

Understanding Photowork

David S. Kirk*, Abigail J. Sellen, Carsten Rother and Kenneth R. Wood

The Mixed Reality Laboratory *
University of Nottingham
Nottingham NG7 2RD, UK
dsk@cs.nott.ac.uk

Microsoft Research Cambridge
7 JJ Thomson Ave, Cambridge, UK, CB3 0FB
{Asellen, carrot, krw}@microsoft.com

ABSTRACT

In this paper we introduce the notion of “photowork” as the activities people perform with their digital photos after capture but prior to end use such as sharing. Surprisingly, these processes of reviewing, downloading, organizing, editing, sorting and filing have received little attention in the literature yet they form the context for a large amount of the ‘search’ and ‘browse’ activities so commonly referred to in studies of digital photo software. Through a deeper understanding of photowork using field observation and interviews, we seek to highlight its significance as an interaction practice. At the same time, we discover how “search” as it is usually defined may have much less relevance than new ways of browsing for the design of new digital photo tools, in particular, browsing in support of the photowork activities we describe.

Author Keywords

Photowork, content-based image retrieval, digital photo albums, searching, browsing, use of images

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

The nature of home photography is changing as the use of digital cameras becomes increasingly pervasive. The shift from paper prints to digital images has sparked a number of changes in people’s practices with their personal collections. One is the increase in size of people’s collections as the costs of film and printing no longer apply and the costs of digital storage decrease. Not only are more pictures being taken, but people are taking more pictures of highly similar things such as the same object or scene from a variety of subtly different views [8]. Other changes include the ability to “tinker” with individual images, such as cropping,

altering the color balance, removing red eye and so on, which often results in multiple copies of similar images. Likewise, tools which allow users to stitch together multiple images (to create panoramas, or montages) create related but new versions of images. Added to this is easy duplication which means the same image can exist in many different locations. Taken together, these changes mean more flexibility in our photo practices, but they also mean more complexity in the kinds of things users can do and the resulting organizational structures with which users work.

Not surprisingly, there has been much research into the design of digital tools to help us cope with these growing collections, especially with regard to more intelligent ways of browsing through and searching for images. Browsing is generally understood to refer to ways of viewing large collections of images where the user’s goal is not necessarily well defined. For browsing, the common strategy has been to maximize the use of screen real-estate and present as many images as possible at any given time [1,7]. Search tools, on the other hand, are provided so that users can do more goal-directed finding of particular images. Most search tools have been based on either text-based searching of users’ annotations of pictures [9,16] or searching on the basis of some other form of appended meta-information [11].

In part, because users are reluctant to annotate their images with text [15], some effort has been made to incorporate Content-Based Image Retrieval (CBIR) algorithms into digital photo albums currently on the market, where search tools do automatic analysis of aspects of the image itself. However, this has largely been restricted to color matching functions (as in Microsoft’s Digital Image Suite or Adobe’s Photoshop). The results of such searches are not easily understood by users, as they often expect search results for similar images to be based on semantic similarities. This may be one reason such features have not found wide appeal [14].

Motivation for the Study

Accordingly, the study we describe was initially motivated by a desire to investigate the potential of more advanced computer-vision-based CBIR tools to address users’ needs to search their collections. It quickly became apparent, however, that there was no substantial research in the CHI literature that could put users’ search needs in context. For

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2006, April 22-27, 2006, Montréal, Québec, Canada.
Copyright 2006 ACM 1-59593-178-3/06/0004...\$5.00.

home users, there were no clear answers to the questions: Do users search through their collections? How often? Why? Using what criteria? Without this understanding of people's current practices, we felt we could not deliver any valid assessments of the potential value of new search tools no matter how technically advanced. Likewise, with regard to browsing, there are clear questions as to how it is integrated into common tasks that users actually perform with their personal digital photos. For example, we need to know what kind of browsing users undertake with their collections, and in the context of what tasks. Without this understanding of what people do, it is difficult to test the utility of any new search or browsing tool in a realistic context. Indeed, there are many research studies designed to test new kinds of software for working with image collections, but the common approach is to contrive a searching or browsing task and then to measure the speed at which users can find target images [e.g., 2, 12].

This led us to design and conduct a field study to put people's practices with their own personal collections into context. Here we wanted to focus on the practices between capture and eventual sharing of photos as the primary "end use" of interest. One reason for focusing on practices oriented toward sharing is that there is now substantial work [e.g. 6] describing the importance of sharing as one of the main goals of taking pictures. Further, Frohlich and others have described many different ways in which people share both paper prints and digital images.

However, this emphasis on sharing in previous work has also meant there has been an emphasis on design of technology (or 'photoware') to support sharing using a broad and grounded understanding of these activities, in both analog and digital media, as its foundation [see 3, 6]. While these efforts may well lead to important innovations, likewise there is much to be learned from studies of current practices around the organisation and preparation of photos (especially digital photos) before they are ready to be shared. These have not yet been mapped out in any great detail. Furthermore, we would suggest that it is the process of working with personal photo collections, including the downloading, selecting, organizing, editing, and filing of them which represents the most significant proportion of the work that users of large digital photo collections must do. If this is the case, then it is within this *photowork* that searching and browsing tools may have the most important role to play in future. It is therefore important we develop an in-depth detailed understanding of the nature of this kind of work so we can build on and leverage these existing skills.

Another reason for focusing on users' behaviour with photos after capture and prior to end use is that, as we will describe, much of the management of a user's personal image collection is not necessarily solely for long term archiving, but may be done for a variety of shorter term purposes involving sharing. These various tasks, as we shall see, shape users' behaviour in managing their collections for tasks

they will do in the future. Likewise, their activities with images in the course of these tasks have a legacy effect on their collections. Management is therefore not a simple matter of filing and archiving but, as we will describe, is a more complex matter consisting of many different activities.

FIELD STUDY

For the field study, we conducted a set of interview and observation sessions with twelve home users of digital photos (ten PC users and two Mac users) to understand people's usual practices with their own personal photos from capture of images to the point at which they might share their photos with others. The simple selection criteria for the study were that participants be routine users of digital cameras and that they have at least 1000 digital photos in their personal collections. Our participants were from a variety of both technical and non-technical backgrounds with an even mix of genders and an age range of 19-61. A natural tension of such qualitative research is the small sample size of our participant group. But as appropriate for such research we felt that the depth of understanding generated, and the variety of differences in photo-handling practices observed, helped to balance this concern.

Interviews were conducted either in participants' own homes or at their places of work. The work-based interviews (of which there were four) accommodated the growing number of people who regularly use laptops at work, and who tend to carry their entire digital photo collections around with them. As such, these participants tended to perform the photowork activities that we were interested in whenever they had time, which for some meant spare moments at work.

Our participants used a mixture of laptops and desktop PCs (with one participant using a Media Centre to work with his photos). Only the Mac users in the study used any kind of digital photo album software, although two of our PC-using participants used camera-specific proprietary software for managing the download of their pictures.



Figure 1. A participant in the study being interviewed in front of her laptop.

During the interview session, participants were first asked to download some recent pictures from their camera so that their immediate download and photo-handling practices could be observed. In front of their computers, the participants were then questioned about their photo archiving hab-

its, their editing practices and their photo sharing behaviours (see Figure 1). They were asked questions such as:

- When and why do you delete your pictures?
- How do you file your photos away (using file structure, folder labels, and so on)?
- Do you use digital photo album software? Why or why not?
- Do you edit your pictures, and, if so, in what ways and when?
- When do you look at your pictures?
- Do you share pictures with other people? If yes, why, which pictures and how?

The participants were also exhaustively questioned about everything that they could remember doing with their photo collection over the last 6 months. Every possible form of sharing pictures was probed and prompted. This was done to ensure that participants were adequately thinking about their history of digital photo use, so that they weren't just focused on how they deal and work with photos they had taken recently, but were also considering how they might use older archived photos. This allowed us to get the participants to generate photo sharing tasks which were relevant to them for the next stage of the interview session.

As stated above the next part of the interview session was given over to getting the participants to *engage* in a "pre-sharing" photo task. In this they were encouraged to think of some photos that they would like to share with others and were then asked to go about demonstrating how they would select, organise and prepare these pictures for sharing. This task was selected as a representative kind of "end use" people often engage in with their personal collections [5]. For only a few of the participants did this task involve the pictures downloaded at the beginning of the session. Most users opted to demonstrate a pre-sharing task they had performed on a previous occasion (in some cases many months before).

The collections that we observed people putting together to share with others were quite varied. We had: participants who had been to weddings and wanted to send their pictures to the newly married couple; several participants who wanted to put together collections of their recent holiday photos or photos from a recent trip to show various people (friends and family); and we had one participant who was very keen on putting together a collection of pictures from his university days as a present for friends. For those who were unsure of a theme to choose for the task (three participants) we suggested that they might want to update a relative about what they've been doing for the last six months.

At the end of the session, participants were asked to complete a formal questionnaire which gauged some of their perceptions of potential CBIR tools. However, this part of

the study and its accompanying analysis will be reported in a separate forthcoming paper.

Analysis

From video recordings of the sessions, a detailed observational analysis was made of the behaviours that our participants engaged in when handling their digital pictures. Notes were made detailing the sequential ordering of their activities as we got them to download their pictures and then engage in a suitable task. This allowed us to observe at first hand elements of the 'life-cycle' of photo use they engaged in. The interview questions about their general habits allowed us to probe for reports of behaviour at other unseen stages of the cycle, and elaborate on any other practices they would undertake. These included uncovering whether or not they would normally print their pictures out and if they did, why were they doing it and how often; or asking about how often they shared pictures with other people and getting them to discuss the variety of ways in which they would normally do this, so as to uncover at which points in the lifecycle they would be working with their pictures and in which ways.

UNPACKING PHOTOWORK

Digital photos offer unique opportunities to the user. They offer certain affordances that print photos do not provide. With digital pictures, the opportunity to edit images and to print them changes the traditional role of the user from one as picture "taker" into a picture editor, developer and printer as well. As one of our participants (participant 12) stated in reference to the ability to print her own pictures:

There's a sort of power thing with it, isn't there, like in a sense you control it, however pathetically, you know you've got the paper and you've got the printer, and you can do it, you can do what you like.

This shift in the balance of power, however, means that dealing with images encompasses work with individual photos as well as work with collections, increasing the potential amount of overall work required. The other side effect of this is added complexity in users' collections. Collections are even more likely to increase in size with sets of highly similar pictures as originals and edited copies are retained. This exacerbates the already increased number of pictures in collections as a side effect of the ease and low cost of digital capture. As participant 2 stated:

When you take a picture, you think, well I can afford to take another two or three more in case one's blurred, because it doesn't cost anything to take pictures, so you end up with thousands.

In order to understand how users work with these complexities in the digital realm, we began by considering the 'life-cycle' of the digital photo from its moment of capture to its subsequent use. During this we noted the practices and the kinds of work in which our participants were engaged at various points within the cycle, noting points of commonality as well as contrast amongst them.

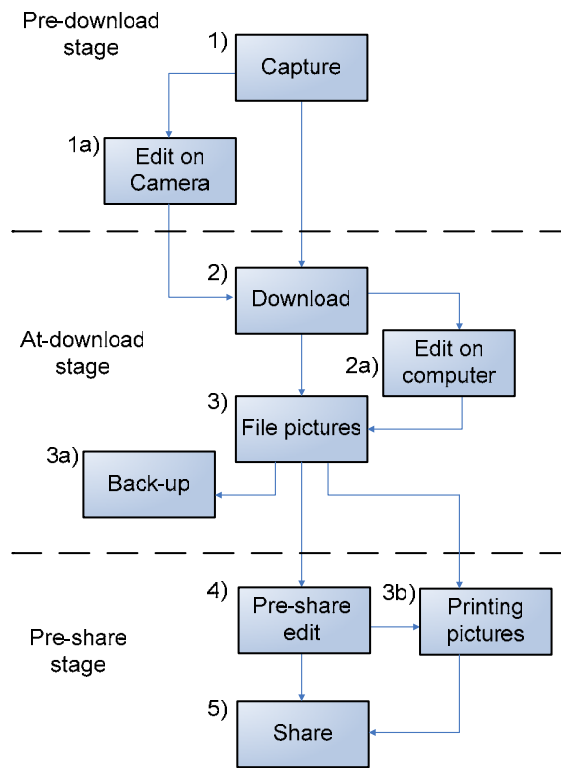


Figure 2. Flow diagram of photowork lifecycle.

Figure 2 illustrates the resulting key activities that are integral to digital photowork. This is further divided into three stages having to do with the capture of images, subsequent download to a computer, and then activities around preparation for sharing. The discussion below strives to articulate how these activities of photowork are actually constituted, considering each stage in turn and the activities within those stages. The discussion then continues with a consideration of some over-arching themes of browsing and searching integral to all of the phases, and ends by drawing out the implications of this analysis of photowork for the design of software to support home users of digital photo collections. We also consider some implications of this work for conducting future research on searching and browsing.

“Pre-Download” Stage

Participants told us that photowork could begin just after the moment of image capture in that they would often engage in a simple editing process on the camera itself (Activity 1a in Figure 2). At that point, seven of our participants said they habitually reviewed and deleted unwanted pictures.

Interviewer: What determines if it gets cut?
 Participant 9: If we look ugly or it’s a bad picture.

Reported motivations for this included a social imperative from others in a picture to delete an unwanted shot, or deleting particularly poor quality shots (such as blurry or out of focus shots). This was a review of pictures just taken.

Equally, however, practices of reviewing and deletion of collections of pictures on the camera were tied to aspects of memory management with memory card size limiting the number of pictures that could be taken. This was especially the case when access to download facilities was restricted, such as when on holiday, as explained by participant 12:

We don’t usually delete when we’re taking them, like if we’re in Australia say, we delete most evenings, because we’ve got a memory stick and we can only take so many.

In these instances the review of pictures went further back and involved the evaluation of all pictures currently on the camera. This editing of collections of pictures did not occur for all users however, especially those for whom storage on the camera was not an issue. When it did, this process was very simple, involving sequential review and deletion of images and little else.

“At Download” Stage

At the time of downloading the images, the kinds of processes participants undertook became richer and more variable. Partly this was a function of why and when people downloaded.

Our participants revealed that there were a variety of reasons for deciding to download images on their cameras onto a computer. For three of our participants, who had large camera storage capacity, the decision was driven by anticipating the amount of time it would take to complete the download, as participant 8 articulated:

But also when I have over 300 pictures, it kind of overwhelms me the time I would spend, so I force myself to deal with them.

For the others however, the decision to download was tied to the nature of the photo capture itself. For many, the capture of images and subsequent download was event-driven. They informed us of how they took pictures at an event and then chose to download their pictures after the event and deal with them in that moment.

Interviewer: What governs when you download?
 Participant 3: When I’ve come back from some particular event or holiday or something like that.

In other instances, participants spoke of needing a specific photo or set of photos for the explicit purpose of passing them to someone else (for example needing pictures for an eBay item). In these instances, the pictures were captured so that they could be downloaded and dealt with straightaway, as described by participant 1, referring to the capture and email of images of a car he wished to sell:

So I grabbed my camera and went out and took some pictures, copied them across in the way that I’ve done and then I emailed them to him.

In sum, then, downloading was either technically-driven, event-driven, or goal-driven.

Editing on the computer (Activity 2a)

When we asked participants to download their photos during our interviews, we observed a variety of practices of filing, selecting and editing. For half of our participants,

downloaded pictures were immediately filed and archived. However, the other half took this opportunity to “work” on their photos by engaging in some form of editing. As the participants reviewed their downloaded pictures, some began the process of reviewing and deleting, this stemming from an acknowledgement that there were often multiple pictures of a highly similar nature, and since not all were needed, only the ‘best’ pictures were retained. All of these participants deleted pictures at this point sometimes based purely on image quality factors (such as blurriness and exposure) but also often based on less concrete properties such as ‘that’s not a very good picture of me’.

Some specific editing of particular pictures also occurred at this point including the re-orientation of pictures to ensure proper viewing, and some ‘red-eye’ corrections, although this tended only to occur where such facilities were made readily accessible (as in the case of the two Mac users). One user also performed more advanced editing at this point by altering a picture’s color balance, but for most it appeared that if this sort of editing occurred, it was left until later as we will describe. Thus the editing and organizing occurring at the download stage can be described as more coarse-grained filtering of pictures, deleting any obvious unwanted pictures, with only occasional work occurring on any particular photo.

Filing on the computer (Activity 3)

Filing and archiving appeared to work along relatively similar lines for most of the participants. This process was heavily influenced by two factors. First, the tendency for digital cameras to store pictures in a ‘folder’ on the camera meant that the process of download became the dragging and dropping of folders between camera and file store. This “pre-packaging” combined with the fact that often pictures in these folders were all of one event meant that the folder was simply re-named to reflect the event, this renaming often conforming to a convention of ‘date - event name’ (as in Figure 3). The second factor was the fact that most users rejected the use of any form of actual digital photo album software. Despite often having such software bundled with the camera when purchased, we saw no evidence of it being used, reasons given including:

- Participant 3: It’s as good as anything else (referring to Windows Explorer)
- Participant 4: It seems like a hassle.
- Participant 6: I don’t know it was just easier; I just use the one that comes with the computer. I just use the most straightforward one.

This meant that users often relied on using Windows Explorer to navigate their picture archive. The obvious exceptions to this were the two Mac users who inevitably used iPhoto, but this decision was clearly more related to the fact that it had been so heavily integrated into the operating system rather than any active choice.

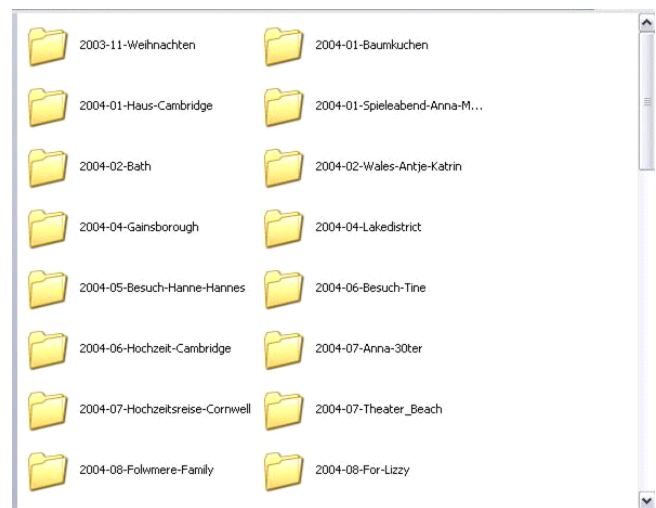


Figure 3. Typical folder structure for a participant.

When participants did have pictures from more than one event on their camera, a common strategy was to copy the folder from the camera over to the file store and name it after the download date, maybe with a broader event name category included. Pictures within the folder would then be split into sub-folders either at the point of download or would be selected later for specific uses in a later review process (which we will return to in the next section).

For many users, the convention of naming folders with date of capture (and with the inclusion of an event related name such as ‘holiday Tunisia’) provided an extremely useful tool for later navigation of the photo archive (as also observed in 17). By grouping pictures in event based folders as a natural part of the downloading process, as a byproduct these pictures would also be grouped according to a common time frame and common relevant locations. Our participants clearly recognized the value of preserving this organizing principle that capture itself and even the simplest software confers. (Likewise, there is evidence that users of paper photo prints do this, preserving prints in the wallets in which they come back from the developers. This same technique preserves aspects of grouping by event, time and location all at once [17].)

Backing-up Pictures (Activity 3a)

Not all of our participants had explicit backing-up strategies. For many, the storage and filing of pictures on their PC remained adequate, for others however there was an explicit desire to take their pictures off their PC and store them somewhere that they felt was safer. For the participants who did this it meant burning pictures onto a CD. One of our participants claimed that the decision to back-up her pictures to CD was based on two key factors, time and PC performance. Her pictures were left archived on her laptop for around 18 months at which point she felt it was of sufficient value to copy them to a CD. This process was combined with what she felt was a decline in performance of her laptop towards the end of the 18 months as the pictures filled up her hard drive. Our participants reported that when

they did back-up they would often again evaluate and delete some pictures before they were committed to the CD, discarding any items that they thought were particularly repetitious or of poor quality.

“Pre-Share” Stage

Having edited and organized photos at download, we then observed another stage during which work was done in preparation for sharing pictures. In our interviews we confirmed that for some participants there could be an imperative to download and immediately prepare pictures for sharing, but for others, once the pictures were downloaded and filed, they were left for a while before undertaking any pre-sharing activities. Participants stated that their preference for either method depended primarily on time constraints or immediate need for the pictures. In fact, because of the interview time constraints, our participants carried out the pre-sharing activities immediately after the download stage.

There are two important observations to make at this point. First, whether our participants said they immediately prepared photos for sharing after downloading, or waited some time, there was a clear point at which they decided to undertake a significant amount of work sorting and preparing photos for sharing. Second, generally speaking, participants reported that the pictures which they used in this process were normally pictures which were recently taken, and the ones most recently downloaded and filed. This has important implications for search and browsing tools as we will discuss.

Our participants spoke of a variety of ways in which they normally shared pictures including viewing pictures on a laptop, viewing on a Media Centre, handing out or posting prints, sending emails containing pictures, burning CDs to send to people and posting pictures to web-pages, either personal blogging pages or public access spaces (such as Flickr), or sometimes both.

Pre-share edit (Activity 4)

Before this process of sharing took place, we noticed several commonly observed editing practices in which all of our participants engaged. This occurred even if such activities had not occurred at other stages, such as at the pre-download or at-download stages. These activities are best described as beginning with a session of photo “triage” where pictures to be shared were sorted and selected from a bigger set of possible pictures. Further, this larger collection of pictures from which pictures were being selected consisted of only a few folders at most (normally the most recently downloaded folders). We observed that participants would look through each folder in turn, would select a picture to share, and copy the picture into a separate folder (leaving the original as it was). Criteria for selection included good image quality, good composition, and more personal preferences, as discussed earlier. We also saw that the natural sequence of the pictures in the “to be shared” folder was often retained as this usually conformed to the time sequence in which they were taken. Occasionally if the

pictures were not in this temporal order (which is a primary structuring device for narration of the pictures), or there was some other pressing reason, the pictures to be shared were re-ordered in a way so as to support some other kind of narrative, for example describing who was at a party rather than the order of events over an evening.

In addition to this triaging process, as pictures were copied into a new folder for sharing, they were often individually edited at the same time. This included re-orienting them (if not already done in a previous edit session), cropping them so as to re-frame the photo and aid composition, and correcting for red-eye (if these facilities were available, which, for ten of our users, they were not). Apart from these kinds of activities, only one participant carried out any more advanced editing, using detailed color balancing on some of his images. However, because of the software they were using to sort their collections, these more technical kinds of jobs often required the user to load the selected pictures into separate picture editing software such as Photoshop, as participant 3 said:

I can do more serious mucking around with Photoshop, for example, to splice pictures together, or if I really wanted to do something serious like edit things out.

For some users (Mac users of iPhoto) there was the option to ‘enhance’ the pictures although it wasn’t entirely clear to the users what this process actually did, so it was rarely bothered with, because of the perceived inconsistent results.

Participant 7: Maybe if I hit the enhance button it might make it look better, sometimes I don’t know if it’s better or not.

Interviewer: Do you know what it’s doing?

Participant 7: Not particularly, no.

The scale of these editing and sorting behaviours was tied to the medium in which they were to be shared. Often if someone only wanted to send a small number of pictures this would prompt the use of email. The process would likely only consist of selecting ‘best’ pictures and re-orienting. If more pictures were being chosen to share with the intention of posting on a web page then more time would be spent selecting and editing those pictures to be posted. There was also usually a stage similar to email wherein the pictures would be suitably compressed for the web format. If a larger number of pictures were to be shared then frequently the CD option was preferred, with users spending equal time editing and sorting pictures but not needing to spend time compressing them. This process of editing / filtering one’s pictures for public consumption was something that participants said they enjoyed. Our participants often commented about the nature of the pre-sharing task that we asked them to engage in, arguing that normally they would spend longer making sure they had exactly the right picture. This is exemplified by participant 12:

Oh that’s not a good one because it’s quite dark. If I’d been on my own I might have taken a bit longer to go through them, to choose some better ones.

A key point to stress is that for triage and detailed editing purposes, most of our participants dealt with a relatively small number of pictures (maybe up to 100 – 150 for any given event) which were stored in only a couple of folders at most. The pictures that were edited for sharing were generally a subset of pictures which had been triaged and copied from the main body of pictures.

Printing pictures (Activity 3b)

Another aspect of photowork that repeatedly appeared in the interviews was reference to the printing of digital photos. Printing had an important relationship to filing, archiving and sharing behaviours. For several of the participants there was a strong desire to print digital photos so that they could be added to existing print albums.

Participant 5: What I'm doing is over there (points to physical photo albums), one to eight and counting, those are my albums, because I prefer to open and look.

Interviewer: Do you normally print your pictures then?

Participant 6: The best ones. Yes.

For some participants this was an ongoing practice but for others this was seen as an eventual desired goal, as participant 12 also stated:

My plan is one day, when I finish all these other things I'm doing, is to print out, to create some photograph albums, from when we've had the digital camera, because we haven't got any albums.

In many cases the decision to print a picture was based on an attempt to share pictures with others who had restricted access to technical resources such as older generation family members.

Again my parents recently said "oh can we have another picture of you?" So ok, print one out and send it off.

Other reasons for printing included the desire to create photo piles that could be left on coffee tables for general perusal (one participant) and also the creation of prints for putting up on wall displays at home (four participants).

Discussion of the practices engaged in when printing pictures demonstrated that similar forms of editing and sorting practice were performed when deciding which pictures to print. Again the triaging of pictures was performed, selecting only the best pictures and removing bad images and repetitious shots from the sequence. And again, further editing of individual photos would be performed. For those printing for wall displays, there was an incentive to enter into an iterative cycle of printing and editing to ensure that color balance and contrast were optimal. However, three of our participants informed us that once the sorting had been performed, the pictures were merely put together into a folder and taken to an automatic printing facility; this was the preferred option for those participants wishing to add prints of digital pictures to existing printed picture albums. The pre-printing activities of home users of digital photo collections therefore appeared to follow similar patterns to the pre-sharing behaviours.

SEARCHING AND BROWSING IN PHOTOWORK

Having elucidated the general practices of photowork, we now consider its relevance to notions of searching and browsing as it is usually understood in image software design. Our primary interest in studying photowork was to understand searching and browsing behaviours in a realistic activity-centred context. As we observed the practices of our participants, discussed their history of photo-handling practices and uncovered the structure of the tasks they performed with their pictures, it became apparent that direct searching per se was not an activity in which home digital photo users *often* engaged. Nor could we discern many opportunities in which participants might *want* to do so, at least within the kinds of tasks that we discussed with them. There are two important points to make here:

- First, when analysing the interviews and asking our participants about the kinds of things they did with their photos, it became apparent that they only infrequently looked back in the past through their collections. Rather, at all of the stages we have outlined, they were more likely to deal with recent images.
- Second, the commonly understood definition of searching for a picture is that this is a goal directed task – there is some picture in mind which is being sought. But our observations and interviews suggest that this was not a common activity. Looking for a particular picture, or even searching for a set of pictures based on some pre-defined criteria simply did not arise as a significant or frequent activity that people did with their home collections.

This is not to say that search activities as they are conventionally defined *never* occur in the course of normal practices with home image collections. Rather it is to say that we saw no strong evidence of such activities in the kinds of tasks we discussed and observed. Furthermore, as our initial motivation had been to study photowork practices with an eye toward the development of better search techniques, this was an issue on which we focussed in our interviews.

On the other hand, one of the most common and time consuming activities we observed was the triaging or sorting of images. This kind of activity was done by considering any one photo against a collection of others (and making decisions about what to keep and what to delete, what to share and what not to share, for example). These processes would seem more relevant to what is typically thought of as browsing than as searching.

Furthermore, the very nature of capture of images with digital photography means that the highly similar pictures which tend to be compared one against the other during sorting are normally ordered next to each other in time series. This already therefore supports an effective way to present pictures for performing these processes. For example, as discussed earlier, triaging and sorting of the pictures generally consists of extracting repetitious shots and shots

of poor image quality, or ones which don't meet personal preference standards. This requires, as stated, the active comparison of multiple pictures in the sequence at any given time.

So how did participants then "find" photos they wanted to use? Here we observed that, when asked to search for a picture to share, in most instances participants could find any given picture just using the date-event naming structure of their files.

Further, with respect to the *structure* of browsing within photowork tasks, for sharing activities, participants said they most often wanted to disseminate pictures from a recent event (e.g., "sharing the photos from our last holiday"). This meant that people were reviewing and browsing pictures from only a few folders of pictures at most. And, within folders, there were only a limited number of pictures, few folders ever containing more than a couple of hundred. So it was rare that our participants felt the need to scan through large numbers of pictures at any given time.

But organising photos around events had more far-reaching consequences. Even in those interviews where a participant was asked to 'create a selection of pictures to update someone about what you've been doing for the last six months' (which by its very nature required the user to browse through multiple folders), there was a common strategy to reduce the complexity of the task. It was observed that users would first think of an event that was relevant and then navigate to that specific folder using the date – event name structure built into their photo archive by their labelling conventions. Having found the relevant folder they would then open that folder and review the pictures within it. Having selected some as suitable for sharing, they would then move onto the next event / folder. This means that the actual 'search-space' or 'browse-space' was actually relatively limited. Again, when placed in context of what people really do, conventional notions of searching and browsing begin to look different.

IMPLICATIONS

In this paper we have introduced and explicated photowork as a key set of interaction practices in which home users of personal digital photo collections engage. On that basis, we have argued that it is difficult to see a natural place within current practice where search tools, in particular, might fit. Whereas some might argue that it is impossible to understand how users *might* work with more advanced searching facilities without deploying such features, we believe that from an in-depth understanding of the existing practices of digital photo users we can uncover ways of introducing new tools and features by leveraging current practices.

The most fundamental implication of our results is that users don't often 'search', at least not in the commonly understood goal-directed target search interpretation of the term, and not in the context of the tasks we studied. Whilst there is clear evidence of the need for those who work with *im-*

personal digital photo libraries to have enhanced photo search capacity [10], it is difficult to find strong evidence of this for people's home collections of photos. Furthermore, while obviously we could not investigate *all* tasks that users do with their home collections, in our interview we were careful to discuss a comprehensive range of things that people do, including the variety of ways in which people share their photos. This was not limited to the tasks we observed but covered their practices more widely through our interview questions. We make no claim that direct target searching never happens. We wish, however, to draw attention to the balance of photo-handling activities and illustrate that it appears that a significantly larger proportion of time is spent in browsing-like filtering and triaging activities than in search activities.

Such considerations cast significant doubt on the ecological validity of the raft of user studies of photo browsing and search tools which assess usability on the basis of the speed at which users can find a target image in a directed search task. Such directed search tasks are simply not representative of the kinds of tasks that we have observed occurring within photowork.

Another implication is that even though we might characterise what people do during photowork more as browsing than searching, the term "browsing" is rather general and thus perhaps not particularly useful as a way of describing what users do with images during these activities. Browsing implies casual "looking" in the context of ill-defined and changing goals. But we have seen that this looking and reviewing is often goal-directed, in the service of the range of activities we have described. The key point here is that many of these activities depend on the scanning and viewing of collections or sub-collections of images where the ways in which those collections are organised is consequential and important.

On a more constructive note, we can consider how highlighting and describing the nature of photowork can suggest ways in which software tools might move forward. Photowork is both time-consuming and complex, and would therefore seem an ideal place where advancements can be made to ease the management of increasingly large photo collections. On the basis of these findings, we suggest that:

- Intelligent ways to support search, such as the ability to search by content, might be better deployed as tools to help users cluster and view large collections of images. Consider, for example, an algorithm which supports search for particular classes of objects (cars, trees, people and so on). Rather than implementing a "search by similar object" feature, such an algorithm could support new ways of viewing collections clustering by object type. Likewise, intelligent "search" tools which look for a particular object, presence of people, similar layouts, similar scenes and so on, could instead allow users to see their collections in new ways through

filtering and grouping along different criteria. This could support the “narrowing down” of collections into sub-groups which then becomes the basis of the more focussed browsing we have seen users do.

- The frequency with which users spent time sorting photos suggests that features to better support these processes may be central to the value of digital image tools. For some users, sorting occurred at up to five points during the life cycle: on the camera after capture, at download, prior to sharing, prior to long term back-up, and prior to printing. This process in its various forms is thus important to understand and design for.
- For example, coarse-grained filtering of photos “to keep” as against those “to delete” happened at all stages. Here, computer vision techniques helping to isolate and cluster poor quality (as in blurry photos) images might considerably speed up and facilitate this process. Likewise tools which might be able to carry out other technical assessments using computer vision techniques (such as automatically highlighting the best technical compositions amongst similar images) might help speed and enrich this process.
- However, sorting in general will mainly be in the user’s hands, suggesting that new techniques offering both flexibility and speed will be of very high value. It is in this area of rapid sorting of pictures or triaging that software could make important contributions. Here, the ability to triage for different user-defined purposes would be particularly useful. Some nice examples of such tools have emerged from the research world within the last few years [4, 9]. Further, the automatic clustering of similar images in support of a first pass through a large collection may support the early filtering process before more fine-grained decisions are made. More generally, clustering similar images may be a more effective way to support browsing in general by displaying only representative images from a cluster when doing large-scale browsing.
- Again in support of sorting, and looking forward into the future, techniques which move beyond the desktop PC to new forms of input and new types of display may find their place. Touch-sensitive tabletop displays seem an ideal context for doing this kind of photowork. In addition, the use of two hands and multiple-touch input on tabletop displays to carry out quick paper-like sorting seems a potential valuable direction in which to pursue new interaction techniques to support photowork.
- Another finding from this research is the fact that the point at which users are most likely to carry

out extensive editing work on individual pictures is just prior to sharing. By its very nature this task is time-consuming, so tools which make it easier and more enjoyable will be valuable. However, we have seen that over-automating features also confuses users. One potential middle-ground is for automatic editing tools to offer up a series of potential “solutions” to a user and to let them choose. For example, an “auto-cropping” tool might be applied to an image, but would present several possible alternatives to a user to select from. This may also give them a better intuitive understanding of what that feature does.

- Related to this issue, the work that goes into preparing sub-collections of images for specific purposes (be they for sharing or not) may well be converted into meta-information that users might want access to. So, for example, they may want to see at a glance all the images they sent to a friend the month before in order to send to another friend. This is similar in notion to the ‘mini-reprint albums’ suggested by Frohlich et al [6]. Or users may want to know which images they posted to the web over the years in order to create an album of favourite photos. These, after all, they may have spent considerable time on in terms of editing and polishing. Thus software tools could leverage the photowork that people undertake for one end purpose to later support them in another, for example by allowing them to filter and view groups of images using tags added automatically during earlier photowork sessions. Interestingly there is already some hardware support for this with some digital cameras allowing the instant tagging of images for specific later use – increasing the potential for on-camera filtering.
- Finally, we have seen that the participants in our study used simple organising principles for their filing of photos, mainly driven by time and event, and relied heavily on these in searching and browsing images. The fact that users do this naturally casts doubt on the need for photo grouping algorithms which are being developed to cluster pictures using GPS data (Naaman et al 2004) as grouping by time and event also tends to group by location as a by-product. But more positively, it provides confirming evidence of the power of both time and event as a way to support effective navigation within large image collections. This fact is beginning to be exploited not only within image collections [7] but in more extensive collections of personal media [13].

In summary, we have drawn out a number of implications to do with people’s use of their *own* personal digital image collections. We must add that there may be many other opportunities which arise for intelligent searching and

browsing tools when considering use of *other people's* photos. For example, tagging images with time or location may well be useful when a friend gives you a set of photos from a recent trip. In addition, people often collectively own image collections (such as in families). Here, searching and browsing tools might have quite different requirements when idiosyncratic filing systems may be less appropriate. These issues are clearly for future research.

A final qualification is that, while we have found no strong evidence for more targeted searching or searching further back in the past in photowork, there may indeed be interesting and compelling opportunities within other kinds of uses for digital images. Here, we can think of intelligent screen savers, the automatic display of images in digital picture frames and slide shows where older images may be sought, and where intelligent search tools may well find their place.

CONCLUSION

In this paper we have attempted to bring to the fore a set of activities we call photowork and to derive a set of design implications from a deeper understanding of what this means. While photowork is best understood in terms of the life cycle of photo management after capture and before the "end use" of images, we have sought to highlight its significance both as a series of effortful and complex practices, but also as practices in service of many kinds of end goals, especially sharing. As such, we hope that this descriptive framework gives a more realistic context against which the potential value of new digital photo management tools can be assessed.

ACKNOWLEDGMENTS

Many thanks to all the people who participated in this project. Thanks also to David Frohlich whose expertise in this area and advice on this paper are much appreciated.

REFERENCES

1. Bederson, B. PhotoMesa: a zoomable image browser using quantum treemaps and bubblemaps. In *Proc. UIST 2001*, (2001), 71–80.
2. Cox, I.J., Miller, M.L., Minka, T.P., Papathornas, T.V., Yianilos, P.N. The Bayesian Image Retrieval System, PicHunter: Theory, Implementation, and Psychophysical Experiments. *IEEE Tran. On Image Processing*, 9, 1, (2000), 20-37.
3. Crabtree, A., Rodden, T. & Mariani, J. Collaborating Around Collections: Informing the Continued Development of Photoware. In *Proc. of CSCW 2004*. ACM Press (2004), 396-405
4. Drucker, S., Wong, C., Roseway, A., Glenner, S., & De Mar, S. *Photo-triage: Rapidly annotating your digital photographs* Microsoft Research Technical Report, MSR-TR-2003-99, December 2003.
5. Frohlich, D. M. *Audiophotography: Bringing photos to life with sounds*. Springer (2004).
6. Frohlich, D. M. Kuchinsky, C. Pering, A. Don, and S. Ariss. Requirements for Photoware. In *Proc. CSCW 2002*, (2002), 166–175.
7. Huynh, D. F., Drucker, S. M., Baudisch, P. and Wong, C. Time Quilt: Scaling up Zoomable Photo Browsers for Large, Unstructured Photo Collections. In *Proc. CHI 2005*, ACM Press (2005),1937-1940.
8. Jaimes, A., Chang, S. and Loui, A. C. Detection of Non-Identical Duplicate Consumer Photographs. *Proc. ICICS – PCM 2003* 16-20.
9. Kang, H., Shneiderman, B., .Visualization Methods for Personal Photo Libraries: Browsing and Searching in the PhotoFinder. In *Proceedings of IEEE International Conference on Multimedia and Expo* (2000).
10. Markkula, M., and Sormunen, E. End-user searching challenges indexing practices in the digital newspaper photo archive. *Information Retrieval*, 1(4), (2000), 259–285.
11. Naaman, M., Harada, S., Wang, Q. Y., Garcia-Molina, H. & Paepcke, A. Context Data in Geo-Referenced Digital Photo Collections. In *Proc. of Multi Media 2004*. ACM Press (2004), 196-203
12. Platt, J. C., M. Czerwinski, and B. A. Field. *PhotoTOC: Automatic Clustering for Browsing Personal Photographs*. Technical Report MSR-TR-2002-17, Microsoft Research, 2002.
13. Ringel, M., Cutrell, E., Dumais, S., and Horvitz, E. Milestones in Time: The Value of Landmarks in Retrieving Information from Personal Stores. In *Proc. of Interact.*, (2003).
14. Rodden, K. and Wood, K. Does Organisation by Similarity Assist Image Browsing? In *Proc. CHI 2002*, ACM Press (2002), 190-197
15. Rodden, K. and K. Wood. How Do People Manage Their Digital Photographs? In *Proc. CHI 2003*, ACM Press (2003), 409–416.
16. Shneiderman, B., and Kang, H. Direct annotation: A drag-and-drop strategy for labeling photos. In *Proceedings of the International Conference on Information Visualisation*, IEEE, (2000), 88–95.
17. Vroegindeweij, S. *My pictures: Informal image collections*. HP Labs Technical Report No. HPL-2002-72R1, (2002).