An Active Browser for Personal Photo Collections

Ankit Gupta
University of Washington
ankit@cs.washington.edu

Aditya Sankar
University of Washington
adityas@cs.washington.edu

ABSTRACT
The size of personal photo-collections has been increasing exponentially due to recent technological advances in camera industry. The current photo-browsers act more as passive querying interfaces where the users can search for their photos and view them. This model relying on human memory to view photos does not work well in the case of large photo-collections and many photos, in-turn significant photographs tend to get lost with time. In this paper, we propose the idea of an *active* photo browser which makes photo-suggestions which could either be picked at random from users' personal collections or related to the current photo being viewed or related to the user's current setting. We first conduct a user survey to identify the need for such a tool and then run a user studies with a prototype photo-browsing tool to get feedback and to understand the users' mental model of connections between the photos. The knowledge of mental model will allow the photo-browser to make more relevant photo-suggestions in the future.

Author Keywords
Photo browser, Active suggestions

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation: User Interfaces

INTRODUCTION

Hardware advances in the digital camera industry have made cameras very accessible to people in the past few years. The cameras are also being integrated into mobile devices like laptops, tablets and phones which allow on-the-fly photography. As a result of this the size of the personal photo-collections has been rising exponentially. In October, 2009 *Flickr*, an exclusive photo-sharing website, boasted of having over 4 billion photos shared by the users [2]. On the other hand, *Facebook*, a social networking website which also allows people to share their photos, reported over 10 billion photos shared by people till October 2008 [1]. These numbers are more than an year old and would almost certainly have grown. This is a clear indication of the increasing number of photos that the users are capturing. Besides sharing their photos online, the users usually create physical prints, or just keep the photos stored on memory devices like the hard disk or CD/DVDs. In this paper, we aim to look at ways by which the people can have better experience with their personal photo-collections.

There are a lot of commercial photo-browsers today like Google's *Picasa*, Microsoft's *Live Photo Gallery*, Adobe's *Lightroom*, Apple's *iPhoto* etc. These tools allow users to tag photos and organize them based on those tags. Users can view all the photos of a desired person, or from an entire trip, or according to the capture time, or located in the same desktop folder. But all these tools are *passive* in nature. The user needs to make a query based on what photograph she/he wants to see and the tool provides a way to do that. This inherent passive nature of the tools does not allow the users to experience photos they might have forgotten about over the period of time, which is quite likely given the large personal photo-collections these days.

In this paper, we propose the idea of an *active* photo-browser. This active photo-browser makes photo suggestions to the users from their personal photo-collections based on the current *context*. We define the context as the user's current surroundings or current photograph that she/he is watching. This is different from random photo-screensavers or desktop photo-gadgets which show a random photo at a time and provide no connections to other photos if one wants to view them triggered by the interest in the current photo. In this work tackle two key questions:

- Is the active photo browser useful? This will help us identify the need for an active photo-browser.
- How do people make mental connections between their photos while browsing them? This will help us build a model to make photo suggestions.

Overview

We first discuss the related research in the field of personal photo-browsers. We then describe our hypotheses for this idea and an initial user survey that we did to validate these. Motivation by the survey-based validation of our hypotheses, we designed and implemented a suggestion-based photo browser. We describe the design issues, implementation and the user study next and then discuss the results of the study.

Project report for CSE510: Advanced topics in Human-Computer Interaction. Spring 2010.

In the end, we summarize the work and discuss some potential shortcomings of the study and ideas for a more comprehensive user study.

RELATED WORK

The problem of photo-browsing is a well understood and dense area of research. Apart from the popular online photo sharing services (Flickr, PicasaWeb, Facebook etc.), research projects such as WWMX [11], MediaBrowser [7] and TimeQuilt [8] use embedded EXIF information to index photos according to various single-dimensions such as time, geolocation etc. More recent work in computer vision [10] enables geo-spatial arrangement of photographs for 3D browsing.

In contrast to the above approaches which pay little attention to the actual mental model behind browsing photos, we conduct a rigorous study to understand the photo browsing behavior of the users and then design a browser prototype which allows people to browse through their photos in a way which is coherent with their mental model, thus leading to a better photo-browsing experience. The most relevant research projects are the PhotoMesa [6] and PhotoFinder [9] projects, which enable the user to express their mental-models via the interface, with semantic tagging. Our approach, however, requires less explicit training from the user and instead infers mental-models via an initial calibration phase.

HYPOTHESES AND SURVEY-BASED VALIDATION

We had the following hypotheses on which we based the idea of an active photo browser:

- The size of personal photo-collections is increasingly large.
- People do not revisit all of their photos often and will enjoy the 'aha' feeling of seeing some older photo that they might have forgotten about.
- People tend to remember photos based on the people in them, and/or the places/trips they are from and/or the approximate time when they happened.

We did a web-based user survey to validate these hypotheses. We created a survey using Google Forms and distributed it in the open through publicity channels like emails and social networking websites. We collected responses from 71 people in all (65 males and 6 females). The users had varying occupation and demographics and all were in the age range of 18-35 years.

Figure 1 confirms that people have an increasing number of photos in their personal photo-collections. We attribute this to the upcoming camera capabilities in the mobile devices as well as the economic accessibility for the digital cameras. Figure 2 suggests that most of these photos lie on the hard disk or are stored on CDs/DVDs. In the survey we did not ask the users how many of the photos from their photo-collection they share online. This can be interesting if we want to study privacy perception models for photos but beyond the scope of this work.

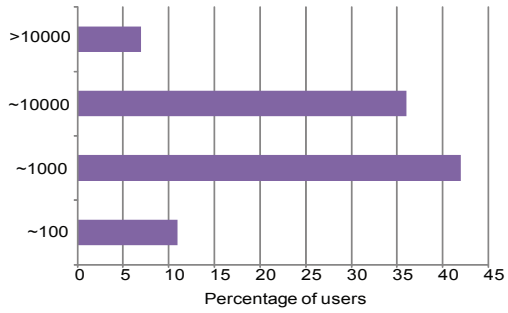

Figure 1. Typical size of personal photo-collections.

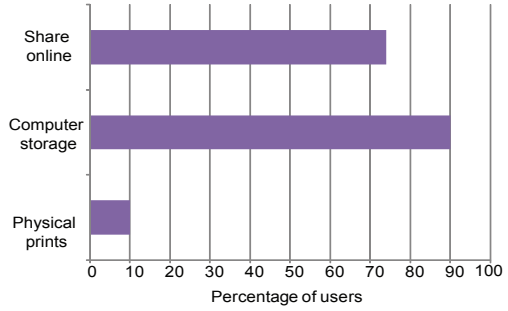

Figure 2. What people do with their photos.

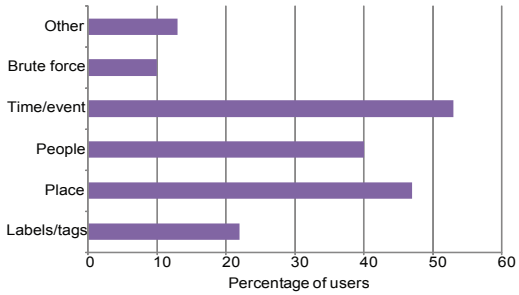

Figure 3. How people search through their photos.

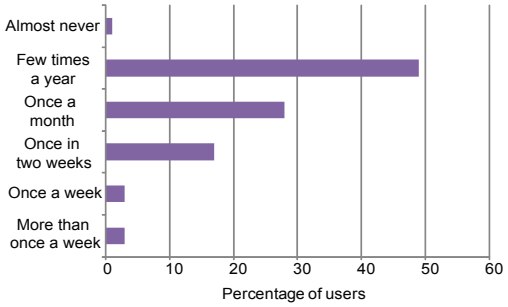

Figure 4. How often people revisit their photos.

From the survey, we wanted to get an idea for the need of an *active* photo-browser. For this we asked the people if they felt they revisited their photos often and what the rea-

sons were when they did so. Figure 4 shows that almost 50% of the users did not visit their old photos more than a few times a year. Some of them do so at being reminded of a photo while talking to friends/family or while looking at some other related photo. But most of the users did not specify any particular reason for revisiting their photos. We asked them if they felt the need to look their photos more often. This was to identify any already existing desires among the users for a suggestion-based photo browsing tool. Around 35% of the users felt that they would like to revisit their photos while the rest said that they were fine as is. We argue that 35% is significant section of users towards whom we this active browser. As might be the case, many of the users who did not feel the need might identify the need once they start using the tool. For this project, we did not do the comprehensive in-person user study with all the survey participants to be able to concretely see this behavior.

A key concept behind the idea of an active photo browser is that of photo suggestions. The browser will make these suggestions using the mental connections that people have between photos and their current context. To get an idea for this, we asked the participants about the properties by which they usually remember the photos. Figure 3 shows that people tend to search for a photo using different properties of a photo such as the people in it, the event, the place, the time, or other ways like the vaguely remembered folder structure or at times a brute force search. The relative weights to these properties give us some idea about what properties play a more important role in remembering photos and we base our in-person user study to analyze transitions between the photos along these properties.

USER STUDY

System Design

The goal of our user study is to infer individual users' mental models of their personal photo collections. The approach taken is to present a user with a few photo-browsing related tasks and collect 'think aloud' feedback at every step of the task. Formulating such a study to accurately capture a mental model, which is inherently an intangible aspect of human behavior is an interesting and challenging design problem. Based on the feedback from the initial user-survey, we proceeded to design an interactive software prototype, to achieve our goal.

Our design process started with an initial paper-prototype of our system. We, the authors, self-validated the design of our prototype by printing out several hundred images from our personal photo collections and running the study amongst ourselves. Figure 5 demonstrates our paper prototype in use.

The final design we propose is in the form of a digital canvas of photographs. The photographs are laid out on a grid with a prominent center photograph surrounded by eight *suggestions* from the users' photo collection. The affordances to the user are as follows:

1. The user may click on the center photograph, revealing a prominent large version of the photograph and dimming the rest of the canvas, illustrated in Figure **??**.

Figure 5. Paper prototype of our system being tested.

2. The user can click on any of the surrounding images to transition to them. This would change the center photograph to the one selected and present new suggestions around the center photograph.

3. During a transition the user is asked to select from a list of options, why he/she transitioned to a particular photograph. Transitions are described in further details in the next section.

Data from the transitions is anonymously collected and can be analyzed post-experiment, to infer why users transitioned between photographs, revealing information about the mental associations between images in their collections.

We hope this simple and minimalistic design serves to direct the attention to the actual visual content of the photographs, rather than captions or metadata. However, adding support for captions and metadata would be trivial.

Transitions

Our prototype displays a photograph in the center, surrounded by eight suggestions. Currently, the suggestions are picked at random from the rest of the users' photo collection. During a transition, the user is asked the reason for transitioning to that photograph. The options we provide allow the user to indicate that they observed a similarity in people, places, events/trips, context and visual appearance. The user may also select 'No reason' for their transition, to indicate that they picked the next image out of whim. An 'other' option allows the user to enter free-form text describing the reason for their change, if it cannot be categorized as any of the above. These proposed transition options are based on feedback we received during our initial survey, asking users how they currently search/browse their photo collections. We take advantage of their prior familiarity with this *dimensional* photo-browsing paradigm, but the key difference being that our design allows for *multi-dimensional* browsing rather than searching or browsing photos across any single dimension, such as the same event.

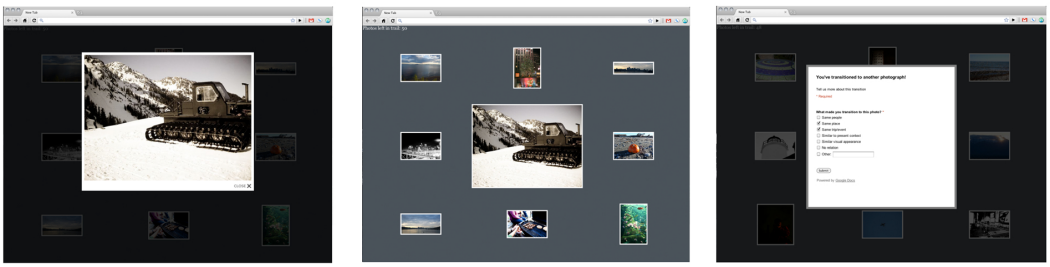

(a) (b) (c)

Figure 6. Screenshots of our prototype system in use: (a) User selects center photograph for an enlarged view, (b) Steady-state of system, waiting for user input, (c) Modal dialog presented to the user during a transition.

Implementation

Based on these ideas we proceeded to implement our digital prototype. The nature of the study required that we work within a few set constraints as described below:

Constraints

1. **The prototype must utilize users' personal photos**: This is critical for a meaningful inference of individual mental models. Running the study on generic photographs would be baseless as individual users tend to have different styles of capturing photographs.
2. **It should be possible to conduct the study remotely**: Since users' personal photo collections might be private, we must enable the study to be run remotely so that user may upload their own collections and run the study without the need for an external observer.
3. **Prototype must be cross platform**: As the study is remote, we have no control on the operating system or applications on the users' computers. Hence the system should be cross-platform and accessible to a wide audience.
4. **Photo collections are generally inconsistently/poorly organized**: Several users across various platforms have different ways of organizing their photo collections. Making meaningful unified use of this data is difficult.

Prototype Implementation

Based on these constrains we elected to implement our system as an online web-application which can run in a standards compliant browser across a variety of platforms and operating systems. Screenshots of our system are shown in Figure 6

Our system utilizes Google's Picasa Web online photo-sharing service to store users personal photos. The photo canvas itself is coded entirely in Javascript and HTML and is hosted online on the authors' website [5]. Transition data is collected via Google's web-form infrastructure and can be viewed only by the authors. Note that our prototype does not store users' photos, which remain on their personal Picasa Web accounts.

Our prototype runs without any server-side processing. Hence all our code can be obtained directly by viewing the source HTML of our online prototype [5].

Conducting The Study

Before the user begins the study, we direct them to a short instructional web-page[4], which describes the task and required tools to get started. It also explains the concept of transitions with the help of visual aids.

The user commences the study by uploading a randomly chosen subset of their photographs, four-hundred in number, to Google Picasa Web [3], a free, online, photo-sharing service. Picasa Web exposes an RSS feed containing links to the photographs and corresponding metadata. By copying the feed URL and supplying it to our web-prototype, we are able to display users' personal photos and display them on the HTML canvas.

The user transitions to photographs he or she would like to view next and provides transition meta-data along the way. The user is free to quit the study at any time, although we recommend that they complete fifty transitions. This is indicated by a counter on the top-left of the screen. The user is also free to provide more than fifty transitions, until a point that they are bored or begin to experience fatigue from the task.

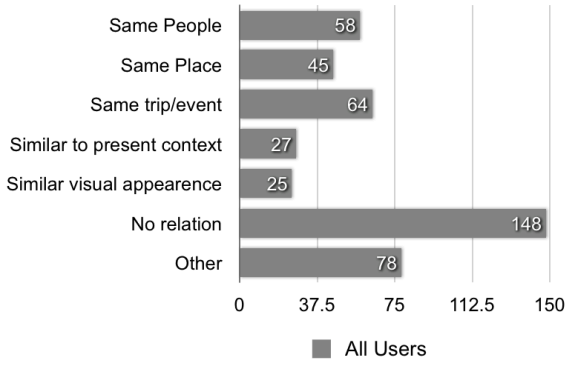

Figure 7. Summary of transition data collected across all users.

RESULTS

Study Demographics and Statistics

We conducted our study over two weeks. We had 5 (3 male, 2 female) participants who volunteered their time and photo

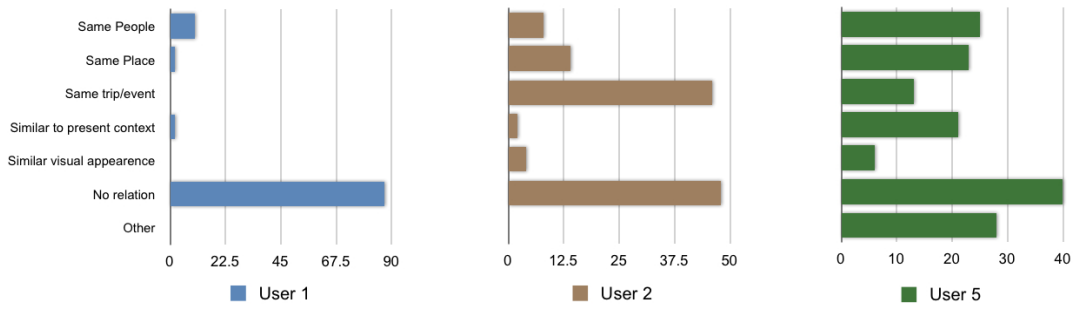

Figure 8. Individual tabulations of user data. We observe User 1 browses photos mostly at random, whereas Users 2 and 5 tend to browse images balanced across all categories.

collections. Each participant provided 400 randomly sampled photographs from their collections. We collected data for a total of 352 transitions ($\tilde{7}0$ per participant). The data was then tabulated and analyzed per-user since mental models can vary between individuals. We did however, also summarize the data across all users to get a feel for the average user behavior shown in Figure 7. The individual results and inferences are discussed in the next section.

Inference

Upon analyzing the transition data, we found that even across a small group of users, the reasons for transition vary widely. From this we infer that users display different behaviors and patterns while browsing photos, based on personal preference and the contents of their collection. The individual results from the experiment are presented in Figure 8. As we can observe User 1 prefers to browse photographs without any particular relation between them. User 2 in contrast prefers to browse images from the same trip or event, though a large number of browsed images also fell under the *no relation* category. The browsing behavior of User 3 is more balanced.

The data collected herein can be fed back into an active photo-browsing system such that the *suggestions* can be better modeled to individual users' preferences. We foresee a system which runs a short *calibration phase* similar to our proposed system, to infer browsing behavior and models itself to mach the user more closely. We hope that this approach will make the photo-browsing experience more enjoyable with the end goal of having users' revisit their old photographs more often and in a much more engaging and stimulating manner. Some anecdotal feedback we have received has encouraged us to believe that this indeed is the case.

Qualitative Feedback

We informally interviewed participants after they had completed the study to collect feedback on their experience. Some of the anecdotal feedback we received was:

- "Just the study tool is interesting and more non-linear and free-form than directory based browsing. Random sample also helps in looking uniformly at all the old photographs."
- "... my mom also joined in and we both had a fun time looking at our photos from different places and times."
- "The suggestions helped me watch other photos from my trips which I had not cared to share online"

Users in general felt that such an active photo-browsing tool would be helpful in re-discovering old photographs and that the user-study by itself was fun and entertaining.

CONCLUSION

In this paper, we proposed the idea of an *active* photo-browser which makes photo suggestions to the users besides displaying the current photo. We did a initial survey to validate the need for this tool. We found that the users have large photo-collections these days and a considerable number of the users did feel that they did not revisit their photos often and an active browser will be useful.

We designed a prototype of an active browser to allow users to actually experience the idea and give feedback on it. A second purpose of the prototype was to understand users' mental connections among their photos which will then allow the tool to make better photo-suggestions. We had some hypotheses connections from which we asked the users to choose one when they made a transition from the current photo to the next chosen one. These included similar people/place/trip/event/time. They could also input another connection that was not in the list. From the user studies, we found that people did make connections across the hypothesized dimensions. To our surprise, random photo suggestions also work quite well which means that just showing random photos to people from their collection can also trigger their interest.

Since we did not have access to all the tagged information in users' photo-collections, our browser prototype just gave random photo suggestions but asked people to input the connections that they saw between the photo-transition they choose. This approach has some potential drawbacks.

- For every photo, we wanted the suggestions to cover all the hypotheses connections so that the user has the op-

portunity to navigate along any of them. Currently, we assume that the random suggestion strategy will cover all the connections in the longer run. Thus, currently the photo-transitions that user makes is biased by the potential insufficient distribution of the random suggestions.

- For every photo-transition, the users have to input the connection manually. This may serve as a disconnect while they are browsing their photos.

Ideally we want to have access to the pre-tagged information about the hypotheses connections so that the suggestion photos can cover all the connections besides some random photos. We think that it is still useful to have the user manually enter the connection they felt to discover any new connections that we did not hypothesize or see if the users' connections sometime differ from the pre-tagged information. Conducting this study will require access to pre-tagged information but will provide more useful and accurate data for analysis.

Once we have the first design of the active photo-browser, we plan to distribute it to users by potential integration with current commercial photo-browsers or in form of a new client-end application. This will allow to experience the tool outside of a "study" mode at their own leisure and will benefit us to capture longer term user feedback. Our grand idea of this active photo-browser is that, when deployed, it will learn on-the-fly users' mental connections as they interact with their photos in the longer run and will customize its active behavior of providing useful suggestions accordingly.

In summary, we conclude that the idea of an active browser is interesting to users as it allows them to see more photos than they usually do. It lets them rediscover photos which they might have forgotten about and hence relive past memories.

ACKNOWLEDGMENTS

The authors would like to thank James Fogarty and Hao Lu for their guidance and feedback on this project. We would also like to thank the participants in our user study who took out time from their busy schedules to run our experiment.

REFERENCES

1. Facebook photo stats. `http://www.facebook.com/note.php?note_id=30695603919`.

2. Flickr photo stats. `http://blog.flickr.net/en/2009/10/12/4000000000/`.

3. Google picasa web sharing service. `http://picasaweb.google.com`.

4. User study instructions. `http://www.cs.washington.edu/homes/ankit/course_projects/510/start.html`.

5. User study prototype. `http://www.cs.washington.edu/homes/aditya/projects/510/`.

6. B. B. Bederson. Quantum treemaps and bubblemaps for a zoomable image browser. In *In Proc. User Interface Systems and Technology*, pages 71–80. Press, 2001.

7. S. M. Drucker, C. Wong, A. Roseway, S. Glenner, and S. De Mar. Mediabrowser: reclaiming the shoebox. In *AVI '04: Proceedings of the working conference on Advanced visual interfaces*, pages 433–436, New York, NY, USA, 2004. ACM.

8. D. F. Huynh, S. M. Drucker, P. Baudisch, and C. Wong. Time quilt: Scaling up zoomable photo browsers for large, unstructured photo collections. In *In CHI 2005 extended abstracts*, pages 1937–1940. ACM Press, 2005.

9. H. Kang and B. Shneiderman. Exploring personal media: A spatial interface supporting user-defined semantic regions. *Journal of Visual Languages & Computing*, 17(3):254 – 283, 2006.

10. N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: exploring photo collections in 3d. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, pages 835–846, New York, NY, USA, 2006. ACM.

11. K. Toyama, R. Logan, and A. Roseway. Geographic location tags on digital images. In *MULTIMEDIA '03: Proceedings of the eleventh ACM international conference on Multimedia*, pages 156–166, New York, NY, USA, 2003. ACM.