

Easy on that Trigger Dad: A Study of Long Term Family Photo Retrieval

Steve Whittaker, Ofer Bergman and Paul Clough

Department of Information Studies, University of Sheffield, S10 2TN, UK.

Tel: +44 (0) 114 6340

Fax: +44 (0) 114 278 0300

s.whittaker@sheffield.ac.uk

Abstract We examine the effects of new technologies for digital photography on people's longer term storage and access to collections of personal photos. We report an empirical study of parents' ability to retrieve photos related to salient family events from more than a year ago. Performance was relatively poor with people failing to find almost 40% of pictures. We analyze participants' organizational and access strategies to identify reasons for this poor performance. Possible reasons for retrieval failure include: storing too many pictures, rudimentary organization, use of multiple storage systems, failure to maintain collections and participants' false beliefs about their ability to access photos. We conclude by exploring the technical and theoretical implications of these findings.

Photography, PIM, Long term retrieval, Digital memories

Introduction

Recent technical developments in information capture, storage, retrieval and distribution have led to huge interest in 'digital memories'. It is now possible to capture important digital information ('Lifelogs') about multiple aspects of our lives for later retrieval, including videos, documents, conversations and even medical sensor data (Gemmell et al., 2006; Tan et al., 2007). Of all these digital resources, photographs are usually regarded as the most potent triggers for past memories. Within the family, photographs are cultural artifacts which document events shaping family life, often telling a story about relationships within, and between, family members (Chalfen, 1987; Frohlich, 2004; King, 1986; Williams, 2000). As a result, people make concerted efforts to generate photographic records of important events (even at the occasional cost of disrupting the event or decreasing the participation of the picture-taker in the event itself). One can't imagine a birth, child's birthday, wedding, or graduation without associated photographs. And increasingly in the age of mobile camera phones, even casual social meetings or meals are accompanied by obligatory visual recording (Cohen, 2005).

There has recently been a revolution in the technology of photography, with the increased availability of digital cameras and cheap storage. Pioneering studies of picture technology revealed that, in the past, there was a relatively high cost associated with generating photos, so that people would usually only have small numbers of pictures associated with an event (Frohlich et al., 2002; Rodden and Wood, 2003). In some cases, these developed pictures would be filtered for quality and transferred to albums to showcase specific events or periods in personal life. Even with these relatively small collections, it is worth noting however, that people often felt their collections to be poorly organized, and the task of maintaining them to be onerous (Frohlich et al., 2002).

Such new digital technology may alter the ways in which people capture, organize and access photos. For example, digital pictures are extremely inexpensive to capture and store compared with their analog counterparts. And new technologies are being developed that allow people to automatically cluster and label their photo collections by automatically identifying events and organizing them into hierarchical structures (Graham et al., 2002; Mulhem and Lim, 2003; Platt et al., 2003; Naaman et al., 2004; Cooper et al., 2005; Adams et al., 2006; Chen et al., 2006). Further, automatic processing of visual content has also been used to label pictures using technologies for face recognition and object detection (Suh and Bederson, 2007; Cui et al., 2007). Finally users can verbally annotate pictures or record accompanying audio (Rodden and Wood, 2003, Frohlich 2004). One question, then, is how these new developments will change people's ability to access their pictures in the longer term?

Various studies of digital picture-taking reveal large increases in the number of pictures that are taken and stored compared with analog capture. Although people exercise some quality control by deleting digital pictures both immediately after they are taken, as well as when they are uploaded from camera to computer (Kirk et al., 2006), many more digital pictures end up being stored. For example, Rodden and Wood (2003) found that in just 6 months, people built up digital collections that were already *half the size* of the analog collections they developed over many previous years. These studies of digital behaviors also show that users employ simple organizational schemes when filing pictures; folders are named

with simple labels ('holiday in Tunisia') or combinations of simple labels and dates ('2002-4-London'). With the exception of Macintosh users, who organize with iPhoto, there is also little use of dedicated photo software for organization (Kirk et al., 2006). In some cases, photos are left in temporary folders for sharing with friends and family, with the intention of later organization or annotation, but such planned later re-organization seldom occurs (Frohlich et al., 2002). And sophisticated annotation and search features are used infrequently after initial experimentation (Rodden and Wood, 2003). These studies also paint a consistent picture of retrieval, showing a focus on *recent* photos. Users commonly upload, edit and organize recent pictures in preparation for sharing these with others (Frohlich et al., 2002; Kirk et al., 2006). This might involve deleting poor quality or unattractive pictures, editing (e.g. red eye removal, cropping), and some relabeling of folder names from those provided by the machine, along with filtering to select best exemplars of the event or activity.

Here we revisit the emphasis of prior research on short term retrieval. For example, in two of the early studies (Frohlich et al., 2002, Rodden and Wood, 2003), the technology was novel, and part of the incentive for participation was that users were given their first digital camera. As a result, those studies necessarily focused on small, recently constructed photo collections. However a key value of analog pictures is long term storage and retrieval (Chalfen, 1987, Williams, 2000). The aim of this study is therefore to discover whether people still take digital pictures for the long term, and if so, how effective they are at retrieving older pictures when the low cost of storage causes a sharp increase in collection size. Of course, it has only just become possible to address this question as people have only recently begun to develop long term digital collections as digital cameras only became widely available in the mid 1990's.

Our study looked at picture retrieval in families with young children. Parents with young families make a good study population; they should be motivated to take and organize many pictures to archive family history for themselves and their children (Chalfen, 1987, Petrelli et al., 2008, Stevens et al., 2003, Williams, 2000). Their age means that they are likely to have been exposed to digital photography and to be reasonably adept in their use of computers. We examined

parents' ability to retrieve pictures relating to important past events in family life, e.g. weddings, vacations, birthdays, or social gatherings - precisely those events that previous studies have shown to be a key reason for taking analog pictures (Chalfen, 1987, Frohlich et al., 2002, Williams, 2000).

Specifically we focus on the following issues:

1. *Archival Value*: People invest effort in taking and storing family pictures. Traditionally this was done for long term purposes, but prior research into digital photography has focused mainly on sharing *recent* pictures with friends and family. So how valuable is *long term retrieval* of digital pictures of family events?
2. *Access*: How *successful* are participants in accessing these older digital pictures? Prior research shows people are adept with recent pictures, but how good are they at accessing pictures of important events from further in the past?
3. *Organizational Strategies*: How do participants store and organize their pictures in the long term? Does participants' organization involve filtering and selection? And if so, which selection strategies do they use? Previous work has shown organization involves some early deletion, with photos organized into folders with simple event or event/time labels. Is this true for longer term organization, and how do organizational methods affect retrieval?
4. *Access Strategies*: What strategies do participants use to retrieve digital pictures? Consistent with previous work, do participants rely on accessing picture folders by topic and scanning within these for relevant information, or do they use different retrieval methods for longer term access?

We conclude by discussing the implications of our results for the design of future photo management and retrieval systems, as well as for more general digital capture techniques.

Method

We interviewed parents regarding their digital family picture collections. We first elicited their views on the value of their archives. The interviews included retrieval assignments in which participants were asked to show the interviewer pictures from important past events. They were then asked to reflect on their retrieval process and organizational strategies.

Participants

Our participants were 18 parents of young families who photograph their families (amongst other things) using a digital camera, storing these pictures on their personal computers. Participants were not related one to another, and were all professionals, aged 38-43. In all cases, we chose as participants those family members who were largely responsible for uploading and organizing family pictures. Only seven (39%) of the participants were women. This seems to be a change from analog photo practices where women are mainly responsible for organizing pictures (Hirsh, 1997).

Participants assessed themselves as having ‘medium to high’ computer skills, and used computers as part of their everyday jobs. All but two had used an analog camera since they were teenagers, and a digital one for a range of 2-12 years ($M = 4.86$, $SD = 2.47$). We also asked them when they had last retrieved a family picture that was more than a year old. This turned out to be 39 days ago on average ($SD = 37$). In other words people had looked at older pictures within the last month or so. The majority of participants (13) reported viewing their family pictures using their computer, 3 mainly via a physical picture album, and two used both methods equally. Our focus was on digital picture retrieval, so although all participants printed out a small minority of pictures, we did not examine retrieval of physical pictures or albums.

Procedure

We attempted to interview all participants in their homes, although for their convenience two were interviewed at work, where we asked them to access personal photos from their own laptops. We saw no obvious differences resulting

from conducting home or work-based interviews. Interviews were audio recorded and transcribed. They included three phases.

1. General motivations for photo archiving and selection of significant past events: After gathering general demographic information, we first asked participants *why* they take pictures of family events, and elicited their views about the value of their photo archives. We used a mixture of open-ended and Likert style structured interview questions. Without explaining the subsequent retrieval task, we then asked them to name significant family events from more than a year ago that they had photographed digitally. To avoid having the participants choosing events that they could easily retrieve, this part of the interview took place away from their computer.

2. Retrieval task: After identifying these key events, the interviewer asked the participants to sit at their computer and show him pictures relating to these events. Sample requests were ‘Find me a picture of your son’s birthday’, or ‘Find me a picture of your holiday in Y’. Participants themselves judged whether or not they found these pictures, and it was very obvious from participants’ reactions whether they thought they had been successful or not. The interviewer was careful not to bias the results by suggesting participants moved on to the next task when they found difficulties retrieving pictures. The participant was solely responsible for determining whether the search had failed.

Retrieval time (i.e. the time that elapsed between the request to find a picture and when participants announced that they had found the picture, or gave up on the task) was measured *after* the interview by analyzing the audio recording. We did this measurement post-hoc so as not to apply pressure or stress the participants into feeling that they were being evaluated. Each participant was given 3-5 retrieval tasks amounting to 71 tasks altogether. Task related events included: birthdays (28 tasks), family trips (18), first pictures of a particular child (7), first day at school or kindergarten (4), public holidays and celebrations (5), and other more idiosyncratic events (9). The majority of the tasks (71%) were suggested by the participants who selected the target event themselves as being a significant past event they would like to revisit. This was to imitate, as far as possible, the

situation where people try to locate specific pictures to share with friends and family, either to commemorate significant past family events or to reminisce about the distant past (Chalfen, 1987). In the remaining cases, when the participant failed to spontaneously generate such events, the interviewer suggested standard events based on the age of the children and general knowledge about what family events might have been recorded (e.g. birthdays, first school days, family holidays or trips). These suggestions were all accepted by participants¹. On average the retrieval events occurred 3.1 years before the interview ($SD = 1.57$ years).

3. Reflection Phase: The participants were then asked to reflect upon the retrieval task and evaluate their performance in terms of speed and ease. They were also asked about their storage, selection and retrieval habits. Many of these questions were of a Likert type: the interviewer read out a sentence and the participant chose the extent to which they agreed to it on a scale that ranged from 1 - “strongly disagree” to 5 - “strongly agree”. More open-ended questions were used to determine users’ reactions to their retrieval efforts, as well as their strategies for capturing, organizing and maintaining their picture collections.

In addition to measuring retrieval time for the search task, we analyzed the interview recordings to identify key reasons for participants’ lack of retrieval success, their reflections on how they organized pictures, their retrieval strategies and attitudes to deletion, as well as general views on their archives. We present quotes to represent their views. We also collected representative screen shots to illustrate different retrieval strategies and where possible information about the size of people’s collections.

Results

Do participants value long term retrieval of their family pictures?

Participants were highly interested in long term retrieval of their family pictures. We analyzed the content of responses to the open question: “Why do you take

¹ There were no statistically significant differences in retrieval success ($t(64)=0.42, p>0.05$), or retrieval time ($t(64)=0.56, p>0.05$), between the tasks suggested by the participants and the ones suggested by the interviewer, so we combine results here.

pictures of your children?” Sixteen out of 18 participants (89%) spontaneously generated answers that referred to long term purposes such as: “*It’s important for me that they’ll have [the picture collection] when they are grown up, so they will be able to leave home with a big box of memories. But also for me – to conserve these moments*” (AC), “*I want to document my children, to eternalize them; so that I will always have these pictures and can always look at them*” (LB), and “*I want to reminisce and show my children [the photos] later on*” (SS). Participants were also given statements referring to long-term access and asked to state the degree to which they agreed with them (on a 1-5 Likert scale). Table 1 provides these statements and participants’ responses.

Table 1: Desirability of Long Term Picture Retrieval – where 1 = ‘strongly disagree’ and 5 = ‘strongly agree’.

Evaluation Statement	Evaluation Mean	SD	One-tailed t value	Significance (17 degrees of freedom)
<i>One of the main reasons for me to take pictures is being able to view them years later</i>	4.72	0.46	43.47	p< 0.0001
<i>One of the main reasons for me to take pictures of events related to my kids is being able to view them years later</i>	4.78	0.43	47.38	p< 0.0001
<i>I would like to easily access pictures from specific past events</i>	4.67	0.59	33.32	p< 0.0001

These results clearly indicate that long term retrieval is a major motivation for taking family pictures. All responses were significantly different from a neutral score (3) when evaluated using a one sample t test, indicating strong agreement with the above statements. Furthermore, longer term retrieval seems to occur reasonably often. Although we didn’t collect data about the frequency of long term access, users reported last accessing older pictures (that were more than a year old) an average of 39 days before.

How well do participants succeed with long term retrieval?

In the retrieval task, participants were asked to show the interviewer digital pictures from 3-5 salient past events concerning their children. Results are presented in Table 2.

Table 2: Success and Retrieval Time for the Retrieval Task

	Retrievals <i>N (%)</i>	Retrieval time <i>Mean</i>	Retrieval time <i>SD</i>
Found	43 (61%)	153 sec.	203 sec.
Not Found	28 (39%)	226 sec.	193 sec.
Total	71 (100%)	182 sec.	201 sec.

In contrast to their expectations, our participants were successful in retrieving pictures in only slightly more than half of the retrieval tasks (61%). In the remainder (39%), participants simply could not find pictures of significant family events. All participants were highly motivated to find the relevant pictures, as indicated by their repeated attempts to find pictures as well as their obvious frustration when they failed to do so. They were allowed as much time as they liked to do this, and all unsuccessful searches were voluntarily terminated by the users themselves.

Of the 28 unsuccessful retrieval tasks, 21 (75%) were pictures that the participants believed to be stored on their computer (or on CDs) but which they subsequently could not find. The remaining 7 were pictures participants initially thought were stored digitally, but during the retrieval process changed their minds into thinking were taken with an analog camera. The average time participants took to find the required pictures was about 3 minutes, with an average of about 2.5 minutes for successful retrievals and nearly 4 minutes for unsuccessful ones.

After the retrieval tasks, participants were also asked to evaluate their retrieval experience along 2 dimensions: speed and ease. They were asked to determine on a 1-5 Likert scale (ranging from 1 “strongly disagree” to 5 “strongly agree”) to what extent they agreed with the following 2 statements which were presented separately: “I think that finding pictures was (fast/easy)”. The results are shown in

Table 3, along with a one sample t-test to test their deviation from the mid range score, revealing that they disagreed with both statements.

Table 3 Participants' Evaluation of their Retrieval Task Performance (1 = 'totally disagree' and 5 = 'totally agree').

Evaluation statement:	Mean	SD	One-tailed t value	Significance (17 degrees of freedom)
'I think that the retrieval task was –'				
'Fast'	2.06	1.11	3.61	p<0.01
'Easy'	2.28	1.18	2.6	p<0.05

Moreover, participants spontaneously reflected about the retrieval process using emotionally-laden language: *"Can I say what I think about that search? It was very difficult. I feel [my picture archive] is a very big mess. I have no idea [where things are]. It has no logic. It has nothing. I am full of despair. It is easier to give up on seeing them [the pictures] altogether"* (LB). Other participants commented: *"I feel like a student who failed a test"* (OB) and *"I am dissatisfied with my organization as photos are scattered everywhere"* (RW). The interviewers felt that it would be unethical to leave the participants with these feelings after the interview. They did their best to reassure them that there was nothing out of the ordinary about their particular picture collections, and that we believed this to be a general problem.

How do participants store their pictures and how does this affect retrieval?

Based on participants' comments and behavior during and after search, we now discuss several potential reasons supplied by participants' to account for their unexpectedly poor retrieval performance: too many pictures, distributed storage, no hierarchy, false familiarity, and lack of maintenance.

Too many pictures: The most frequent explanation participants gave for their retrieval difficulties was that they had large numbers of pictures to search. It was difficult for obtain accurate estimates about the exact number of digital photos each person had, due to a lack of organization. Photos were often distributed across multiple storage devices and machines. However we were able to obtain accurate estimates for 4 users, yielding an average of 4475 pictures but with huge

amounts of variation (SD 3039). This restricted dataset made it hard to statistically test the effects of collection size on retrieval performance.

Consistent with previous work (Frohlich et al., 2002, Kirk et al., 2006, Rodden and Wood, 2003), participants felt that they were taking many more digital pictures than they had with analog equipment. All participants pointed to the low cost of capturing large numbers of digital pictures. However, during retrieval they realized that having too many pictures has its price when this mass of pictures competes for their attention, making it hard to locate specific ones. One participant put it in these words: “*Once I used to take a picture or two, now there are 20 pictures for each occasion. All of a sudden you have thousands of pictures because there is no economic constraint. This creates overload. It’s hard to find our favorite pictures. It’s not like it used to be when we had a single album*” (OS). So although there may be 30 digital pictures of a wedding where in the past there might have been only 3, this does not seem to make retrieval easier, as users have to find these 30 from collections of thousands of pictures spread across multiple folders that may each contain hundreds of items.

This is an interesting finding, because, consistent with other research (Kirk et al., 2005), participants all made definite efforts to reduce the overall number of pictures by *filtering* and *negative selection*. For example they deleted poorly focused or unwanted pictures, both when pictures were first taken, as well as at upload. Participants were asked to estimate the percentage of the pictures that they delete on their camera and as they transferred them to their computer. They estimated that they deleted 10% of pictures on the camera (SD = 17%), and 8% on the computer (SD = 13%). This amounted to 17 % altogether (SD = 18%).

Distributed storage: In addition to having many pictures, some participants also struggled with finding pictures from past events because their collections were not all located in one place, or a single consistent filing system. One participant (PH) stored photographs on two separate computers, on CDs/DVDs (as backup) and in physical albums. The same participant noted his use of inconsistent storage organizations across different media - which made re-finding photos hard. To rectify these inconsistencies, he had started to make passes through his archive

(when time-permitted), to organize into new folders, adding tags and a picture title. However, this detailed level of organization required considerable effort and as a result, it only existed for part of his photo collection. Another participant (RW), an IT technician at a local school, had set up a network with a file server in his house, photos being stored in different network drives (as well as folders), making the search task one of locating the correct drive, and then identifying the correct folder. Another participant (PD) stored digital videos and images on five external hard drives (labeled as drives 1 to 5), as well as their computer (with some recent photos still residing on a memory stick waiting to be downloaded). Having to locate the correct hard drive or location made finding photos a major problem: *“it was difficult to find the pictures I wanted because I first had to find the correct hard drive.”*

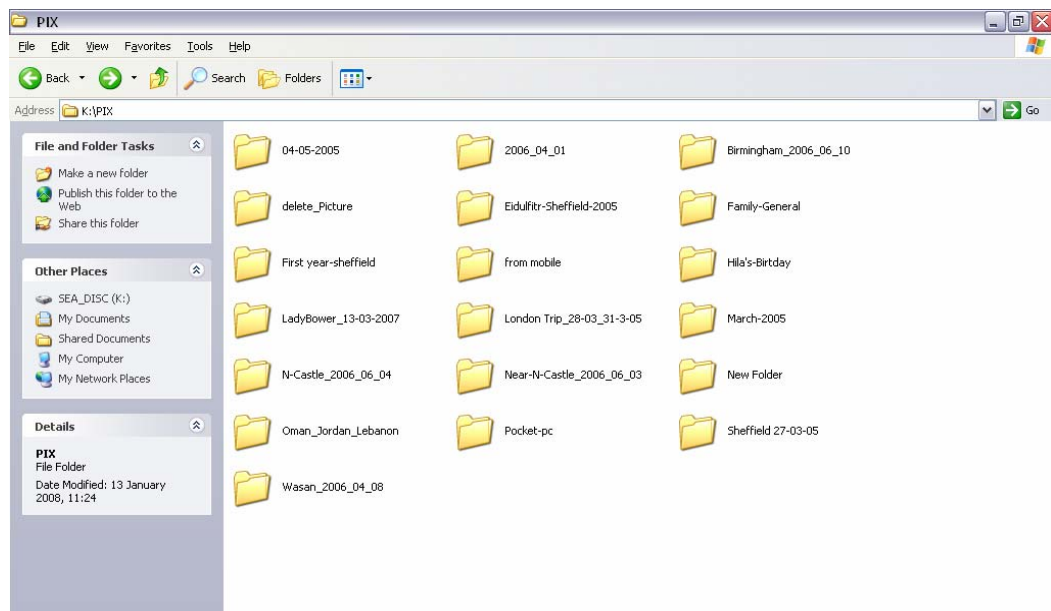


Figure 1: Typical Folder Organisation: note the use of negative selection in the delete_Picture folder, as well as heterogeneous folder names

Minimal Hierarchical Organization: Participants typically relied on a single main picture storage location (such as the “My Pictures” folder). For participants with multiple storage devices (computers and hard drives) there was usually a single main storage location for each device. They usually stored their pictures at that location as multiple folders in a single flat hierarchy with minimal subfolders (see Figs 1 and 2). As a result, when they began scanning inside that main folder, numerous irrelevant folders competed for their attention. Furthermore, a given folder might contain pictures that related to multiple events (possibly because they

were uploaded at the same time). This made identification of the correct folder and picture hard.



Figure 2: Typical Folder Organisational scheme showing heterogeneous folder names

Typical folder structures are shown in Figs 1 and 2. Participant 1 (Figure 1) has no subfolders. But even when participants did use subfolders, they were often inconsistent in how they used them. Participant 2 (Figure 2) uses subfolders occasionally ('famliy' 'freinds', 'scotlan-july 2005'), but not all the time. Furthermore, he is also inconsistent in the *level* of organization that he applies. For example 'mypics' is a subfolder with a label that suggests a superordinate folder covering his entire collection. In addition for both participants: there is a mixture of very general labels ('my old photos', 'from mobile'), place-based labels ('scotland', 'notigham-photos), time-based labels ('2006-04-23'), people

(‘Family and freinds’) and mixed labels (‘chatsworth-8-07-2006’). Note too, that some of these dates are computer generated and based on *upload* time which may not correspond with the time when the events actually took place.

Only three participants constructed an organizational hierarchy that included systematic use of *subfolders*: AC organized her pictures in multiple subfolders within higher level folders labeled with the year in which they were taken; AR organized his picture subfolders within folders representing a time period that started with the printing of pictures from the previous folder and ended with its own printing; OB had an idiosyncratic hierarchy: her family pictures were in a subfolder called “family” (which meant “family 2006” for her) within another folder called “family”, however she seemed to know her way fairly well around it. The 3 participants who systematically used subfolders had a higher proportion of successful retrievals on our long-term retrieval task than those with more rudimentary organization ($t(16) = 2.38, p < 0.05$), although there were no differences in their average retrieval times to access pictures on that task ($t(16) = 1.51, p > 0.10$).

Overall 7/18 (39%) participants had experimented with photo software including Picasa (3 participants), Photoshop (2 participants), Pixier, Kodak album, ACDC, and Google gadgets. Of these, only one participant used software on a regular basis. The others relied instead on operating system folders for organization and access. Some picture organization software (Picasa, Photoshop) automatically organizes users’ folders by time. Overall there was no evidence that experience using photo software led to a greater proportion of successful retrievals on our long-term retrieval task ($t(16) = 0.15, p > 0.05$), indeed there was a suggestion that those people who used a dedicated software program took longer to retrieve their photos on that task, than those who did not use such software ($t(16) = 1.94, p = 0.06$). A possible explanation for this was given by OB: “*Although software like Picasa organizes your pictures for you, the only organization I remember is the one I create*”, adding with a bitter smile “*and when I don’t create it, I can’t remember anything*”.

There are various reasons why time-based representations supported by software may not be enormously helpful. Firstly, these time representations are only accurate if the camera date is correctly set. While 15/18 people (83%) had the correct time on their camera at the time of our study, they pointed out that automatic dating could be misleading: as some programs labeled folders based on *computer* upload date (as opposed to when the picture was actually taken). Further they noted camera dates could be wrong, for example when batteries had been changed or in different time zones. Finally for older pictures, participants were sometimes unable to remember even an approximate date when the picture was taken, so that even accurate system dates would have been of limited use. We return to this issue when we discuss retrieval strategies.

In addition to these general organization schemes, and consistent with previous work (Bergman et al., 2008, Kirk et al, 2006), participants engaged in *positive selection*, identifying favorite pictures to increase their visibility and availability. Eleven out of the 18 participants (61%) reported using positive selection, for an average of about 9% of pictures in total. As the operating system does not offer dedicated support for this, users applied various workarounds to achieve it. They stored their favorite pictures in special folders and then: printed them, emailed them, used them as screen savers, or (in the case of one person) posted them on them as albums on the Internet. However when people retrieved, they focused on their entire collection and not on this favorite subset. OZ was the only one to exploit an explicit software design feature to privilege particular pictures: in Picasa he added a 'star' to his favorite pictures. He was able to see at glance how many 'stars' each folder contained, and also view his favorite pictures from all folders together.

Another potential way to improve retrieval might be to annotate pictures. Consistent with other studies (Rodden and Wood, 2003, Frohlich et al., 2002), we found very little evidence for this. Only two users did any form of annotation. One user who had only recently begun organizing pictures using a computer, tried annotating digital photos using the same method she used for her physical albums: by manually annotating paper lists of her pictures. She admitted that this method "makes no sense" because annotation and picture are stored in two separate

locations, and therefore need to be retrieved separately with nothing to connect them (see The Subjective Context Principle in (Bergman, Beyth-Marom, & Nachmias, 2003, 2008)). Eventually she lost the paper containing her annotations, after which she abandoned this strategy. The other participant was highly experienced, but annotated only intermittently when he had time.

False familiarity: Previous work has highlighted how participants are able to exploit their familiarity with recently taken pictures to quickly scan, sort and organize materials for sharing with others (Frohlich et al., 2002, Kirk et al., 2006). Possibly as a result of these experiences with recent pictures, our participants expected themselves to be very familiar with their entire picture collection. After all, it was their family, and they remembered taking part in the events and taking pictures at those events. They had even downloaded the pictures from their camera to their personal computer themselves. Despite this, however, their attempts to access older pictures indicated they were often unfamiliar with the way their pictures were organized. In most cases, it seemed that they hadn't accessed the vast majority of their pictures since they were uploaded. We saw evidence of this during the retrieval task, when pictures appeared in the "list" view. Participants universally preferred to view pictures in the thumbnail view for easier scanning. Had the participants previously opened these folders, the thumbnail view would have remained at the interview. So while participants may initially have been very familiar with their pictures, this may have decayed over time. However, contrary to our expectations, there was no evidence that the specific age of the target picture affected success on our long-term retrieval task ($r(60) = 0.06$, $p > 0.10$) or the time take to find it in that task ($r(60) = 0.04$, $p > 0.10$). But the absence of a direct correlation may occur because all pictures older than a certain date are equally unfamiliar.

Some participants attempted to account for their poor retrieval by arguing that they hadn't given folders meaningful names. However 67% of participants made efforts to apply meaningful labels instead of relying on software defaults. Such changes did not seem to guarantee they could find their pictures, possibly because naming schemes were inconsistent (see Fig. 2). People who used meaningful labels were neither more successful on the long-term retrieval task ($t(16)=0.28$,

$p > 0.05$), or faster to retrieve pictures on that task ($t(16) = 0.16$, $p > 0.05$). Participants' comments and behaviors also suggested that the meaning of such names was forgotten over time. Finally, participants commented on difficulties in remembering changes over the years in organizational schemes they had enacted or software they had used.

Lack of organization and maintenance: How can we explain this rudimentary organization and false familiarity? Parents typically have very little spare time to organize their photos (even though they would like to do this more). There always seems to be something more urgent to do. Consistent with this, one participant commented that his attitude to photos was “*collect now – organize later – view in the future*” (SS). The difference between home and work information was noteworthy here. Two participants took pictures occasionally as part of their jobs. After the retrieval tasks which of course involved personal data, we asked those two participants how they thought they would fare at retrieving their professional, (as compared with personal), photo collections. Although we did not test their retrieval of work-related pictures, they expected to do much better with professional pictures contrasting the differences in organization and maintenance as follows: “[*with my personal pictures*] *I need to delete the bad pictures, put everything in place and give meaningful names. Look, that’s what I’m doing to my pictures at work. However because I have to do it [organize my personal pictures] in my “spare time”, something that doesn’t really exist, I don’t do it*” (OB).

What strategies did participants use to retrieve their pictures?

Consistent with studies of autobiographical memory (Brewer, 1988; Linton, 1978; Wagenaar, 1986), 6 participants tried to use knowledge of related events to remember the *approximate date* when the target event occurred and then navigate to the folders they thought might contain these pictures. Specific folders were chosen because their name (if there was a meaningful name) contained a date close to the guessed date, or because the name was thought to relate to it. After opening the candidate folder, they changed the view to “thumbnails” and scanned the pictures having first confirmed that they were the right ones. If it was the wrong folder, they navigated to an alternative folder using the above criteria and repeated the process.

Another 2 participants tried to remember the *exact date* when the event had occurred and to find folders from that date. This worked when folders had been labeled with correct dates although in many cases, labels were purely textual. We have already noted problems with this strategy. First participants may be unable to accurately remember the date of the target event. Second the date label itself may be inaccurate either because of problems with camera settings or the folder date represents the upload date as opposed to when the picture was actually taken.

The retrieval strategy for the remaining users seemed to resemble *trial and error*: users would cycle through their entire collection accessing folders to see whether they contained promising pictures, moving on to other folders if they did not.

We also asked people about whether they ever used other ways to find information, such as search. Two participants said they had used the search option to retrieve pictures. Others claimed this to be impossible, as they didn't name individual pictures. However, when during one retrieval task a participant's wife suggested that he should search for the folder's name, both participant (and interviewer) were surprised by the immediate positive results.

Discussion

Much of the user-focused literature on digital photos has looked at people's behavior with relatively small collections of recent events, examining the practices by which people process and share collections with others. While such recent activities are clearly important, here we found strong evidence that *long term archives* are also highly valued. Furthermore, people experience problems in accessing such long term archives, with almost 40% of accesses being unsuccessful. This lack of success seemed to occur for a variety of reasons. Because of the ease of capture and storage, participants now have larger collections of digital pictures. However, these digital photos seem to be organized in a rudimentary manner - arising partly from the time and effort involved in maintaining large collections that are sometimes distributed across multiple storage locations and media. A related factor is false familiarity: participants have a strong (but apparently misguided) belief that their involvement in the initial

events will guarantee that they will be able to successfully retrieve photos relating to those events, without subsequent efforts to systematically organize those pictures. And even when participants worked to generate meaningful labels for their folders, these labels were sometimes forgotten – detracting from their usefulness at retrieval. At the heart of these problems is that, despite their perceived value, participants don't spend much time accessing or maintaining their collections so that organization and access difficulties are often undiscovered.

While the above findings are somewhat different from previous picture studies, some of our other results are consistent with prior work. For example, we found that users tended not to use bespoke picture programs - relying instead on folders provided by the operating system (Kirk et al., 2006). And, as in other studies (Rodden and Wood, 2003), we observed that annotation was highly infrequent - being limited to 2 participants one of whom was a relatively newcomer to digital photography. Participants' failure to return to, and later reorganize what were intended to be temporary organizations, has also been documented previously (Frohlich et al., 2002).

Our findings also link to more general research on Personal Information Management (PIM). We observed problems arising from people's inability to *organize* and *maintain* long term personal collections, as well as their inability to *determine which new information is likely to be of long-term value*. Other PIM research has documented how the low cost of keeping digital information (such as documents, emails and web bookmarks) has implications for how much gets kept (Jones, 2008; Marshall, 2007). As in those PIM studies, we found that people kept large numbers of photos, without entirely anticipating the consequences of large collections for future retrieval. Our participants also had little spare time to filter and organize their pictures, which meant that necessary re-organization was seldom carried out – which is a perennial and well documented problem in PIM (Boardman and Sasse, 2004; Jones, 2008; Whittaker and Sidner, 1996; Whittaker and Hirschberg, 2001). There is a close parallel here with people's web-bookmarking habits. It is well known that people generate large collections of bookmarks, but that these are infrequently used to re-access the web (Tauscher

and Greenberg, 1997; Abrams et al., 1998). And because bookmarks are seldom accessed, users often fail to discover that these are in urgent need of reorganization (Aula et al., 2005).

One key difference between pictures and PIM, however, is that picture collections are less well organised than emails, paper or bookmarks; but seemingly of very high subjective value. There are also major differences between the organization practices observed here, and other aspects of PIM, such as how often people access their stored collection. It was clear from our study that participants accessed picture archives very infrequently. This contrasts with certain parts of the file system which are accessed on a frequent basis, e.g. when people access documents or emails related to a current project (Bellotti et al., 2005; Bergman et al., 2008; Whittaker, 2005). There are two immediate consequences to such frequent access. The first is that participants can often clearly *remember* where frequently visited files are stored making them easier to access (Bergman et al., 2008). Second, this frequent exposure provides participants with opportunities to discover whether their organization is adequate, and to make necessary modifications to an inoperative organization.

There are also various new empirical questions that this study points up. Here, we tested only the parent who organized the pictures. Future research could examine retrieval tasks with other members of the household who are also interested in these collections. We expect that these others will be much less familiar with the picture collection organization and so their success rate will be significantly lower. And when we have digital collections that stretch over decades, future research could compare longer term digital picture retrieval, with the many studies of long term analogue access (Chalfen, 1987, King, 1986). Yet other issues concern people's ability to access and exploit pictures taken using new generation automatic devices, which take pictures when the user moves or there is a change in their environment (Sellen et al., 2007). How will people view, access and retrieve from these new types of collections?

Our results also suggest various design implications. Participants often tried to retrieve pictures based on the approximate time of the event, or by remembering a related event. One interesting link to explore would be between pictures and

calendars to support *event-based retrieval*. Photos might be viewed and accessed in relation to the *activities* the user was carrying out at the time; activities that could be inferred from the user's calendar, allowing users to locate pictures based on what they remember they were doing around the time that the picture was taken. This approach has been shown to be useful in other work (Eldridge et al, 1991; Czerwinski and Horvitz, 2002). Ringel et al. (2003) took a similar approach with their Landmarks interface, integrating representations of personal (calendar) events with public events and linking these to search for desktop files. Users found it beneficial to be able to see their personal information organized on a timeline that had been populated with private and public events (e.g. disasters, public holidays) that served as landmarks for retrieval.

Another much simpler design implication concerns the operating system's "default view." People prefer to browse and scan pictures using thumbnails. However the prevalent operating system we looked at here, Windows XP, does not recognize that a folder contains only pictures, and always presents the default list view. Users changed this default view repeatedly to thumbnails for each picture folder they accessed. A simple system change would be to present folders containing only pictures as thumbnails by default - preventing the need for this.

It was also clear that dedicated photo software didn't seem to benefit most users. Many had briefly experimented with such software and used it for editing pictures (e.g. cropping, red eye removal), but seemed to be overwhelmed by the vast number of features on offer. In addition there were various useful organizational features that seemed to be missing from such programs. Consistent with other work (Frohlich et al., 2002), users wanted ways to identify, select and sort key pictures. All users engaged in *negative selection*: deleting 17% of generally poor quality pictures that they didn't want in the collection. However a more promising approach might be to devise tools that encourage *positive selection* (such as the 'star' feature in Picasa) to privilege important or preferred pictures, enabling users to retrieve from a much smaller collection. In addition, we might be able to infer picture importance automatically based on user actions such as printing emailing, screen saving, web casting, or direct editing (cropping/redeye removal), which all suggest a picture is critical to the user. Software that analyzes such user actions to

rank pictures could help users identify valued parts of their collection (Frohlich et al., 2002). Such action-based techniques have already proved useful for retrieval in settings where single or multiple users access multimedia materials (Bergman et al., 2000, Kalnikaite and Whittaker, 2007, 2008, San Pedro et al., in press).

Our results also suggest a need to further explore techniques that help users organize and maintain picture collections. Combining metadata about *when* and *where* a picture was taken, with low level information about its contents could be used to cluster pictures (Graham et al., 2002; Mulhem and Lim, 2003; Platt et al., 2003; Naaman et al., 2004; Cooper et al., 2005; Adams et al., 2006; Chen et al., 2006). Users could then supply an appropriate event label for that cluster. This could support the retrieval strategies we observed where participants would try to access a particular photo by association, first recalling related events they were confident that they could locate in time. Another very simple modification would be to ensure that automatic date labelling refers to the *capture* date rather than the upload date. Other possibilities involve face recognition and object detection (Suh and Bederson, 2007; Cui et al., 2007) or lightweight annotation (Shneiderman and Kang, 2000, Kang, Bederson, and Suh, 2007). There are three important caveats associated with these technologies however: they must be lightweight, accurate and should be integrated into the current folder structure. Our participants made only rudimentary attempts to organise their collections - showing that if software-aided organization requires extensive effort then people are unlikely to carry it out. Techniques that misclassify information may also exacerbate problems in finding information in already poorly organized collections. Finally, our participants tended not to use any dedicated software to retrieve their pictures. To guarantee wide deployment, future techniques therefore need to be integrated into the existing folder structure, instead of attempting to replace it. An important topic for future research is to explore whether and how new techniques might be harnessed given these user constraints.

At a more general design level, one construal of our findings is that user practices associated with digital pictures have yet to catch up with what the technology offers. In the analog domain, users had smaller numbers of photos which they would share with friends and family, once these had been developed. Albums

might later be created containing favorite pictures, by careful sifting through this relatively small collection. In contrast people take many more digital pictures which tend to be uploaded privately to a computer for storage, and viewed relatively infrequently. As yet there are no equivalent digital practices for the social sharing of a recently developed set of analog prints. New technologies such as digital photo frames have not been embraced by users in part because of the effort involved in setting these up (Petrelli et al., 2008). Furthermore, few of our participants mentioned creating digital albums, although some printed digital pictures to mount in *analog* albums. One area that may be changing is the practice of uploading to the web. Only one of our participants stored pictures on the web and they generally weren't positive about general picture sites such as Flickr, because they perceived these sites to lack privacy controls. However they felt there may be future utility for storage and sharing for web-based sites that support strictly controlled access for friends and family.

Finally, most of our participants (possibly because of the larger numbers of digital pictures they had taken) viewed organization and maintenance of their digital collection to be onerous. If we are to avoid digital archives being *out of sight and out of mind* (Petrelli and Whittaker, 2008), we need to think more carefully about what new tools might allow participants to better share older digital pictures with others and what might motivate them to access and organize their collections. Many of the problems that participants experienced here might have been reduced if they had more exposure to the extent and organization of their collections, exposure that is achieved in the analog domain by the practices of social sharing and album preparation.

Our results are not only relevant to studies of photowork, they have more general implications for lifelogging and PIM. While there has been great recent interest in technologies that make it possible to store vast amounts of personal information (Gemmell et al., 2006, Tan et al., 2007), there has been far less study of the *value* of such archives, or of *people's ability to retrieve such information* from long-term digital stores. Our results suggest that the emphasis that Lifelogging places on *capture* may be misguided. Instead, it seems that participants have problems in accessing, maintaining and using such collections. If such collections are to

realise their potential, we need to focus on new tools that allow participants to filter, evaluate, maintain and share the huge digital collections that they are now accumulating.

One could argue that digital camera users are quite happy with the way they currently view their pictures – opening picture folders almost at random. While accidental finding has its advantages (e.g. finding pictures of a forgotten event), our results indicate that people often use this strategy out of necessity rather than by choice. Our participants clearly indicated they would like to be able to retrieve pictures of specific events, although the price they need to pay (in terms of time and organizational effort) for such controlled access might force them to compromise and instead focus on casual browsing. Moreover, digital photography is rather new. Our participants had owned digital cameras for around 5 years and the events they searched for took place 3 years before on average. As their digital picture collection continues to grow in size, their ability to retrieve pictures of a certain family event might be expected to decrease: both because the users' memory for its location degrades and because each new folder they add distracts them from their target. If we fail to develop effective new tools by the time their children mature, our participants may be able to give them “a big box of memories” but not the key to find specific pictures within this box.

Conclusions

In this paper, we have investigated people's ability to retrieve personal photos related to personal events from more than a year ago to better understand the ways that people store and access photos for long term retrieval. Through an empirical study involving 18 parents of young families, we found that people failed to find almost 40% of pictures in a total of 71 retrieval tasks, despite most participants indicating the importance to them of carrying out such tasks. This work contributes to existing user-based studies of people's photowork by addressing long term retrieval, a subject which has received little attention to date.

Despite recent technical advances in the field, this paper highlights a more fundamental problem related to the way in which people organize their personal information, particularly for long-term access. It is clear that for many people, the

ability to collect more digital information is not matched by a similar ability to organise and maintain such information. We analyzed people's organizational and access strategies and discuss the reasons for this poor performance. These include: storing too many pictures, rudimentary organization, failure to maintain their photo collections and false beliefs about their ability to access photos.

It is clear that technical advances could assist with organization, but only if people are able (and willing) to use them. In further work, we plan to experiment with technologies for improving long term retrieval, developing and testing with a range of users, prototype systems for photo management. Only by developing lightweight, accurate, new tools, will we allow users to regain control of, and access to, their increasingly unwieldy collections.

Acknowledgements

We thank our participants for their time and efforts. This research was funded by the European Union Marie Curie Grant, TOK 30008 (Memoir). The authors are solely responsible for the content of this paper. It does not represent the opinion of supporting funding agencies, and the supporters are not responsible for any use that might be made of data appearing therein.

References

- Abrams D, Baecker R, Chignell M (1998) Information archiving with bookmarks: personal web space construction and organization. In: Proceedings of ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 1998), Los Angeles, California, New York: ACM Press.
- Adams B, Phung D, Venkatesh S (2006) Extraction of social context and application to personal multimedia exploration. In: Proceedings of the 14th Annual ACM international Conference on Multimedia (MULTIMEDIA '06), Santa Barbara, CA, USA, New York: ACM Press, pp 987-996.
- Aula A, Jhaveri N, Kåki M (2005) Information search and re-access strategies of experienced web users. In: Proceedings of WWW 2005, pp 583-592.
- Balabanović M, Chu LL, Wolff GJ (2000) Storytelling with digital photographs. In: Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '00), The Hague, The Netherlands, New York: ACM Press, pp 564-571.
- Baxevanis A (2006) The Journey of a Digital Photograph, MSc dissertation, UCL.
- Bellotti, V.; Ducheneaut, N.; Howard, M. A.; Smith, I. E.; Grinter, R. E. Quality versus quantity: e-mail-centric task management and its relation with overload. *Human-Computer Interaction*. 2005; 20 (1-2): 89-138.
- Bergman, O., Beyth-Marom, R., & Nachmias, R. (2003). The user-subjective approach to personal information management systems. *Journal of the American Society for Information Science and Technology*, 54(9), pp. 872-878.
- Bergman, O., Beyth-Marom, R., & Nachmias, R. (2008). The user-subjective approach to personal information management systems design - evidence and implementations. *Journal of the American Society for Information Science and Technology*, 59(2), pp. 235-246.

Bergman O, Beyth-Marom R, Nachmias R, Gradovitch N, Whittaker S (2008) Improved Search Engines and Navigation Preference in Personal Information Management, *ACM Transactions On Information Systems*, 26(4), pp. 1-24.

Bergman O., Beyth-Marom R., Hadar D., & Dekel A., (2000). From “learning-by-viewing” to “learning-by-doing”: A video annotation educational technology tool. *Proceedings of ED-MEDIA 2000, World Conference on Educational Multimedia Hypermedia & Telecommunications* (pp. 1555-1556). Montreal, Quebec, Canada.

Boardman R, Sasse MA (2004) Stuff goes in the Computer but it doesn't come out, In: *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'04)*, New York: ACM Press, pp 583-590.

Brewer W (1988) Memory for randomly sampled autobiographical events, In: U. Neisser & E. Winograd (Eds.), *Remembering Reconsidered*. New York: Cambridge University Press, pp 21-90.

Chalfen, R. (1987). *Snapshot Versions of Life*. Bowling Green Ohio, Popular Press.

Chen C, Oakes MP, Tait J (2006) Browsing Personal Images Using Episodic Memory (Time + Location), In: *Proceedings of European Conference on Information Retrieval (ECIR 2006)*, pp 362-372.

Cohen KR (2005) What does the photoblog want? *Media, Culture and Society*, 27, pp 883-901.

Cooper M, Foote J, Girgensohn A, Wilcox L (2005) Temporal event clustering for digital photo collections, *ACM Trans. Multimedia Comput. Commun. Appl.*, 1(3), 269-288.

Cui J, Wen F, Xiao R, Tian Y, Tang X (2007) EasyAlbum: an interactive photo annotation system based on face clustering and re-ranking, In: *Proceedings of the ACM SIGCHI conference on Human factors in computing systems (CHI'07)*, San Jose, California, USA, New York: ACM Press, pp 367-376.

Czerwinski M, Horvitz E (2002) Memory for daily computing events. In: *Proceedings of HCI 2002*, pp 230-245.

Eldridge M, Lamming M, Flynn M (1991) Does a video diary help recall? *EuroPARC Technical Report*, No. EPC-1991-124.

Frohlich D, Kuchinsky A, Pering C, Don A, Ariss S (2002) Requirements for photoware, In: *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work (CSCW'02)*, New Orleans, Louisiana, USA, New York: ACM Press, pp 166-175.

Frohlich D (2004) *Audiophotography: Bringing Photos to Life with Sounds (The Computer Supported Cooperative Work Series)*, ISBN 1402023316, Springer-Verlag: Berlin.

Gemmell J, Bell G, Lueder R (2006) MyLifeBits: a personal database for everything, *Communications of the ACM (CACM)*, 49(1), pp 88-95.

Graham A, Garcia-Molina H, Paepcke A, Winograd T (2002) Time as essence for photo browsing through personal digital libraries. In: *Proceedings of the 2nd ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL '02)*, Portland, Oregon, USA, New York: ACM Press, pp 326-335.

Hirsch M (1997). *Family Frames: Photography, Narrative and Postmemory*. Harvard University Press.

Jones W (2008) *Keeping Found Things Found: The Study and Practice of Personal Information Management*. Burlington, MA: Morgan Kaufmann Publishers.

Kalnikaitė, V & Whittaker, S (2007). Software or Wetware? Discovering When and Why People Use Digital Prosthetic Memory. *Proceedings of CHI07*, 71-80, New York, ACM.

- Kalnikaite V., and Whittaker, S. (2008). Social Summarization: Does Social Feedback Improve Access to Speech Data? In Proceedings of Conference on Computer Supported Cooperative Work, 9-12, New York. ACM Press.
- Kang, H., Bederson, B., Suh, B. Capture, Annotate, Browse, Find, Share: Novel Interfaces for Personal Photo Management. *International Journal of Human-Computer Interaction*, volume 23, issue 3, 315-337.
- King, G. (1986). Say "Cheese". The Snapshot As Art And Social History London, Collins.
- Kirk DS, Sellen A, Rother C, Wood K (2006) Understanding Photowork, In: Proceedings of the ACM SIGCHI conference on Human factors in computing systems (CHI'06), Montreal, Canada, New York: ACM Press, pp 761-770.
- Linton M (1978) Real world memory after six years. In: M. Gruneberg, P. Morris, & R. Sykes (Eds), Practical Aspects of Memory. London: Academic Press.
- Marshall CC (2007) How People Manage Personal Information over a Lifetime, In: Personal Information Management (Jones and Teevan, eds.), University of Washington Press, Seattle, Washington, pp 57-75.
- Mulhem, P., Lim, J. H. (2003) Home photo retrieval: Time matters. In: *International Conference on Image and Video Retrieval, Lecture Notes in Computer Science*, vol. 2728, Springer, 308-317.
- Narayanaswami C, Raghunath MT (2004) Expanding the Digital Camera's Reach, *IEEE Computer*, 37(12), pp 65-73.
- Petrelli D, Whittaker S, Brockmeier J (2008) Autotopography: What can Physical Mementos tell us about Digital Memories? In Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'08), New York: ACM Press, pp 53-62.
- Petrelli D, Whittaker S (2008) The Use of Photos in the Home and in the Computer: Similarities and Differences, In: Proceedings of Workshop on "Collocated social practices surrounding photos" at CHI 2008.
- Platt JC, Czerwinski M, Field BA (2003) PhotoTOC: Automatic Clustering for Browsing Personal Photographs, In: Proceedings of the Fourth IEEE Pacific Rim Conference (ICICS-FCM 2003).
- Ringel M, Cutrell E, Dumais S, Horvitz E (2003) Milestones in time: The value of landmarks in retrieving information from personal stores. In: *Proceedings of Interact 2003*, pp 228-235.
- Rodden K, Wood KR (2003) How do people manage their digital photographs? In: Proceedings of the ACM SIGCHI conference on Human factors in computing systems (CHI '03), New York, NY, USA: ACM Press, pp 409-416.
- San Pedro, J., Kalnikaite, V., Whittaker, S. (in press). You Can Play That Again: Exploring Social Redundancy to Derive Highlight Regions in Videos. To appear in *Intelligent User Interfaces*.
- Sellen, A., Fogg, A., Hodges, S., Rother, C. and Wood, K. (2007). Do life-logging technologies support memory for the past? An experimental study using SenseCam. In *Proceedings of CHI '07*, New York ACM Press.
- Shneiderman, B. Kang, H. (2000). Direct Annotation: A Drag-and-Drop Strategy for Labeling Photos. *Proc. International Conference on Information Visualisation (IV2000)*. London, England.
- Stevens, M, Abowd, G., Truong K., Vollmer, F. (2003). Getting into the Living Memory Box: Family archives and holistic design. *Personal and Ubiquitous Computing Technology*, Vol. 7 (3-4), 2003, pp. 210-216.

- Suh B, Bederson BB (2007). Semi-automatic photo annotation strategies using event based clustering and clothing based person recognition. *Interact. Comput.* 19(4), pp 524-544.
- Tauscher L, Greenberg S (1997). How people revisit web pages: empirical findings and implications for the design of history systems, *International Journal of Human-Computer Studies*, 47(1), pp 97-137.
- Tan, D., Berry, E., Czerwinski, M., Bell., G., Gemmell, J., Hodges, S., Kapur, N., Meyers, B., Oliver, N., Robertson, G., and Wood, K. (1997). Save Everything: Supporting Human Memory with a Personal Digital Lifetime Store. In W. Jones and J. Teevan (Eds.), *Personal Information Management*, University. Of Washington Press, Seattle.
- Wagenaar W (1986) My memory: A study of autobiographical memory after six years. *Cognitive Psychology*, 18, pp 225-252.
- Williams C (2000) The meaning of family photographs, [Available online: <http://homepage.mac.com/williamszone/dostal/research/meaning.html>]
- Whittaker, Steve (2005): Supporting Collaborative Task Management in E-mail. In *Human-Computer Interaction*, 20 (1) pp. 49-8888
- Whittaker S, Sidner C (1996) Email Overload: Exploring Personal Information Management of Email. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'96)*, pp 276-283.
- Whittaker S, Hirschberg J (2001) The character, value and management of paper archive. *ACM Transactions on Computer Human Interaction*, 8, pp 150-170.