Beyond the Shoe Box: Foundations for Flexibly Organizing Photographs on a Computer

Karen D. Grant, Adrian Graham, Tom Nguyen, Andreas Paepcke, and Terry Winograd

Computer Science Department
Stanford University
Stanford, CA 94305 USA
{kgrant, aag, tomnguyen, paepcke, winograd}@stanford.edu

ABSTRACT

As a foundation for designing computer-supported photograph management tools, we have been conducting focused experiments. Here, we describe our analysis of how people initially organize collections of familiar images. We asked 26 subjects in pairs to organize 50 images on a common horizontal table. Each pair then organized a different 50-image set on a computer table of identical surface area. The bottom-projected computer tabletop displayed our interface to several online, pile-based affordances we wished to evaluate. Subjects used pens to interact with the system. We highlight aspects of the computer environment that were notably important to subjects and others that they cared about less than we had hypothesized. For example, a strong majority preferred computer-generated representations of piles to be grid-shaped over several alternatives, some of which mimicked the physical world closely and others that used transparency to save space.

Keywords

digital photographs, piles interface, tabletop display, user study, personal digital library, pile representations, pile manipulation, collaboration, spatial organization system, clusters, interaction design

INTRODUCTION

Countless consumer-level photographers "archive" the fruit of their labors in stacked shoe boxes. The boxes themselves overflow with envelopes from photographic development laboratories. This ad hoc storage scheme renders these photos of trips and celebrations difficult to share and enjoy. The growing acceptance of digital cameras merely transfers this problem to the computer and its file system. In fact, the absence of incremental costs for digital shots can exacerbate the problem.

Tools for organizing digital images lag behind the

Figure 1: Subjects organizing photographs on the computer table

quickening stream of uploaded snapshots. Most recreational photographers continue to experience the organization of photographs as onerous, preventing them from browsing, retrieving, and sharing photographs with pleasure. There are a variety of applications for photo organization on a workstation [16, 17]. We were particularly interested in organization as a collaborative effort on a shared collection of photographs. We chose to experiment with a "computer table" that combines the informality of multi-person, shared tabletop activity with the advantages of computer support. The table has the appearance of a common dinner table with a glass surface that serves as a bottom-projected computer screen. Replacing the mouse are eBeam [18] ultrasonic pens.

As with other technology that supports collaborative work, tabletop displays have been rarely deployed, particularly as they have been expensive and difficult to build at an appropriate resolution. However, with new display technologies emerging for computer and television screens, the opportunities for larger tabletop displays especially suited for collaboration are increasing.

When people work together sorting photos on a physical table, the style is fluid and informal -- very different from the structured file system hierarchies and grids that appear in typical computer programs. We set out to explore software that is better suited to this less formal task style, beginning with studies of how the widely known "pile" metaphor can be deployed in service of a digital photo-organizing tool on a computer table. The interface metaphor of piles adapts to the computer desktop the widely employed method of organizing related material by physically stacking it into little towers -- or piles. The tendency to generate piles and the potential for using this organizing scheme in computer interfaces have been documented, for example, in work by Malone [9], Mander, et al. [10], and Rose, et al. [12].

To our knowledge, however, the use of piles for organizing personally familiar images on a traditional table, or on a computer table, has not been studied in detail. Experimentally verified answers to a variety of questions can impact the design of related computer-based tools. When do users construct piles while organizing pictures? Can we find patterns by which users shape individual piles and arrange piles with respect to each other? How do users physically interact with piles? Can we identify some small number of manipulations that accomplish the rearrangement, image placement, and miniature searches for images that comprise the photo-organizing task? Would users follow the same organizing strategies on a computer table as they do on a regular table?

In search for these answers, we conducted two studies:

1. We observed pairs of users organizing a set of conventional photographic prints (photos) on a conventional table. The photos showed still frames from the movie "The Wizard of Oz" (1939), which all subjects had watched at some point in their lives. Prior to the organizing task, each subject read a summary of the movie to refresh his or her recall of the movie. This enabled us to match in a multi-subject experiment the normal task situation, in which people organize photos of personally familiar people and events, in the company of others familiar with them.
2. We had the same subjects organize a different set of still frames from the movie on our computer table, for which we had implemented a simple pile-based organizing system. This tool was not a complete photo-organizing system dealing with input and output, files, etc. It was built specifically to test the use of piles and the structures naturally created by the subjects.

We analyzed videos of the sessions, logs of captured interactions, and surveys that each subject completed.

Related Work

Several projects have examined the strengths and shortcomings of computer tables [15]. Elliot and Hearst [3] compared user preference and performance for two architecture-related tasks on a tablet, a desktop monitor, and a tilted computer table. They did not consider the pile metaphor, and their focus and experimental setup differed considerably from ours. Nevertheless, we were able to confirm some of their results where our studies intersected with theirs.

Other related studies have considered the domain of organizing Web site bookmarks. Although this is a different task and object domain from photos, there is overlap in the focus on storing and organizing collections of objects. A few studies explore a visual representation approach and have considered spatial location. Data Mountain [11] represents Web sites as thumbnail images and allows the user to move and group images in a "3D" plane. Although, the issue of object occlusion is examined in this research, the researchers did not seem to consider transparency, orientation, or the other visual representation characteristics that we are investigating. The Lifestreams system [5] uses a time-ordered approach in which the documents on the user's workspace are represented visually as icons and then organized strictly chronologically without the expression of semantic relationships.

As a domain area, digital photography has prompted a rich field of research. The traditional metaphor for organizing electronic information is a file-manager system such as those used in the popular commercial programs ACD Systems' ACDSee [16], Apple iPhoto [17], and Microsoft Window XP's My Pictures [19]. Much work on digital photograph management has emphasized the tasks of indexing, searching and browsing, which are complementary to our task of supporting users in the initial organization of a photo collection. We share with Kuchinsky et al. [8] and Shen et al. [13] the belief that for personal photographs, storytelling is a powerful organizational mechanism. Our focus, however, was not to present a particular organizational support, but to explore the variety in how people organize photographs with a storyline in mind. We aimed to uncover methods that support maximal flexibility in supporting this interaction style.

Like Kuchinsky et al.'s scraplets [8] and Shipman et al.'s objects [14], our subjects clustered photographs. However, we found that our subjects, when unencumbered, occasionally used all 360 degrees of orientation to indicate relationships among photographs, not just vertical or horizontal aligned positions. One of our goals was to explore how important this flexibility is for the subjects. Some indexing approaches automatically extract features such as color [4] or face recognition [8]. PhotoFinder provides a search system that focuses on specifying queries and presenting results that allow specific items of interest

Figure 2: Subjects organizing photographs on the traditional table

to be found efficiently [7]. Browsing systems have been developed to enable navigation over a large set of previously organized photos, which deliberately facilitate "serendipitous discovery" [1, 6, 7, 13]. Our focus in this research was on the earlier stage of initially categorizing and sorting a set of personally meaningful photographs.

EXPERIMENT & SYSTEM DESIGN

Experimental Setup

Figures 1 and 2 show, respectively, the computer table with its bottom-projected PC color screen tabletop and the traditional table, which was the same table with a cardboard cover with black tape on the cardboard delineating a 1.24m x .94m (49" x 37") work area. This area was identical to the screen below it.

Two pens were provided, each of which, when touching the screen, operated like a physical mouse. The hardware could track only one pen at a time, so subjects alternated use (see discussion).

The computer screen showed the images on a plain light-blue background. We did not expose a traditional Macintosh/PC desktop or any menus, scrollbars, etc. In our interface, the photographs were the dominant visible objects on which the user performs direct manipulations. A screen shot of the main interface can be seen in Figure 3c. An overhead digital camera recorded the contents of the table surface. A video camera on the sidelines recorded video and all sound in the room.

Experimental Procedure

The 26 subjects (14 women, 12 men) ranged in age from 15 to 51 and averaged 26. They were drawn from a diverse range of professions, including art, education, library sciences, food services, law, business, and computer science. Fifty-six percent still used chemical film as their primary photographic technology, while 44% used digital cameras. When asked how they currently store their photographs, 69% replied that they initially keep their photos either in the original envelopes in a shoe box or in a bag.

The experiment consisted of four parts. In Part I subjects organized a set of physical photographs without computer support, constrained only by the workspace surface area. In Part II the same subjects repeated the organizing task with (different) images on a computer table. In Part III we showed subjects mock-ups of eight designs for how the computer might display automatically generated piles of photographs. In Part IV we tested recall of photo placement. The results of Part IV are not covered in this paper.

Our study was designed to mimic users' (future) transition from physical interactions with photographs to a new computer table supported interaction style. We wanted to observe this transition when it was made under controlled conditions of keeping the environmental and task factors as comparable as possible.

We also wanted to make sure to learn as much as possible about users' spontaneous, unpolluted manipulations of personally familiar photographs in a purely physical setting. We therefore exposed all subjects to the physical table first and then to the computer table. Note, however, that in this setup, learning effects along the duration of each subject session are not compensated for. Comparisons of behavior *between* tables are therefore of limited accuracy. We did control the study *within* each table type and report results from those statistics only. Our observations between table types are offered informally.

After completing each part, each subject filled out a questionnaire with 41 Likert-scale and other multiple-choice questions as well as an invitation to fill in free-form comments. All subjects made use of the free-form fields, and most did so extensively.

During Parts I and II, subjects worked in pairs. The partners had been acquainted with each other an average of 5.9 years. All pairs had known each other at least a few months.

Part I

We gave subject pairs 45 photographs, each showing one of 650 frames extracted from the movie "The Wizard of Oz." The same set was used for all subjects. We gave them 10 minutes to organize the photos on the traditional table. Although we wanted to learn how users would freely organize photographs personally known to them, we did give our subjects the framework of a task. We instructed them to work towards a scrapbook or online photo journal for Dorothy, the movie's central character. Subjects were instructed to stay within the area outlined by the black tape.

We asked them to talk aloud during their process. After they had arranged the 45 photos to their satisfaction, we provided an additional five images to add to their implicit structure. At the end of this time, we took several photographs of the photo arrangement on the tabletop.

Part II

In Part II, we again instructed the subjects to keep the scrapbook or online photo journal in mind. We exposed the computer tabletop screen that showed a perfectly edge-aligned stack of 25 images on the lower right corner. The images were again frames from the movie, but different from those used in Part I. Piles were indicated by an eight-pixel color frame around the (often irregular) outline of each pile. For reasons of resolution and consequent viewability, the photographs on the computer table had a diagonal 15% larger than the paper photos on the traditional table.

The software enabled subjects using the pens to perform several operations:

Moving photos: A photo could be dragged by pressing the pen down on the photo, moving the photo, then lifting to release it. If the photo was released overlapping any image in a pile, it was added to that pile. If the photo was initially in a pile and released outside the pile, it was removed from the pile.

Moving piles: An entire pile could be moved by a drag-and-drop operation beginning in its colored border. Lifting the pen from the surface when the color frame of a dragged pile touched or overlapped the color frame of the target pile merged two piles. Otherwise, when the color frames were not touching, piles remained overlapped, yet kept their separate identities and distinct colored borders.

Enlarging photos: A pen tap on an image enlarged the image to about one fourth the size of the total work area. Tapping again anywhere on the tabletop returned the image to its regular size.

There was no affordance for rotation of the images, which remained ortholinear.

At the end of 10 minutes we introduced the subject pair to a new way of using the computer. Tap-holding the pen on a pile border popped up a menu with three alternative options for displaying the images of an entire pile all at once or collapsing a pile into an icon. Data on the use of these options is not included in this paper. Once subjects understood these new facilities, we asked them to continue to organize an additional 20 images. After they had arranged all 45 photos to their satisfaction, we provided an additional five images to add to their structure, as in Part I. Our software logged data such as all pile locations, sizes of piles, etc.

Part III

For Part III we separated the subject pairs. Subjects were shown 20cm x 28cm (8" x 11") mock-ups printed at a resolution comparable to that of the computer table (~48 dpi) and mounted together on a poster board. (See Figure 6.) They were asked to indicate their first, second, second-to-last, and last choices.

HYPOTHESES AND RESULTS

Shapes that Make a Pile

Based on our initial observations and our review of the literature, we hypothesized that subjects would create a wide variety of pile shapes but that certain core patterns would naturally emerge. We found this to be true on both the traditional table and the computer table. During our pilot studies, we found organizational layouts similar to Shipman, et al. [14]. However, our goal was to identify optimal support for organizing photographs on a computer. We needed more detailed data on the relative frequency of the organizational structures and on the frequency of change among the structures. Our emphasis was on photographs that have meaning to the organizer. We also were interested in how people interacted with the structures while building them. We wanted to find what people did with both physical- and computer-based images that were specifically photographs. We found that among the wide variety of pile representations created in the study, five distinctive patterns emerged: *Grid, Column, Row, Overlap,* and *Stack*.

Comments made by users of the *Grid* representation (Figure 3a) indicated that subjects were concerned with avoiding overlap and with making the photos align flawlessly.

Subjects also created orderly *Column* (Figure 3b) and *Row* (Figure 3c) piles. Even when the subjects had infinite freedom to orient the photos in any direction they wanted on the traditional table, when using the orderly arrangements of *Grid, Column,* or *Row*, they universally chose a perpendicular orientation, aligned with the table edges.

Stack (Figure 3d) and *Overlap* (Figure 3e) both contain photographs on top of each other. A *Stack* involves nearly complete overlap, while in an *Overlap,* the photographs are more spread out. Some subjects aligned photos in *Stacks* so precisely that only the topmost image was visible.

On the traditional table, in preparation for constructing piles, all subjects would initially lay the photographs on the table with little or no overlap, including those who later used *Overlap* piles.

3a) Grid

3b) Column

3c) Row

3d) Stack

3e) Overlap

Figure 3 (a-e): Pile Shapes created on the traditional and computer tables.

Comparison of Shapes

Our initial hypothesis was that because we had tried to provide photo-manipulating actions that mimicked those of the real world, subjects would use the same shapes on the two tables. Indeed, the subjects created similar shapes on both the traditional table and the computer table. However, the prevalence of certain patterns that were created was different—due to learning effects, or preference related to the computer table.

For example, on the traditional table, the percentage of subjects using *Grid* as their primary pile type (23%) was not significantly different from the percentage using *Overlap* (15%). However, on the computer table, *Overlap* (54%) was significantly more popular than *Grid* (0%) ($p < 0.025$). We cannot reliably compare the usage of *Grid* on the traditional table versus the computer table because our experiment was not set up to explore this question. We informally note that the change in *Grid* usage between the two table types was not statistically significant.

Sixty-seven percent of the subjects who used some of the more regular rectilinear representations such as *Column* and *Grid* on the traditional table switched to the looser configuration of *Overlap* on the computer table. Furthermore, when moving from the traditional table to the computer table, significantly more subjects switched from *Column* and *Grid* to the *Overlap* representation than switched from *Overlap* to one of these two more rigid representations ($p < 0.025$). The subjects reported that they were less concerned about obscured images while working

Figure 4: Ambiguous pile membership on the traditional table.

Figure 5: Piles on the computer are surrounded by colored borders.

on the computer table because they could easily select even an almost covered photograph with a click to enlarge it.

Distinction of Grouping

In laying out groups of physical photos, there can be ambiguity in where one pile ends and the next one begins, as shown in Figure 4. In designing the computer version, we included color borders around piles (Figure 5). Our hypothesis was that this would make it easier to distinguish pile boundaries.

This hypothesis was confirmed at a 95% significance level ($p < 0.025$), but the difference between perceived ease of pile distinction between the two tables was not as strong as expected. Only 69% found a pile grouping easier to distinguish from other pile groupings on the computer compared to the traditional table. However, 35% of the subjects explicitly commented in their written surveys that one of their favorite features with our application was that the piles had background colors, whereas none listed this as a negative feature.

Eighty-five percent of the subjects reported that the computer's concept of what constituted a pile matched their own notion.

Spatial Relationships

One of our interests in developing a pile-based computer interface was the hypothesis that people use the spatial arrangements among piles in meaningful ways.

This hypothesis was supported by the visible results and subjects' comments. Eighty-five percent created spatial relationships among piles on either the traditional table or computer table. Some ordered according to the narrative chronology, placing the earliest photographs of the story in the upper left portion of the work surface and then ordering them either across and then down or down and then across. Some used spatial location to indicate thematic similarity. For example, one pair of subjects placed their Scarecrow pile next to their Dorothy pile and created an area for "character piles," explicitly giving a name to this larger grouping.

Our software did not easily support ordering of individual photos within piles. Yet despite this difficulty, the majority of subjects did manage to do some ordering within piles, even on the computer, with uses of both horizontal and vertical ordering.

For example, as demonstrated in Figure 3c, some subjects ordered the photos within a pile from left-to-right even though the experimental setup did not explicitly support this activity. Subjects varied in which axis they preferred to order along. This difference among subjects was particularly pronounced in the *Grid* formation on the traditional table. When the set of photos had a within-pile order, some subjects chose a left-to-right, row ordering while some subjects chose an up-to-down, column ordering. This is relevant for future designs because we cannot assume a universal understanding or preference for the row or column ordering in a representation.

Pile Display Alternatives

We produced static mock-ups of eight alternative methods for displaying a pile on a computer table, which included four variants of an *Overlap* arrangement, three of a *Stack* arrangement, and one of a *Grid* arrangement, as shown in Figure 6.

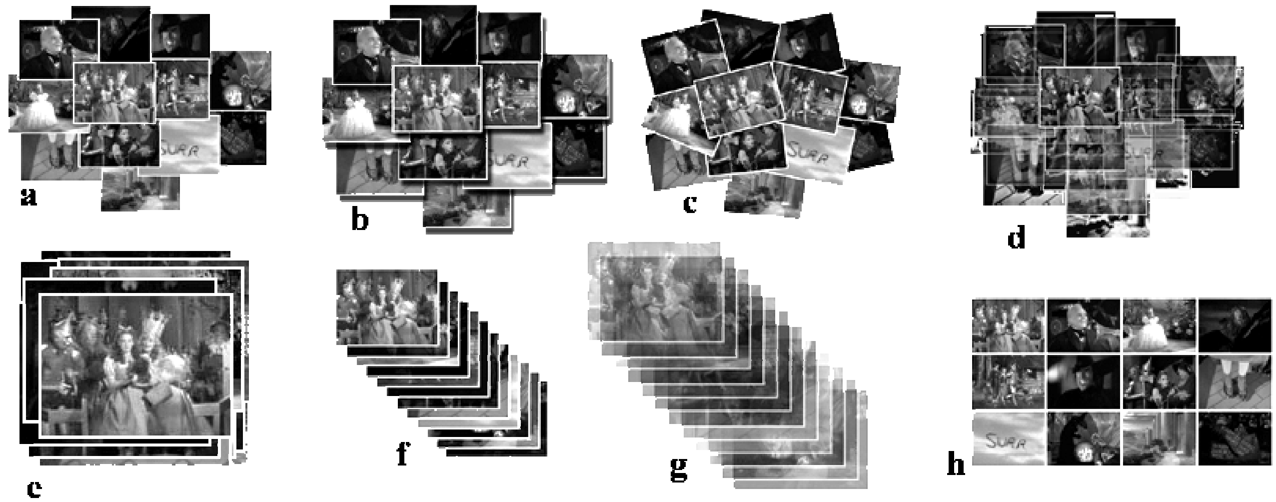

Figure 6: Eight choices for Pile Representations: (a) Simple Overlap (b) Drop Shadow Overlap (c) Tilting Overlap (d) Transparent Overlap (e) Informal Stack (f) Accordion Stack (g) Transparent Accordion Stack (h) Grid

(a) Simple Overlap: Images are overlapped so that enough of each image is visible for identification.

(b) Drop Shadow Overlap: Like *Simple Overlap*, with 7-pixel drop shadows added to each image to create the illusion of three dimensions.

(c) Tilting Overlap: Like *Simple Overlap*, with images randomly tilted by 0 to 45 degrees away from the vertical axis in either direction.

(d) Transparent Overlap: Like *Simple Overlap*, with all images other than the top image at 50% transparency. Each image is framed by a 2-pixel white border.

(e) Informal Stack: Each successive image is randomly shifted up, down, left, or right by a distance between 0 and 10 pixels.

(f) Accordion Stack: Each image is shifted to the right and down by 5 pixels.

(g) Transparent Accordion Stack: Like *Accordion Stack*, with all images set to 70% opacity, rather than the 100% in the schemes above.

(h) Grid: Non-overlapping rows and columns.

Our hypothesis was that for computer-based images, subjects would prefer renderings that were more “aesthetic” and less "formal," taking advantage of features such as transparency and arbitrary rotation. Subjects were asked to indicate their first, second, second-to-last, and last choices, with the results shown in Figure 7.

Our hypothesis was not confirmed inasmuch as the less conventional representations with rotation and transparency were not highly rated. The top preferences, *Grid*, *Simple Overlap*, and *Drop Shadow Overlap,* are all rectilinear and expose maximal surface area. Moreover, *Grid* and *Simple Overlap* have one-to-one counterparts in the physical world. The drop shadows of *Drop Shadow Overlap* attempt to mimic the real world in using a drop shadow to indicate three-dimensional depth.

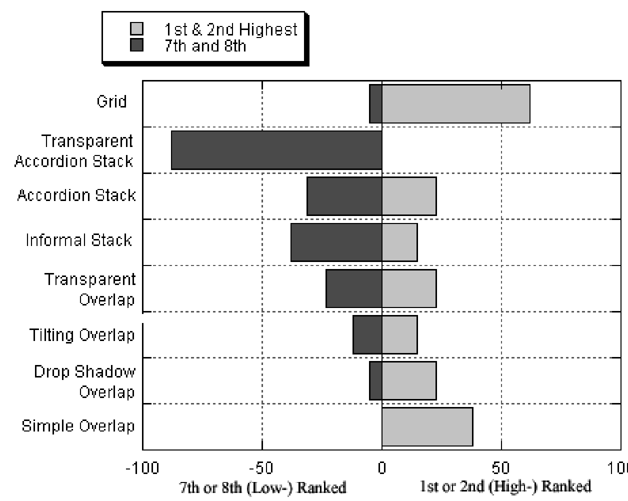

Figure 7: Subjects' Evaluation of Pile Representations. Percentages to the right of zero are the sum of the percentages of subjects who chose the representation as first or second choice; those on the left are the sum of the percentages of subjects who chose the representation as last or next to last.

The widespread preference for the *Grid* might suggest that maximum area exposure is an advantage despite the negative effect on the use of space. Several subjects who favored *Grid* acknowledged this resource disadvantage.

Consistency between Preference and Action

We hypothesized that there would be a familiarity effect in which people would express preference for the forms they had actually used rather than others that our software did not enable them to use directly.

Data indicated that these stated preferences matched subjects' behavior. Using physical movement or standard drag/drop, subjects could realize *Grid*, *Accordion Stack*, *Simple Overlap*, and *Simple Stack* on both the traditional and the computer table. Even with this restriction, 54% of the subjects had used their most preferred representation during the previous work with their partner on the traditional table. On the computer table, 73% realized their most preferred representation.

Ambivalence: Machine vs. Physical Universe

Through the experiments, our subjects gave evidence of ambivalent expectations for how they expected the computer to behave. On the one hand, subjects wished the computer to mimic the physical world as much as possible. For example, two subjects rejected *Transparent Overlap* on the grounds that this representation did not "look natural." Several subjects regretted the inability to pick up individual photos to look at and the loss of physical contact with the photographs. On the other hand, several subjects judged *Tilting Overlap*, the pile with photos arranged informally, as not "computer-like" enough for their taste: "Computers should help add neatness. The disarray is unnecessary. I might as well not use a computer." Or, "Too messy for a computer. [I] expect more."

Beyond an expectation of neatness, one subject even praised the more rigid computer setup because it "forces good categorization from the start." Others requested snap-to-grid at both the image and pile levels. We observed this strong desire for strict rectilinearity also on the traditional table, although to different degrees among subjects. For example, partner A of one subject pair exhibited strong initiative in the organizing process. The pair was following a *Grid* discipline. Partner A made the decisions along the way and retained control over photo placement. Partner B, meanwhile, trailed behind A, straightening each and every photo to make the grid perfectly rectilinear. Not surprisingly, B requested computer support for lining up collapsed piles as well.

At the same time, subjects recognized and appreciated the flexible aspects of the computer setup. For example, the ease of moving piles as a unit on the computer table was widely acknowledged, as was the ease of temporary image enlargement. Several subjects, including B, requested the ability to scale individual photos to different sizes: "Different pictures have different emotional value and should therefore be given different sizes/priority."

It seems as if this task caused subjects to draw on two mindsets at once. The task resonated with subjects' informal, domestic, creative tendencies. At the same time, however, the task spoke to their structured, disciplined streaks that are more widely associated with computers.

CONCLUSION

Flexibility and fluid interaction in general are important characteristics of tools for early phase work, such as initial photo organization. Piles offer this kind of flexibility, but they do not have the more highly developed structures of systems with file hierarchies, tagging, timelines, and the like. We do not expect to completely replace these, but want to find an appropriate mix in which the early informal stages are effected with simple, natural mechanisms and then integrated seamlessly with the more fine-grained, reflective work to which highly structured systems are geared.

Our study suggested several lessons for the design of systems to help consumer-level users organize their personal digital photographs.

First, it is important to provide a variety of pile patterns to support user diversity. The rectilinear family of patterns (*Row, Column, Grid*) and the more space-efficient occlusion-tolerant family (*Stack, Overlap*) were both widespread.

At the same time, the computer "freed" some of the most structurally rigid *Column* and *Grid* subjects to change their habits. Sixty-seven percent of them moved to the looser *Overlap* pattern. User interface affordances that support this loosening were enthusiastically welcomed, such as the use of overlap, the use of one-click temporary enlargement of images, and the use of background color to enable messy pile placement.

On the other hand, some computer-only interface elements, such as transparency and layer-by-layer pile inspection (flip), were not successful.

We saw a preference for the *Grid* pattern as the favorite method for the computer in displaying automatically generated piles. Given the relatively small number of photographs in the experiment, we were not able to see how these preferences would change if large numbers made *Grid* layouts unfeasible. There are further options, including reduced thumbnails that could be explored to keep the full visibility and regularity of grids while reducing the use of space.

Like Elliott et al. [3], we were surprised that image resolution was less of an issue than we had feared. Some subjects did realize that the computer images were not as attractive as the physical photographs. But except for a computer graphics engineer, no one listed this in his or her responses to questions about the shortcomings of the computer system despite a resolution of ~21 dpi compared with a printed photograph (~200 dpi).

Finally, although we did not systematically compare our tabletop system with the same software on a conventional workstation, we learned several facts about the use of tabletop displays. People adapted quickly to using the pen as a pointing device, and the parallax problems observed by Elliott et al. [3] were not an issue for us. We don't know whether the tilted table of their study, different pens, or perhaps the size of the manipulated objects account for this difference. On the other hand, we were quite clearly told (by 69% of our subjects) that the pens of both partners who collaborate on tasks need to be operable simultaneously. All subjects succeeded in working around this hardware limitation, but a completely concurrent dual-pen operation is clearly a future necessity. New hardware [2], which we plan to incorporate in future work, addresses this problem.

The software prototype for the computer table was designed to measure particular parameters. It was far from a best-design prototype and can be improved in many areas to develop a useful system. We were strongly encouraged that computer support, and in particular support on a large display, can add value in this domain. We realized that flexibility is a requirement for aspects such as pile shape. But at the same time, flexibility does not seem to be needed throughout the system. For example, none of the subjects requested control over the colors of the pile backgrounds. An understanding of where degrees of freedom are truly required is particularly important in a system that is intended for lay people.

We found that our study supported the use of a simple computer interface and direct manipulation. Nevertheless, certain computer features could add convenience without compromising the simplicity principle. For example, several pairs liked to initially spread out all photos to view them. This took time on a physical table; a computer tool could easily help with this. Also in our next phase of testing, we would consider offering a mechanism for sorting within piles and the ability to add captions. An open question that we did not address was movement based on flinging and acceleration. Reach and accessibility issues on a large tabletop might suggest such additional features.

On a broader level, computer-supported interaction with personal photographs straddles a 'fault line' between domestic, private activity on one side and efficiency-oriented workplace activity on the other. Sometimes the resulting emotional and practical requirements are in contention. But this treacherous design ground, when carefully navigated, promises to yield insights that will be of use on both sides of that line.

ACKNOWLEDGMENTS

We wish to thank the subjects who participated in this study for their time and enthusiasm. This work was supported in part under NSF Grant IRI-9817799.

REFERENCES

1. Bederson, B.B., PhotoMesa: A Zoomable Image Browser Using Quantum Treemaps and Bubblemaps. in *UIST '01*, (Nov., 2001), 71-80.
2. Dietz, P. and Leigh, D., DiamondTouch: A Multi-User Touch Technology. in *UIST '01*, (Orlando, FL, 2001), ACM Press, 219-226.
3. Elliott, A. and Hearst, M.A. A Comparison of the Affordances of a Digital Desk and Tablet for Architectural Image Use Tasks. *International Journal of Human-Computer Studies*, *56* (2), 2002. 173-197.
4. Flickner, M., Sawhney, H., Niblack, W., Ashley, J., Huang, Q., Dom, B., Gorkani, M., Hafner, J., Lee, D., Petkovic, D., Steele, D. and Yanker, P. Query by image and video content: The QBIC system. *IEEE Computer*, *28* (9), 1995. 23-32.
5. Freeman, E. and Gelernter, D. Lifestreams: A Storage Model for Personal Data. *ACM SIGMOD Bulletin, March 1996*. 80-86.
6. Graham, A., Garcia-Molina, H., Paepcke, A. and Winograd, T., Time as Essence for Photo Browsing Through Personal Digital Libraries. in JCDL 2002, (July, 2002).
7. Kang, H. and Shneiderman, B., Visualization Methods for Personal Photo Collections: Browsing and Searching in the PhotoFinder. in *ICME 2000*, (July, 2000).
8. Kuchinsky, A., Pering, C., Creech, M.L., Freeze, D., Serra, B. and Gwizdka, J., FotoFile: A Consumer Multimedia Organization and Retrieval System. in *CHI '99*, (Pittsburgh, PA, 1999), 496-503.
9. Malone, T.W. How do people organize their desks?: Implications for the design of office information systems. *ACM Transactions on Information Systems (TOIS)*, *1* (1), 1983. 99 - 112.
10. Mander, R., Salomon, G. and Wong, Y.Y., A 'pile' metaphor for supporting casual organization of information. in *CHI '92*, (1992), ACM Press New York, NY, USA, 627 - 634.
11. Robertson, G., Czerwinski, M., Larson, K., Robbin, D.C., Thiel, D. and Dantzich, M.V., Data Mountain: Using Spatial Memory for Document Management. in *UIST '98*, (San Francisco, CA, 1998), ACM Press, 153-162.
12. Rose, D.E., Mander, R., Oren, T., Poncéleon, D.B., Salomon, G. and Wong, Y.Y., Content awareness in a file system interface: implementing the 'pile' metaphor for organizing information. in *SIGIR : ACM Special Interest Group on Information Retrieval*, (Pittsburgh, PA, USA, 1993), ACM Press New York, NY, USA, 260 - 269.
13. Shen, C. and Neal Lesh, B.M., Paul Beardsley, Ryan S. Bardsley, Personal Digital Historian: User Interface Design. *Proceedings of Extended Abstracts of Human Factors in Computing Systems,* 2001. 29 - 30.

14. Shipman, F.M., Marshall, C.C. and Moran, T.P., Finding and using implicit structure in human-organized spatial layouts of information. in *CHI '95*, (Denver, CO, 1995), ACM Press/Addison-Wesley Publishing Co., 346-353.

15. Wellner, P. Interacting with paper on the DigitalDesk. *Communications of the ACM*, *36* (7), 1993. 86-97.

16. http://www.acdsystems.com

17. http://www.apple.com/iphoto

18. http://www.e-beam.com/

19. http://www.microsoft.com/windowsxp