

# PhotoMagnets: Supporting Flexible Browsing and Searching in Photo Collections

Ya-Xi Chen  
Media Informatics  
University of Munich, Germany  
yaxi.chen@ifi.lmu.de

Michael Reiter  
Media Informatics  
University of Munich, Germany  
reiterm@cip.ifi.lmu.de

Andreas Butz  
Media Informatics  
University of Munich, Germany  
andreas.butz@ifi.lmu.de

## ABSTRACT

People's activities around their photo collections are often highly dynamic and unstructured, such as casual browsing and searching or loosely structured storytelling. User interfaces to support such an exploratory behavior are a challenging research question. We explore ways to enhance the flexibility in dealing with photo collections and designed a system named PhotoMagnets. It uses a magnet metaphor in addition to more traditional interface elements in order to support a flexible combination of structured and unstructured photo browsing and searching. In an evaluation we received positive feedback especially on the flexibility provided by this approach.

## Categories and Subject Descriptors

H.5.2 [Information Interfaces and Presentation]: User Interfaces.

## General Terms

Design, Experimentation, Human Factors.

## Keywords

Browsing, search, photo collection, flexibility.

## 1. INTRODUCTION

The digitalization of media has influenced our life, and in particular it has changed our ways of dealing with media collections. Existing software for managing media collections mainly offers linear structures corresponding to the underlying technical concepts (e.g., folders and events), or metaphors borrowed from the analogue counterparts (albums, stacks, etc.). These interfaces support certain activities very well, such as reorganizing collections along different criteria, or quickly locating a specific item within a large set. However, other important activities with media collections are more dynamic and often unstructured, such as casual browsing and searching or storytelling [3, 13]. This explorative behavior and flexible association within media collections is poorly supported by today's interfaces. We believe that applications supporting these activities in a flexible way will be appealing and lead to more powerful UIs. We explore how to support the flexibility with photo collections and designed a prototype named PhotoMagnets, which uses a magnet metaphor in addition to more traditional and

well-known interface elements such as a tree view. In an evaluation we received positive feedback and the results confirmed the effectiveness of PhotoMagnets in supporting flexible and unstructured activities.

## 2. EMERGING UI CONCEPTS FOR MEDIA COLLECTIONS

In the field of Information Retrieval, researchers have already realized the limitation of query-based search, especially in situations in which users have no efficient information to formulate queries [16, 22]. Beyond query-based search, Marchionini proposed the concept of Exploratory Search [16], in which querying and browsing strategies are combined to foster learning and investigation. Some researchers also addressed the exploratory aspect of browsing. Bawden identified three categories of browsing activities [4]: purposive browsing, capricious browsing and exploratory browsing. Cove and Walsh proposed a similar taxonomy [8]: search browsing, general purpose browsing and serendipity browsing.

Exploratory behavior is very common with media collections. Finding a target in a directed search is not a typical activity within photowork [15]. People rarely look for one specific item and instead, a significantly larger proportion of time is spent in browsing-like filtering rather than in explicit search activities [10, 15]: People are satisficing rather than optimizing their selection. They only attempt to find an acceptable alternative rather than the optimal solution. During browsing and searching, they are easily sidetracked by serendipitous discoveries. People often change their mind and they might look for a particular item but in the end select something different. People also tend to interact more often with recent items than with older ones [15, 20].

For photo collections, capture time and folder structures are the main ordering principles [9, 20]. As people tend to keep digital photos organized around the same event, some work further explored automatic event detection, which groups photos into different events and selects one representative photo for each. TimeQuilt [14] and Phototoc [18] are two representative applications. Besides ordering photos chronologically, it has recently become a fashion to automatically display photos on a map based on their geo-coordinates [17, 23]. However, the semantic gap between the physical coordinates and the high-level perception cannot be narrowed very well **Error! Reference source not found.** Most existing systems rely on pre-defined categories (such as where, when and who), and flexibility is not efficiently supported. PhotoArc [1] and MyLifeBits [11] are two of the few systems allowing users to freely create a sequence of photos by drag-and-drop. Shen et al. also investigated collaborative interactive presentations on large displays[21].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICMI-MLMI'10, November 8-12, 2010, Beijing, China.

Copyright 2010 ACM 978-1-4503-0414-6/10/11...\$10.00.

### 3. SUPPORTING FLEXIBILITY WITH PHOTO COLLECTIONS

As motivated above, people’s activities around photo collections require flexible interfaces. Applications with fixed and pre-defined structures cannot efficiently support the flexibility observed in user behavior. In general we believe that an interface supporting flexibility should follow some common characteristics:

*Encourage exploration:* It is essential to support exploratory activities beyond explicit query-based search. On the other hand, People often switch back and forth between different activities, such as browsing and directed search. A seamless and tight integration of these functionalities will actively encourage this.

*Facilitate serendipity:* People easily get sidetracked by serendipitous findings and inspirations. The interface should encourage users to discover such serendipitous findings, and make it easy to revisit older parts of a collection to rediscover forgotten photos. An anchor point should always be kept on track, to allow the user to recover easily from distractions.


*Keep the user in control:* As House et al. [13] found, people prefer interaction and feedback over full automation. Active interactions can help users to make more sense of the underlying system logic and enhance their feeling of control and enjoyment.

### 4. PROTOTYPE DEVELOPMENT

Based on these general design considerations, we developed a prototype named PhotoMagnets to support flexible combinations of browsing and search in photo collections.

#### 4.1 INITIAL CONCEPT DEVELOPMENT

Initially we proposed two interface concepts for browsing and one for search, in which the user can actively manipulate the layout of the collection. Since event is a fundamental concept for organization, we employ it as the basic principle for browsing and searching. BubbleBrowser is a circular layout (see Figure 1a), in which events are represented as bubbles. Other events sharing the same contexts, such as location or person, can be opened in other bubbles, which are connected with lines to illustrate their logical relations. Similarly, events in TreeBrowser are displayed by their browsing order, and the browsing history is shown as a tree in an extra panel (see Figure 1b).



**Figure 1. The sketches of the initial concepts: a) BubbleBrowser. b) TreeBrowser. c) Magnet metaphor.**

Inspired by Dust&Magnet [24] and Dynamic Query [2], we proposed a magnet metaphor to facilitate flexible search (see Figure 1c): Multiple search criteria act as magnets, attracting matching photos while repelling others. We expected the magnet metaphor to facilitate more flexible query formulation and to offer an intuitive visualization for the retrieval results. To facilitate serendipitous browsing, the user can search for visually similar photos by using example photos as magnets. Semantically “related images” are retrieved based on tags.

We sketched these concepts on paper and tested them with 8

subjects who were regular photo consumers. We first introduced our concepts with these sketches, and then illustrated them with an example collection of 90 print photos. The results showed that both bubble and tree layouts were appreciated. The magnet metaphor and “related images” were commented as useful, albeit with overall slightly lower grades. Although the connecting lines illustrate the browsing history, the connected bubbles seemed confusing especially when the connected events were topically irrelevant. TreeBrowser was commented as similar to the common hierarchical browsing, and 5 out of 8 participants preferred a chronological layout over the browsing history.

Based on the participants’ feedback, we refined our concepts: The connecting lines between bubbles were excluded from the BubbleBrowser. A chronological order was preferred in TreeBrowser, therefore, all events are now ordered by time and the detailed hierarchical structure of the collection is shown on a separate timeline. The visualizations of the retrieval results can be saved in the magnet view. We conducted a second round of tests with 8 more participants and all refinements were confirmed. The chronological structure in TreeBrowser was appreciated, and BubbleBrowser and the magnet metaphor were considered useful to support unstructured activities.

#### 4.2 THE PHOTOMAGNETS UI

Our preliminary tests showed that both chronological and loosely structured representations have their respective advantages in supporting structured and unstructured activities. Therefore, we decided to employ TreeBrowser for structured browsing, while BubbleBrowser and the magnet metaphor were integrated into one view to support browsing and searching in a more flexible way. The PhotoMagnets prototype was implemented based on the Piccolo toolkit [5] for Zoomable User Interfaces.

##### 4.2.1 TreeBrowser

TreeBrowser consists of three components (see Figure 2): a scrollable timeline (top), a zoomable presentation canvas (center) and a tag panel (left). The user can organize, browse and tag photos in a hierarchical event structure. While importing a collection, events are automatically detected and the collection is segmented. By default, the first photo of each event is chosen as the event representative photo (RP), and the first event RP is selected as the collection RP. All imported collections are displayed chronologically on the timeline: Each collection is shown with its RP, accompanied by general information, such as the automatically calculated time span, collection name and location defined by the user. Events are visualized as peaks in a temporal histogram (see Figure 2). The width indicates the event’s time span and the height illustrates the number of photos taken in each event. To allow exploration of each event, with mouse hovering on one peak, the collection RP is temporally replaced by the corresponding event RP, and a tooltip displays its name and the number of photos in it.

The user can load a collection from the timeline onto the presentation canvas by dragging its RP from the timeline. As Figure 2 shows, the loaded collection is led by its RP, and photos are shown chronologically in respective event blocks. Through a shortcut menu triggered by the right mouse button, the user can set a photo as the collection or event RP which is labeled with “R”. Above the collection RP and the first photo in each event, there are two text fields, in which the user can input the name and location. Clicking on one photo creates a slideshow window, in which loaded photos are shown chronologically in full size.

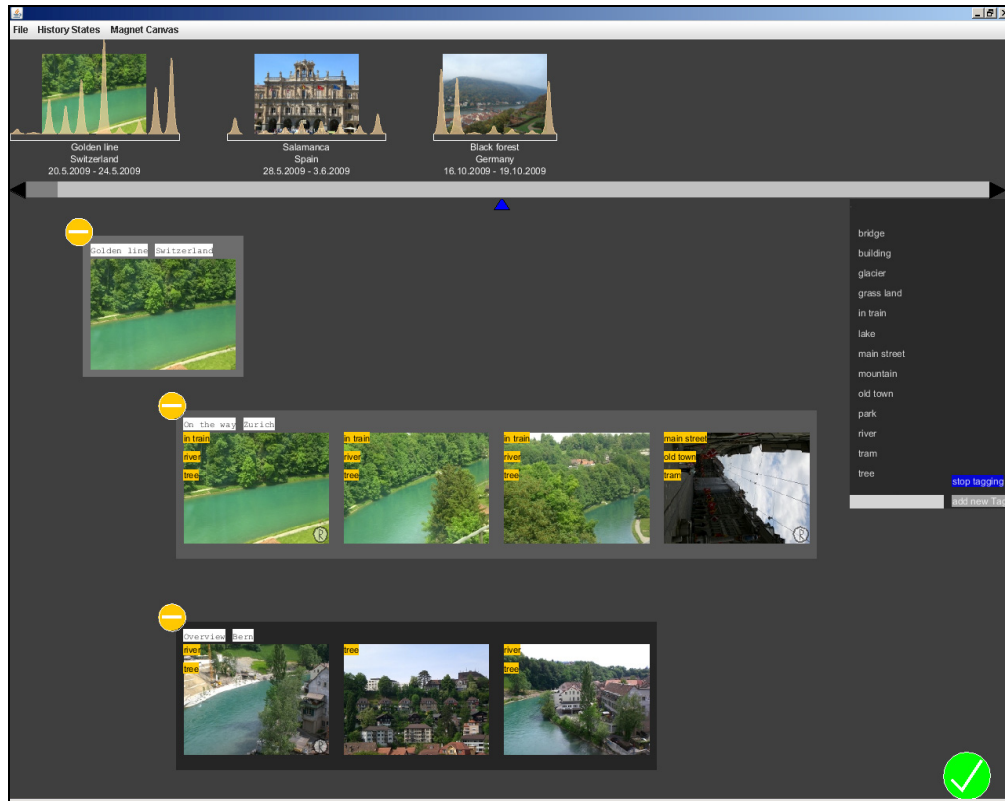


Figure 2. The interface of TreeBrowser.

Event refinement is allowed by drag-and-drop, and the chronological order is automatically maintained. Tagging is another activity supported in TreeBrowser. The user can generate new tags, select multiple tags and assign them to multiple photos by clicking on each of them once. Tags can be removed by clicking each of them once. All changes can be saved by the green button in the bottom right corner.

#### 4.2.2 BubbleBrowser

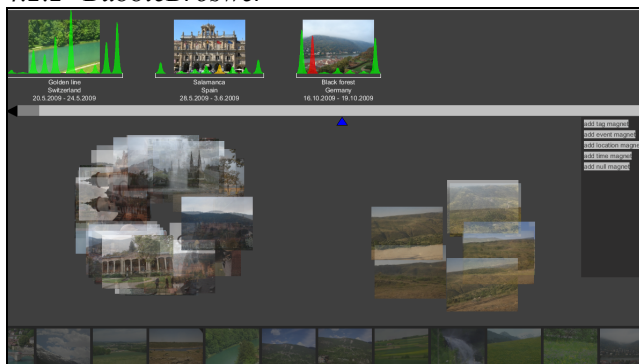


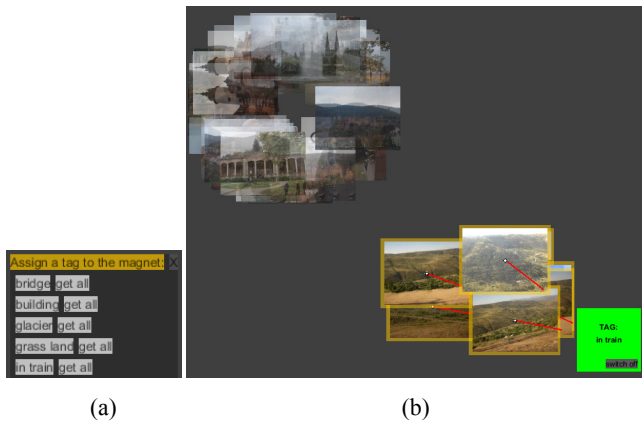
Figure 3. The interface of BubbleBrowser.

The user can switch between TreeBrowser and BubbleBrowser with the Tab key. BubbleBrowser is designed to support unstructured activities, such as flexible browsing and searching. It consists of four panels (see Figure 3): a timeline (top), a presentation canvas (center), a magnet list (left) and a panel for related images (bottom). The loaded photos are displayed in event

bubbles, ordered clockwise by time. Other photos sharing common tags are displayed as “related images” in a flow animation. When dragging one related photo onto the canvas, the user can choose between loading this photo only, or loading photos with same tags from the same event or collection. Similar to TreeBrowser, double clicking on one photo leads to the extra slideshow window. A photo can be deleted with the Del key. Through the shortcut menu with mouse right clicking, the user can delete the corresponding event, or view it in TreeBrowser.

Five types of magnets are offered in the magnet panel: tag-, event-, location-, time- and null-magnet. The first four types of magnets can be used to search photos matching certain attributes, for example, containing same persons, from a specific event or location, or captured in a certain time period. The null-magnet groups photos that do not match any other magnets. After selecting a tag-, event- or location-magnet, relevant attributes are shown in a pop-up window. Figure 4a is an example for a tag-magnet. A time-magnet provides four sliders for the definition of the year, month, day of a seed date and a time span. There is no pop-up window for null-magnet.

Three usages of magnets are allowed: Magnets can be used to attract matches from loaded photos, from all collections, or from certain events. In the first case, the user can first load certain events. By choosing one of the attributes, such as a tag “in train” in Figure 4a, the cursor changes to a magnet, a green square showing the magnet type “tag” and corresponding attribute “in train”. After dropping the magnet onto the canvas, the matching photos are automatically attracted to it. The red line(s) illustrate by which magnet(s) each photo is attracted.



**Figure 4. Using a magnet to attract matches from the loaded photos: a) Setting the magnet attribute. b) The magnet automatically attracts matches from loaded photos.**


Events in the timeline are rendered differently according to the number of matches loaded from them: Green for none, yellow for some and red means that all photos from this event are loaded. The magnitude of a magnet is visualized as its size and can be adjusted by mouse wheeling while pressing the ALT key (see the two magnets with different sizes in Figure 7). The position of a magnet can be changed freely by drag-and-drop and the corresponding matches will be replaced accordingly. A magnet can be activated/deactivated by clicking the “switch on/off” button on its bottom right corner. A magnet can be deleted by selecting “delete” in its shortcut menu.

The second usage of magnet is to attract all matches from the timeline by choosing “get all” in Figure 4a. Matches outside the loaded photos are loaded and form a new cluster (see the cluster on the right hand side in Figure 5).



**Figure 5. Using a magnet to attract all matches .**

In the third case, the user can use a magnet to load matches from certain events in the timeline. By choosing “sweep timeline” in the shortcut menu of a magnet, the cursor becomes a magnet, and events containing matches are highlighted (see Figure 6). The user



**Figure 6. Using magnet to attract matches from certain events.**


can sweep the magnet over the timeline, which has the same effect as normal mouse hovering. Events can be selected by clicking on them, which will be highlighted in red and the number of total matches will be shown in the magnet. After dropping the magnet back on the canvas, matches from the selected events are loaded.

Besides the five fundamental types of magnets, there is another type of similarity-magnet. Through its shortcut menu, a photo can be converted to a similarity-magnet, which will attract other visually similar photos. The user can conduct a complex search based on multiple or duplicate magnets. Figure 7 illustrates the retrieval results of 1 location-magnet, 1 time-magnet and 1 similarity-magnet (rendered with a green frame). Photos matching multiple magnets are placed between the corresponding magnets.



**Figure 7. A search based on multiple magnets.**

When a new magnet is added in, the current state of BubbleBrowser will be automatically saved in an extra history sidebar. The user can also save one manually through the "history states" menu. The history sidebar can be shown/hidden with the H key. The yellow frame indicates the current screen (see Figure 8), and the user can switch to another status by selecting one snapshot.



**Figure 8. The history sidebar holds all saved snapshots.**

### 4.2.3 Underlying algorithms

#### 4.2.3.1 Event segmentation

Similar to Phototoc[18], events are detected based on the photos’ capture time. All photos are sorted chronologically, and the event boundary is defined as a noticeable time gap.  $g_i$  is the time difference between two successive photos  $i$  and  $i+1$ , and  $g_N$  is considered a gap if it is larger than a local average of all gaps on a logarithmic scale (empirically,  $K$  is chosen as  $\log(17)$  and  $d$  as 10):

$$\log(g_N) \geq K + \frac{1}{2d+1} \sum_{i=d}^d \log(g_{N+i}) \quad (1)$$

#### 4.2.3.2 Image analysis

Image similarity is calculated based on the FCTH feature extractor, which includes color and texture information in one quantized histogram [7]. Initially the image is segmented into blocks. For each block, a 10-bin histogram is derived for 10 preselected colors in the HSV color space. This histogram is then expanded to 24-bins by including hue-related information for each color. For each image block, a Haar Wavelet transform is applied

to the Y component (Luminosity in YIQ color space), and then the histogram is expanded to 192-bins by integrating the extracted texture information. A quantization is applied to limit the final length of the feature descriptor to 576 bits per image. The Tanimoto coefficient is used to measure the similarity between

two images, which is then normalized to a value between 0 and 1.

When a photo is converted to a similarity-magnet, the similarity between this example photo and other loaded photos will be calculated dynamically. Photos with a similarity higher than a threshold (empirically chosen as 0.5) will be included in the retrieval results (see Figure 9).



Figure 9. A similarity-magnet with the attracted photos.

#### 4.2.3.3 Model for magnetism simulation

Force directed layout is a popular approach to implement the magnet magnetism. The distance between each pair of items is controlled by the attracting force. In order to achieve a non-overlapping layout, a repulsive force is used to keep a certain distance between items. This method is easy to realize and has already been included in existing toolkits, such as Prefuse [12] for interactive Information Visualization design. However, it has one main shortcoming: Due to the co-existence of attracting and repulsive forces, the stability of the layout decreases enormously with increasing numbers of items.

Therefore, we propose another method to simulate the magnetism. Instead of introducing both attracting and repulsing forces for each pair of items, we only define the attracting force for a magnet and its matches. By pressing the left mouse button on one magnet, the attracting force is activated, and the corresponding matches are gradually attracted to it. Items attracted by multiple magnets will be placed between them (see Figure 7). The attracting force is canceled when the left mouse button is released or matches reach the boundary of the magnet.

For each match, its overall movement  $\vec{v}$  is calculated with Formula 2:  $n$  is the number of corresponding magnets,  $\vec{L}_p$  is the location of this item. For each magnet,  $\vec{L}_i$  is its location,  $m_i$  is its magnitude (set to 1.5 by default) and  $a_i$  is its attraction force.

$$\vec{v} = \sum_{i=0}^n \frac{m_i * a_i}{|\vec{L}_i - \vec{L}_p|} * (\vec{L}_i - \vec{L}_p) \quad (2)$$

For a tag-, event-, location and null-magnet,  $a_i$  is set as 1. For a similarity-magnet,  $a_i$  is the similarity between the match and the example. For a time-magnet,  $a_i$  is calculated with Formula 3:  $\Delta t$

$$a_i = 1 - \frac{\Delta t}{t_{\max}} \quad (3)$$

is the time difference between the photo's capture time and the seed date of the magnet.  $t_{\text{span}}$  is the time span of the magnet.

To avoid instability and oscillation, we do not introduce an automatic repulsive force to each pair of items. Instead, a repulsing functionality is offered on demand: By holding the

space key, photos repulse each other, temporarily and gradually transforming the current layout into a non-overlapping layout.

## 5. EVALUATION

We conducted an evaluation of PhotoMagnets, aiming to assess its effectiveness in supporting flexible browsing and searching. Considering the novelty of the interface components such as magnets, instead of comparing with another quite different system, we chose an in-depth study, in which a usability test and an explorative study were conducted. We were specifically interested in the performance of the combination of our two concepts in supporting different browsing and searching activities.

### 5.1 Participants

We recruited 18 participants, 6 female and 12 male. Their age ranged from 19 to 37 with an average age of 27 years. The participants were regular PC users and hobby photographers who own large photo collections. Only 4 participants had tagged their photos (no more than 30% of their collections).

### 5.2 Settings and procedure

Each participant was required to provide three photo collections (with an average size of 150 photos for each, and the time span ranged from 1 to 15 years). During the study the participants were equipped with a laptop. The study was conducted at the participants' home and lasted on average about 2 hours per participant. It was recorded on video using the Think-Aloud protocol. All the scores were rated on a 5-point Likert-scale where 5 represented the highest score.

The participants first filled out a pre-questionnaire with demographic information and the general experience with their photo collections. After a brief introduction of PhotoMagnets, the participants played around with a collection of 149 photos in a 30-minutes trial. Then they joined an interview, which consisted of two sessions. In the first session, the system usability was tested based on a set of pre-defined tasks covering the main functionalities, such as importing collections, refining events and tagging photos in TreeBrowser; loading photos from certain events and conducting search based on multiple magnets in BubbleBrowser. To offer the participants a chance to use PhotoMagnets more freely and to give us a chance to observe their behavior more closely, the participants were asked to execute two open tasks in the second session: Introducing their photos to the interviewer (storytelling) and creating a photo album as a gift for a friend who will move to another city (album creation). After completing all the tasks, they filled out a post-questionnaire concerning the overall impression of PhotoMagnets.

### 5.3 Results

#### 5.3.1 System usability

The timeline was commented as an important component ( $M=4.61$ ,  $SD=0.50$ ) that offers general information about the collections, such as time span and photo distribution. The color coding of the event histogram was appreciated ( $M=4.06$ ,  $SD=1.16$ ). Loading photos from the timeline by drag-and-drop was also appealing ( $M=4.72$ ,  $SD=0.57$ ).

In TreeBrowser, the automatic event segmentation was satisfactory ( $M=4.83$ ,  $SD=0.38$ ). The manual refinement was appreciated ( $M=4.78$ ,  $SD=0.55$ ) but seldom used. Although tagging is a tedious and time-consuming task [6, 20], the tagging functionality in TreeBrowser was highly appreciated ( $M=4.94$ ,

SD=0.24) and commented as very easy to assign multiple tags to multiple photos. Although no quantity of tags was required, the majority of our participants assigned multiple tags to 80% of their photos. Three participants expected the tagging functionality to be included also in BubbleBrowser.

Sweeping a magnet to load photos in BubbleBrowser was appreciated (M=4.22, SD=0.94). The red lines on photos were commented as helpful for deriving the relations between magnets and photos (M=4.17, SD=1.15). Figure 10 shows the ratings of magnet performance in the aspects of usefulness and ease of use, with error bars representing  $\pm$  standard error of the mean. We can observe that the tag-magnet received highest scores. The null-magnet was considered useful in maintaining a neat display, and the similarity-magnet was commented as entertaining. A more functional time-magnet was desired, for example, the matches could be automatically ordered by time.




Figure 10. Ratings for each magnet type.

### 5.3.2 Observation of user behavior

During the two open tasks of storytelling and album creation, the tag-magnet was the most frequently used magnet type (on average 4 times per participant). Event-, similarity- and null-magnets were used less and nobody used location- or time-magnets. Although the history sidebar was considered useful, it was only used by one participant who deleted some photos by accident. This illustrated that the history also represents an undo functionality for operation mistakes in addition to its original function of tracking the exploration history. Considering the overall positive feedback and the limited size of the tested collections, we expect that these functionalities will be used more often in a long-term usage with larger collections.

Related images were generally appreciated. 2 and 5 participants respectively were sidetracked by the suggested images. Regarding the performance in supporting these two tasks, magnets were appreciated in general and received higher scores with album creation (see Figure 11, with error bars representing  $\pm$  standard error of the mean).




Figure 11. Ratings for magnet in two open tasks.

The participants' practical behavior was closely observed in the two open tasks, and the following results were derived:

#### 5.3.2.1 Storytelling

The storytelling behavior was studied based on Balabanovic's storytelling modes [3]: In the photo-driven mode, "...the subject explains every photo in turn, the story prompted by the existing sequence of pictures". The story-driven mode is defined as "...the subject has a particular story in mind [...], then gathers the appropriate photos and recounts the story".

To introduce their photos, 5 out of 18 participants started with TreeBrowser. The photo-driven mode dominated and they told their stories in the slideshow window by explaining photos sequentially. The other 13 participants started with BubbleBrowser. The story-driven mode was prominent: Among the loaded candidates, they either manually selected certain photos, or used magnets to group the desired ones. They then told their stories based on these selections. Figure 12 shows a story created by a participant.

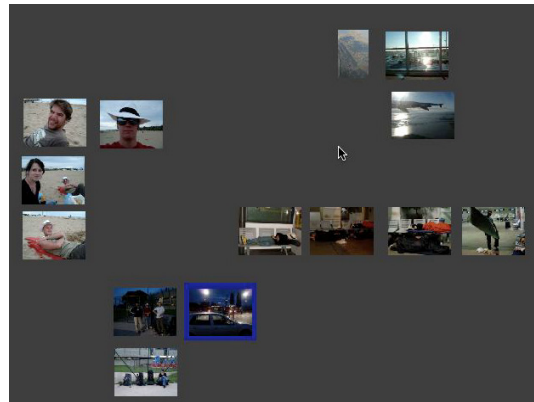


Figure 12. A story created by a participant.

Specifically, the participants spontaneously discovered some creative usages of magnets. Besides searching and clustering, magnets were also helpful for deriving implicit contextual information. For example, participant 9 used two tag magnets to attract photos containing two friends. The results reminded him of their relationship as they always appeared in a pair in those photos. Participant 8 even discovered a new way to exclude undesired photos: By reducing the magnitude of a magnet below zero, a negative magnet was created, which repulsed instead of attracted photos. This functionality was not foreseen in the interface design, but showed that the interface was appropriated by the user, which is a good sign regarding its flexibility.

In general, with a clear chronological order in TreeBrowser, the story lines were fluid and the quality of the stories was generally high. Storytelling in TreeBrowser was thus commented as easier than in BubbleBrowser.

#### 5.3.2.2 Album creation

Considering the flexible activities allowed in BubbleBrowser, all participants created the album in this view. Most of them first loaded some candidates with magnets, from which they made their final decisions, and then manually repositioned these selections as a collage.

Magnets were commented as useful for clustering and comparison. Although similarity-magnet performed not always correctly as expected, participants still considered it very

interesting and some of them even got inspirations from the unexpected results, which is consistent with Toms' suggestion of introducing serendipity via poor similarity [22]. For example, participant 7 used a similarity-magnet to choose one photo from a set of embarrassing photos of one friend. We expected that users may get more inspirations and discover more interesting findings with a larger data set.

## 5.4 Expert interview

Besides hobby photographers, we also discussed PhotoMagnets with four experts, three employees in digital media press and one professional photographer. In general, PhotoMagnets was commented as useful in supporting their professional workflows. They especially appreciated its value to facilitate their routine work of searching with loosely defined goals and dynamically producing different draft layouts, which is essential to enhance the discussion with their clients. Magnets, especially the similarity-magnet, were highly appreciated for grouping photos with a feeling of progression. Beyond the currently applied overall similarity, more advanced visual features were desired, such as similarity in a certain color theme and an adjustable similarity threshold. To introduce PhotoMagnets into the professional field, some modifications are necessary, for example, introducing different organization principles beyond time and events.

## 5.5 Discussion

During the evaluation we received an overall positive feedback on PhotoMagnets. The main functionalities, such as event segmentation, tagging, and loading of photos from the timeline, were generally appreciated. The magnets were especially appealing in facilitating search, clustering and comparison.

Although the magnet metaphor is a rather novel concept, the participants seemed to understand it and demonstrated some creative usages. Magnets allow the definition of areas of interest by loading, repositioning and grouping photos. The interactive nature of the interface enables the user to execute complex tasks with simple and active control, and thus facilitates complex query formulation and refinement. Generally the participants were immediately able to use magnets to accomplish the tasks of storytelling and album creation, which required a high degree of flexibility. As the results of the study confirmed, the magnet metaphor is consistent with the common sense of physical magnetism and illustrates the underlying system logic. The new uses of magnets and the different strategies developed by the participants show a high degree of interface appropriation, which speaks strongly for its flexibility. The stability of the visualization was one of the main concerns, but was successfully established using our specific method for the magnetism simulation.

To improve reproducibility of one usage session, the arrangement of the canvas can be automatically or manually saved in a history sidebar, which makes it easy to reproduce the exact same representation. The reproducibility needs further improvement, as the exact same layout for two subsequent applications of the same given set of magnets can not be promised due to the randomness of automatic repositioning.

The size of the test set in the evaluation was rather small and scalability needs further investigation towards realistic collection sizes. Occlusion is solved by the repulsing functionality, which should be further combined with a tight coupling method to avoid waste of space. Another possible solution may be grouping matches automatically and offering semantic zoom-in on demand.

Regarding the evaluation methodology, free exploration was encouraged in the two open-ended tasks, and the observation of user behavior yielded a number of insights on the users' practical usage patterns. In accordance with other studies of user behavior [9, 15], we noticed that users seldom conduct explicit search with well-defined targets, but instead often navigate through the collections.

Magnets help query formulation and ease the grouping and comparison of relevant items. They also let users understand the underlying relations and discover semantic relationships. The Magnet metaphor greatly benefits from a convenient tagging functionality, specifically with photos, which initially lack meaningful names. Serendipity can stimulate curiosity and in the photo case, we suggest to utilize similarity and semantic metadata (such as tags) to induce serendipity. Although the content analysis is not always as precise as expected, in some cases it can also practically stimulate serendipity by its poor quality. Due to the time constraint in the laboratory evaluation, we expect to learn more about the user's practical usage in a long-term experiment. Since our expert interview illustrated a general appreciation among professional users, a more elaborate study with experts will bring more insights on the applicability of PhotoMagnets in professional user groups.

We received encouragement as well as valuable suggestions for improvement. The most frequently required functionality was tracking time in BubbleBrowser, especially among the retrieval results of time-magnet. A selection area was desired to exclude selected photos from the influence of certain magnets. In addition, a tagging functionality in BubbleBrowser was requested. Tag-magnets were frequently used, so we expect that a well-organized structure of tags will further facilitate tagging. Semantic analysis of tags can be introduced to enhance the efficiency of tagging, for example by automatically suggesting relevant tags. Currently both structured and unstructured activities are supported in the two respective views, and the user can use the loaded photos as anchor points to switch between them. A semantic coordination can facilitate smooth switching between the two browsing styles and avoid losing track of the context. Since these features do not conflict with the essence of PhotoMagnets, we believe that they can be conveniently integrated in future versions.

## 6. SUMMARY

People's activities around media collections are often unstructured and explorative. Today's linear interface models cannot efficiently support this exploratory behavior. We explored how to support flexibility with photo collections and proposed PhotoMagnets to support flexible photo browsing and searching. Both structured and unstructured activities are supported and the user can actively manage their photo collections. In the evaluation PhotoMagnets received overall positive feedback, especially for the magnet metaphor. The open tasks in the evaluation allowed an in-depth observation of user behavior, from which interesting information about the practical usage patterns were derived. The high degree of flexibility was appreciated in supporting dynamical activities such as flexible browsing and searching, storytelling and album creation. Users also showed a high degree of interface appropriation. Further improvements are possible and in our future work we will continue the exploration of supporting an exploratory experience with media collections.

## 7. ACKNOWLEDGMENTS

This research was funded by the Chinese Scholarship Council (CSC) and by the German state of Bavaria. We would like to thank the participants of our user study.

## 8. REFERENCES

- [1] Ames M. and Manguy L. 2006. PhotoArcs: a tool for creating and sharing photo-narratives. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Ext. Abstracts (Quebec, Canada, April 22-27, 2006). CHI '06. ACM Press, 466-471.
- [2] Ahlberg C. and Shneiderman B. 1994. Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays. In Proceedings of the SIGCHI conference on Human factors in computing systems (Boston, Massachusetts, USA, April 24-28, 1994). CHI '94. ACM Press, 313-317.
- [3] Balabanović M., Chu L. L. and Wolff G. J. 2000. Storytelling with digital photographs. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Ext. Abstracts (The Hague, The Netherlands, April 1-6, 2000). CHI '00. ACM Press, 2657-2660.
- [4] Bawden D. 1986. Information Systems and the Stimulation of Creativity. *Journal of Information Science*. Vol. 12, No. 5, 203-216.
- [5] Bederson B. B., Grosjean J. and Meyer J. 2004. Toolkit design for interactive structured graphics. *IEEE Transactions on software engineering*, 30:535-546.
- [6] Bentley F., Metcalf C. and Harboe G. 2006. Personal vs. commercial content: the similarities between consumer use of photos and music. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Ext. Abstracts (Quebec, Canada, April 22-27, 2006). CHI '06. ACM Press, 667-676.
- [7] Chatzichristofis S. A. and Boutalis, Y. S. 2008. FCTH: Fuzzy Color and Texture Histogram – A Low Level Feature for Accurate Image Retrieval. In Proceedings of the 2008 Ninth International Workshop on Image Analysis for Multimedia Interactive Services (Klagenfurt, Austria, May 7-9, 2008). WIAMIS '08. IEEE Press, 191-196.
- [8] Cove J. and Walsh B. 1988. Online text retrieval via browsing. *Information Processing & Management*. Vol. 24, No 1, 31-37.
- [9] Drucker, S., Wong, C., Roseway, A., Glenner, S. and De Mar, S. 2004. Mediabrowser: reclaiming the shoebox. In Proceedings of the working conference on Advanced visual interfaces (Gallipoli, Italy, May 25-28, 2004). AVI '04. ACM Press, 433-436.
- [10] Frohlich D., Kuchinsky A., Pering C., Don A. and Ariss S. 2002. Requirement for photoware. In proceedings of ACM Conference on Computer Supported Cooperative Work (New Orleans, Louisiana, USA, November 16-20, 2002). CSCW '02. ACM Press, New York, NY, 166-175.
- [11] Gemmell J., Aris A. and Lueder R. 2005. Telling stories with mylifebits. In proceedings of IEEE International Conference on Multimedia & Expo (Amsterdam, The Netherlands, July 6-8, 2005). ICME '05, 1536-1539.
- [12] Heer J., Card S. and Landay. J. A. 2005. Prefuse: A Toolkit for Interaction Information Visualization. In Proceedings of the SIGCHI conference on Human factors in computing systems (Portland, Oregon, USA, April 2-7, 2005). CHI 2005. ACM Press, 421-430.
- [13] House N. V., Davis M., Takhteyev Y., Good N., Wilhelm A. and Finn M. 2004. From what to why: social use of personal photos. In proceedings of ACM Conference on Computer Supported Cooperative Work (Chicago, Illinois, USA, November 6-10, 2004). CSCW '04.
- [14] Huynh D., Drucker S., Baudisch P. and Wong C. 2006. Time quilt: Scaling up zomable photo browsers for large, unstructured photo collections. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Ext. Abstracts (Quebec, Canada, April 22-27, 2006). CHI '06. ACM Press, 1937-1940.
- [15] Kirk D., Sellen A., Rother C. and Wood K. 2006. Understanding photowork. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Quebec, Canada, April 22-27, 2006). CHI '06. ACM Press, 761-770.
- [16] Marchionini G. Exploratory search: from finding to understanding. *Communications of the ACM*. Vol. 49, No. 4 (April 2006), 41-46.
- [17] Naaman M., Song Y.J., Paepcke A. and Garcia-Molina H. 2002. Automatic organization for digital photographs with geographic coordinates. In proceedings of ACM/IEEE Joint Conference on Digital Libraries (Portland, Oregon, USA, July 14-18, 2002). JCDL '02. ACM Press, 53-62.
- [18] Platt J., Czerwinski M. and Field B. 2002. Photoc: automatic clustering for browsing personal photographs. Microsoft Research Technical Report, MSR-TR-2002-17 (2002).
- [19] Preece J., Rogers Y. and Sharp H. 2002. Interaction design: beyond human-computer interaction. John Wiley.
- [20] Rodden K. and Wood K. R. 2003. How do people manage their digital photographs. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Fort Lauderdale, USA, April 5-10, 2003). CHI '03. ACM Press, 409-416.
- [21] Shen C., Lesh N. B., Vernier F., Forlines C. and Frost J.: Sharing and building digital group histories. In Proceedings of the conference on Computer supported cooperative work (New Orleans, Louisiana, USA). ACM Press, 324-333.
- [22] Tomes E. G. Serendipitous Information Retrieval. 2000. In Proceedings of the First DELOS Workshop on Information Seeking, Searching and Querying in Digital Libraries (Zurich, Switzerland, December 11-12, 2000).
- [23] Toyama K., Logan R. and Roseway A. 2003. Geographic Location Tags on Digital Images. In Proceedings of the 11th international conference on Multimedia (Berkeley, CA, USA, 2003). MM '03. ACM Press.
- [24] Yi J. S., Melton R., Stasko J. and Jacko J. Dust & Magnet: multivariate information visualization using a magnet metaphor. In Proceedings of the international conference on Information Visualization (London, UK, 2005). IV '05.