5–14.
- [Black et al., 2004] Black, L. J., Carlile, P. R., and Repenning, N. P. (2004). A dynamic theory of expertise and occupational boundaries in new technology implementation: Building on barley’s study of ct scanning. *Administrative Science Quarterly*, 49(4):572–607.
- [Brooke, 1996] Brooke, J. (1996). SUS-A quick and dirty usability scale. *Usability Evaluation in Industry*, 189(194):4–7.
- [Buchanan et al., 2007] Buchanan, T., Paine, C., Joinson, A. N., and 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.
- [Carlile, 2004] Carlile, P. R. (2004). Transferring, translating, and transforming: An integrative framework for managing knowledge across boundaries. *Organization science*, 15(5):555–568.
- [Casey and Olivera, 2011] Casey, A. J. and Olivera, F. (2011). Reflections on organizational memory and forgetting. *Journal of Management Inquiry*, 20(3):305–310.
- [Connerton, 1989] Connerton, P. (1989). *How societies remember*. Cambridge University Press.

- [Conway, 2009] Conway, M. A. (2009). Episodic memories. *Neuropsychologia*, 47(11):2305–2313.
- [Corbett, 1997] Corbett, J. M. (1997). Towards a sociological model of organisational memory. In *System Sciences, 1997, Proceedings of the Thirtieth Hawaii International Conference on*, volume 2, pages 252–261. IEEE.
- [de Holan, 2011] de Holan, P. M. (2011). Agency in voluntary organizational forgetting. *Journal of Management Inquiry*, 20(3):317–322.
- [Deeds and Decarolis, 1999] Deeds, D. L. and Decarolis, D. M. (1999). The impact of stocks and flows of organizational knowledge on firm performance: An empirical investigation of the biotechnology industry. *Strategic management journal*, 20(10):953–968.
- [Dierickx and Cool, 1989] Dierickx, I. and Cool, K. (1989). Asset stock accumulation and sustainability of competitive advantage. *Management science*, 35(12):1504–1511.
- [Dobberkau et al., 2016] Dobberkau, O., Goslar, J., Gsedl, I., and Wolters, M. (2016). ForgetIT Deliverable D10.4: Organizational Preservation Report.
- [Douglas, 1986] Douglas, M. (1986). *How institutions think*. Syracuse University Press.
- [Easterby-Smith and Lyles, 2011] Easterby-Smith, M. and Lyles, M. A. (2011). In praise of organizational forgetting. *Journal of Management Inquiry*, 20(3):311–316.
- [Feldman and Feldman, 2006] Feldman, R. M. and Feldman, S. P. (2006). What links the chain: An essay on organizational remembering as practice. *Organization*, 13(6):861–887.
- [Fernandez and Sune, 2009] Fernandez, V. and Sune, A. (2009). Organizational forgetting and its causes: an empirical research. *Journal of organizational change management*, 22(6):620–634.
- [Fernández-Molina and Peis, 2001] Fernández-Molina, J. C. and Peis, E. (2001). The moral rights of authors in the age of digital information. *Journal of the American Society for Information Science and Technology*, 52(2):109–117.
- [Ferraris, 2010] Ferraris, M. (2010). *Documentality: Why it is necessary to leave traces*. Fordham University Press.
- [Gawande, 2010] Gawande, A. (2010). *The checklist manifesto: How to get things right*. Profile Books.
- [Hacking, 1998] Hacking, I. (1998). *Rewriting the soul: Multiple personality and the sciences of memory*. Princeton University Press.
- [Halbwachs and Coser, 1992] Halbwachs, M. and Coser, L. A. (1992). *On collective memory*. University of Chicago Press.

- [Hart and Staveland, 1988] Hart, S. G. and Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Advances in Psychology*, 52:139–183.
- [Hecker, 2012] Hecker, A. (2012). Knowledge beyond the individual? making sense of a notion of collective knowledge in organization theory. *Organization Studies*, 33(3):423–445.
- [Jackson et al., 2006] Jackson, S. E., Chuang, C.-h., Harden, E. E., and Jiang, Y. (2006). Toward developing human resource management systems for knowledge-intensive teamwork. In *Research in Personnel and Human Resources Management (Research in Personnel and Human Resources Management)*, volume 25, pages 27–70. Emerald Group Publishing Limited.
- [Jedlowski, 2001] Jedlowski, P. (2001). Memory and sociology themes and issues. *Time & Society*, 10(1):29–44.
- [Kansteiner, 2002] Kansteiner, W. (2002). Finding meaning in memory: A methodological critique of collective memory studies. *History and theory*, 41(2):179–197.
- [Kaufman and Rousseeuw, 1990] Kaufman, L. and Rousseeuw, P. J. (1990). *Finding groups in data*. Wiley, New York NY.
- [Kim and Ross, 2008] Kim, Y. and Ross, S. (2008). “the naming of cats”: Automated genre classification. *International Journal of Digital Curation*, 2(1):49–61.
- [Kleiner et al., 2012] Kleiner, M. M., Nickelsburg, J., and Pilarski, A. M. (2012). Organizational and individual learning and forgetting. *Industrial & Labor Relations Review*, 65(1):68–81.
- [Koutstaal et al., 1998] Koutstaal, W., Schacter, D. L., Johnson, M. K., Angell, K. E., and 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.
- [Levine et al., 2002] Levine, B., Svoboda, E., Hay, J. F., Winocur, G., and Moscovitch, M. (2002). Aging and autobiographical memory: dissociating episodic from semantic retrieval. *Psychology and Aging*, 17(4):677–689.
- [Loftus and Palmer, 1974] Loftus, E. F. and Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of verbal learning and verbal behavior*, 13(5):585–589.
- [Logie et al., 2014] Logie, R., Niven, E., Wolters, M., and Mayer-Schönberger, V. (2014). ForgetIT Deliverable D2.2: Foundations of Forgetting and Remembering - Preliminary Report.
- [Logie et al., 2015] Logie, R., Niven, E., Wolters, M., Mayer-Schönberger, V., and Allasia, W. (2015). ForgetIT Deliverable D2.3: Foundations of Forgetting and Remembering - Intermediary Report.

- [Lyles and Schwenk, 1992] Lyles, M. A. and Schwenk, C. R. (1992). Top management, strategy and organizational knowledge structures. *Journal of Management Studies*, 29(2):155–174.
- [Mahler, 2009] Mahler, J. G. (2009). *Organizational learning at NASA: the Challenger and Columbia accidents*. Georgetown University Press.
- [Maus et al., 2016] Maus, H., Schwarz, S., Jilek, C., Wolters, M., Rhodes, S., Ceroni, A., and Gür, G. (2016). ForgetIT Deliverable D9.5: Personal Preservation Report.
- [Mills, 1940] Mills, C. W. (1940). Situated actions and vocabularies of motive. *American sociological review*, 5(6):904–913.
- [Nissenbaum, 2004] Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington law review*, 79(1):101–139.
- [Olivera, 2000] Olivera, F. (2000). Memory systems in organizations: an empirical investigation of mechanisms for knowledge collection, storage and access. *Journal of Management Studies*, 37(6):811–832.
- [Rowlinson et al., 2009] Rowlinson, M., Booth, C., Clark, P., Delahaye, A., and Procter, S. (2009). Social remembering and organizational memory. *Organization Studies*, 31(1):69–87.
- [Schulenberg and Melton, 2008] Schulenberg, S. E. and Melton, A. M. A. (2008). The Computer Aversion, Attitudes, and Familiarity Index (CAAFI): A validity study. *Computers in Human Behavior*, 24(6):2620–2638.
- [Spender, 1996] Spender, J.-C. (1996). Organizational knowledge, learning and memory: three concepts in search of a theory. *Journal of organizational change management*, 9(1):63–78.
- [Walsh and Ungson, 1991] Walsh, J. P. and Ungson, G. R. (1991). Organizational memory. *The Academy of Management Review*, 16(1):57–91.
- [Weick and Quinn, 1999] Weick, K. E. and Quinn, R. E. (1999). Organizational change and development. *Annual review of psychology*, 50(1):361–386.
- [Wolters et al., 2013] Wolters, M., Niven, E., Logie, R., Mayer-Schönberger, V., Maus, H., Dobberkau, O., Damhuis, A., Niederée, C., Kanhabua, N., Nilsson, J., Mezaris, V., Gallo, F., Chen, D., and Greenwood, M. (2013). ForgetIT Deliverable D2.1: Evaluation Plan.
- [Wolters et al., 2015] Wolters, M. K., Niven, E., Runardotter, M., Gallo, F., Maus, H., and Logie, R. H. (2015). Personal Photo Preservation for the Smartphone Generation. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, pages 1549–1554. ACM.
- [Wynn and Logie, 1998] Wynn, V. and Logie, R. (1998). The veracity of long-term memories: did Bartlett get it right? *Applied Cognitive Psychology*, 12(1):1–20.

Acronyms

IQR Inter-Quartile Range. 25, 26

OR Odds Ratio. 24

A Sample Callback Sheet From 2014 Festival Study

Photo 1 (IMG_20140810_124705)

1. How much does this photo mean to you?

A lot

Nothing at all

<input type="checkbox"/>									
--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------

2. How long would you be willing to take to find out where this photo is?

Don't care as long as I know

I definitely have it

Immediately

<input type="checkbox"/>									
--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------

I don't care about finding this photo again

3. How long would you be willing to wait to get this photo once you've said you want it?

Don't care as long as I know

I definitely have it

Immediately

<input type="checkbox"/>									
--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------

I don't care about finding this photo again

4. Imagine that your photo management software has suggested these 18 words (descriptors / tags) to describe the photo.

Please tick the descriptors you want to keep.

<input type="checkbox"/> Building	<input type="checkbox"/> scotland	<input type="checkbox"/> Cityscape
<input type="checkbox"/> fringe	<input type="checkbox"/> beautiful	<input type="checkbox"/> hhgh
<input type="checkbox"/> unbored	<input type="checkbox"/> jamesplays	<input type="checkbox"/> edfringe
<input type="checkbox"/> Daytime Outdoor	<input type="checkbox"/> hmfc	<input type="checkbox"/> Apartments
<input type="checkbox"/> edinburgh	<input type="checkbox"/> iraq	<input type="checkbox"/> edfringe2014
<input type="checkbox"/> indyref	<input type="checkbox"/> News	<input type="checkbox"/> Urban Scenes

Photo 2 (IMG_20140810_124811)

1. How much does this photo mean to you?

A lot

Nothing
at all

<input type="checkbox"/>									
--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------

2. How long would you be willing to take to find out where this photo is?

Don't care as long as I know

I definitely have it

Immediately

<input type="checkbox"/>									
--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------

I don't care about finding this photo again

3. How long would you be willing to wait to get this photo once you've said you want it?

Don't care as long as I know

I definitely have it

Immediately

<input type="checkbox"/>									
--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------

I don't care about finding this photo again

4. Imagine that your photo management software has suggested these 18 words (descriptors / tags) to describe the photo.

Please tick the descriptors you want to keep.

hmfc

hhgh

edinburgh

scotland

edfringe2014

Furniture

edfringe

unbored

Male Person

iraq

beautiful

Background Static

Table

indyref

Text

fringe

jamesplays

Person

B Schedule for ForgetIT Semi-Structured Interviews

B.1 Preamble

- Thank interviewee for time
- Introduce Participant Information Sheet
- Explain background to project and key research aims
- Explain interview process and time expected to take:
 - Estimated to take no more than 90 minutes
 - Transcription and quotation will be anonymous but institutions, institution type, and type of role of the interviewee in the organisation may be referred to in outputs
 - Interviewee able to stop and/or withdraw at any time
 - Confirm permission to record
- Ensure Consent Form is signed and Participant Information Sheet is given to participant

B.2 Interview Topics and Questions

Common Knowledge

- Tell me about the organisation: What is its mission?
 - History?
 - Who are the users?
- Whats the story behind you getting appointed to this role/project?
 - How long have you been part of the organisation?
 - How did you end up in this role?
 - Previous work experience?
 - What relevant knowledge did you bring with you from previous work?

Tacit Knowledge

- How did you learn the skills and acquire the expert knowledge required for this role?
 - Education?
 - Training?
 - In house training?

- Reading manuals/documents?
- Talking to previous holders of your current role?
- Talking to managers?
- Learning by performing the role?
- What does a general day or project span look like for you?
 - Do you plan and organise your own work schedule or is this determined by a protocol/manager/external events such as emails arriving?
 - Key milestones?
 - Reporting process?
 - How do you deal with interruptions?
- What are the key activities you are engaged with?
 - Who do you work with?
 - How are decisions arrived at?
 - How flexible is this process?
 - When are other experts called in?
- What happens if staff leave?
 - What's the staff turnover rate?
 - Is there a handover process?
 - What records are kept and used?
 - Passing on of process knowledge?
 - Passing on of organisational knowledge?

Encoded Knowledge

- What sort of admin requirements are there in your role?
 - What is recorded and how?
 - Who is this aimed at?
- What are the systems for recording what the organisation does?
 - In the memory of individual staff?
 - Databases?
 - Committee/group minutes?
 - Audio video resources?
 - Paper documents?
 - Where are these located?
 - Who is responsible for these?

- Who has access?
- Has there been a significant shift from paper-based to IT systems over the last 10 years? How was this managed?
 - Are the paper systems still used/archived?
 - What is still not electronic?
 - Was the shift associated with a change in the info recorded?
- What external resources do you draw on?
 - Company-wide databases?
 - External & commercial data?
 - Where are these located?
 - Who is responsible for these?
 - Who has access?
 - Free resources/internet?
 - Knowledge networks?
- What contingency plans are in place to retain knowledge?
 - What happens if the KM systems aren't available?
 - Are there backups?
 - What if funding/budgets change?
 - What if strategic collaborations were to fail?

Conclusion

- Who are the other key people we should talk to?
- Any questions?

B.3 Afterward

- Agree any further action (e.g. follow-up interview)
- Remind participant of contact details
- Thank interviewee for time and help

C Material for the Organisational Evaluation

C.1 Participant Information Sheet

PROJECT TITLE

Creating a test bed for evaluating organisational preservation strategies

INVITATION

In this study, we are seeking to create a test environment for a new approach to managing and preserving web site content. It is based on the enterprise content management system TYPO3.

In order to create this environment, our project partner dkd have created five instances of a business website for Fake's Famous Fish Shop. You will be asked to edit one of these web sites and visit the four remaining web site with specific browsing tasks in mind. You will also be asked to share your thoughts on the usability of the TYPO3-based annotation and editing environment that dkd have created.

This research is being conducted under the supervision of Professor Robert Logie and Dr Maria Wolters at the University of Edinburgh as part of a European grant titled 'ForgetIT'. We were not involved in the development of the system created by dkd and are only contributing test data and usability feedback. This project has been approved by the Psychology Research Ethics Committee (9 – 1516/1).

WHAT WILL HAPPEN

The session is split into four parts. Part 1 is the intake interview, we will ask you about your background in web design. This will be audio recorded.

In Part 2, you will receive a brief TYPO3 tutorial, explore the TYPO3 website you will be editing through a dedicated editor interface, and fill in a brief survey about your experience. After the introduction, you will receive a tutorial on the annotation process, and we will ask you to extend, edit, and annotate existing content on the web site. You will be asked to fill in a short survey about your experience.

Your interaction with the TYPO3 web site in Part 2 will be audio and video recorded. We would like to encourage you to think aloud while navigating and editing the web site.

In Part 3, you will be asked to browse and like pages on the remaining four web sites. The liking mechanism does not require access to your social media pages. You will not be recorded as you do this, and page visits and likes will be logged anonymously.

Finally, in Part 4, we will ask you to reflect on your experience with TYPO3 and the annotation process in a short semi-structured interview. This interview will be audio-recorded

and transcribed for analysis.

TIME COMMITMENT

Overall the study is expected to last around 2 hours.

PARTICIPANTS' RIGHTS

You may decide to stop being a part of the research study at any time without explanation. You have the right to ask that any data you have supplied to that point be withdrawn/destroyed. You will still be paid for your contribution.

You have the right to omit or refuse to answer any question that is asked of you.

You have the right to have your questions about the procedures answered (unless answering these questions would interfere with the study's outcome).

If you have any questions as a result of reading this information sheet, you should ask the researcher before the study begins.

BENEFITS AND RISKS

There are no known benefits or risks for you in this study.

COST, REIMBURSEMENT AND COMPENSATION

You will receive \$20 in return for your participation. The payment can be made via One4All gift cards which are redeemable at a wide range of retailers (see www.one4allgiftcard.co.uk or in cash.

CONFIDENTIALITY/ANONYMITY

The data we collect do not contain any personal information about you. Video/ audio recordings will be anonymised as far as possible and stored on a secure, password protected disk to which only members of the project have access. No one will link the data you provided to the identifying information you supplied (e.g., name, address, email). The data collected in these sessions may form the basis of conference presentations or publications in academic journals. You will not be identified in any dissemination of research findings.

FOR FURTHER INFORMATION

Professor Robert Logie or Dr Maria Wolters will be glad to answer your questions about this study at any time. You may contact Prof. Logie at rlogie@staffmail.ed.ac.uk / 0131 651 1394 or Dr Wolters at maria.wolters@ed.ac.uk / 0131 650 3425. If you want to find

out about the final results of this study, please tell the researcher.

<https://www.forgetit-project.eu/>

INFORMED CONSENT FORM

PROJECT TITLE

Creating a test bed for evaluating organisational preservation strategies

PROJECT SUMMARY

In this study, you contribute to the creation of a set of web sites and associated editing and usage data that will be used to test intelligent archiving strategies.

Please tick the boxes below to indicate that:

- You have read and understood the Participant Information Sheet.
- Questions about your participation in this study have been answered satisfactorily.
- You are taking part in this research study voluntarily (without coercion).

PERMISSION TO USE VISUAL/ AUDIO RECORDINGS

As part of this study we use video recordings to assess use of the software and audio recordings to transcribe interview responses, but this is not compulsory.

- I give permission for audio recording to be used during the sessions
- I give permission for video recording to be used during the sessions

The recordings are anonymised and will only be accessible by researchers on the project. If you are happy for us to use your data in presentations please indicate below:

- I give permission for my data (video or audio recordings) to be used in academic presentations (fully anonymised).

Participant's Name (Printed)⁸

Participant's signature

Date

Name of person obtaining consent (Printed)

Signature of person obtaining consent

C.2 Intake Interview

Part 1: Demographics questionnaire

- Age

⁸Participants wishing to preserve some degree of anonymity may use their initials (from the British Psychological Society Guidelines for Minimal Standards of Ethical Approval in Psychological Research)

- Gender
- Current occupation (part time student, full time student, unemployed, employed part time, employed full time, homemaker, retired)
- Highest educational qualification (Highers, A-Levels, Bachelor, Master, vocational qualification)

Part 2: Interview

In the advert, we asked for people who have experience with editing, maintaining, and creating web sites. Could you tell me:

- What is/was the context in which you edited, maintained, or created web sites? [name web sites]
- How long have you done / did you do this?
- Were you employed by anyone? (Self-employed? / For free? / Own site?)
- What are the key activities you undertake/undertook?
- What skills are /were required?
- Are/Were you part of a team? If so, what size is that team?
- Have you worked formally in web development? For how long? [List all relevant posts]

Experience with Web Development Environments

What are the web development environments you use at the moment?

Have you used other web development systems prior to this?

Had you heard about TYPO3 before you saw our ad?

- If yes, what did you know about the system?
- Have you worked with it previously?
- Do you know people who have worked with it?

Skill Development

Now, I would like to ask you about how you learnt the skills you used to edit, create, and maintain web sites: [link back to skills mentioned earlier]

- Did you undertake formal training, for example at university or professional courses?
- Have you received any in-house training from your employer?
[If yes to any of those two, which ones were in the last 12 months?]

- Do you share skills between colleagues?
- What information resources do you draw on to develop your job skills and learn new things? (Wordpress.org, Stack Overflow, YouTube, Online course, Printed Manuals, etc.)

Team Work

When working on a project, do the briefs you are working to change (for example, because of customer issues, technical difficulties, changes in management priorities)?

If this happens, how is this handled?

When organising your work, what kind of things are written down or done through project management software?

What sort of things are decided face to face or in a meeting and never formally recorded?

Tagging / Metadata / SEO

Have you done any search engine optimization work? If so, please describe it with examples.

Have you been involved in tagging content or creating metadata (either automatically or manually)? If so:

- What was the aim of tagging or adding metadata for these projects?
- What were the key elements or categories that you labelled with metadata?
- Who (or what) was intended to make use of that metadata?
- How was it used?

C.3 Debriefing Interview

What is your impression of TYPO3?

How does it compare with the other web content management systems you know?

(Encourage people to be specific, ask about details)

What did you like best?

What did you like least?

What did you think about the annotation mechanism?

What did you like best? Why?

What did you like least? Why?

If you could change one thing about the annotation tool, what would that be?

How useful were the automatic annotations?

What was particularly impressive?

What was missing?

Show screenshots of web1, where the fish from the dkd demo is annotated.

Get feedback on the way multiple instances of the same concept are displayed

Get feedback on the names of the buttons - better wording?

Could you see a system like this being of use in your current work environment?

How would you use this tool in practice?

Function of the index

Additional prompts:

I noticed you had problems with / it took you some time to do X. Could you tell me a bit more about your problems?

C.4 Interview Schedule for Odd Participant Numbers

Intro to TYPO3

Videos:

<https://www.youtube.com/playlist?list=PL1wplB0LBgiaNwPGA4aVxp7SdmRZRCSGK>

1. Log into the TYPO3 backend in the first tab of the browser (username: writer, password: writer).
2. Spend a few minutes browsing through the page tree from the Page View, and opening individual pages.
3. Go to Aquarium Fish Accessories - Aquarium and do the following:
 - (a) In the description of one of the aquarium products, add at the end Recommended fish: Goldfish, Oscar
 - (b) Add links to the pages of the recommended fish.
 - (c) Save and view the results in the frontend.
4. For 1-3 pages from Fish - Freshwater Fish, do the following:

- (a) Replace the pictures with another picture from the folder Images/new/Freshwater
 - (b) Add 1-2 other pictures from Images/new/Freshwater.
 - (c) Save and view the results in the frontend.
5. Go to the News folder (NOT the Fish News) and do the following:
- (a) Add a new piece of fish news by choosing its contents from the News tab, and add to it any image from TYPO3 from Images/new/News.
 - (b) Edit an additional 1-3 pieces of fish news (e.g. shorten them, add some text, etc.).
 - (c) Save and view the results in the frontend.

Questionnaire time - please fill in the SUS questionnaire about TYPO3

Annotations

Videos:

<https://www.youtube.com/playlist?list=PL1wplB0LBgiYlQihHcE0RDfj1RV6IngQJ>

1. Annotate a random selection of 4-6 fish subpages of the Freshwater Fish page:
 - (a) Perform automatic annotation
 - (b) Edit the annotations. For the first page you work on, do all of these steps, for the rest, choose what you would prefer to do.
 - i. Go through the annotations and delete any that are not useful.
 - ii. Select a term which was not already annotated and annotate it. For locations and place names, there is a resource file in the tabs.
 - iii. Select a term that has already been annotated and that has a subregion that can also be annotated, select the subregion, and annotate it.
 - (c) Index the annotations.
 - (d) Save the page.
2. Use Mimir Search to:
 - (a) display all annotations.
 - (b) search for all mentions of entities of type Fish.
 - (c) search for all countries.
3. Go to the News folder (NOT the Fish News), and for any 2-4 of the news items:
 - (a) Perform automatic annotation.
 - (b) Edit the annotations as you did for the Fish subpages, doing 1-3 steps per page.
 - (c) Index the annotations.
 - (d) Save the page.

Questionnaire time - please fill in the SUS questionnaire about the annotation process.

Surfing the Fish Web

You are now done with editing your website, and its time to look through the other websites.

1. Go to the website in the second tab. Look through the freshwater fish, go to the pages of fish you like, and fishlike the pages (choose at least one fish, no other restrictions).
2. Go to the website in the third tab. Look through the seawater fish, go to the pages of fish you like, and fishlike the pages (choose at least one fish, no other restrictions).
3. Go to the website in the fourth tab. Go through the news, read a few pages, and fishlike all news items that you find interesting.

4. Go to the website in the fifth tab. Go through the whole website, explore all pages that are not related to news and fish, and fishlike any that you find interesting.

C.5 Interview Schedule for Even Participant Numbers

Intro to TYPO3

Videos:

<https://www.youtube.com/playlist?list=PL1wplB0LBgiaNwPGA4aVxp7SdmRZRCSGK>

1. Log into the TYPO3 backend in the first tab of the browser (username: writer, password: writer).
2. Spend a few minutes browsing through the page tree from the Page View, and opening individual pages.
3. Go to Aquarium Fish Accessories - Decoration and do the following:
 - (a) In the description of one of the aquarium products, add at the end Recommended fish: Goldfish, Oscar
 - (b) Add links to the pages of the recommended fish.
 - (c) Save and view the results in the frontend.
4. For 1-3 pages from Fish - Seawater Fish, do the following:
 - (a) Replace the pictures with another picture from the folder Images/new/Seawater.
 - (b) Add 1-2 other pictures from Images/new/Seawater.
 - (c) Save and view the results in the frontend.
5. Go to the News folder (NOT the Fish News) and do the following:
 - (a) Add a new piece of fish news by choosing its contents from the News tab, and add to it any image from TYPO3 from Images/new/News.
 - (b) Edit an additional 1-3 pieces of fish news (e.g. shorten them, add some text, etc.).
 - (c) Save and view the results in the frontend.

Questionnaire time - please fill in the SUS questionnaire about TYPO3

Annotations

Videos:

<https://www.youtube.com/playlist?list=PL1wplB0LBgiYlQihHcE0RDfj1RV6IngQJ>

1. Annotate a random selection of 4-6 fish subpages of the Seawater Fish page:
 - (a) Perform automatic annotation
 - (b) Edit the annotations. For the first page you work on, do all of these steps, for the rest, choose what you would prefer to do.
 - i. Go through the annotations and delete any that are not useful.
 - ii. Select a term which was not already annotated and annotate it. For locations and place names, there is a resource file in the tabs.
 - iii. Select a term that has already been annotated and that has a subregion that can also be annotated, select the subregion, and annotate it.
 - (c) Index the annotations.
 - (d) Save the page.
2. Use Mimir Search to:
 - (a) display all annotations.
 - (b) search for all mentions of entities of type Fish.
 - (c) search for all countries.
3. Go to the News folder (NOT the Fish News), and for any 2-4 of the news items:
 - (a) Perform automatic annotation.
 - (b) Edit the annotations as you did for the Fish subpages, doing 1-3 steps per page.
 - (c) Index the annotations.
 - (d) Save the page.

Questionnaire time - please fill in the SUS questionnaire about the annotation process.

Surfing the Fish Web

You are now done with editing your website, and its time to look through the other websites.

1. Go to the website in the second tab. Look through the freshwater fish, go to the pages of fish you like, and fishlike the pages (choose at least one fish, no other restrictions).
2. Go to the website in the third tab. Look through the seawater fish, go to the pages of fish you like, and fishlike the pages (choose at least one fish, no other restrictions).
3. Go to the website in the fourth tab. Go through the news, read a few pages, and fishlike all news items that you find interesting.

4. Go to the website in the fifth tab. Go through the whole website, explore all pages that are not related to news and fish, and fishlike any that you find interesting.

C.6 List of Themes and Subthemes

C.6.1 Themes and subthemes for the TYPO3 evaluation

1. Notes on the video tutorials
2. User memory of the functionality
3. Technical problems
4. Debriefing conclusions
 - (a) General opinions
 - (b) Comparison with other CMSs
 - (c) Liked best
 - (d) Liked least
5. Functionality problems
 - (a) Technical issues
 - (b) Other wanted features
6. Design problems
 - (a) Interaction problems
 - (b) Navigation problems
 - (c) Presentation problems
 - i. Page layout
 - ii. Button presentation
 - iii. Icon presentation
 - iv. Menu presentation
 - v. Tab presentation
 - vi. Presentation of other interface elements
7. Support problems
 - Absence of a help

C.6.2 Themes and subthemes for the evaluation of the annotation tool

1. Notes on the video tutorials
2. User memory of the functionality
3. Technical problems
4. Debriefing conclusions
 - (a) General opinions
 - (b) Liked best
 - (c) Liked least
 - (d) Change one thing
 - (e) Practical usefulness
 - (f) Automatic annotation
 - i. General opinions/practical usefulness
 - ii. Particularly impressive
 - iii. Missing elements
5. Problems which violate Nielsen's first heuristic: "Visibility of system status"
6. Problems which violate Nielsen's second heuristic: "Match between the system and the real world"
7. Problems which violate Nielsen's fourth heuristic: "Consistency and standards"
 - (a) Consistency problems between the tool and user expectations from other systems
 - (b) Consistency problems between the tool and the rest of TYPO3
 - (c) Consistency problems within the tool
8. Problems which violate Nielsen's fifth heuristic: "Error prevention"
9. Problems which violate Nielsen's sixth heuristic: "Recognition rather than recall"
10. Problems which violate Nielsen's eighth heuristic: "Aesthetic and minimalist design"
 - (a) Contrast problems
 - (b) Alignment problems
 - (c) Proximity problems
 - (d) Other problems
11. Problems which violate Nielsen's ninth heuristic: "Help users recognise, diagnose and recover from errors"
12. Problems which violate Nielsen's tenth heuristic: "Help and documentation"

13. Utility problems

- (a) Unuseful features
- (b) Other wanted features