

Lost in Memories: Interacting With Large Photo Collections on PDAs

Susumu Harada, Mor Naaman, Yee Jiun Song, †QianYing Wang, Andreas Paepcke

Stanford University

{harada, mor, yeejiun, paepcke}@cs.stanford.edu †wangqy@stanford.edu

ABSTRACT

We developed two browsers to support large photo collections on PDAs. Our first browser uses a traditional, folder-based layout that utilizes either the user's manually created organization structure, or a system-generated structure. Our second browser uses a novel interface that is based on a vertical, zoomable timeline. This Timeline browser does not require users to organize their photos, but instead, relies solely on system-generated structure. Our system creates a hierarchical structure of the user's photos by applying time-based clustering to identify subsets of photos that are likely to be related. In a user experiment, we compared users' searching and browsing performance across these browsers, using each user's photo collection. Photo collection sizes varied between 500 and 3000 photographs. Our results show that our timeline browser is at least as effective for searching and browsing tasks as a traditional browser that requires users to manually organize their photos.

Keywords

Handheld devices, mobile computing, pen and tactile input, photo browser

INTRODUCTION

As the computational power and storage capacity of PDAs increase, photo browsers are emerging as a feasible and important application for these devices. Like calendars, photographs are most useful when available at a moment's notice. Personal photographs in particular are often shown spontaneously during the course of conversations with friends and family ('co-present sharing', [3]). Effective search capabilities are essential in these situations, in addition to instant access to all of one's photos, as the need for any particular photograph arises unexpectedly. This fact poses an interface challenge: it is difficult to design a system that allows instantaneous search and retrieval of

photographs, particularly on a device that has limited screen space. This problem is exacerbated by users' increasing capacity to create digital photographs, which results in rapidly growing photo collections.

We explore one avenue towards solving the problem of searching and browsing through large collections of personal photographs on a small screen. In this paper we are not concerned with the problem of transferring photographs to the PDA, but instead focus on the interface issues that arise, assuming that all of a user's photographs are already available on the PDA.

Interacting with personal photograph collections on a small display presents a vexing problem. Most promising approaches for compensating for the lack of screen real estate involve the use of textual metadata, such as the location at which photos were taken, and what they depict. However, casual photographers often lack the time and inclination to create such textual metadata. Image analysis techniques that generate such data are still immature. In addition, the entry of search terms on PDAs is cumbersome at best, making the interaction with metadata frustrating.

Our approach to this problem involves the extraction and processing of metadata from the headers that digital cameras prepend to the image files they generate. In [5], Graham et al. explored the use of photo creation time as a foundation for desktop browsers that arrange and manage photos automatically. We report now on our experience with the usage of time as an organizational principle on PDAs. As is usually the case, even if underlying algorithms transfer in part to the PDA, user interfaces need to be rethought from scratch.

In this work we limit ourselves to collections of personal photos; that is, images that were photographed or collected by the person who then interacts with the images on the PDA. Chronology has been shown to be a very important factor in users' interactions with their own photo collections [11]. Furthermore, Graham et al. confirmed in [5] that single-photographer photo creation time can be related to events by well-known mathematical methods, which we summarize below. Work in [3,11] confirms that events are users' natural way of thinking about their photos.

Time, of course, is not the only, and reportedly not even the most powerful memory cue. Stronger memory-to-image associations seem to be, in order of power, the factors



Figure 1. Opening-game screen for the Timeline browser

‘who’ and ‘where’ [14]. In addition to time references, then, visual clues in the form of representative photos must have a pervasive screen presence as well.

On these principles, we are constructing a series of time-informed interfaces, the first of which we call the Timeline browser (Figure 1). We compared this novel interface to a baseline of more traditional, folder-based interfaces (see, for example, [8,10]). To this end, we pitted the Timeline browser with its automatically generated hierarchical structure against a traditional browser under two conditions. For the first condition, we had subjects operate the traditional browser over their own photo collections as they had organized them on their personal computers. For the second condition, we replaced the owner’s storage organization with one that was created by the clustering algorithm. This allows us to test the efficacy of both our automatic organization and the time-informed interface.

The next section provides a short summary of the clustering algorithm used. For more detail, see [5].

Clustering Algorithm

Suppose, for example, that a person’s photo-taking activities during a three-month period exhibit the following pattern. She first takes pictures during a one-week trip to California. This is then followed by five days when no pictures are taken. That period of inactivity is followed by a flurry of photos during her daughter’s birthday.

The clustering algorithm would partition the resulting collection into two clusters, A and B. Cluster A would contain the trip photos, and cluster B would comprise the birthday shots. Within each cluster, picture taking will usually again be ‘lumpy,’ this time at a finer time granularity than ‘trip’ and ‘birthday.’ The algorithm will therefore recursively decompose A and B into smaller clusters. The California trip might be partitioned into three sub-clusters that include, respectively, images taken on a one day visit to San Francisco, a series of photos shot during a three-day road trip, and photos of the last day’s travel home. All this clustering is based purely on the time when images were taken, not on image recognition.

More technically, the initial clusters are created by detecting gaps of more than 24 hours between two consecutive photos. The intuition behind this first pass is that picture-taking events are usually more than 24 hours apart from each other. Then, the photos in each cluster are examined further to detect “outliers” - irregular gaps between two consecutive photos. When such an outlier is found, the cluster is divided further at the detected time gap. This process is applied recursively for each cluster until each cluster holds less than 30 photos and spans less than 8 hours. Thirty is, not coincidentally, the number of thumbnails we can fit on one PDA screen.

Of course, this algorithm is based on heuristics and is not perfect. If the birthday event took place immediately upon the return from California, the two will fall into the same high-level cluster. However, in most cases the clustering algorithm will detect these as distinct events when analyzing the cluster during the next step of the algorithm.

The Timeline browser we report on here is based on the above algorithm. We included in the Timeline browser interface only time-related features, because we wished to test these facilities before proceeding to add other useful machinery, such as auto-rotation, continuous scaling, or cluster summarization.

The following sections (i) introduce the three browsers, (ii) describe the goals and design of our experiment (iii) present and discuss results, (iv) provide leads to related work, and (v) conclude the paper.

BROWSERS

As mentioned above, we implemented two browser interfaces: the traditional Baseline browser, and a novel Timeline browser (TL). The Baseline browser was implemented in two variants, which we call *Baseline Manual* (BM) and *Baseline Automatic* (BA).

All the browsers offer the user three views that correspond to three main phases of search, as used in [2,15]: an ‘opening game’ during which the user navigates to the general vicinity of the target photo, a ‘middle game’ where the search is narrowed, and an ‘end game’ where users examine an individual photo.



Figure 2. Opening-game screen for the Baseline Manual browser

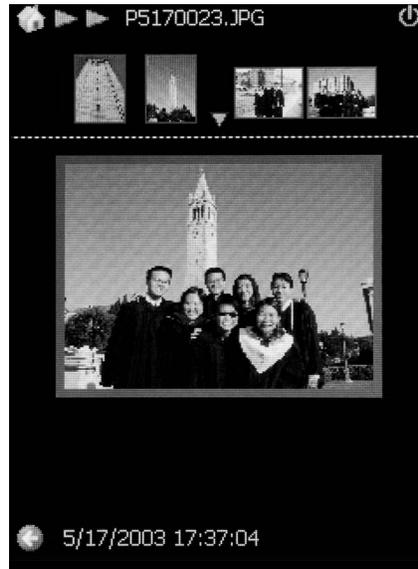


Figure 3. End-game ("Preview") screen for the Baseline browsers



Figure 4. Middle-game ("Thumbnail") screen for the Timeline browser

The Baseline Browser

The Baseline browsers are designed in a fashion similar to folder-based photo browsers available today for the PDA [8,10]. BM is based on the user's own organization, while BA is based on system-generated organization. We will present the BM interface briefly, and then describe the difference in photo organization between BM and BA and its effect on the interface.

The opening game for BM, shown in Figure 2, lists all folders and photos under the current folder. Instead of showing a folder icon, each folder is represented as a button control (the first three rows in Figure 2), each consisting of three sample photos, the name of the folder, number of subfolders (if any), and the total number of photos contained in this folder. Tapping on a folder will drill down into it, showing its subfolders and photos in the same fashion. Below the rows of the folders, we display a grid of thumbnails of all photos within the current folder.

At all times, the top of the screen displays a series of red arrowheads to indicate how deep the current folder is situated in the hierarchy. The Back and Home buttons carry the usual meaning of climbing one level up and returning to the top of the hierarchy respectively.

In the BM browser, the middle game is just a special case of the opening game, where there are no more subfolders in the current folder. In this case, all the thumbnails of the images in the folder are presented in a grid. The user can scroll down the grid if it spans more than one screen.

At any point, tapping on one of the photos in the grid will take the user into the end game (Figure 3). In this screen ("Preview Screen"), the selected photo is enlarged. Four other photos appear in a row at the top. These are the next and previous two photos, included for context and quicker navigation. Tapping any of the four replaces the enlarged

photo with an enlargement of the tapped photo. The context changes accordingly. Tapping on the enlarged photo cause that photo to fill the entire screen. Another tap will return the user to the Preview Screen.

Baseline Automatic (BA) is a variation on BM that uses the system-generated structure instead of the user's own photo organization. There are just a few differences between the two browsers. First, in BA the photos appear only in the lowest level of the hierarchy. Therefore, we have a distinct opening game where users only choose between subfolders to drill down into. When there are no further subfolders (middle game), users see a grid of photos in the current folder. The second difference between BA and BM is the lack of folder names in the opening game. Instead we show the time span of the folder (the times when the first and last photos in the folder were taken).

The Timeline Browser

The Timeline browser (TL) presents a novel photo-browsing interface that relies heavily on the system-generated hierarchical organization described above. There are three main concepts behind the Timeline design. First, we enhance the role of time in browsing of the collection by providing the user with maximal time context. Second, we reduce the number of photos shown on the screen, using the additional screen space to provide enhanced time orientation instead. Third, we do not overload the screen, deliberately preserving as much negative space as we can. Again, we hope this allows the user to focus on the time context and the few photos as clues. At least one study [11] concludes that on desktops, browsers should display as many thumbnails as possible. We depart from this advice in our PDA browser on the basis of page layout principles that warn of clutter. Our sense is that clutter is at least as detracting on a small screen as, for example, on print media. Our experiments were not designed to answer this

question, and we will investigate the issue further in future work.

The opening game for the Timeline browser is shown in Figure 1. The screen is partitioned into three columns. In the middle is a vertical timeline, and to the right and left of the timeline is one column of pictures each. Each picture represents a system-generated cluster (“album”). The time range that an album covers is shown next to the picture. You also see how many photos are in the album. For a quick overview there is also a row of dots - the more dots, the more pictures are in the album.

The timeline in Figure 1 spans 4 months (the application actually starts in full year view). At any screen, we display the time range that is covered by this screen at the top. There are *never* more than 10 albums shown on the screen. If the clustering algorithm generates more than 10 albums, we merge the closest ones until we have 10 albums. Again, keeping a limited number of clusters is geared towards maximizing the time context, trading it off with narrower representation of albums and photos on the screen.

A white line runs from each album to the timeline. The point at which the line touches the timeline corresponds to the time when the photos in that album were taken. If there is a thick vertical line at the point of their intersection, the height of the vertical line indicates the time span during which the photos in the corresponding album were taken. For example, you can easily see the top left album in Figure 1 spans the first few days of April.

Tapping on an album will drill into a finer grained time range – the best range (in integral number of month, weeks or days) this album fits in. We are still in the opening game, only now we see the lower-level albums inside the tapped album and possibly other adjacent albums that fall into this time range, only the former are highlighted while the latter are not. Another way to display a finer grained time range is to drag the stylus on the timeline to select the desired time range. The selected time range will expand to fill the screen, showing in more detail the albums contained in it.

We implemented a way to quickly “peek” into the photos in an album during the opening game. The user can hold the stylus on a photo, and then start circling the stylus around the photo. This allows the user to flip in place through all the photos in that album. Users can circle forwards or backwards, by moving the stylus in a clockwise or counterclockwise direction.

Eventually, the user navigates and taps on a “leaf album” that contains just photos. Hopefully, she had done so when getting very close to the photo she was searching for. The final tap takes the application into the middle game. In the middle game (Figure 4), the user is presented with a grid of thumbnails. The thumbnails of the tapped album are aligned on the top, and are highlighted. From this point, the user can navigate back and forth (using the arrow icons at the bottom of the screen) through thumbnails of her entire

collection, ordered by time. To aid in the navigation, dates are displayed above the thumbnails wherever the date changes. The range of dates represented on the screen appears on the top of the screen; the current position within the entire collection appears on the bottom.

There is also a shortcut that allows the user to switch quickly from the opening game into the middle game. From any level of the hierarchy in the opening game, tapping the “thumbnail” icon (second button from the left on the bottom row of Figure 1) will take the user to the middle game, showing thumbnails beginning from the start date of the previous screen.

We decided not to implement the end game for the Timeline browser as we thought it might not be required. Thus, in the TL middle game, tapping a photo will simply cause it to fill the entire screen. Tapping again will return the user to the middle game view of thumbnails.

DESIGN OF EXPERIMENT

We recruited 17 subjects for our experiment. The most important criterion for subject selection was that each subject needed to have shot and stored a sizeable number of digital photographs. Each subject interacted with her own photographs throughout the experiment. Subjects’ ages ranged from 19 to 48, with the highest representation in the 20s. Twelve subjects were male, and five were female. One of the photo collections comprised of 474 photographs, but all others were significantly larger. Nine exceeded 1,000 photographs, and one was 28 photos shy of 3,000. The average collection size was 1,200 images.

We loaded each subject’s existing collection onto a Hewlett-Packard H5500 Pocket PC. Its relevant hardware specifications are 128MB of RAM, a 400MHz Intel XScale processor, and a 240x320/64K-16 bit (65K) color display with a viewing area of 57.6mm x 76.8mm (WxH). The operating system was Microsoft Windows Mobile 2003.

We reduced all photos to thumbnails such that the long edge comprised 40 pixels; the short edge was generally 30 pixels long. The large image in the Baseline browser Preview Screen (Figure 3) was a 190x143 rendition.

Our experiment followed a within-subject design. We exposed each subject to three experimental conditions: the Baseline browser with the subject’s own, manually arranged photo organization (BM), the Baseline browser with the time cluster based organization that our system generated automatically (BA), and the Timeline browser (TL), also with the system generated organization.

Each subject completed two tasks on each browser. The first was a Search Task. We showed the subject one of her own photos on a laptop screen and asked her to find that photograph in her collection by navigating on the PDA. We set a three-minute time limit for this task, and asked subjects to work as efficiently as they could.

The second task was a Browsing Task. We asked the subject to select ‘good’ pictures for a collage that represented some portion of the subject’s life. For example, we asked for pictures that would make a collage of the subject’s friends. Other collage assignments were family, trips, special events, and scenery. We asked subjects to select photos from as broad a time span and set of occasions as possible. For this task we imposed a time limit of two minutes.

For each browser we had subjects complete the photo Search Task four times, having each subject find a different photo each time. We asked subjects to perform the Browsing Task once for each browser. The collage target (friends, events, etc.) was different for each browser.

Each time subjects completed both tasks under one of the conditions, they were asked to complete a questionnaire. We asked questions such as the helpfulness of the photo organization, the subject’s degree of satisfaction, the amount of frustration, and adequacy of the allowed time. Answers were encoded on a 10-point Likert scale.

In order to minimize experimenter bias during the selection of photos for the Search Task, we had a computer randomly select photos from each subject’s collection. The computer presented one random photo after the other to one of the experimenters. The experimenter accepted or rejected each photo based on the following criteria: a photo was rejected if (1) the picture was taken at the same event as one that had already been chosen, or (2) the photo did not display any recognizable context, and the subject was therefore not likely to identify it in her collection; all other photos were accepted. The study in [9] followed a similar procedure and reports positive experience with this approach.

Notice that if we had simply run subjects through the three conditions, BM, BA, and TL, the subjects would have been exposed to the Baseline user interface twice, while interacting only once with the Timeline interface (recall that BM and BA differ only in the organization of the folders, not in their interface). To avoid disadvantaging the Timeline interface we had subjects perform the Search and Browsing tasks twice on the Timeline interface. In addition to the fairness issue, this double exposure (TL1 and TL2) allowed us to study learning effects on the Timeline interface. The two exposures were never administered in immediate succession. We counterbalanced BM, BA, TL1, and TL2 by rotating the order in which successive subjects were exposed to these conditions.

Throughout the experiment we recorded quantitative data, such as the amount of time subjects spent in each view of the respective browser, the number of view switches they initiated during the tasks, and the number of times the subjects made use of special features, such as the photo flipping facility. We also recorded the time taken for each trial of the Search Task and the number of photos collected within the allotted two minutes in the Browsing Task.

RESULTS

We first list results concerning the speed and performance of the different tasks. We then present the users subjective evaluation of the browsers and the tasks. We conclude this section with statistics about the usage of the features of our Timeline browser.

Speed/Performance

Search Task

On average, subjects showed the best search performance with the Timeline interface ($M=67.18$, $SD=33.5$), followed by Baseline Automatic (BA) ($M=80.4$, $SD=46.8$) and Baseline Manual (BM) ($M=80.8$, $SD=45.8$). However, the differences in search time among these three interfaces were not significant ($p=0.3$). The good news is that as users got familiar with TL, their searching performance improved. The learning effect between the two TL conditions was significant. The average completion time for the first set of TL trials (TL1) was 76 secs; for the second set (TL2) it was 58 secs ($t(16) = 2.663$, $p<0.05$).

When comparing the Search Task completion times for all four conditions, TL2 had shorter average search time ($M=58.3$, $SD=28.7$) than both BM ($t(16)=2.18$, $p<0.05$) and BA ($t(16) = 2.27$, $p<0.05$). However, when we compared the search time between BM and TL2 only among those nine subjects who worked the search task first with BA and then with BM (we call this set BM2), the difference was not significant¹.

With TL, subjects found more photos within 3 minutes than with BA or BM. The success rate with TL was 91%, significantly higher than BM ($M=80\%$) and BA ($M=75\%$). There was a marginally significant learning effect between TL1 and TL2 for success rate as well ($t(16) = 2.07$, $p=0.06$). In TL1, 87% of the search tasks were completed. This number rose to 96% in TL2. Condition TL2 was also better than BM2 (83% success rate, $p<0.1$).

As a measure of navigation performance, we counted the number of backtracking (“Back” or “Home”) operations during the search task. The fewer backtracking moves, the better the subject’s orientation. The BA interface required 56% more backtracking than BM, and 62% more than TL. There was no significant difference between BM and TL. We again see improvement from TL1 to TL2. The difference in backtracking between TL2 ($M=4.0$, $SD=3.9$) and BM ($M=7.4$, $SD=5.5$) is significant ($t(16)=2.26$, $p<0.04$). Moreover, TL2 required less backtracking than BM2 ($p<0.01$).

Browsing Task

The measurable parameters of the browsing task (e.g., number of photos found) were not significantly different

¹ For all search results, we charged the maximum search time of 3 minutes for failed search attempts.

across the different conditions. The average number of photos found was between 4 and 5 for all conditions.

Subjective Measures

Search Task

The perceived easiness of the search task was significantly higher for TL2 compared to TL1 ($t(16)=2.705$, $p=0.01$). TL2 was rated easier than BA (marginally significant, $t=1.86$, $p=0.07$). There was no other significant difference in perceived ease of search among the conditions. There was also no difference between the conditions in terms of satisfaction or the helpfulness of the photo organization.

We asked the subjects how well they knew where to look for each picture when they started the search. With BM, subjects claimed to have known best, at the outset, where the picture was. The subjects ranked BM first followed by TL2, TL1 and BA. The difference between BM and BA was significant ($p<0.05$), and the difference between TL2 and BA was marginally significant ($p=0.07$).

Browse Task

Most users claimed that the BM organization (i.e., their own) was more helpful than the system organization for the Browsing Task. However, there was no difference between the conditions in terms of satisfaction or perceived completeness of the Browsing Task.

Usage Of Timeline Browser

Figure 5 shows the average total time users spent at each level of the Timeline application during the four Search Tasks. For example, users spent an average aggregate of 44 secs in the year level (11 secs per task, 17% of the time). In total, users spent an average of 58% of the time in “timeline view” (the opening game on its different levels) and the rest (42%) in “thumbnail view” (the middle game). The numbers were a little different for the Browsing Task, where users spent relatively more time (48%) in the thumbnail view than during the Search Task. The relative times spent in day and year views are similar to the Search Task, while the relative times spent in week and month view are reduced.

The subjects used the drag-zoom feature on average 3.8 times during TL1, and 2.9 times during TL2. Recall this feature enables drilling into finer grained time range by dragging the stylus on the timeline to cover the desired time range.

The “flipping” feature, where users can tap the representative photo of the album and circle around it to see other photos from the album, was not extensively used. Only seven of the subjects tried to use this feature.

DISCUSSION

The experiment’s results informed our future work along three main dimensions. We learned (i) how manual and automatic photo organization and the time-informed user interface affect search and browsing performance. We

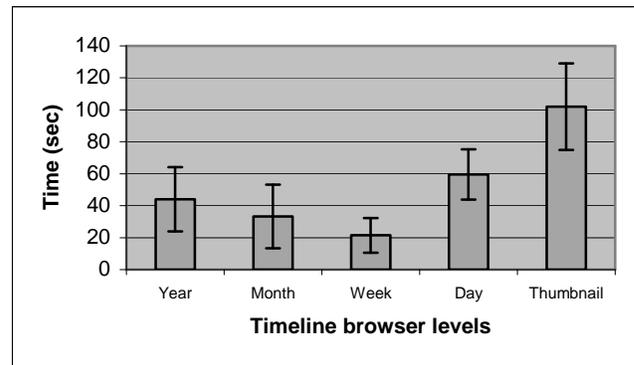


Figure 5: Average time subjects spent in each level of the Timeline browser hierarchy

learned (ii) that we could count on continued exposure to the Timeline browser to improve user performance even more. And (iii) we gained insight into how much the interfaces encourage zooming behavior.

Time and Browsing Performance

Recall that one of our central questions was whether the automatic generation of organization coupled with an appropriate interface can compete with manually organized collections underneath a traditional interface. The most important answers emerged from the task of finding known photographs.

As we showed above, comparing TL2 with BM, that is discarding the first exposure to TL as a learning experience, brought TL2’s search performance to a significant 29% improvement over BM. Even when we allow learning of BM as well (by counting only the trials for which BM followed BA in the counterbalancing rotation), TL2 can still hold a position of equivalence. Similarly, when we separate the success rate of finding photos from the search speed², TL’s speed was again equivalent to BM’s, while TL’s success rate was significantly (11%) higher than BM’s.

The manually organized Baseline, BM, was also equivalent in the search time measure to the automatically organized Baseline, BA. Given the identical interfaces of BM/BA, their search time equivalence suggests that the automatic organization works well.

The questionnaire data adds that subjectively, as well, all conditions were perceived as equivalent. None of the photo organizations or interfaces was perceived as being more ‘satisfying’, ‘helpful’, or as enabling a more ‘organized search process’, except for a preference for BM with the Browse Task.

² We obtained this measure by treating search failure as missing data, rather than substituting the maximum time allowed in place of failed search attempts.

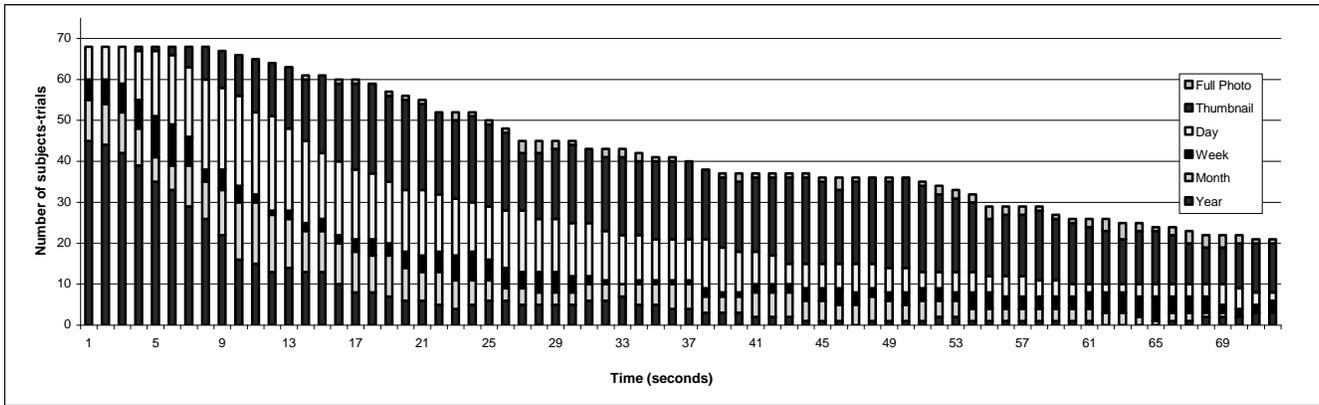


Figure 6. View dwelling time distribution for TL2

Learning Effect

A significant learning effect was evident for the Timeline interface. On the other hand, despite the fact that TL2 (the second exposure to the Timeline condition) yielded a significant speed-success improvement of 18% over TL1 and a marginal improvement in success rate, subjects did not feel in TL2 that they were better able to predict ahead of the search where they would find a given photo.

Navigation

The small PDA screen requires frequent screen changes, either through scrolling or through ‘drilling’ down. Both of the interfaces in question emphasize drilling. The time subjects spent at each level of depth is therefore crucial for understanding user behavior.

View Depth Dwelling Time

Not surprisingly, subjects spent most of their Baseline time in the ‘thumbnail’ view. We see a more informative result in the Timeline browser. Recall that the ‘levels of depth’ in that interface correspond to time granularities, such as ‘year’, ‘month,’ etc. The distribution of time spent at these various levels shows sharp preferences, which are strongest for the Browsing Task. As the results show, by far predominant in either task were extended stays in the ‘day’ and ‘thumbnail’ views.

In Figure 6 we illustrate exactly where in the time granularity and full vs. thumbnail views subjects spent time as they progressed through the Search Task in TL2. The x-axis shows the advancing search time. Every bar is partitioned into six portions, each of which corresponds to one time granularity or view (see legend). The relative size of each portion reflects the number of subjects that were browsing at that granularity during the respective bar’s time range. The y-axis is the number of subject-trials (17*4=68 for the TL2 Search Task). Notice that at the outset (first click) most subjects stayed in the year view. By the 20th second the majority of subjects had switched to day or thumbnail view. The height of the bars decreases as subjects find their photos. Half of the photos were found by the 47th second. This pattern suggests either that the year view was very effective (allowing users to focus quickly on

a few days), or that the intermediate views need to be improved.

Backtracking

The frequency of descending and backtracking through the view hierarchy is a good indicator of false trails. Those, in turn, are usually costly in terms of time and satisfaction. The results show that BA exhibits significantly higher backtracking rates than either BM or TL. The lowest backtracking rate is evident in the TL2 condition, an encouraging finding. This result is one possible factor that accounts for the superior TL2 time performance.

Subject Comments and Feedback

At the end of each experiment session with a subject, we conducted a debriefing session where we explained the purpose of the experiment and solicited general feedback regarding their experience with the system.

For a majority of the subjects the thumbnails were too small. We had purposely pushed the thumbnail size down to the very limit of possible acceptability. The rationale was that the subjects would be interacting with their own collections and thus should be able to recognize their photos with minimal level of detail. We will need to revisit this issue.

There were quite a few suggestions made by the subjects for improving the interface. The specific screen that received the most number of suggestions was the Baseline browser end-game screen (“Preview Screen”). One of the suggestions was to provide more context about the thumbnails on the top row by increasing the number of preceding and succeeding thumbnails, and also to indicate the current index of the center photo (e.g. “12 of 47”). The subjects also expressed the desire to have the ability to scroll to the next screen-full of photos in the Preview Screen, instead of being restricted to scrolling at most by two photos, as is currently the case.

We also received suggestions to add support for the circular flipping gesture in the Preview Screen so that the user can flip through the medium sized images in that view. There were also requests for a full-screen slideshow feature.

RELATED WORK

There are a number of commercial photo browser applications available for the various handheld device platforms. We reviewed six of the most popular commercial photo browsers on a popular PDA software website (<http://www.handango.com>). None of the browsers support automatic organization of photos, or allow browsing through the photos based on creation time. The features supported by the commercial browsers include thumbnail views, Windows-Explorer-style folder views, and slide show support. Two of the commercial browsers [8,10] also provide an intermediate screen similar to our Baseline browser Preview Screen. Most browsers are geared towards browsing a small collection of pre-selected photos, rather than the user's complete image collection.

On the other hand, quite a few research projects have studied photo browser applications for the desktop [1,6,7,9, and others] or other platforms [12]. In [1], for example, the authors try to maximize screen utilization for displaying images in a hierarchical organization.

Some of these photo browser projects use timestamps of photos as a hint for photo organization [4,9]. However, most of the browsers rely primarily on the manual annotation/grouping of photos, and use time only as a secondary tool. Some commercial desktop photo browsers, such as Adobe's Photoshop Album, have begun to use time as a primary organizational tool, but often use time only as a hint for the ordering and density of photos, and do not use automatic tools to group photos based on time.

Additional time-based clustering techniques for photos have been suggested in [4,9,13]. The interface we present in this paper is orthogonal to time clustering techniques, and in fact, can be adapted to take advantage of different clustering methods.

ACKNOWLEDGMENTS

We thank the subjects of our experiment, and poor Matt Wang and Wendy Tobagus who helped parse our log files.

CONCLUSION

Building on existing methods for automatically organizing digital photographs, we developed a user interface for browsing thousands of photographs on a personal digital assistant. In a user experiment, we compared this new system against a more traditional interface using each subject's own, manually organized photo collections. We found that the automatically organized system performed at least as well in search and browsing tasks as the traditional system, which requires that users invest a significant amount of time and effort in manually organizing their photos. We observed that users show a decreased need for search backtracking and an increase in search success using our new time-based system. We were therefore convinced

that our automatic photo organization scheme and timeline interface provide a solid foundation for further development.

REFERENCES

1. Bederson, B.B. PhotoMesa: a zoomable image browser using quantum treemaps and bubblemaps. *Proc. UIST 2001*.
2. Buyukkokten, O., Garcia-Molina, H., Paepcke, A., and Winograd, T. Power Browser: Efficient Web Browsing for PDAs. *Proc. CHI 2000*.
3. Frohlich, D., Kuchinsky, A., Pering, C., Don, A., and Ariss, S. Requirements for Photoware. *Proc. ACM CSCW 2002*.
4. Gargi, U. Consumer Media Capture: Time-Based Analysis and Event Clustering. *HP Tech Report Number HPL-2003-165*, 2003.
5. Graham, A., Garcia-Molina, H., Paepcke, A., and Winograd, T. Time as Essence for Photo Browsing Through Personal Digital Libraries. *Proc. JCDL 2002*.
6. Kang, H. and Shneiderman, B. Visualization Methods for Personal Photo Collections: Browsing and Searching in the PhotoFinder. *Proc. IEEE International Conference on Multimedia and Expo*, 2000.
7. Kuchinsky, A., Pering, C., Creech, M.L., Freeze, D., Serra, B., and Gwizdka, J. FotoFile: a Consumer Multimedia Organization and Retrieval System. *Proc. CHI 1999*, pp. 496-503.
8. PhotoPals Photo Album. <http://photopals2002.com>.
9. Platt, J.C., Czerwinski, M., and Field, B.A. PhotoTOC: Automatic Clustering for Browsing Personal Photographs. *Microsoft Research Tech Report Number MSR-TR-2002-17*, 2002.
10. Resco Picture Viewer for Pocket PC. <http://www.resco-net.com/picview.asp>.
11. Rodden, K. and Wood, K.R. How do people manage their digital photographs? *Proc. CHI 2003*.
12. Shen, C., Lesh, N., and Vernier, F. Personal digital historian: story sharing around the table. *interactions*, 10(2):15-22, 2003.
13. Stent, A. and Loui, A. Using event segmentation to improve indexing of consumer photographs. *Proc. SIGIR 2001*. ACM Press, 2001, 59-65.
14. Wagenaar, W.A. My memory: A study of autobiographical memory over six years. *Cognitive psychology*, 18:225-252, 1986.
15. Yee, K.P., Swearingen, K., Li, K., and Hearst, M. Faceted metadata for image search and browsing. *Proc. CHI 2003*. ACM Press, 2003, 401-408.