Assisted Metadata Propagation: Visualizing Contextual Metadata to Reveal Groupings

Carrie Burgener, Benjamin Hill, David Hong
{carrie, bhill, dkhong} @ sims.berkeley.edu
School of Information Management and Systems
University of California Berkeley

Table of Contents

Abstract

Given the high cost of labor involved in annotating large collections of images, the goal of this project is to create an 'ideal' layout of a photo collection for the eventual purpose of quick tagging and event identification. The arrangement is dynamically laid out based on who took the picture, who was around when then picture was taken, and when the picture was taken (date, time of day, and day of week, and where the photo was taken). The layout of the photographs is determined by a calculated value of "closeness" between any two photographs based solely on the metadata associated with each of the photos.

Keywords

Metadata, Personal Photography, Visualization, Context, Cameraphone, Clustering

Introduction

The gap between what computers know about digital photos and what people think about when they mentally classify their photo collections impedes the ability of photo collection management software to accurately sort and retrieve photos without extensive user interaction. Computers work in terms of Cell Tower ID, exact time of photo capture, nearby Bluetooth IDs, and other specific information, while people think in terms of events, collections, people that might be in the photo, and terms that the person might later use to retrieve the photo group. The difficulty of the visual medium lies in the inherent semantic gap between text and images: text has the advantage of semantics and syntax, while images, barring major breakthroughs in both AI and computer-vision, do not share the same expressive ability. One image or photograph could have different meaning and interpretations, especially in the case of any possible automated annotation.

Our visualization's goal is to help bridge the semantic gap when organizing and annotating a photo collection by creating the ideal “cognitive” spatial layout of a photo collection. The visualization attempts to arrange the photo thumbnails in such a way as to match people's underlying cognitive process of identifying groups of related photos (or as close to ideal as we can get). The data behind this layout is a dynamically calculated value of “closeness” between any two photos, based on similarities in the photos' various types of computer-understandable metadata. The relative attractive weighting of the various types of metadata will be easily and interactively adjustable through the user interface, allowing a live exploration of the photo collection's visualization with a minimum of user effort.

Our hope is that a force-directed graphing and clustering of photos will provide an interface intuitive enough as to allow people to adjust the weights to maximize the similarity between the on-screen grouping of photographs and how an individual would classify or tag their photo collection. Using the Mobile Metadata Media 2 (MMM2) framework and data, we explore this type of grouping on personal photo collections to help users understand some implied semantics that may emerge from the metadata collected. First, we will look at other visualizations of photo collection management and

then discuss the outcome of the interface tested by users in follow-up testing. We hope that this technique of clustering could be applied to the larger problem of photo management and may possibly provide an efficient method for photo tagging and annotation.

Related Work

Background

Christel and Consescu discuss the two paradigms of access and retrieval of images and videos: concept-based versus content-based visual indexing. Concept-based indexing provides more expressive means of communicating an image to an end user but requires more manual processing. On the other hand, content-based indexing looks to analyzing lower level properties of an image such as individual pixels and differences in color, texture, and shape. Content-based indexing can be automated but often is too low-level for most end-users.

However, the challenge of effective retrieval and organization still exists, especially with video retrieval. Christel and Consescu talk about some of the problems that arose related to search and retrieval of video with the TRECVID 2003 collection. In particular, the authors wanted to assess the value of visual information in contrast to accompanying speech narrative in closed caption video. Participants of this study were given the task of searching for videos either in a system that had an accompanying information channel (closed caption) or a purely visual one. The results demonstrated that the system with closed caption outperformed the purely visual search. In effect, text search was still predominant and preferred.

Rodden, Basalaj, et al discuss the usefulness of having thumbnails arranged according to mutual similarity. Prior research found that textual documents were easier to browse when giving a representation of some notable structure by clustering similar items. This understanding motivated the development of an interface similar to our own based on creating an undirected graph of similar photos according to some textual description. Images would be arranged spatially closer or further away from one another according to mutually similar items. One key difference in the interface from our own is the puzzle-like fashion in which the thumbnails are arranged according to each photo's orientation.

Amazon Create Your Own Ring

This is an example of an interface that is used to filter for a few objects in a sea of many, and is an interface that is almost completely driven by sliders as a filtering mechanism. This multiple-slider approach allows the current “state” of the system (in their case the ring, in our case the formula that defines closeness between two photos) to be adjustable along many dimensions at the same time, without excessively confusing the user. Both the “range” sliders, where the user is bounding a range of possible values, and the “volume” like graphic, indicating increase towards the right of the slider range, have been

incorporated into our revised screen mockup and align well with the information gathered in our user surveys.

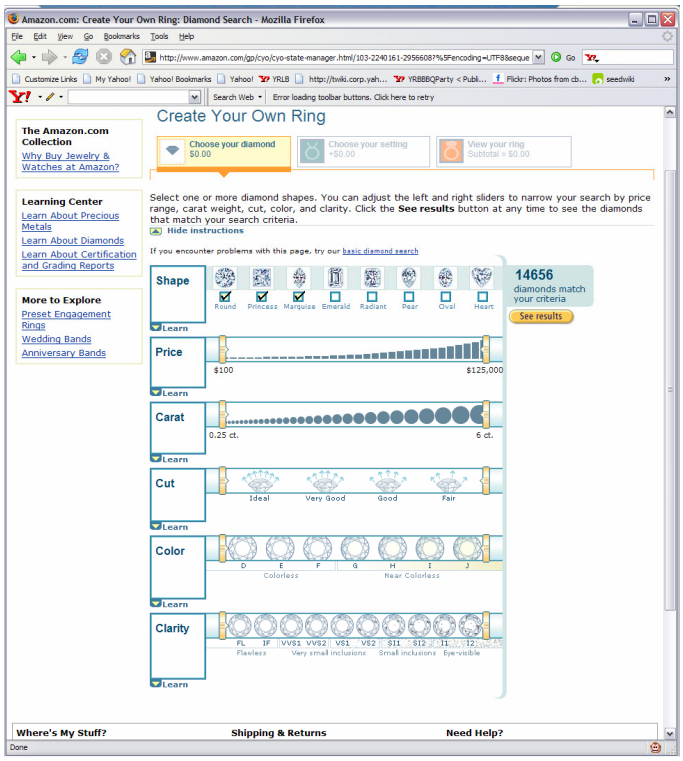

Flickr

The Flickr photo sharing website is used by many people as a convenient way to upload and tag images. The portion of the Flickr interface that was most inspiring was the filtering timeline mechanism in the Organizr section of the website. It is useful to see how long a series of photos is and where in time photos were taken. The group used this idea of a filtering timeline in its later sketches.

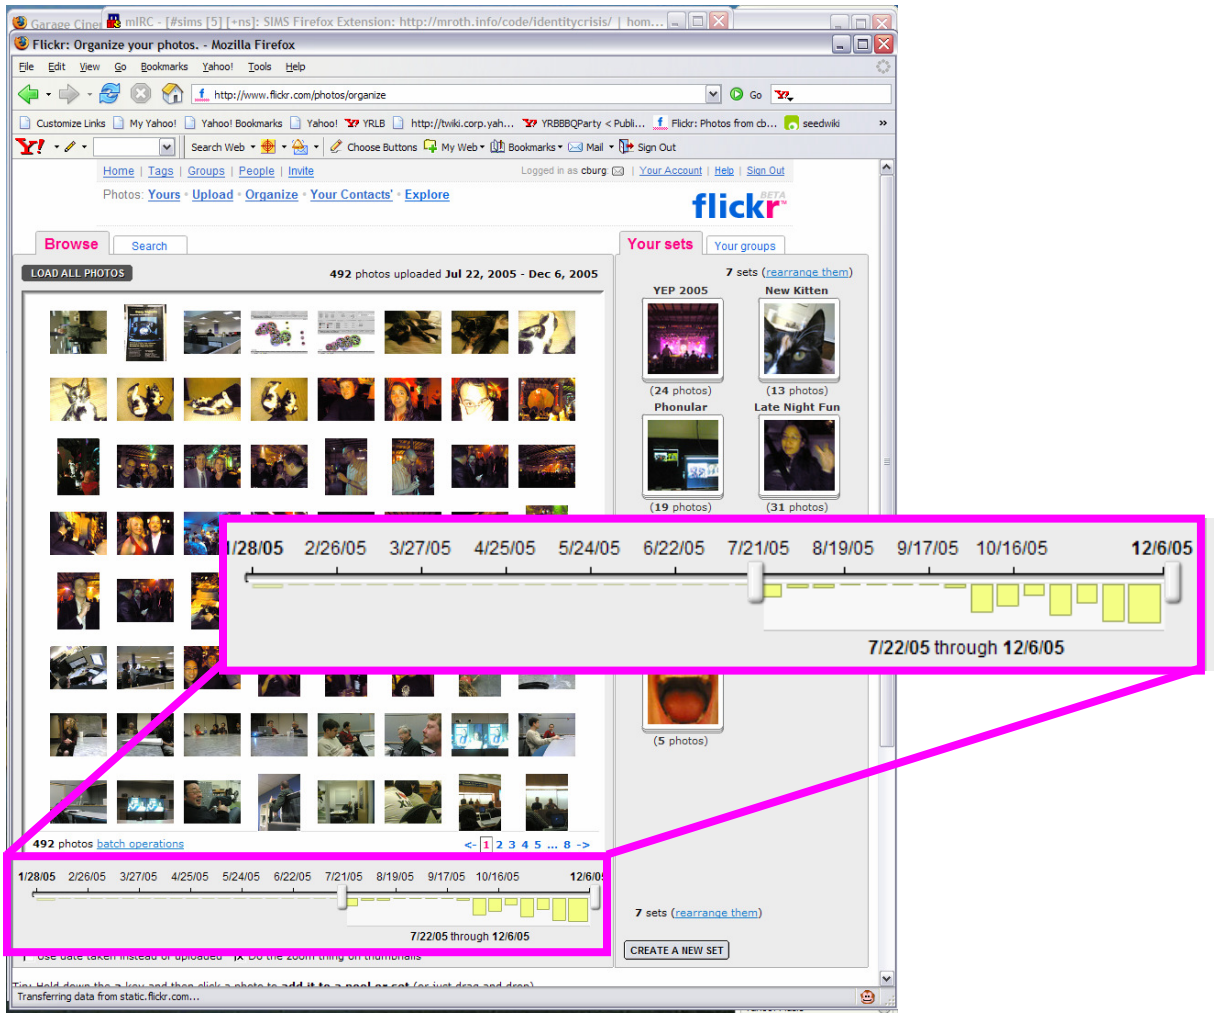

Picasa

Picasa is a popular photo-management tool released by Google. Picasa provides an easy-to-use tool to organize, edit, and share pictures. The structure of the program is largely influenced by the folder layout in the Microsoft Windows file tree (along with EXIF data extracted from the images). One of the key features of Picasa is the timeline that allows you to scroll through images in a chronological fashion. The timeline includes a micro/macro overview of the images by noting the number of images per month and year along the bottom of the screen while allowing the user to scroll through each folder and display a representative image from it. The interface is smooth and fluid and encourages users to flip through and get a sense of possible significant events. However, the major drawback is that the program is largely driven by how a user manually organizes photos and folders. In our user testing, several users had seen Picasa and were able to provide us with comparisons between our prototype and the experience that Picasa provides.

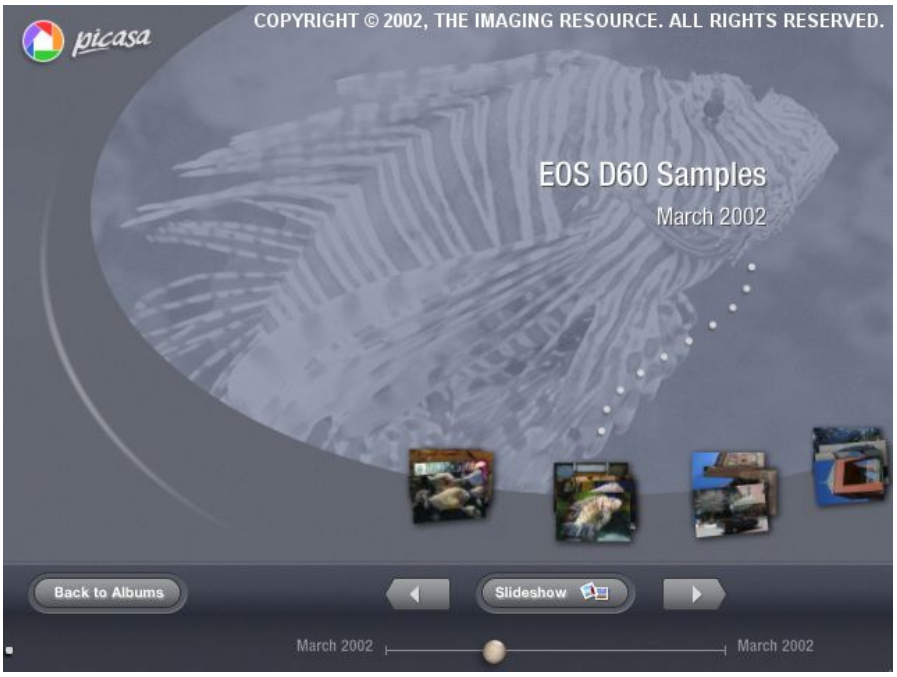

Lessons Learned

Our group found that sorting by time was the most common way of partitioning photos, as almost all software has access to the sequence that the photos were taken in, if not the exact time that the photos were taken. This allowed for a rough match to cognitive groups, but was still lacking in several respects. The highly interactive sites like Amazon Ring have us ideas into how to create an interface that had continuous feedback with respect to every action the user took. This also informed the method that we are currently using for the slider control, and correlated with some of the feedback we got in our user survey: it is best if every action had an immediately visible effect. Our current use of "threshold" values in the sliders, where if a user is outside the threshold there is no visible effect, is less useful than a continuous range where changes always have an effect.

Our Project

Visualization Properties

Our primary goal when designing the AMP visualization was simplicity: we believed that the user should have a minimum of clutter on the screen for both controls and feedback. Because of this goal, we limited the AMP interaction points with users to the top controller portion and the large photo table area. The main point of interaction is through the top sliders, changing the available force-of-attraction variables to control grouping of the photos. The user can also interact directly with the thumbnails through drag-and-drop in the large bottom thumbnail display to re-arrange and untangle interesting clusters of photos.

Layout

The "attractive force" (edge weights) are determined by the contextual metadata associated with every pair of photographs displayed, with each potential contributing factor interactively defined by the user's position of all of the sliders. The dynamic (and sometimes organic looking) clustering of the thumbnails is accomplished through a simple force directed layout, with each thumbnail repelling every other thumbnail a constant amount, and each thumbnail potentially attracted to the other thumbnails with similar metadata.

Our hypothesis was that with the right balance of forces, the layout would cluster into groups that made cognitive sense. An earlier prototype with visual edges (see close-up below), while more cluttered, more clearly represents the desired clustering behavior. See the usability analysis for details as to how useful the clusters actually were.

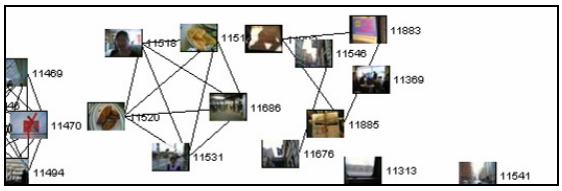

Color

Color in AMP is used simply to prevent overlapping and occlusion of groups that may be in near spatial proximity. Since clusters represent mutually exclusive events, we need to highlight the categorical differences in each grouping. As such, using colors that are different in hue is important. This rule is contingent on the spatial layout of each cluster. Clusters that overlap should use differing hues. However, clusters that may be farther apart could simply change the saturation level.

Size

Currently size is used to highlight the thumbnail that the mouse is hovering over, in order to vie a closer view of the currently selected photo without disrupting the rest of the visualization. In a future revision of AMP, we hope to make the thumbnail sizes dynamic with respect to the number of thumbnails on the screen. This is very important to aid in recall: while our users seemed able to identify photos when they were part of a cluster, we found the threshold to be around 35 pixels before the users started to have problems identifying isolated pictures. Making the thumbnail size dynamic with respect to the number of images displayed on the screen would make better use of the available screen size, while minimizing overlap and maximizing recall.

Properties of Interactive Graphics

While current paradigms of browsing photo collections involve looking at grids of images, this method of organization does not invite users to explore the data. In addition,

if users have large sets of photos, the grid interface along with more serialized views like a timeline do not allow for the investigation for emergent patterns.

Our interface emphasizes the principle of interactivity and exploration of one's photo collection. More importantly, it provides a fun and easy-to-use tool to adjust the visual presentation. In any highly interactive graphic environment, users must be able to manipulate the data by selecting query parameters. The AMP interface accomplishes this through the controller bar section.

AMP also provides a quick way to easily interpret the data. Clustering immediately draws the attention to the underlying similarity of photos that might not have been apparent at first glance. Given the continuous and rapid feedback of the interface, the grouping of photos helps users validate any multitude of associations.

Finally, the AMP interface enables for a wide but meaningful range of viewing images. This exploration of clustering gives immediate feedback and encourages incremental and iterative loopbacks so as to satisfy individual cognitive organization.

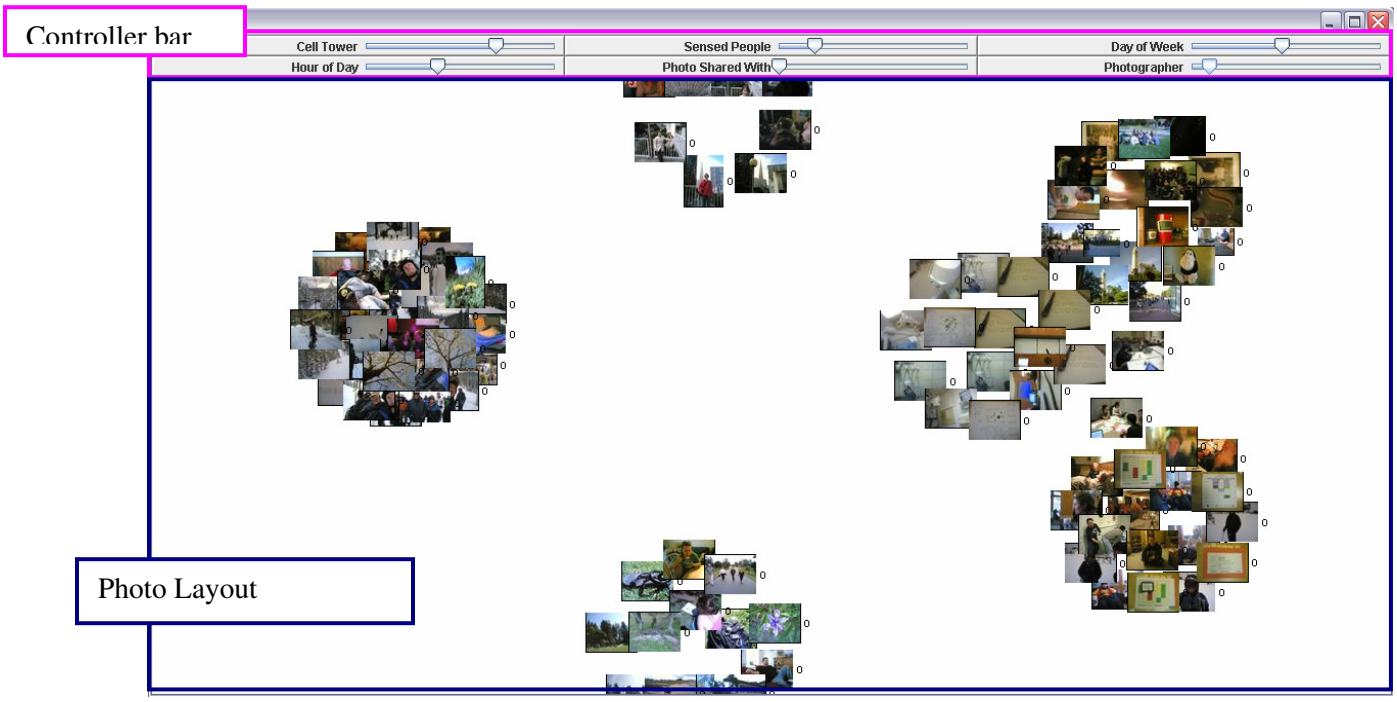

Our Dataset

All group members and usability participants were test subjects for the Mobile Media Metadata 2 project. For this project we used extensible markup language (XML) outputs from the MMM2 database. MMM2 was a piece of software developed for Nokia 7610 cameraphones. The software detected when a participant took a picture, in addition to enabling the user to share the pictures with friends and family, the application also stored contextual metadata associated with each photograph. This contextual metadata included, who took the picture, individuals (with Bluetooth phones) that were nearby when the picture was taken, where the camera was when the photo was taken, what date

it was, and what time it was. The unique data provided by MMM2 was invaluable for this visualization.

Insights Provided by AMP

Simple Interface

AMP provides a very simple input interface that allows users to explore their photo collections, without having to understand the majority of the data before they input. This interface then presents photos laid out visually that enables users to view relationships in their photo collection. In the future we believe that this type of software could be used for tagging, more will be discussed in the future directions section.

Browsing Vs. Search

Like other popular photo albuming software and websites we have approached the problem of finding photos as a browsing task. Our examples of Picasa and flickr highlighted two pieces of software that enable users to move through their photographs in a visual way. We have extended this concept a bit further, moving beyond a grid layout of photographs, to show interrelatedness.

Filtering

Although the sliders are useful in adjusting the visualization, they do not filter the pictures that are being displayed. This filtering can prove to be very important. For example some MMM2 users have more than 1000 photos, it would be extremely difficult for a human to process that many photographs at the same time, the screen would either be cluttered or the photographs would have to be too small to be viewable. Additionally, current computing power begins to hit limits when using a force-directed layout with too many nodes at the same time, extending the need for intelligent filtering.

It is likely that the sliders will have to be augmented in future versions with filtering mechanisms, or they will need to be replaced with another type of input widget.

Grouping Data

Looking at a series of clusters of photos on a computer screen may give enough clues to how the clusters were created, it also may not. Doing other labeling including displaying metadata with the photos, and clustering data is being considered.

Usability Analysis

For our Usability Analysis we recruited two groups of people, one group of individuals were new to the AMP interface (Beginners), and one group was more familiar with the AMP interface concept as well as the MMM2 data (Experts). We also found that testing the visualization with the owners of the photos is essential for accurate feedback. People are able to recognize photos that they are familiar with at a much smaller thumbnail size, and are much faster to draw associations between groups of photos. Unfortunately, this limited our test population to those that had taken 100 or more photos using MMM2, a necessary restriction to ensure a sufficiently rich dataset for the visualization.

Tasks

The tasks created for our usability analysis were fairly generic mostly because the team members had no way of determining the types of photographs each of the participants would have. At the beginning of the test users were asked to identify what each of the labels on the sliders of the interface meant, participants were also asked to identify what the sliders in general did. After the simple identification task users were asked to manipulate the sliders to find meaningful groups of photos, when they found a meaningful cluster they were asked to identify what the cluster was. At the end of the test participants were asked to react to the interface, reporting what they liked, what they did not like and any changes they would suggest.

Beginners

The beginner participants were recruited from the MMM2 user pool. All of the participants were second-year graduate students in the School of Information Management and Systems at the University of California Berkeley.

Beginner 1

- Berkeley Graduate Student
- This user used to use Picasa to organize their collection but stopped because of the effort.

Feedback

- Cell tower was initially thought to mean "how easily one can upload photos"
- Hour of Day was though to be a 0 to 24 hour slider. "Sensed" had no meaning
- After playing with the sliders, this user was able to identify photos taken at home vs. taken at South Hall and photos taken with a 202 design group
- The photos scrolling off the screen was distracting
- Wanted a "clear all" button, and the ability to have bigger icons

Beginner 2

- Berkeley Graduate Student
- This user didn't have a favorite photo manager program

Feedback

- All sliders were through to be filters – hours for the hour of the day, day of week represented Sunday – Monday
- Sliding sliders part-way confused the user when there was no visible effect on the clustering
- It was unclear what time range of photos were being displayed, and why some of the sliders didn't have the expected effect.
- User noted that it would have been useful to be able to view the metadata associated with photos to find out why some sliders had no effect when the user thought they should

Experts

The expert participants were recruited from the MMM2 designers. All of the expert participants have worked on similar photo sharing and tagging problems in the past.

Expert 1

- Berkeley Undergraduate Student
- Part of the design team of the MMM2 system

Feedback

- Generally did not understand what the sliders were associated with on first look
- Did not know if the sliders related to each other
- Did understand after playing with the interface that all the way to the left was turned off and all the way to the right is turned on all the way
- Pointed out that there were two clusters of pictures from the same place and they were in different parts of the screen because they must have a different Cellular ID, wished they would have been closer together
- Felt weird mixing the sliders
- Could identify a picture cluster was on a Sunday because he could recognize his parents in the picture (and he does not regularly hang out with them during the week
- Reported that the constant animation was distracting
- It was weird that all of the photos sort of animated off into space and stuck to the edges of the screen
- Would like to be able to zoom in on a particular section of the screen
- It does make more sense as you spend more time with it
- This participant was particularly interested in being able to dive deeper seeing details about photographs, larger representations of the photos and then being able to move quickly back to the visualization

Expert 2

- Berkeley Faculty Member
- Involved in the design, development and direction of MMM2 system

Feedback

- Would have liked to be able to 'lasso' a selection of pictures to allow for zooming, and quick bulk processing
- Thought the interface could be improved by placing the sliders vertically then the team could allow the sliders to be reorganized the top slider being the most important hierarchically
- It would be nice to have a feature that clustered by closeness in time, so unlike day of the week, and time of day it would allow to pick a grouping of pictures that happened at a similar time
    - This might allow for more 'event' like clustering
- This participant thought that the current interface looks too 'techy'

Expert 3

- Berkeley Graduate Student
- Involved in interviewing and analysis of the usage of MMM2 system

Feedback

- It was not obvious what the sliders are supposed to be doing right away
    - Guessed it probably increased the importance of a category
    - At one moment stated "oh I see how this works"
- "It is a cool viz as far as playing with the data, I can't see myself going back to it multiple times"
- "It is not so meaningful to me in the way I think about my photographs"
- I always end up with the big blob in the middle of the screen
- Combining the sliders is confusing
    - Might be useful to make clusters with one of the sliders, then zoom in on that cluster and re-cluster with one of the other sliders
- Might be easier to have an on and off switch for each of the variables that effect the cluster…since it seems like on/off is enough
- The constant animation gets really distracting after awhile

Expert 4

- Berkeley Graduate Student
- Involved in developing other areas of the MMM2 system

Feedback

- Any additional information next to the thumbnail sis distracting
- The sliders were seen to be continuous, but had confusion as to how that mapped to discrete values like day of week and cell tower ID.
- People Sensed was though to be a slider showing number of people
- Overall, initial cognitive map of the sliders was all based on filtering
- Moved cell slider half-way, noted that it had no effect
- Was able to identify clusters after putting cell tower to 100%
- South Hall, Home, Berkeley South Side
- Tried to combine clusters through brush and link (possible future feature?)
- Noted that photos taken in short bursts are already related
- Too hard historically to organize photos across downloads form a camera – but would very much like to!
- Noted that the tool was fun in itself – wanted to just let people play with it!

Expert 5

- Berkeley Graduate Student
- Involved in developing other areas of the MMM2 system

Feedback

- Using a trackpad with the interface was very difficult – would have preferred a mouse for the test
- "Photos shared" was confusing – did it mean photos shared with me?
- Was able to locate clusters of meetings initially
- Wanted to be able to increase icon size
- Comments of the visualization included "Small and globby" and "one blob is attacking another blob!"
- This user doesn't organize photos in flickr, mostly uses it purely for a broadcast platform

Consolidated Takeaways from Usability Analysis

Sliders

- Sliders are difficult for users to understand – because people are used to thinking of "Day of Week" as discrete values, they almost always associated the left side of the slider with Sunday.
- The majority of participants' immediate reaction is to think of the sliders as filters, something which they are more used to seeing.
- Combining sliders only rarely creates more visually interesting or useful clusters
- It is too easy to get a big glob of photographs in the middle of the screen and not know what actually is going on – beyond a certain threshold, the clustering looses value.
- Most of the participants turned the slider all the way on or all the way off and did not play with the various settings
- Terminology on the sliders needs to be improved for understandability

Work Area

- Provide users a way to turn of constant animation, create an algorithm that evens out over time
- Provide a way to keep the photographs and clusters from moving all the way to the edge of the screen
- It would be useful to allow a user to see the data associated with a picture or a cluster, but also enable them to turn the display off when it isn't needed

Implementation Tools

In developing AMP we used the Prefuse Graph Visualization Toolkit, running in the Java 5.0 development environment. All coding was accomplished in Eclipse.

The group utilized the J2EE(JSP) API tool for exporting XML metadata from the MMM2 database, provided by the MMM2 "metadata tagging" API.

Future Work

Based on previous work, background research, and our usability testing results, we identified several key features that we believe would help to move AMP from an interesting prototype to a viable product. The biggest improvement would be improving accessibility – retrieving the data necessary for driving the visualization is a several step process. Optimally, this could be consolidated into a single step, or switch to using a platform that is more commonly used.

The following mockups take into account our research and user testing. They include the following key features, clearly separating "filtering" from "clustering."

- Time range slider
    - Although not identified in the usability analysis providing a time range slider will be crucial to AMP in the future. This allows a user to select a

range of contiguous time, it also works like a filter narrowing the number of photographs displayed on the screen as a single time.

- History "Bookmarks"
    - Allowing a user to save some preferred settings, creating a history of their favorite clusters was identified as important
- Filtering mechanisms
    - It is useful to provide filtering mechanisms based on critical feedback from class, beginners and experts. All filters except for the selection of the time range are provided in a drop-down advanced display. The advanced display is the second figure shown below.
    - Cell Tower
        - The advanced options provided with the cell tower slider allow for a user to select specific towers. The team is aware that selecting a cell tower number is exceptionally not user friendly and that a city, state or even zip code would be more understandable than a cell tower id number. Currently that information is unavailable to us.
    - Hour of Day:
        - The advanced options provided with hour of day allows the user to select specific times they would like photos to be displayed from.
    - Day of Week:
        - Like hour of day the day of week filter allows for specific days of the week to be selected and deselected.
    - Sensed People, Photographer, Photo Shared With
        - Use a single set of filters. A user can choose the people they would like included or not included then use any of the sliders to change the attraction of the remaining photos.
- Coloring of groups
    - It could be useful to utilize basic light coloring behind pictures that are being grouped/clustered with using the same variable.
    - There also might be a use for creating colored groups based on the selection of data or strength of attraction used by the sliders. For example, if Cell Tower is red and hour of the Day is blue photographs grouped by two variables would have purple behind them. (NOT DISPLAYED)
- Metadata display
    - Textual display of data on the work area could be useful, as long as there is an option to turn it off and on, since the amount of data will most likely become overwhelming (NOT DISPLAYED)
- Slider feedback
    - Giving some feedback to what the sliders are doing, and what the setting is will be of great importance in the future as AMP is developed further. Currently a shadowed indicator has been added. This may not work well since during some usability tests participants thought that the sliders were signifying quantity and not pull (and the ramp will not resolve that misunderstanding)
- Multiple Views/Windowing

- It might be useful to have an overview window to display the entirety of the dataset and then allow the user to zoom in on a section, this would allow for context to always be retained and larger views of photos to be accommodated without having to go to a separate screen.

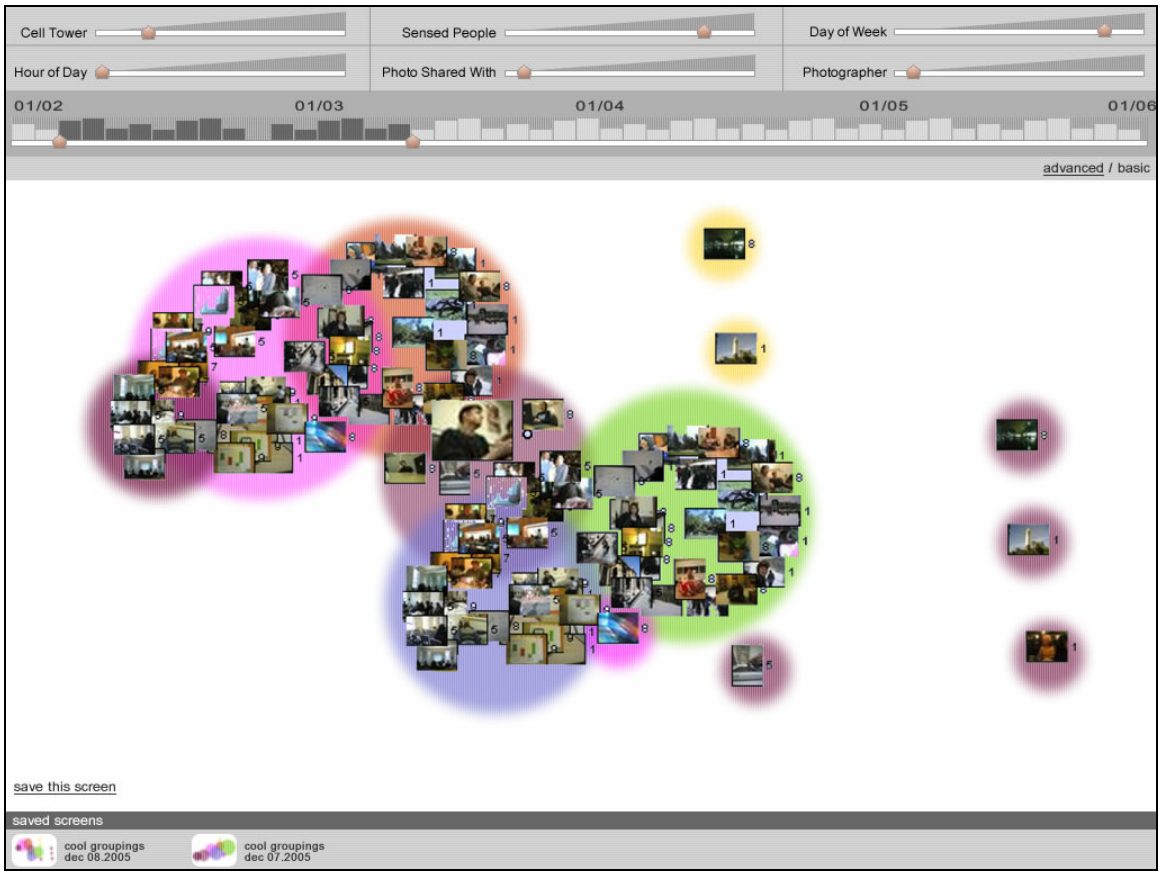

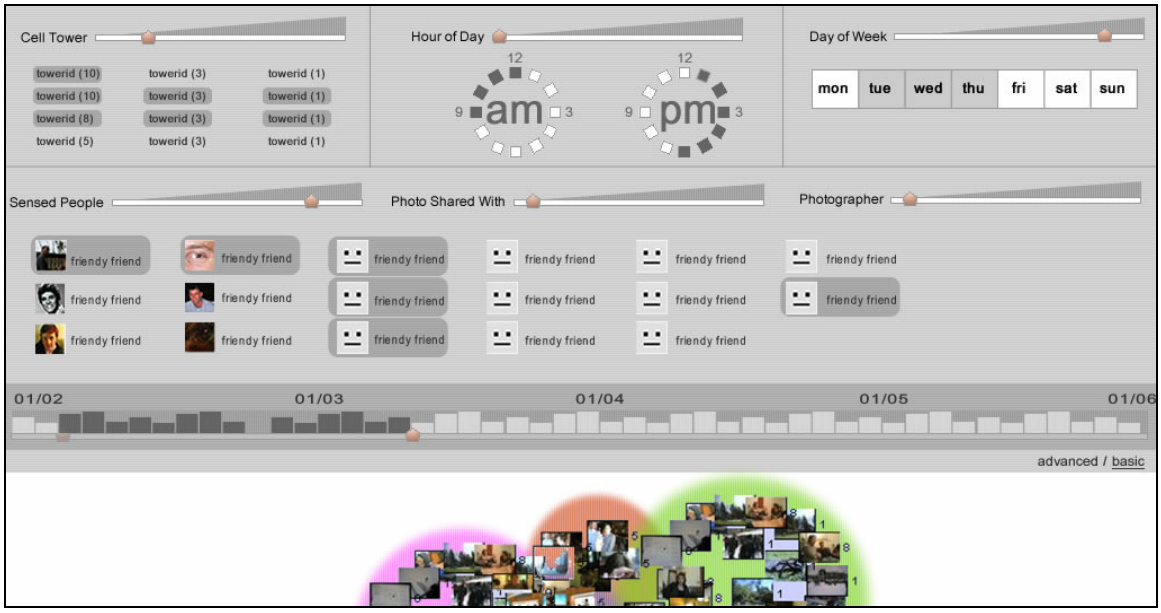

In addition to the features visible in the mockups, the following features are being considered:

- The thumbnails would dynamically resize to fit the available screen space as the filters changed the number of visible thumbnails.
- Wall-forces could keep thumbnails from edging off the screen
- A "tagging" feature assigns one or more key words to a colored thumbnail cluster, identifies a group of photos for later recall.
- Ways to represent the sliding categories hierarchically allowing users to specify which one of the variables has precedence
- Time sliders that allow selection of contiguous times, this would allow a user to cluster all photos that were taken close together, the current interface does not support such a cluster.
- A simpler representation of the slider interface, using a binary on and off switch instead of changing springiness via the sliders the software would define the springs behind the scenes.
- Integration with popular photo repository sites that support metadata, including flickr.

Summary

Photo management challenges many of the traditional notions of browsing and searching. Looking at the available tools like Google Images or Picasa, organization and retrieval of photos is still a fairly labor intensive and manual process. However, by including principles of filtering and grouping by similarity, we can get closer to potentially creating the optimal cognitive spatial layout of a photo collection.

In summary, AMP has been designed to dynamically create an 'ideal' layout of photographs for photo viewing and tagging. By utilizing the contextual metadata associated with every pair of photographs displayed, we are able to generate a calculated value of closeness between any two photos. Users are given a chance to explore the data and effectively touch the variables by manipulating slider controls to generate varying views of the photos. Clustering of the photo thumbnails is accomplished through a simple force directed layout, with each thumbnail repelling one another while potentially attracted to the other thumbnails with similar metadata. The interactivity of the interface provides quick and rapid feedback along with continuous scaling of photo attractiveness.

We have conducted a preliminary usability analysis to test if our interface is understandable, and is actually helpful in the photo organization process. We have found that the clustering of photos is understandable and users can pick out days, events and people, but the sliders are not well defined and are thus confusing for even expert users. While feedback is currently provided in terms of clustering of photos, ideally, providing some transparency behind what each slider control influences will be necessary. Also, the ability to filter the various slider inputs would be useful as well.

Acknowledgements

We would like to thank Marti Hearst, Jeff Heir and our Fall 2005 is247 class.

We would also like to thank all of the individuals that volunteered to participate in our usability analysis.

Thank you to Marc Davis and Garage Cinema Research for the use of MMM2 phones.

If you would like to try the AMP interface and were part of the MMM2 project, please contact one of the AMP team members for a compiled version of the code and instructions on how to import your personal photo metadata file.

Bibliography

Amazon.com, Amazon Ring Selector. http://www.amazon.com/gp/cyo/cyo-state-manager.html/ . Accessed December 6, 2005.

M. Baldonado, A. Woodruff, and A. Kuchinsky. Guidelines for Using Multiple Views in Information Visualization. Proc. Advanced Vis. Interfaces 2000, Palermo, Italy, May 2000, pp. 110-119. http://www.sims.berkeley.edu/courses/is247/f05/readings/Baldonado_MultipleViews_AVI00.pdf

Cynthia A. Brewer. Color Use Guidelines for Data Representation, Proceedings of the Section on Statistical Graphics, American Statistical Association, Alexandria VA. pp. 55-60. 1999. http://www.personal.psu.edu/faculty/c/a/cab38/ColorSch/ASApaper.html

Stuart Card. “Degree of Interest” http://www2.parc.com/istl/projects/uir/pubs/items/UIR-2002-11-Card-AVI-DOITree.pdf

Michael G. Chirstel, Ronald M. Conescu. “Addressing the challenge of visual information access from digital image and video libraries” From: Proceeedings of the 5th ACM/IEEE-CS joint conference on Digital Libraries in Denver, Colorado, ACM Press, 1-58113-876-8.

Marc Davis, Nancy Van House, Jeff Towle, Simon King, Shane Ahern, Carrie Burgener, Dan Perkel, Megan Finn, Vijay Viswanathan, and Matthew Rothenberg. "MMM2: Mobile Media Metadata for Media Sharing." In: Extended Abstracts of the Conference on Human Factors in Computing Systems (CHI 2005) in Portland, Oregon, ACM Press, 1335-1338, 2005.

Stephen G. Eick, Graham J. Wills. High interaction graphics. European Journal of Operational Research, 81:445-459, 1995. http://www.sims.berkeley.edu/courses/is247/f05/readings/Eick_HighInteractionGraphics_EJOR95.pdf

Google. Picasa Photo Browser. http://picasa.google.com/features/index.html

Alfred Inselberg. Multidimensional Detective. IEEE InfoVis 1997. http://www.sims.berkeley.edu/courses/is247/f05/readings/Inselberg_MultiDimDetective_InfoVis97.pdf

K Rodden, W Basalaj, D Sinclair, KR Wood. "Does organization by similarity assist image browsing?"

Ben Shneiderman. Dynamic queries, starfield displays, and the path to Spotfire. February 4, 1999. http://www.sims.berkeley.edu/courses/is247/f05/readings/Shneiderman_DynamicQueriesSpotfire_99.pdf

Ben Shneiderman. PhotoMesa http://www.cs.umd.edu/hcil/photomesa/

Yahoo!, Flickr Photo Browser. http://www.flickr.com

Ping Yee, Kirsten Swearingen, Kevin Li, and Marti Hearst. Faceted Metadata for Image Search and Browsing. ACM CHI 2003. http://www.sims.berkeley.edu/courses/is247/f05/readings/Yee_Flamenco_CHI03.pdf

Ben Shneiderman. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. Proc. 1996 IEEE Visual Languages. http://www.sims.berkeley.edu/academics/courses/is247/f05/readings/Shneiderman_EyesHaveIt_VL96.pdf
"Overview: Gain an overview of the entire collection.
Zoom : Zoom in on items of interest
Filter: filter out uninteresting items.
Details-on-demand: Select an item or group and get details when needed.
Relate: View relationships among items.
History: Keep a history of actions to support undo, replay, and progressive refinement.
Extract: Allow extraction of sub-collections and of the query parameters."
Multi-dimensional: most relational and statistical databases are conveniently manipulated as multidimensional data in which items with n attributes become points in a n-dimensional space. The interface representation can be 2-dimensional scattergrams with each additional dimension controlled by a slider"
Network visualization is an old but still imperfect art because of the complexity of relationships and user tasks."

Chris Stolte, Diane Tang, Pat Hanrahan. Polaris: A System for Query, Analysis and Visualization of Multi-dimensional Relational Databases. IEEE Transactions on Visualization and Computer Graphics, Vol. 8, No. 1, January 2002. http://www.sims.berkeley.edu/academics/courses/is247/f05/readings/Stolte_Polaris_TVC

G02.pdf
* *Sorting, Filtering, Brushing, Tooltips*

Jeffrey Heer, Stuart K. Card, and James A. Landay. prefuse: A Toolkit for Interactive Information Visualization. ACM CHI 2005.
http://www.sims.berkeley.edu/academics/courses/is247/f05/readings/Heer_Prefuse_CHI05.pdf

S Harada, M Naaman, YJ Song, QY Wang, A Paepcke, Lost in memories: interacting with photo collections on PDAs, JCDL'04.
http://www.sims.berkeley.edu/academics/courses/is247/f05/readings/harada04.pdf
"Our system creates a hierarchical structure of the user's photos by applying time-based clustering to identify subsets of photos that are likely to be related."

Ivan Herman, Guy Melançon, M. Scott Marshall. Graph Visualization and Navigation in Information Visualization: A Survey. IEEE Transactions on Visualization and Computer Graphics, 2000.
http://www.sims.berkeley.edu/academics/courses/is247/f05/readings/Herman_GraphVizNavigation_TVCG00.pdf
* *Would be a great guide to future directions to work on the layout of the icons other than using a force directed layout*

Jeffrey Heer, danah boyd. Vizster: Visualizing Online Social Networks. InfoVis 2005.
http://jheer.org/publications/2005-Vizster-InfoVis.pdf
* *Future work, to use the clustering method that they used to show groups of photos graphically.*