Time Quilt: Scaling up Zoomable Photo Browsers for Large, Unstructured Photo Collections

David F. Huynh[†], Steven M. Drucker[‡], Patrick Baudisch[‡], Curtis Wong[‡]
dfhuynh@csail.mit.edu, {sdrucker, baudisch, wong}@microsoft.com
[†]MIT Computer Science and Artificial Intelligence Laboratory,
The Stata Center, Building 32, 32 Vassar Street, Cambridge, MA 02139, USA
[‡]Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA

ABSTRACT

As automatic image indexing technologies have yet to mature, large digital photo collections still require a correspondingly large manual organization effort. Without such an effort, only the camera-recorded metadata, specifically the creation date, and the photos' visual contents are available. Current zoomable photo browsers (e.g., PhotoMesa) are designed to support visual searches but, in trying to maximize screenspace usage, they cannot effectively convey temporal order. Conversely, a linear timeline layout makes poor use of screenspace and requires excessive panning. We propose *time quilt*, a layout that conveys temporal order while making better use of space than a linear timeline. In an experimental comparison of space-filling, timeline, and time quilt layouts, participants carried out the task of finding photos in their personal photo collections averaging 4,000 items. They performed 45% faster on time quilt.

Furthermore, while current zoomable photo browsers are designed for visual searches, this support does not scale to thousands of photos: individual thumbnails become less informative as they grow smaller. We implemented semantic zooming using representative photos to provide an overview for visual searches in place of the diminishing thumbnails and found a subjective preference for this feature.

Categories & Subject Descriptors: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General Terms: Human Factors, Design.

Keywords: Photographs, timeline, space filling, time ordering, representative thumbnail, zoomable user interface, semantic zooming.

INTRODUCTION

Digital photography has a lower cost for taking and storing pictures than film photography. Consequently, digital photos tend to accumulate much more quickly than film photos. From IDC's 2003 U.S. Consumer Digital Imaging Survey, 17% of the respondents took 50–100 photos per month (or 3,000–6,000 images over a 5 year period). 8.9% of the respondents took more than 100 photos per month.

As a photo collection grows, it becomes harder to manage. Particularly, it takes time and effort to create structures that will facilitate browsing and searching that collection (e.g., tagging the photos with keywords, filing them into a hierarchy of file folders, annotating them with comments). But the very aspect of the collection that demands organization—its size—also defies manual organization: Frohlich et. al. [6] reported that very few families participating in their user study systematically organized their digital photo collections. While users cannot always be counted on to invest time and effort to create complete, coherent structures, most photo browsers rely on the presence of such user-provided metadata to support effective searching and browsing.

There are means for automatically indexing photo collections by visual content analysis, but Kuchinsky et. al. [10] expressed concern that such content-based indexing and retrieval approaches focused on directed search but neglected users' information seeking behaviors. Rodden and Wood [17] reported that the content-based image retrieval feature offered in their user study's test system was rarely used and its perceived utility was low.

In the absence of user-input metadata, one must rely on the metadata recorded by cameras, specifically the creation date, and the photos' visual contents. Not only is the creation date recorded automatically, it is also an essential factor by which people browse their photos [7]. With only creation date and visual content, people can still perform useful tasks such as:

- reminiscing over a past period of time, to answer such questions like, 'what pictures do we have from the last five years?'; and
- finding photos associated with some memorable events (e.g., 'that photo of you hanging from a cliff that I took on our trip to the Grand Canyon two years ago').

We characterize such tasks as *time-based visual searches* as they rely on time information as well as the visual contents of photos.

Zoomable photo browsers (e.g., PhotoMesa [2]) are designed for visual searches by maximizing screenspace usage, thus requiring minimal panning. By adopting a space-filling layout (e.g., quantum treemap [2]), however, they don't convey temporal order very well: although they allow users to cluster photos by creation date, the clusters are not laid out in such a way that communicates their order in time. Changing the layouts of existing zoomable photo browsers to convey time will often require sacrificing screenspace. In the extreme case, a linear timeline trades all space optimization for a straightforward presentation of temporal order.

We propose a layout designed to combine the benefits of timeline and space-filling approaches: our *time quilt* layout wraps photo clusters in columns (see screenshot in Figure 1), effectively compressing the timeline horizontally.

Not only do existing zoomable photo browsers need to convey temporal order, their support for visual searches must also scale to the order of thousands of photos. These browsers currently do not implement any form of semantic zooming to compensate for the loss of visual information when thumbnails become too small to be recognizable and distinguishable from one another. Figure 1 shows how representative photos can be used to provide a more informative overview than a rendition of scaled down thumbnails.

Contributions

In this paper, we make 3 contributions:

- We propose a novel layout called time quilt that combines time-ordering and space-filling layout criteria to support time-based tasks in the absence of user-provided organizational metadata.
- We propose the use of representative photos for implementing semantic zooming in zoomable photo browsers to support visual searches in place of diminishing thumbnails.
- We present a user study that compares the various layout approaches for zoomable photo browsers and tests the usefulness of semantic zooming based on representative photos.

RELATED WORK

The work presented in this paper builds on three fields of prior work: photo browsing and visualization, time information visualization, and semantic zooming.

Photo Browsing and Visualization

Personal digital photography started to attract the attention of researchers in 1999 as Kuchinsky et. al. [10] introduced *FotoFile*, a consumer-oriented multimedia organization and retrieval system. This work recognized consumers' lack of economic incentives for up-front time and effort devoted to annotating their own photo collections for the purpose of directed, keyword-based searches later. Kuchinsky et. al. also noted that automatic content-based indexing and retrieval were not the complete solution, either, as they did not map to consumers' information-seeking behaviors. In light of these observations, *FotoFile* was designed to support both directed searches and exploratory browsing.

In the same year, Combs and Bederson [3] built the first zoomable image browser and found that it performed as well as a conventional 2D image browser, but better than a 3D browser for collections of up to 225 images. However, they did not test their browser on larger photo collections which are more common today. Bederson followed up in 2001 with *PhotoMesa* [2], a second zoomable photo browser that used space-filling layouts as discussed above. *PhotoMesa* was implemented in the *Pad++* UI framework [1] but it did not at the time make use of the semantic zooming support that the framework already offered.

Kang and Shneiderman [9] introduced *PhotoFinder* in 2000, equipped with a set of visual Boolean and dynamic query interfaces. Users can search for photos based on multiple criteria, including date, people, event, location, color, keyword, rating, etc. *PhotoFinder* was designed primarily to address searching rather than browsing.

In 2002, Platt et. al. [15] performed one of the first formal user studies of personal photo collections (averaging 850 images). Their photo browser, PhotoTOC, introduced representative photos to create an overview, table-of-content summary of photos clustered by date and time. Their contributions included the user of representative photos, new algorithms for clustering similar photos by date and other attributes, and methods for selecting representative photographs. Their browser rendered the representative photos in a separate overview pane; while in our work, we render representative photos *in the place of* their clusters—that is, we implement semantic zooming. At the time, the study found that choosing an inappropriate thumbnail to represent a collection caused great difficulty with the users for finding a photo.

In 2003, Rodden and Wood [17] conducted a six month-long study for their work titled, "How do people manage their digital photographs?" They concluded that two of the most important features to support in photo browsers were (1) automatically sorting photos in chronological order and (2) displaying a large number of thumbnails at once. Their participants most commonly wanted to browse their personal photos by event rather than by querying them based on more specific properties.

In 2004, Drucker et. al. [4] presented a careful selection of many previous research concepts integrated into a single browser, the *MediaBrowser*. By integrating temporal clus-

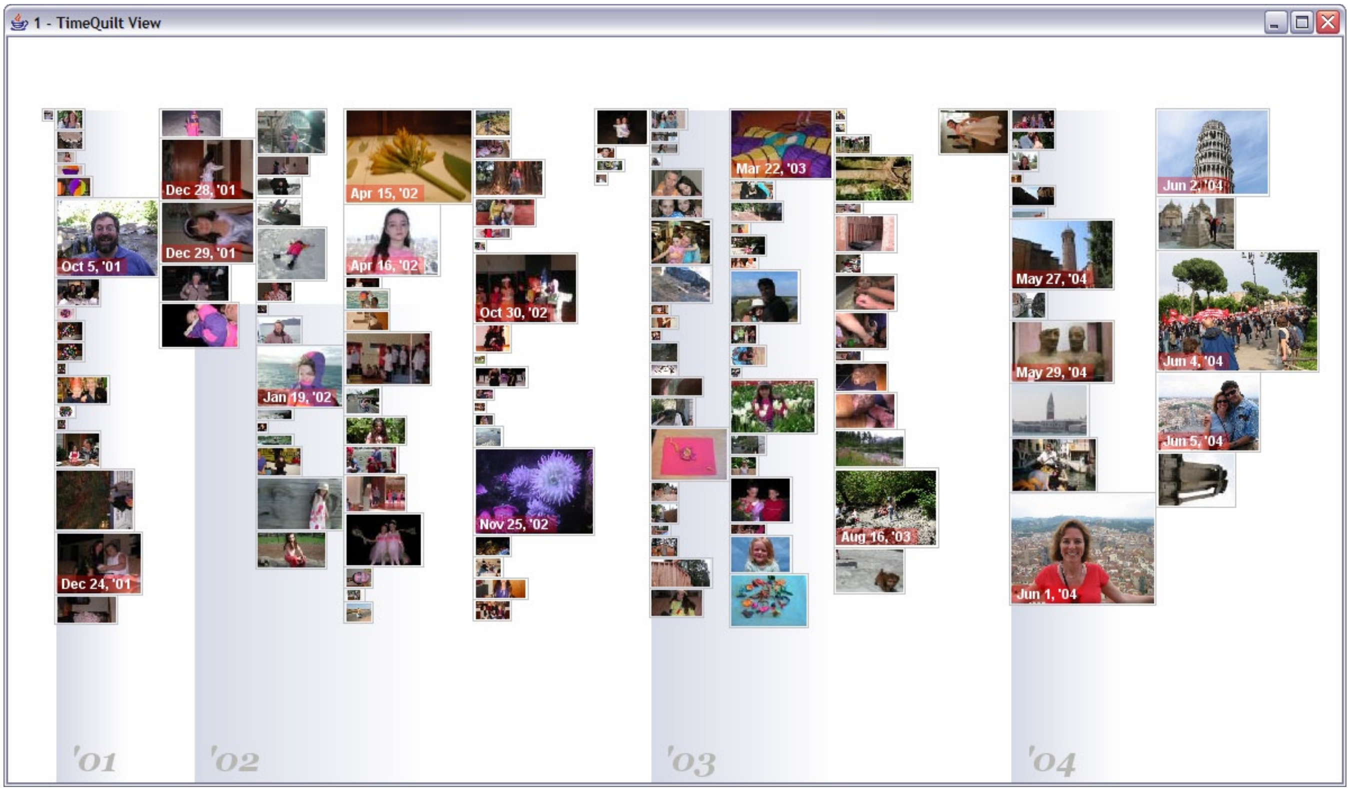

Figure 1. Time quilt – a layout designed to convey temporal order while making better use of screenspace than a timeline, showing approximately 5,500 photos with representative photo overview

tering with rapid selection, they were able to make it easier for users to annotate their collection. In an informal study, they found that loading *MediaBrowser* with more than 500 or 600 objects rendered individual thumbnails hard to distinguish by eye, thus identifying scalability as a topic for future work.

Time Visualization

In the field of visualization, there is also much prior work on presenting temporal information. In particular, Plaisant et. al. [14] presented *LifeLines* as a simple and tailorable visualization environment for showing personal histories in multiple facets. Fertig et. al. [5] presented the system *Lifestreams* for showing a user's personal file system in a timeline format. Ringel et. al. [16] investigated a timeline layout to present search results from personal information corpora. They used photo thumbnails as landmarks on their timeline visualization although the main information being visualized is not photographs but e-mail messages, appointments, and other documents. In other older work on time visualization, Mackinlay et. al. [11] developed the *Spiral Calendar* for rapid access to an individual's daily schedule by integrating details and context using a 3D spiral layout. The software industry has also introduced time-based browsing of personal photo collections. For example, Adobe *Photoshop Album*, Microsoft *Digital Image Library* and *Picasa* provide a linear timeline while *ACDSee* and *iPhoto* support a monthly calendar view.

Semantic Zooming

Semantic zooming was introduced by Perlin and Fox [13] in the *Pad* system in 1993 and later supported pervasively in the *Pad++* system by Bederson and Hollan [1].

To address the problem of shrinking thumbnails, Suh et. al. [18] proposed a method of automatically cropping a photo to keep only its most salient region such that a more recognizable thumbnail can be generated from it. This solution works to some relatively small size, but beyond this, no recognizable thumbnail can be generated regardless of the original image's saliency. This is why we chose to investigate the use of representative photo overview to replace the shrinking thumbnails.

THE TIME QUILT LAYOUT

As mentioned in the introduction, space-filling layouts do not convey chronological order effectively to support time-based visual searches. Figure 2 shows a space-filling layout (quantum treemap) of a sample collection of 5,500 photos clustered by dates and times. While May 8th, 2004 and May 15th, 2004 are closer dates than May 8th, 2004 and June 3rd, 2003, they map to clusters that are further apart. This mapping makes the task of finding the next cluster in time difficult, especially when the visualization is zoomed in.

At the other end of the spectrum, laying out clusters of photos in a timeline conveys temporal order in the most straightforward and familiar way but completely disregards screenspace usage. As illustrated in Figure 3b, a space filling layout maintains a reasonable overall aspect ratio such that the whole visualization can fit on one screen while a timeline yields a severe overall aspect ratio and thus requires excessive panning to browse through consecutive clusters.

We believe that a compromise between the two layout criteria, space-filling and time-ordering, would result in a layout more effective for time-based visual searches than either criterion alone. This compromise can be made by taking a long timeline stretching along only one dimension (as in Figure 3b) and "weaving" it in the other dimension as well, thus effectively compressing it to achieve a better overall aspect ratio (as in Figure 3c). In Figure 4, we propose one particular "weaving" layout called *time quilt*: clusters of photos are wrapped in vertical columns (similar to how text is wrapped in horizontal lines).

Incidental Whitespace as Design Element

In the time quilt layout, there is some whitespace in between the clusters, incidental from the way the clusters are aligned in columns. We believe that this inter-cluster whitespace, absent in space-filling layouts by design, makes the whole visualization more aesthetically pleasant to view. In fact, good designs are those that make effective use of whitespace as a design element to unite and separate other design elements, communicating structure and flow [12].

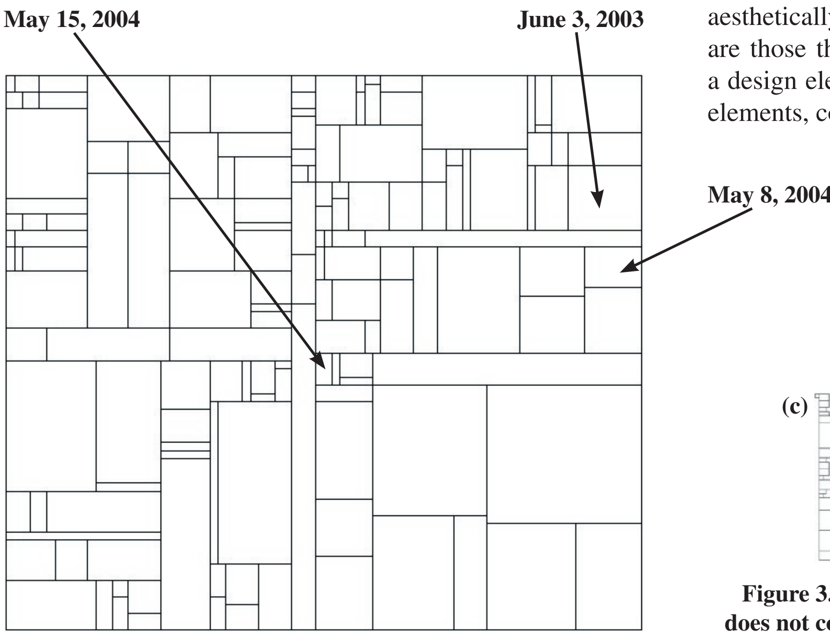

Figure 2. A space-filling layout of a sample collection of 5,500 photos does not convey chronological order: closer dates might map to far-apart clusters

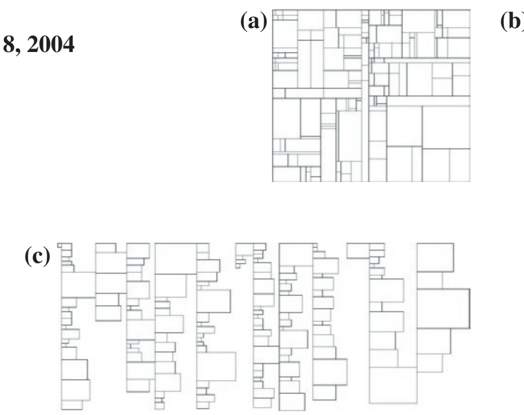

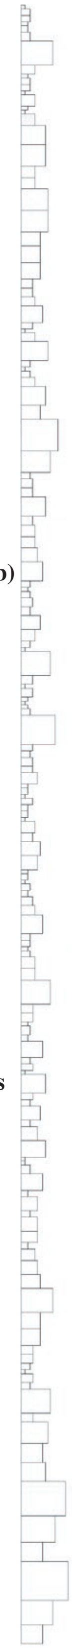

Figure 3. A space-filling layout (a) saves space but does not convey time; a timeline (b) wastes space and requires excessive panning due to severe aspect ratio; weaving the timeline along both dimensions (c) conveys time and achieves reasonable aspect ratio.

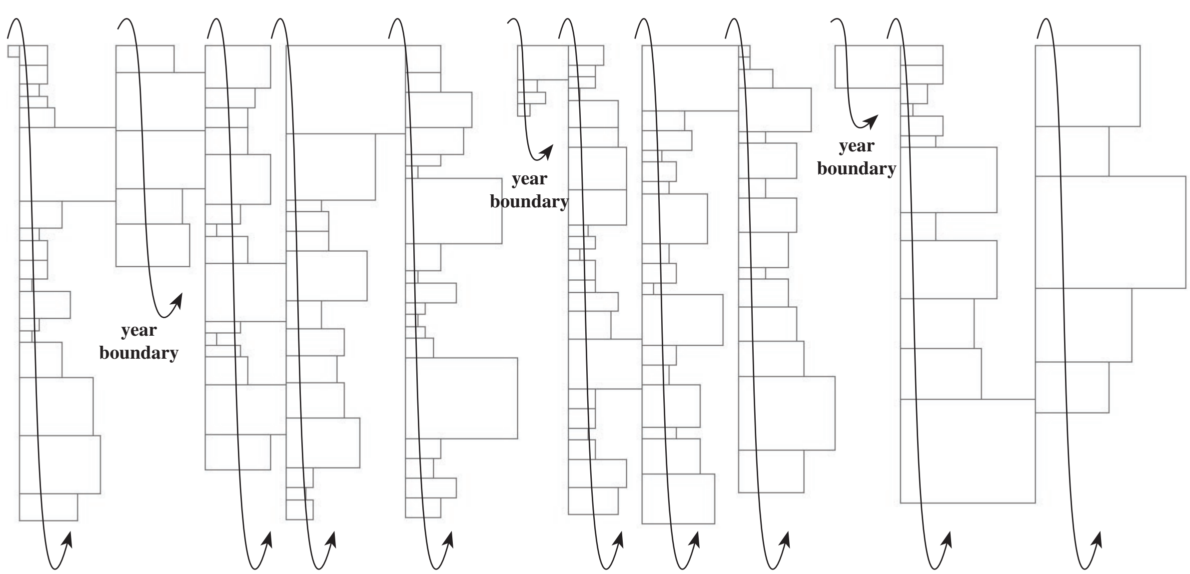

Figure 4. The time quilt layout "weaves" a linear timeline into columns of photo clusters, breaking at a specified maximum column height or at year boundaries

In conjunction with the incidental inter-cluster whitespace, the clusters arranged in columns form a non-uniform pattern that may facilitate spatial memory for re-finding tasks. We believe that time quilt layouts, being sparser than space-filling layouts, are easier to commit to memory.

Cluster Aspect Ratios

Time quilt not only trades screenspace usage for temporal order, it also trades the overall aspect ratio for individual clusters' aspect ratios. In a space-filling layout, all clusters must be laid out together so that the overall aspect ratio of the whole visualization can be controlled, but this strategy results in severe aspect ratios for some of the clusters (see Figure 2).

Time quilt, on the other hand, lays out photos within each cluster independently at first, and then wraps the resulting clusters into columns. This strategy allows us to make each cluster's aspect ratio match that of the screen—this choice makes maximum use of the screen's space when the cluster is zoomed in, which reduces the need for panning.

Such trade-off can be justified as follows: An owner of thousands of photos accumulated over several years may not always want to look at the whole collection all at once. S/he might only want to focus on a few years' worth of photos and does not mind panning to those years before diving down into some clusters within them. Once s/he has zoomed into a cluster, the cluster's aspect ratio becomes more important than the overall aspect ratio of the visualization.

Changes to Photo Collection

The global strategy of a space-filling layout not only yields severe cluster aspect ratios, it also causes major shuffling in the clusters' positions and sizes as a photo collection changes—an inevitability. This behavior is detrimental to spatial memory. Time quilt, on the other hand, "grows" into more columns and leaves the old columns unchanged. (This is like how appending text to an unjustified paragraph leaves the existing lines unchanged.) Although the overall aspect ratio worsens as the collection grows, as mentioned before, the user is likely to focus on a few years' worth of photos anyway.

Other Ways of Weaving The Timeline

There are other ways to weave the timeline in 2 dimensions such that the layout is resistant to changes, the individual cluster aspect ratios can be controlled, and whitespace helps toward the effectiveness of the visualization. For instance, the clusters can be laid out in a spiral or in a zig-zag pattern. However, we chose the saw-tooth pattern in creating time quilt because people are very familiar with the similar concept of text wrapping and it is straight forward to compare the temporal order of any two clusters: if a cluster is located to the right of another cluster, it is later in time, or if both clusters start at the same horizontal position, the lower one is later.

SEMANTIC ZOOMING USING REPRESENTATIVE PHOTOS

To support time-based visual searches, we must show both time information as well as visual contents of photos. Presenting photos as thumbnails works only for some number of images (Figure 5a), beyond which the thumbnails become too small to carry information individually (Figure 5b).

We propose the use of representative photos to implement semantic zooming in zoomable photo browsers: when individual thumbnails become too small to be recognizable at a certain zoom level, a representative photo from each cluster is rendered in place of the diminishing thumbnails (Figure 5c).

Note that the use of representative photos is independent of the layout, although representative photos are more useful when the aspect ratios of individual clusters are reasonable, as can be achieved with either the timeline layout or the time quilt layout.

Figure 6 shows how zooming out from a cluster replaces individual thumbnails with a representative photo.

Other Ways of Using Representative Photos

There are also other ways to implement semantic zooming. For example, several representative photos can be chosen and rendered for each cluster. As the visualization is zoomed out and each cluster becomes smaller, fewer and fewer photos will be rendered in its space. For another example, a mosaic combining several photos in each cluster can be composed and used to represent the cluster. Each approach has its own downside. The first sample approach might not effectively indicate when a cluster is being semantically represented and when it is rendered with individual thumbnails. The second sample approach is difficult to implement and the details in each mosaic might be too small to recognize.

Our choice of rendering the representative photo of a cluster within the space occupied by the cluster also has drawbacks. At any particular level of zooming, only the representative photos of clusters above a certain size will be recognizable. Furthermore, the size of a cluster might not correspond to the importance of the event.

USER STUDY

We conducted a user study to compare 3 layouts: a space-filling layout (quantum treemap [2]), the timeline layout, and the time quilt layout, as well as to test the effectiveness of semantic zooming using representative photos.

Interfaces/apparatus

For this user study, we implemented 4 zoomable photo browsing interfaces in Java, three of which used representative photos (classified in Table 1 and shown in Figure 8): a space-filling browser (SF), a time quilt browser (TQ), a timeline browser (TL), and a space-filling browser without representative photos (SF–).

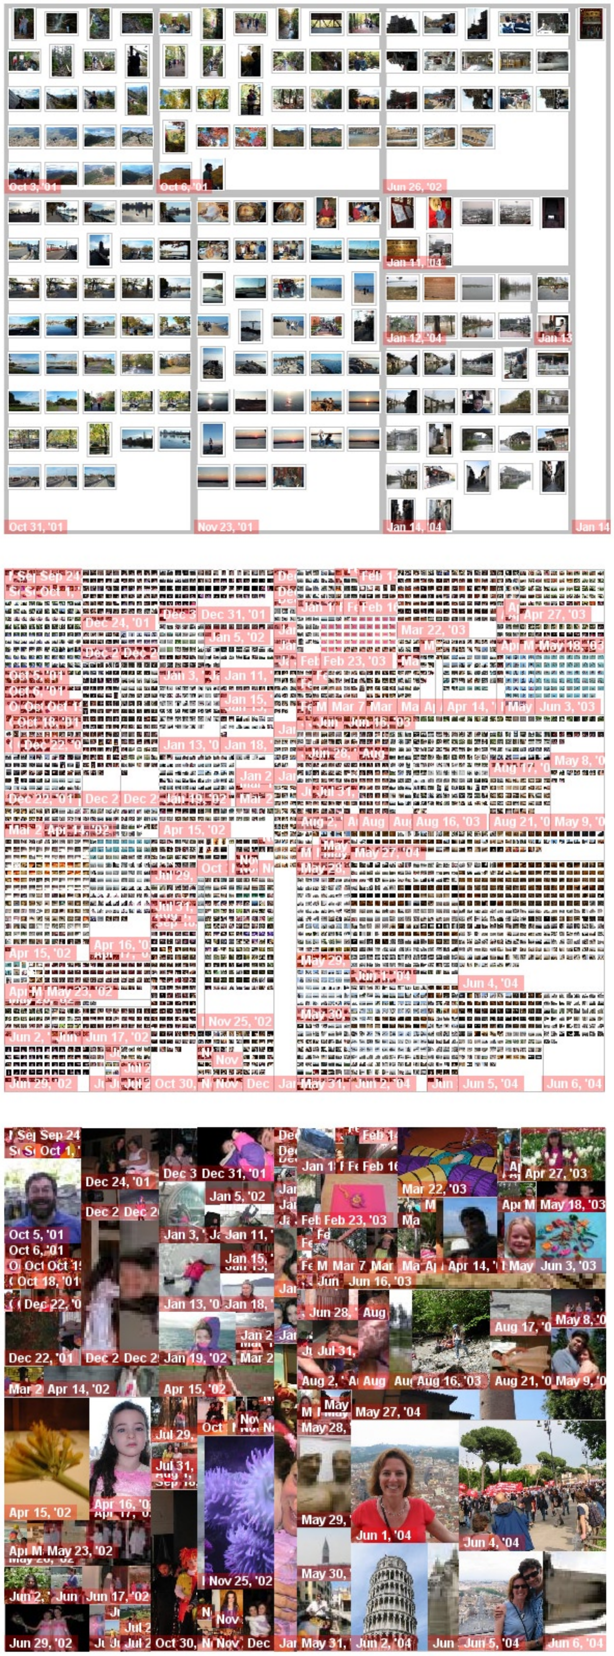

(a)

(b)

(c)

Figure 5. Thumbnails are recognizable at 180 images in a 700 x 600 pixel screen area (a) but became hard to recognize at 5,500 images in the same area (b); representative photos can be used to provide overview in such a case (c).

All interfaces supported zooming through the mouse-wheel. Scrolling the mouse-wheel forward zooms into where the mouse cursor points, and scrolling backward zooms out. Panning is done through dragging with the left mouse button. Selection of a photo can be performed by left-clicking. We limited the zooming levels and the panning range to prevent the users from getting lost in whitespace (famously known as the desert fog problem [8]).

We used Platt's adaptive clustering algorithm [15]. We chose the middle photo of each cluster to be the representative photo.

We ran these interfaces in Windows XP Pro, SP1, on a Pentium IV 3.2 Ghz HT, 2GB RAM machine with an NVidia Geforce FX 5200 128m video card, connected to two 19" CRT monitors, each running at 1280 x 1024 pixels, 96 dpi resolution. The mouse used was a Microsoft IntelliMouse Optical USB and PS/2 Compatible mouse (with a scroll-wheel).

The monitor on the left showed the instruction window (Figure 9) while the monitor on the right showed one of the four interfaces.

Task

The participants' task was to find a shown in the instruction window in the interface window and select it. Each participant was instructed to "go for speed" while maintaining reasonable accuracy. When faced with multiple photos resembling the target photo, the participant was instructed to pick one that was reasonably similar to the target photo rather than pinpointing the exact match. If the target photo was a modified version of another, original photo, s/he could select the original photo instead of the modified version. If a participant was unable to find a photo for several minutes, s/he was instructed to abandon the task.

Procedure

10 people participated in the study: 2 females, 8 males, 28 to 42 years old, with an average age of 35. They were those who responded to a user study advertisement on the Photography Enthusiast forum within our company's intranet, and who were willing to share with us large personal photo collections of over 3,000 images. According to their estimates, they took anywhere from 600 to 25,000 photos a year (mean = 6,980, stdev = 8,410) of topics including: family and friends, children, pets, items of interest, documentation, nature/sceneries, architecture, vacations/trips, school events, work events, outdoor events, sports events, and special events. None of these participants were familiar with any of the interfaces before the study.

Before their user study sessions, the participants shared with us their photo collections ranging from 2,863 to 5,708 photos (mean = median = 3,994, stdev = 928). Each of them also selected 28 favorite photos from his/her shared collection, almost all from different events. These favorite photos were later divided randomly into 4 groups to test the four brows-

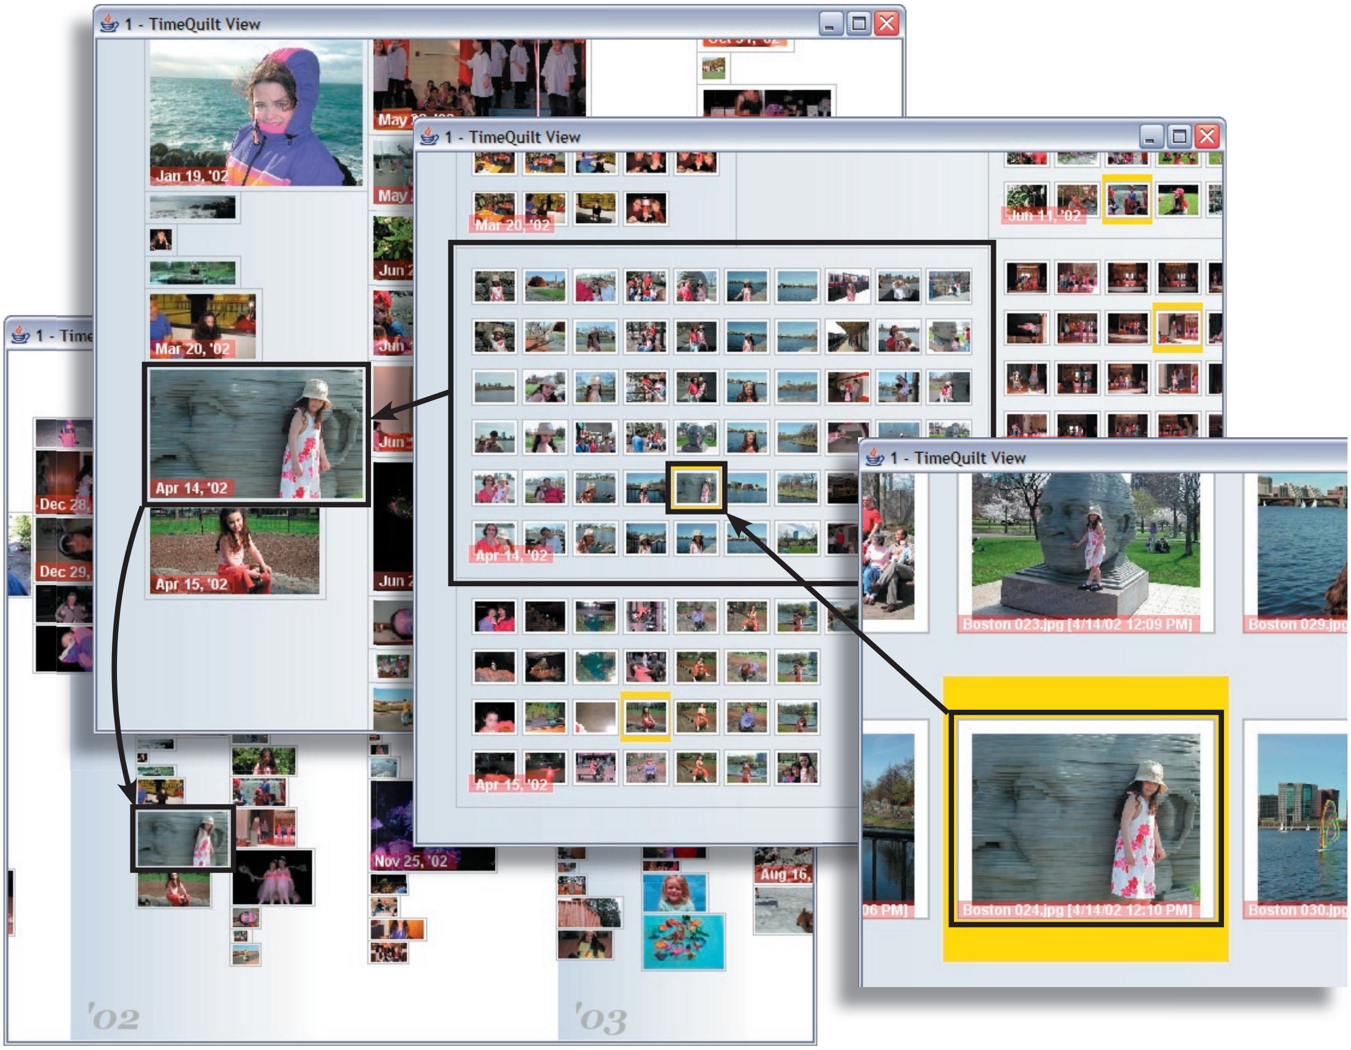

Figure 6. Semantic zooming replaces a grid of thumbnails of a cluster with a single representative photo as the visualization is zoomed out and the thumbnails become too small to be informative

ers, each group consisted of 2 training photos and 5 photos for the actual test. (In a pilot study, we tested the retrieval of randomly selected pictures from a user's collection and found that this task was often impossible since the randomly selected picture would often not be recognized at all.)

We used a within-subject experimental design: each subject carried out the task 5 times (for 5 different target photos selected from the 28 favorite photos) on each of the 4 testbed browsers. In order to avoid sequence effects, the browser order was counterbalanced between subjects.

Each participant received verbal instruction when the study session started and an introduction to each interface before the tasks on that interface were performed. After the introduction, s/he was encouraged to explore the interface by finding 2 "training" target photos.

Upon completing the tasks for each interface, the participant answered questions regarding his/her experience with that interface. At the end of the study session, s/he completed a survey to specify their overall preference. All Likert-scale questions were on a scale from 1 (disagree) to 7 (agree). The whole session took around one hour.

Hypotheses

We had 2 hypotheses:

- Participants would complete the task faster in TQ than in SF and TL because TQ conveyed time better than SF and required less panning than TL.
- Participants would complete the task faster in SF than in SF– because SF provided a more comprehensive overview with the use of representative photos.

Results

Task completion time – Table 2 and Figure 7 summarize the average time participants required to complete the task. Confirming our first hypothesis, they achieved better task completion time on TQ than on TL and SF (44.2% and 45.1% faster respectively). A one-way ANOVA test yielded $p = 0.002$ ($F(3,173) = 5.314$).

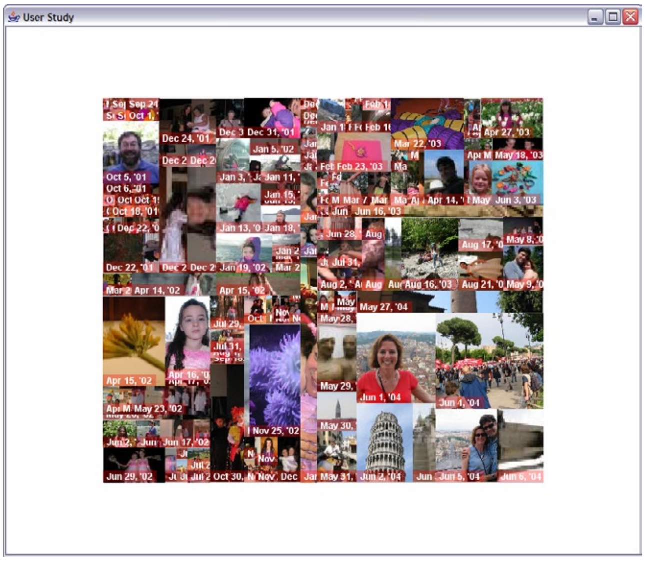

(a) Space-filling interface (SF)

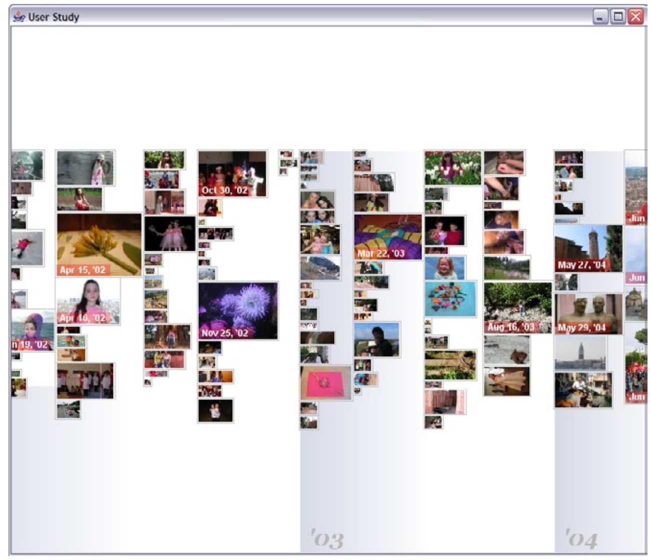

(b) Time quilt interface (TQ)

(c) Timeline interface (TL)

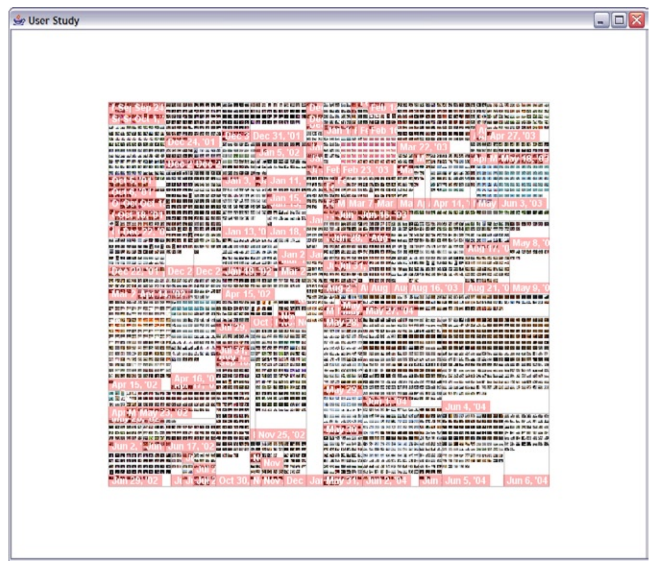

(d) Space-filling interface without representative photos (SF–)

Figure 8. Zoomable photo browsing interfaces for user study

		Layout		
		Space-filling (q. treemap)	Combo (time quilt)	Time-ordering (timeline)
Representative Photos	Yes	SF	TQ	TL
	No	SF–		

Table 1. Four photo browsing interfaces for user study

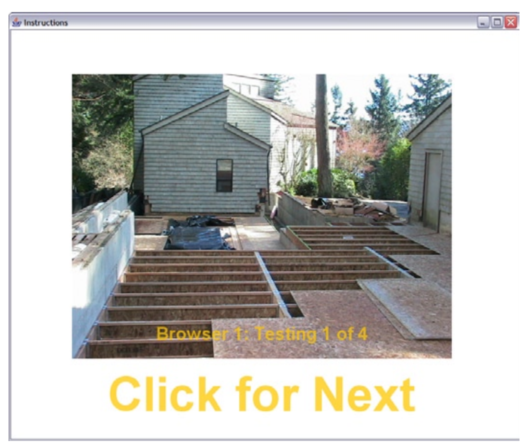

Figure 9. Instruction window on second monitor in user study

Note that we did not count trials with time exceeding 3 minutes. In such trials, participants did not recall enough information to locate the target photos and simply resorted to panning over the entire visualizations.

The participants' comments on the two space-filling interfaces provided some indication for why the space-filling layout made it harder for them to perform the task:

- "Dates are right but locations are completely wrong."
- "We have August and December next to each other... that's totally confusing."
- "Can't always go to the next day—not always next to where a given day is."
- "With the timeline, I have a linear order by which I can tell how far to go."

The difference between the task completion times on SF and SF– was not significant, so we were not able to confirm our second hypothesis.

Subjective preference – Confirming our first hypothesis, the participants specified that they "preferred the arrangement of photos by dates and times over the space-filling arrangement" (Likert: M = 6.20, SD = 1.476, t-test = .009). A Friedman test on the participants' rankings of the four browsers showed that TQ was best preferred: TQ ranked 1.40, TL 2.10, SF 2.80, and SF– 3.70 (N = 10, Chi-square = 18.125, df = 3, Asymp. Sig. = .000).

The participants also specified that they "preferred the use of representative photos" (Likert: M = 5.70, SD = 1.636, t-test = 0.009) even though they did *not* find that "the representative photos were accurate representation of the clusters" (Likert: M = 4.30, SD = 1.418, t-test = 0.520). Their comments also indicated that representative photos were helpful:

- "With representative photos, I have more clues of where things are."
- "There might be a representative picture of that... it'll get me there much quicker."
- "Bad key photos are better than nothing."

One participant tried to explain her worse performance on SF as compared to SF– as follows: The representative photos gave her a false image of the size of her photo collection, so when she zoomed in, the appearance of individual thumbnails overwhelmed her. She also mentioned that this sense of being overwhelmed only occurred for SF but not for TL or TQ, although they all used representative photos. Other participants wanted an intermediate zoom level where both representative photos and individual thumbnails were shown simultaneously.

		Layout		
		Space-filling (q. treemap)	Combo (time quilt)	Time-ordering (timeline)
Representative Photos	Yes	SF 61.7 (40.5)	TQ 33.9 (20.2)	TL 60.8 (41.2)
	No	SF– 56.7 (41.5)		

Table 2. Task completion time in seconds (standard deviation)

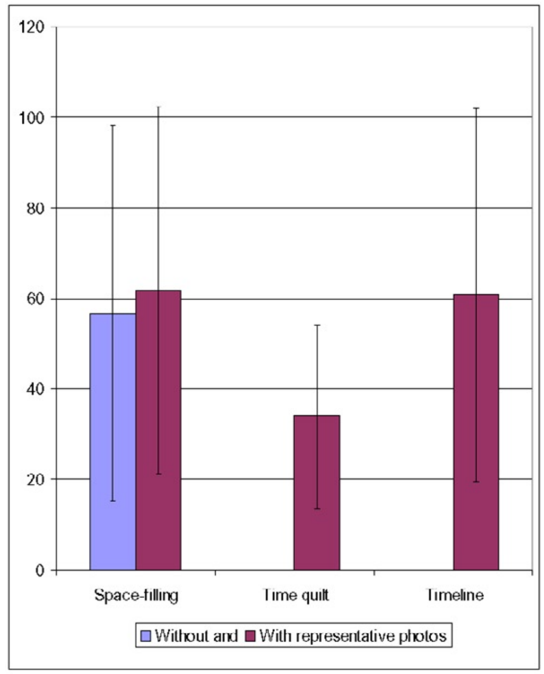

Figure 7. Average task completion time across different layouts (error bars represent standard deviations)

Discussions

Although every digital photo always carries a creation date, the dates of the images shared by the participants turned out less reliable for the following reasons:

- Some participants included scanned images in their photo collections and in the 28 favorites that they selected. The dates on these images were not the dates of the corresponding events.
- Some participants' cameras had faulty dates. Some pictures taken in 2004 were dated 2003.

In addition, participants included photos taken by someone else. As a result, they found it hard to remember approximately when those photos were taken.

CONCLUSIONS

In this paper, we presented time quilt, a novel layout for zoomable photo browsers that conveys temporal order while making better use of space than a linear timeline. We conjectured that this layout was more effective on large photo collections than the space-filling layouts as used in existing zoomable photo browsers. For a large collection, without zooming first into the approximate date range when a photo was taken, visual scanning for that photo over the entire visualization is inefficient. For this reason, time quilt emphasizes temporal ordering over space filling. But space filling is also important for maximizing the amount of visual information the user can see at any time, thus minimizing the amount of panning and facilitating faster visual searches.

We observed that the participants of our user study performed the photo finding task 45% faster on a time quilt-based interface (TQ) than on a space-filling interface (SF) and a timeline (TL). We attributed this difference in task completion time to TQ's better visualization of temporal order over SF and better use of space over TL.

As a second, complementary method for scaling up zoomable photo browsers, we proposed the use of representative photos to implement semantic zooming in order to provide a more informative overview in place of the diminishing thumbnails. We found a subjective preference for this method.

In future work, we plan to verify whether representative photos improve task completion time quantitatively. We would also like to better understand the trade-off between the time-ordering and the space-filling layout criteria: we plan to pinpoint the collection size at which time-ordering makes gains over space-filling.

The time quilt layout itself has much room for exploration and improvement. For example, changing the wrapping columnar layout to a zig-zag layout eliminates column breaks and reduces the amount of panning as the user scans over photo clusters sequential in time. For another example, enforcing a minimum size for representative photos makes sure that every representative photo is recognizable regardless of whether its cluster is large or small.

We believe that the seemingly non-systematic patterns in time quilt layouts facilitate spatial memory in re-finding photos from a large collection. We wish to compare time quilt's effectiveness in inducing spatial memory to that of space-filling layouts, hoping to show that whitespace is not necessarily wasteful.

Finally, we would like to apply this idea of trading off space-filling for time-ordering to laying out visual information from domains other than digital photographs.

ACKNOWLEDGEMENTS

Many thanks to those who participated in our user study and to Ben Bederson, Mary Czerwinski, Julie Guinn, and John Platt for their thoughtful discussions on our work.

REFERENCES

[1] Bederson, B.B. and J.D. Hollan. Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. *Proc. UIST 1994*, 17–26.

[2] Bederson, B. PhotoMesa: a zoomable image browser using quantum treemaps and bubblemaps. *Proc. UIST 2001*, 71–80.

[3] Combs, T. and B. Bederson. Does Zooming Improve Image Browsing? *Proc. Digital Libraries 1999*, 130–137.

[4] Drucker, S., C. Wong, A. Roseway, S. Glenner, and S. De Mar. MediaBrowser: reclaiming the shoebox. *Proc. AVI 2004*, 433–436.

[5] Fertig, S., E. Freeman, and D. Gelernter. Lifestreams: An Alternative to the Desktop Metaphor. *Proc. SIGCHI 1996*, 410–411.

[6] Frohlich, D., A. Kuchinsky, C. Pering, A. Don, and S. Ariss. Requirements for Photoware. *Proc. CSCW 2002*, 166–175.

[7] Graham, A., H. Garcia-Molina, A. Paepcke, and T. Winograd. Time as essence for photo browsing through personal digital libraries. *Proc. Joint Conf. Digital Libraries*, 2002.

[8] Jul, S. and G. W. Furnas. Critical Zones in Desert Fog: Aids to Multiscale Navigation. *Proc. UIST 1998*, 97–106.

[9] Kang, H. and B. Shneiderman. Visualization methods for personal photo collections: Browsing and searching in the PhotoFinder. *Proc. IEEE Intl. Conf. on Multimedia and Expo*, 2000.

[10] Kuchinsky, A., C. Pering, M. Creech, D. Freeze, B. Serra, and J. Gwizdka. FotoFile: a consumer multimedia organization and retrieval system. *Proc. SIGCHI 1999*, 496–503.

[11] Mackinlay, J., G. Robertson, and R. DeLine. Developing calendar visualizers for the information visualizer. *Proc. UIST 1994*, 109–118.

[12] Mullet, K. and D. Sano. Designing Visual Interfaces – Communication Oriented Techniques. Sunsoft Press, Prentice Hall, 1995.

[13] Perlin, K. and D. Fox. Pad: An Alternative Approach to the Computer Interface. *Proc. SIGGRAPH 1993*, 57–64.

[14] Plaisant, C., B. Milash, A. Rose, S. Widoff, and B. Shneiderman. LifeLines: Visualizing Personal Histories. *Proc. SIGCHI 1996*, 221–227.

[15] 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.

[16] Ringel, M., E. Cutrell, S. Dumais, and E. Horvitz. Milestones in Time: The Value of Landmarks in Retrieving Information from Personal Stores. *Proc. INTERACT 2003*.

[17] Rodden, K. and K. Wood. How Do People Manage Their Digital Photographs. *Proc. SIGCHI 2003*, 409–416.

[18] Suh, B., H. Ling, B. Bederson, and D. Jacobs. Automatic Thumbnail Cropping and Its Effectiveness. *Proc. UIST 2003*, 95–104.