

Visualizing Social Photos on a Hasse Diagram for Eliciting Relations and Indexing New Photos

Michel Crampes, Jeremy de Oliveira-Kumar, Sylvie Ranwez, and Jean Villerd

Abstract—Social photos, which are taken during family events or parties, represent individuals or groups of people. We show in this paper how a Hasse diagram is an efficient visualization strategy for eliciting different groups and navigating through them. However, we do not limit this strategy to these traditional uses. Instead we show how it can also be used for assisting in indexing new photos.

Indexing consists of identifying the event and people in photos. It is an integral phase that takes place before searching and sharing. In our method we use existing indexed photos to index new photos. This is performed through a manual drag and drop procedure followed by a content fusion process that we call ‘propagation’. At the core of this process is the necessity to organize and visualize the photos that will be used for indexing in a manner that is easily recognizable and accessible by the user. In this respect we make use of an Object Galois Sub-Hierarchy and display it using a Hasse diagram. The need for an incremental display that maintains the user’s mental map also leads us to propose a novel way of building the Hasse diagram. To validate the approach, we present some tests conducted with a sample of users that confirm the interest of this organization, visualization and indexation approach. Finally, we conclude by considering scalability, the possibility to extract social networks and automatically create personalised albums.

Index Term—Information visualization, Hasse Diagram, indexation, social photos, formal concept analysis, Galois sub-hierarchy.

1 INTRODUCTION

Digital cameras and mobile phone cameras are now in common use for the capture of pictures of people during social events such as friend’s parties, anniversaries and weddings. We call this class of photos ‘social photos’. In order to navigate through large collections of such photos and share them it is necessary to organize and visualize them, normally requiring indexation beforehand. Indexing social photos consists of tagging them with the event (date, place, nature of the event, etc.) and the names of the people. In traditional systems, tagging and organizing are separated. However even when they are indexed, it is not clear how to organize and visualize social photos. In this paper we propose a strategy whereby visualization, organization and indexation support each other.

Although indexing photos with the name of people can be performed manually (a typical example being Facebook) it becomes excessively complicated and time consuming, if not impossible, when one has to manage even tiny collections of photos representing small groups of people. As the state of the art below will show, there are no sufficiently satisfactory systems or supporting tools for this task. Conversely, face recognition techniques, although promising, are still far from providing concrete results for social photos that are not taken in controlled environments. The method we propose is still manual; however it is assisted by a particular visualization and organization strategy. It is based on three key features:

Indexing through ‘propagation’. The content of already indexed photos is used to index new photos. This method requires that photos used for indexing are clearly visible and organized.

Visualization utilising a Hasse Diagram. All existing indexed photos are organized and displayed using classification and

visualization techniques from Formal Concept Analysis (FCA) [1], among which a Hasse Diagram.

Incremental integration. Photos that are newly indexed are automatically incrementally integrated and visualized into the Hasse diagram. These photos then form a part of the basis for further indexing.

All the functionalities described have been integrated in an application called PhotoMap that was tested with volunteer students. According to the users and according to measures that are presented at the end of the paper, the methods and tools that are proposed perform equally or better (albeit with some limitations) than present solutions that are used by applications such as Facebook or Flickr.

Our method is based upon the analysis of social photos. However, it may be extended with care and consideration to other photo collections, or even other item sets.

2 STATE OF THE ART

Social photos are part of personal photos. A user study in [2] shows that most people do not want to bother with indexing. However indexing is sometimes necessary, in particular when the intention is to automatically build and share personalized albums as it will be presented at the end of the paper. Many methods are currently being explored for managing digital photos, particularly social photos both with events and people’s names. Large surveys can be found in [3] and [4].

Photo library management consists in two major tasks: indexing (annotating) photos and then organizing indexed photos in order to ease forthcoming browsing and retrieval. Concerning the indexation step with people’s names, various approaches have been explored to acquire useful metadata. Among them, much research focuses on automatic extraction. Face recognition systems have significantly improved in recent years particularly for controlled positions and illumination [5]. Some are even available as commercial or public applications such as Apple’s iPhoto or Google’s Picasa. Their overall performance is improving though they still require controlled environments with favorable light and orientation (see for instance a summary of the state of the art in this matter in [6] and [7]). This is far from being the regular kind of contexts that we find in social events where light, positions and even peoples’ faces are not controlled as can be seen in this paper’s figures.

• Michel Crampes and Sylvie Ranwez are with LGI2P/EMA Research Center, France, E-Mail: {Michel.Crampes, Sylvie.Ranwez}@ema.fr.

• Jean Villerd is with INRIA-LORIA, France, E-Mail : jean.villerd@loria.fr.

• Jeremy de Oliveira-Kumar is with UNSW, Sydney, Australia, E-Mail: jdok706@cse.unsw.edu.au.

Manuscript received 31 March 2009; accepted 27 July 2009; posted online 11 October 2009; mailed on 5 October 2009.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

In [4] a less ambitious approach is proposed. The authors use automatic recognition of clothes' coloring to identify individuals within an event. They report interesting results after tests even if the recognition process produces many errors. Applying this method to social photos is difficult because the examples that are provided in [4] show people in limited number, with visible upper bodies, under good lighting and with clearly discriminating clothing. As a result manual indexation using computer assistance is the general practice for most social photos, and it is apparent that it will remain this way in the near future.

Manually typing in people's names is both time consuming and tedious. To overcome this drawback several tools for assisting users in manual indexing are proposed, some being described in [3]. The seminal works of Ben Shneiderman [8] on photo annotation led to the concept of direct annotation, which consists in selectable name labels that users drag-and drop directly on the photo, following direct manipulation principles rather than filling in text captions. They are also applied to social tagging which is used in applications like Flickr or explored deeper in the project LabelMe [9] with later expectations for learning and identification of objects in images. Similarly Facebook applies dragging and dropping of individuals' names at precise locations within each photo. However this method limits productivity. We also use dragging and dropping though in a converse way because we apply it to photos as will be presented in section 6. A simple method such as typing in the first letters of a name and automatically completing it is also an efficient approach such as in Flickr as we shall see in section 7.

Bulk annotation [10] improves productivity and efficiency; it consists in automatically or manually selecting photos that are similar according to some criteria and indexing them in one go with existing or new labels, such as in Flickr. We use a similar method when gathering photos in a container for indexing as it is described in section 6. Bulk annotation is also proposed in [4] where event identification uses temporal metadata to build an event hierarchy from annotated photos. In our approach, users select or create an event before indexing with people's names. When a photo is indexed with names, it also inherits the event and using an event hierarchy is not necessary. But in the case where the event is unknown the method presented in [4] could complement our method which only focuses on people.

At the end of the indexation step, photos are indexed with one or several meaningful metadata, such as date, location, event name, people names. The next step deals with the visual organization of photos for supporting future browsing. PhotoMesa's interface [11] dynamically groups photos based on photo metadata, allowing users to combine nominal (places, people) metadata with ordinal metadata (date). Semantic regions [12] have been introduced to let users customize their information space. A region is a set of photos, grouped by the value they share for a specified metadata. Regions can be displayed in a tree-map manner. When new indexed photos are added, they are automatically placed in the appropriate regions, following a metaphor called "fling-and-flock". Sub-regions can be defined but the resulting structure is a single-inheritance hierarchy, so a photo belonging to two regions will appear twice. Our approach concentrates on one particular metadata (people's name) and merges the indexation and organization step in the sense that the same structure is used for both steps in an incremental manner. The overall idea is to allow the user to index new photos by propagating and reusing annotations from existing photos. To do so we propose a visual organization that groups photos by subsets of people, following a lattice structure that allows multiple-inheritance. As a result, each photo belongs to only one group and each group is represented by a unique photo that summarizes it. A group specifies a unique subset of people which can be visually identified and whose names can be propagated to new photos showing the same subset people.

Indexing through information propagation may be performed when inserting contextual information such as in [13] where the

authors use time, geodesy and social information to recover the event and the people. However there is still the need to provide each photo or group of photos the names of the people they contain.

We propose to index photos with people's names through photo content propagation as an original and alternative method that may be used on its own or that may complement face recognition or context information propagation. It is based on the core idea that the best way to index a collection of photos is to recover previously indexed photos organized in a particular display where one can incrementally capitalize on existing information. Moreover, indexing, visualizing and organizing are inherently intertwined.

We use a particular incremental organization and visualization strategy supported by a Galois sub-hierarchy and a Hasse diagram. This kind of display is well known in Formal Concept Analysis where it is used to represent concept hierarchies. FCA based photo browsing methods are described in some applications such as in [14] and [15]. However the authors do not apply these methods to indexation and do not propose the use of an incremental Hasse diagram.

3 INDEXING PHOTOS WITH PHOTOS

We introduce our indexing method with a real example of a person who wants to index a set of photos within our application. Mary, our test subject, went to an anniversary party and took approximately 40 photos that she aims to index with peoples' names in order to share them.

She can display all her photos on the screen that we refer to as the 'photo-board'. First of all, she may index all photos taken during the party with place, date and type of event. All this data is declared under a unique 'event' that she can create and that appears on top of the photo-board. The event she selects will automatically be used as part of the indexation of new photos and become part of the entire event's data.

The most time consuming element of the indexation process concerns the people that are present in the photos. Any photo may show one or more people in different situations during the party. Our indexing method and algorithms increase the speed of this process and improve the accuracy of the result. They can be activated after a few photos have already been indexed or where past events already exist with photos that have been indexed using a subset of the same people. This initialization phase which involves creating a character list will be presented below. We consider for the moment that a set of indexed photos is already available as well as a character list of the people within these photos in alphabetic order. The photo-board is presented with the photos to be indexed on the right and the character list on the left in Fig. 5.



Fig. 1. Indexing a photo with two indexers

To index new photos Mary will make use of propagation, i.e. she will use already indexed photos to index new ones. We call 'indexers' those photos that have already been indexed and that will be used for indexing. We call 'indexees' those photos that are yet to be indexed. In Fig. 1 and Fig. 6 the user is dragging with the mouse an indexee which is surrounded by a blue square. She may visually identify the characters in the picture using a zoom or fish eye lens [16] if this is necessary. Mary can drag the indexee near one or several indexers that contain the people in the photo. Releasing and pressing the button on the mouse selects the nearest indexer which gets surrounded by a blue circle and consequently shows its content. In Fig. 1 the indexee contains three characters. The indexer at the top has already been selected and the indexer at the bottom that is being selected reveals its content. Double-clicking on the mouse button triggers propagation that works in the following way.

Formally let N be the set of characters that are visible on the indexee, S_1 the set of characters that are visible on indexer 1 and S_2 the set of characters that are visible on indexer 2 (for instance Jeremy and Maria as can be seen in Fig. 1). The indexee's index is built through the union of the selected indexers' contents, i.e. $N = S_1 \cup S_2$ in our example.

Since S_1 contains {Jeremy and Willey} and S_2 contains {Jeremy and Maria}, therefore the set {Jeremy, Maria, Willey} is propagated to N . There may be cases when some characters in the indexee are not present in the contents of the photos that are used for indexing. These characters can easily be created and added to the character list and to the content of the indexee at any time. There may also be cases when there are 'intruders', i.e. more characters in the union of the indexers' contents than desired. Consequently the indexee would receive more content than desired through propagation. However the 'intruders' are necessarily in the character list because they were created when indexing some photos that are now the indexers. Clicking on the intruders' names in the character list on the left will subtract them from propagation.

The efficiency of this process may be analyzed in comparison to a basic methodology whereby the user only needs to click on the characters belonging to N in the character list. Several tests conducted have proved that the method and the tools proposed were simpler and faster as will be shown in the evaluation, though with some limitations.

As can be seen above, a preliminary phase consists of organizing the indexers on the photo board so that Mary can easily retrieve those she can make use of. Our method of organization proposes the use of a Hasse Diagram to achieve this result.

4 ORGANIZING INDEXED PHOTOS AS A HASSE DIAGRAM

4.1 The Concept Lattice of indexed photos

In this section, we describe how to organize indexed photos in a Hasse Diagram located on the photo-board for further indexation of new photos through propagation. A Hasse Diagram is a well known representation of a Concept Lattice (or Galois Lattice). In Formal Concept Analysis [1], a set of objects with properties (or attributes) can be organized in a lattice of concepts that contain these objects according to their property commonality. In our case, we consider the photos as objects and the names of the characters in the photos as properties. The organization process starts with the formal context, i.e. a table with the objects as rows and the properties as columns. Any entry is marked (e.g. valued 1) if the corresponding object possesses the corresponding property, and is not marked (e.g. valued 0) if the object does not possess such an entry. Formally, a formal context is a triple (G, M, I) where G is a set of objects, M a set of attributes and I is a binary relation between the objects and the attributes, i.e. $I \subseteq G \times M$. Table 1 presents a formal context taken from our example of the set of already indexed photos O and the set of properties M that are the names of the characters present in the photos.

Table 1. A context with photos and characters

Objects/Properties	Maria	Willey	Jeremy	Peter
Photo P1	1	0	0	1
Photo P2	0	0	1	1
Photo P3	1	1	1	1
Photo P4	1	1	0	0
Photo P5	0	0	0	1
Photo P6	0	0	0	1

The next step in the construction of the Concept Lattice is to define concepts. A concept is defined by a pair of subsets: a subset of objects (called the extent) and a subset of properties (the intent) that the objects share. In our case, a concept is a subset of photos that share the same subset of characters' tags. For the context presented in Table 1, the concepts are nodes in the graph shown in Fig. 2, such as the concept $(\{P2, P3\}, \{\text{Jeremy}, \text{Peter}\})$ where $\{P2, P3\}$ is the extent and $\{\text{Jeremy}, \text{Peter}\}$ is the intent.

Following the process of concept identification, the next goal is to build a lattice whose elements are the concepts. A partial order on formal concepts is defined as follows.

$(O_2, A_2) \leq (O_1, A_1)$ iff $O_2 \subseteq O_1$ (and consequently $A_1 \subseteq A_2$).

The ordered concepts form a complete lattice called a "concept lattice". The set of concepts L is completed if necessary by a top concept that contains all objects (see Fig. 2) and a bottom concept that contains all attributes.

4.2 The Hasse Diagram

Fig. 2 shows the photo Concept Lattice of the example as a Hasse diagram. It is a graph whose vertices are the concepts, ordered from top to bottom according to their order in the lattice, with their intent and extent shown within brackets (bold items will be explained later); the edges of the graph are drawn between concepts when two concepts are directly ordered (without transition through another concept). In our example, each concept is a group of photos P_i with their common characters. As can be seen, a photo as well as a property may appear in several concepts.

4.3 The Object Galois Sub-Hierarchy

In order to simplify the display it is possible to only show reduced extents (bold items in Fig. 3). The reduced extent of a concept (O, A) is the set of objects that belong to O and do not belong to any lower concept, i.e. objects that do not have any attributes other than those in A .

For each object (photo), there exists a unique object-concept that represents the most specific concept that contains the object. Since we want a photo to be associated with a unique and most specific description, only object-concepts are kept in the lattice. This act of pruning is the one proposed in [17] under the name PCL/X. It is illustrated with our example in Fig. 3.

It is now a particular case of what is called a Galois Sub-Hierarchy [18]. To be more explicit we call it an Object Galois Sub-Hierarchy (OGSH). It is a lighter visualization of the data when only focusing upon the objects, in our case the photos.




Fig. 2. Concept lattice associated with the context in Table 1




Fig. 3. The photo Object Galois Sub-Hierarchy

4.4 Building incremental Hasse Diagrams

Many algorithms have been proposed for building a Concept Lattice and its Hasse Diagram (see a quick survey in [19]). More precisely [20] compares three main applications that only consider GSHs that are also of our concern. These algorithms begin by computing the topology of the graphs and then generate the Hasse Diagram, or at most incrementally build the Hasse Diagram whilst computing the topology of the graph. Consequently they can hardly be used in our context since the topology of the graph is always evolving.

In photo indexing, every time a new photo is indexed using the Hasse Diagram of indexers, it must be integrated into the diagram that must consequently be updated. Since the diagram is made of groups of indexed photos, i.e. concepts that are used to index new photos, it is of the utmost importance that the user easily finds the groups of people she is interested in. With this in mind, every time a new photo is added to the diagram, there must be a minimum change of the display with all concepts staying in or near their current location. The overall objective of this strategy is to maintain as much as possible the user's 'mental map' [21]. All diagram display algorithms that require a preliminary computed concept lattice are not applicable here since our Hasse Diagram must be incremental. All other algorithms that build the Hasse Diagram incrementally could be interesting if they would keep all concepts in place or move them as little as possible. However these algorithms are with good reason targeted towards two objectives: speed and aesthetic criteria, where a limitation on the number of edge crossings plays an important role [22]. Although these objectives are still interesting to us, we find a dilemma between edge crossing limitation and an incremental display in order to maintain a stable mental map. This last issue is the most important for us, contributing to the reasoning behind why we will use the links in a way that will accept edge crossings to a certain extent as is explained below.

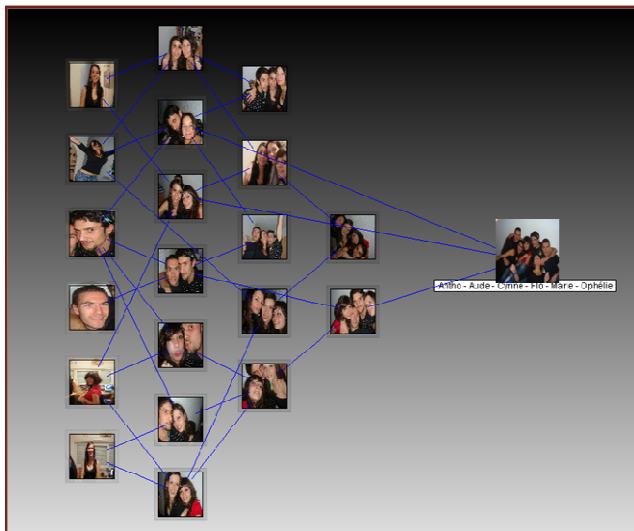


Fig. 4. Self organizing photo Hasse Diagram; for practical reasons, the Hasse Diagram is organized from left to right.

The new algorithm we propose relies upon two techniques. Firstly it makes use of Force Directed Placement (FDP), a well known algorithm in graph drawing [23], which is applied for

instance to lattice drawing in [24][25]. However, in the latter case the lattice is computed first which is against our goal of an incremental display while maintaining a stable mental map. Secondly we take the novel approach of building the Hasse diagram first through FDP and then the OGS emerges as a result of this display.

The algorithm proceeds as follows. Similar to [24], a rank attribute is added to any new photo after it has just been indexed. It is an integer value equal to the number of properties in the object, i.e. the number of characters. A horizontal force is applied to all indexed photos, with an intensity that is proportional to the difference of rank value. As a result, the photos move and settle down at a precise x position corresponding to their rank. They form well spaced columns. Following this they are fixed in x and a force is applied in each column along the y coordinate. The value of this force between any pair of photos is calculated on the basis of the Hamming distance between them, i.e. an integer which is calculated as the number of attribute differences between two attributed objects. As a result of this force, all photos that have the same characters pile up to form groups. These groups are the object concepts, and all groups are separated vertically. In order to present a better display, a third force is applied, which separates the groups at a certain distance, providing a well-spaced separation between them. Every time a new photo is indexed, it either joins an existing concept with the same group of characters or it creates a new group in a column and reorganizes the column vertically. Alternatively it may create a new column, pushing horizontally the existing columns. When a new group is created the algorithm identifies a new concept which is represented by the photo. In practice, a new indexed photo interacts with existing concepts rather than the photos that are part of the concepts. This happens in order to limit the number of interactions. When the new photo creates a new concept links are added to the nearest upper and lower concepts that contain common characters. Fig. 4 shows a result that was built during the tests described in our evaluation. For practical reasons, the Hasse Diagram is organized from left to right. All links are visible between concepts, but as will be shown below, a more desirable strategy is to show only some links.

5 USING THE HASSE DIAGRAM FOR INDEXING

Endowed with a Hasse Diagram, it is now possible for a user to index new photos. As an example we observe Mary indexing a photo with a demonstrative complex interaction. Generally interactions are much simpler and straightforward.

5.1 Adding tags through propagation

Fig. 5 is a snapshot of the screen after some photos have already been indexed. To index a new photo, Mary drags it from the pile on the right and may apply a zoom effect to better identify the people that are present. Fig. 5 shows the photo which is to be indexed at the top of the screen with its standard EXIF data (it is a code containing technical data, the date, etc.). Mary visually identifies four characters with which she must index the photo: Helena, Jeremy, Karol and Maria. In the column on the left, one can see in round pictures those characters that have already been identified from indexed photos. They are in alphabetical order. Updating this list is described below. The photos that are already indexed are organized in a Hasse Diagram in the centre.

Mary can easily spot Jeremy in a group of two people. She drags the photo on top of it. Keeping the right button down, the group becomes circled in green and shows its characters: Jeremy and Richard. Mary may either finalize the indexing using a right-click menu, or just obtain the characters of the selected photo. However, in this particular case, Mary does not want to select the characters of the group because a line shows up indicating that there is another group to the right that contains the same characters of the selected groups plus another character: Helena. This line is a hint for going straight to a bigger group if it is interesting. Since Mary is interested both in Jeremy and Helena, she keeps the button down, leaves the two people concept which loses its green circle and she moves to

the three people concept that now turns to green (Fig. 6). In order to “catch” the three characters, she releases the button. The circle now turns to blue, the three characters are added to an indexing bag for the photo and their pictures are visually extracted from the character list on the left (visible in Fig. 7). She now has Helena, Jeremy and an intruder (Richard) in the indexing bag. She may drag the photo again onto another group of three people just above. This time she catches Karol, Maria and Richard again (Fig. 7).

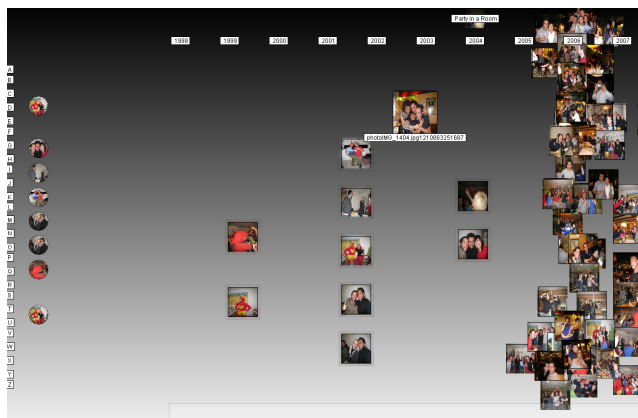


Fig. 5. A new photo to index

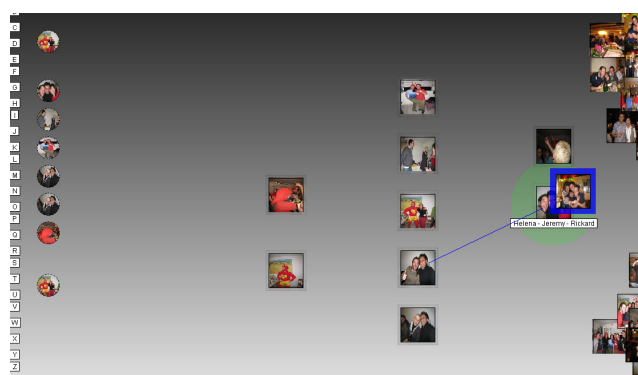


Fig. 6: Gathering three characters

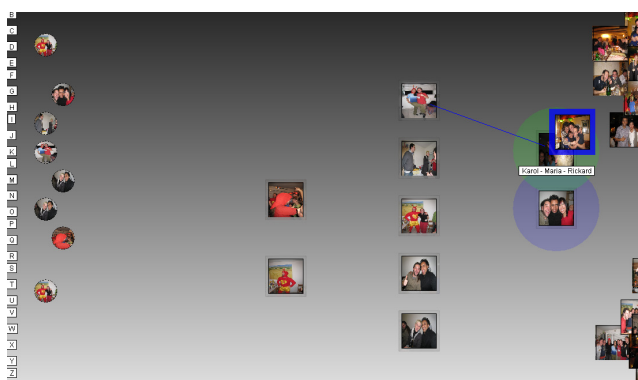


Fig. 7. Getting two more characters

Mary now has five characters in her indexing bag including Richard that she does not want. She drags the photo to the character list over the picture of Richard who is consequently withdrawn from the list. She can right-click and the photo is indexed with the four expected characters. The incremental display algorithm is triggered. If the photo had been indexed with a unique existing concept, it would have joined the corresponding photo pile. If Mary had formed a new group with an existing number of people, the photo would have created a new concept in an existing column. In the current case, since the four character column does not exist, the photo

creates it on the right in the Hasse Diagram. Finally the four characters slide back within the character list (Fig. 8).

6 ENHANCED VISUAL FUNCTIONS AND INTERACTION

The Hasse Diagram display proposes other functions in order to be completely efficient when indexing. Some are traditional (creating new characters and bulk indexing), some are more specific

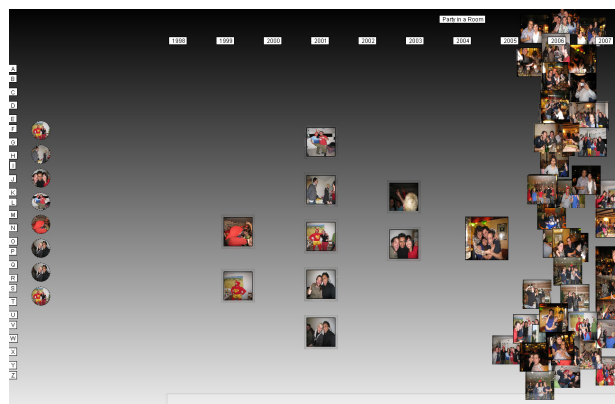


Fig. 8. Updating the Hasse diagram

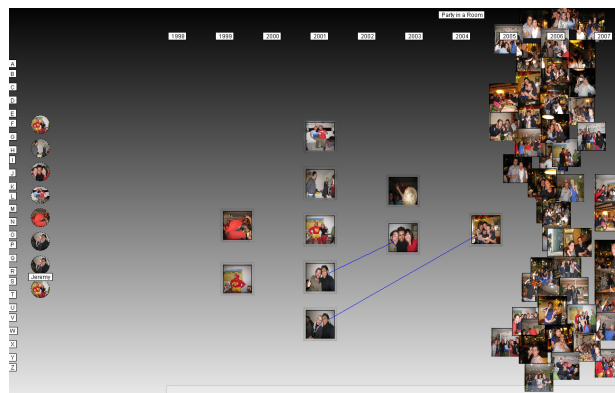


Fig. 9: Looking for a person in groups using links

Creating new characters. When indexing a new photo, there may be an individual in the group who is neither present in existing concepts, nor in the character list. Her name can be typed in like in Flickr and the photo being indexed may be used to represent her in the character list.

Bulk indexing. Indexing several similar photos simultaneously is very efficient. Mary can create a container and drag into it several indexees that show the same group of people. She may then drag the container with its photos over concepts in the Hasse Diagram and index the collection as a unique photo. When indexed, the photos in the container pile up and are included in the Hasse Diagram either in the same concept or by creating a new concept.

Searching ‘Who’s with whom’. We have already seen how the lines between concepts help with navigation during indexing. They can also be used to navigate among the piles of indexed photos to see who is with whom. Hovering the pointer over a concept shows the other concepts that contain the same group of people with others. It is also possible to move over an individual in the character list and the concepts she is part of are revealed within a sub-hierarchy. In Fig. 9 Jeremy appears in four concepts, which means in four different groups of people. Each concept could contain several piled up photos. He is alone with Maria in a pile of photos, while in another pile he is with Maria as well as Helena and Karol. In another pile he is alone with Richard and in pile number four he is with Richard and Helena. The names are not shown because we think that

at this point it is more important for the user to clearly observe the links. Moving over a pile may reveal the names whilst still showing local links from this pile.

Pile spreading. Indexed photos are hidden in piles within concepts and cannot be seen by the user, excluding the photos representing the concepts. If Mary wants to browse and select a particular photo, she can switch on a tool called 'the spreader'. When hovering over a concept with the mouse the hidden photos separate and can be explored. When leaving the concept the photos return to their pile.

7 TESTS

7.1 Organization

Two tests have been conducted. The objective was to compare PhotoMap's performances during social photo indexing with Facebook's and Flickr's performances, the most popular photo indexing applications. They also propose the most common ways of manual indexing. The subjects of the tests were volunteer students, between the ages of 20 and 22 years.

First test. It involved three males and one female. The objective was to provide qualitative opinions on PhotoMap and to fine-tune the assessment method for the second test. Considering the goals of this tests its results will not be presented below.

Second test. It involved seventeen students, among which 12 males and 5 females. They declared being interested in social photos, but none of them declared intensively practicing particular photo software. Only two of them knew Flickr, and 60% occasionally indexed photos with Facebook. They were given a USB key containing 40 social photos that had been taken during a friends' party that they had attended a few days beforehand. Consequently all of them knew the people in the photos. However, none of them knew the photos that had been taken by another student. The photographer took his photos as usual since he did not know at that time that he would be involved in this test. The experiments had to be performed on the same set of photos with Facebook, Flickr and PhotoMap in random order. Five photos were used for a short training course on each application (section 2 describes the indexing process with these applications). Two supervisors observed the subjects during this process, measuring, taking notes and asking open-ended questions at the end of the test.




Fig. 10: Time measures (in seconds) for indexing 15 photos with (1) Facebook, (2) PhotoMap, (3) Flickr

After training, the test subjects were asked to index 20 photos to have a core set of indexed objects. The first experiment was navigation. They were asked to search a particular photo with 3

people that was presented to them among the previously indexed 20 photos. Their time performance was measured. The second experiment consisted of measuring the amount of time taken to index the remaining 15 photos within the collection. They finally were asked to rank the three applications according to the words 'simplicity', 'quality of results', 'assistance' and 'fun' (1 for 'very good', 2 for 'good' and 3 for 'not so good').

7.2 Results.

Results with navigation. In test 2 it took a mean average of 11 seconds with PhotoMap and 20 seconds with the other two applications to find the particular photo with 3 people. The subjects particularly appreciated the human grouping organization and the spreader to find the photo in a stack.

Results for indexing. The first result of test 2 was unexpected and specific. Seven of the seventeen subjects gave up indexing the 40 photos with Facebook or/and Flickr before they had completed because it was too boring. Since this was not planned, we let them abandon indexing further. However, all of them indexed the photos with PhotoMap completely. The answer to an open question about this fact was that the incremental Hasse Diagram was pleasant to interact with and they found it was good for revealing the different groups making indexing immediately rewarding. Considering the 10 people that did not abandon in test 2, the time for indexing 15 photos showed a mean average of 3'10" with Facebook, and no meaningful difference between Flickr (2'30") and PhotoMap (2'35") (see Fig. 10).

Results for questions. Fig. 11 shows only results for PhotoMap and Flickr since Facebook scores less than these two applications for all questions. Each bar represents the percentage of answers 'very good' or 'good'. PhotoMap compares equally with Flickr for the first three questions (bars n° 1 to 3). It is particularly pleasant to use compared to Flickr (bars n° 4: 'fun'). This last result is interesting since it may explain why nobody abandoned indexing with PhotoMap.




Fig. 11: Answers to questions. From left to right (1) simplicity, (2) quality, (3) assistance, (4) fun.

7.3 Synthesis of results

The Hasse diagram on the one hand and its incremental construction on the other hand seem to perform well when users organize, visualize and index social photos in a limited number and on their own diagram. Performances are significantly better when navigating with PhotoMap than with the other two applications. Time measurements for indexing are not significantly better on PhotoMap than on Flickr. However it appears that users get less bored and appreciate the fact PhotoMap gives a dynamic global view of the photos and the people during the indexing process. Our main result is

that PhotoMap turns indexing into a creative visual activity, at least for a limited number of photos which was the case for this test.

Two main issues have still to be investigated in another future test: the case where many photos have to be displayed and indexed (scalability) and a situation where a user is bound to index new photos on a Hasse Diagram that she did not initially build.

8 LESSONS AND PERSPECTIVES

Improvements. Another version is being prepared; as expected it will take into account feedback from the tests. Speed of updating is stressed by users. A lot of other improvements have also been proposed such as a wider use of links for retrieving groups, a different organisation of the columns on the screen, and the suppression of the first column that could be included in the character column.

Scaling up to hundreds of social photos is interesting for managing a whole collection. PhotoMap can merge several events and a more efficient pan and zoom tool will be available in the next version to help visualizing many concepts together on the screen. In saying this, the current problem is that the pictures become very small on top of concepts, and it is difficult to identify a particular group. In practice, some use case studies make us think that people don't want to visualize a whole collection, but instead look for groups of photos containing one or two particular persons. It may be used as a filter to limit the number of concepts to be displayed and the diagram will become smaller and more manageable. Another complementary strategy will be to create a dynamic Hasse Diagram that may be displayed in a scrollable grid and that will be limited to six columns, with the last column containing groups of six people or more. This is possible because other investigations that are not presented in this paper show that there are few groups of four people or more. Conversely there are often many groups of two or three people. Consequently there may be problems with the size of these two columns. When columns are too big in the Hasse Diagram, they could be displayed on two or more columns in the grid. This idea requires more research because it will also be valuable for organising other photos than social photos and indexing them through propagation. The objective is to be able to clearly display about 100 concepts containing more than 200 photos in the Hasse Diagram. Beyond this limit people will use the scroll bar or the pan and zoom facilities.

the different concepts. Extracting social networks from social photos is also described in [26]; however the author's method is based on the analysis of individual photos and not on the analysis of concepts. In another paper we will explain why we think that concepts are a better support for extracting social networks. From these social networks, it is also possible to automatically build personalised photo albums that may be sent individually to people. These other possibilities were presented to the test subjects after their participation. The spontaneous reaction we observed was extremely positive; however other user tests must be performed to obtain more concrete results about their perception. Fig. 12 shows the photo album suggestion tool. When clicking on a button at the bottom of the screen, all photos slide and pile up on the right. The characters come out from the character list on the left. They then form a cloud with proximities depending upon their co-occurrences in the concepts. Each character is given a personalised album that includes closest friends. When moving over a particular character, links with other close friends appear, as well as her photo album; the photos in the album are visible at the bottom of the screen and the characters that are present in the album are seen in the character list. From Fig. 8 to Fig. 12 and back, there is a fluid animation that again aims at maintaining the user's mental map.

This automatic album creation functionality is not only something added to PhotoMap. It is at the core of the application since, as it was explained at the beginning of the paper, people are ready to perform indexing only if it is necessary or if there is a measurable and substantial reward. Unveiling social networks and automatically building personalised album are some such necessary rewards.

9 CONCLUSION

The contributions of this work are a method and an application for intertwining organization, visualization and indexation of social photos. Hierarchical visualization, interaction and an incremental Hasse Diagram-building algorithm are central to the entire process. The tests that we conducted showed that indexation productivity is improved when people use the Hasse Diagram they built through indexing. Navigation is also more efficient. The main result is that PhotoMap turns indexing into a pleasant activity at least for a limited number of photos and on a Hasse Diagram which is well known by the user because she created it. However speed and scalability require some improvements in the Hasse Diagram organisation.

In the future, if face recognition techniques become more efficient, our visual indexing method may be less interesting. However, for the time being, the majority of social photos cannot be indexed accurately by these existing techniques and our method which is based upon information visualization is a good supplement. It may also be reinforced by techniques such as face and clothes' colour detection to suggest groups. Beyond indexing there is still the problem of organizing social photo collections and navigating within them. This paper showed that a Hasse Diagram is a good tool if it is accompanied with helpful assistance such as zooming, local links and a 'spreader'. Moreover, the application opens up many interesting perspectives. In particular, we introduced the possibility of extracting and visualizing social networks as well as the ability to automatically build and display personalized albums. The key issue is to propose different displays that manage complexity and fluidity between the different applications.

REFERENCES

- [1] B. Ganter and R. Wille, "Formal Concept Analysis: Mathematical Foundations", Springer, 1999.
- [2] K. Rodden and K. Wood, "How do People Manage Their Digital Photographs?", Proceedings of ACM CHI 2003.
- [3] H. Kang, B.B. Bederson and B. Suh, "Capture, Annotate, Browse, Find, Share: Novel Interfaces for Personal Photo Management",




Fig. 12: Building social relations and personalized photo albums

Social networks. A very promising approach is to extract various social networks from the observation of who is with whom in

- International Journal of Human-Computer Interaction (IJHCI), volume 23, issue 3, pp. 315-337, Elsevier, Dec. 2007.
- [4] B. Suh and B.B. Bederson, "Semi-automatic photo annotation strategies using event based clustering and clothing based person recognition", *Interacting with Computers* Volume 19, Issue 4, July 2007.
 - [5] P.J. Phillips, P.J. Flynn., T. Scruggs, K.W. Bowyer and W. Worek, "Preliminary Face Recognition Grand Challenge Results", *Proceedings of 7th International Conference on Automatic Face and Gesture Recognition*, Pages 15-24, 2006.
 - [6] NSTC report, "Face Recognition", NSTC, Committee on Technology, Committee on Homeland and National Security, Subcommittee on Biometrics, 2006.
 - [7] W. Zhao, R. Chellappa, P. J. Philips and A. Rosenfeld, "Face recognition: A literature survey", *ACM Computing Surveys*, 35(4), 399-458, 2003.
 - [8] B. Shneiderman and H. Kang, "Direct Annotation: A Drag-and-Drop Strategy for Labeling Photos", *Proceedings of International Conference on Information Visualization (IV2000)*. London, England, 2000.
 - [9] B. Russel, A. Torralba, K. Murphy and W. Freeman, "LabelMe: A database and web-based tool for image annotation", MIT AI Lab Memo AIM-2005-025, September 2005.
 - [10] A. Kuchinsky, C. Perring, M.L. Creech, D. Freeze, B. Serra, and J. Gwizdka, "FotoFile: a consumer multimedia organization and retrieval system", *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 99)*. ACM Press, pp. 496--503, 1999.
 - [11] B.B. Bederson., "PhotoMesa: a zoomable image browser using quantum treemaps and bubblemaps", *Proceedings of ACM Conference on User Interface and Software Technology*, ACM Press, pp. 71-80, 2001.
 - [12] H. Kang and B. Shneiderman, "Exploring personal media: A spatial interface supporting user-defined semantic regions", *Journal of Visual Language and Computing*, 17, pp. 254--283, 2006.
 - [13] F. Monaghan and D. O'Sullivan, "Leveraging Ontologies, Context and Social Networks to Automate Photo Annotation", *2nd International Conference on Semantics and Digital Media Technologies*. p. 5-7. Genova, 2007.
 - [14] P. Eklund, J. Ducrou and T. Wilson, "An Intelligent User Interface for Browsing and Search MPEG-7 Images using Concept Lattices", *Proceedings of the 4th International Conference on Concept Lattices and Their Applications*, LNAI, Springer-Verlag, 2006.
 - [15] S. Ferré, "CAMELIS: Organizing and Browsing a Personal Photo Collection with a Logical Information System", *Proceedings of the 5th International Conference on Concept Lattices and Their Applications*, pp. 112-123, 2007.
 - [16] G.W. Furnas, "Generalized Fisheye Views", *Proceedings of Human Factors in Computing Systems (CHI'86)* Mantei, M., and Orbeton, P. (Eds), ACM, pp.16-23, 1986.
 - [17] R. Godin, G. Mineau and R. Missaoui, "Incremental structuring of knowledge bases", *Proceedings of the International Knowledge Retrieval, Use, and Storage for Efficiency Symposium (KRUSE'95)*, Santa Cruz, pp. 179-198, 1995.
 - [18] R. Godin and H. Mili, "Building and Maintaining Analysis-Level Class Hierarchies using Galois Lattices", *Proceedings of OOPSLA'93*, volume 28, pages 394-410. ACM Press, October 1993.
 - [19] P. Valtchev, R. Missaoui and P. Lebrun, "A Fast Algorithm for Building the Hasse Diagram of a Galois Lattice", *Proceedings of the Colloque LaCIM*, Montréal 2000.
 - [20] G. Arévalo, A. Berry, M. Huchard, G. Perrot and A. Sigayret, "Performances of Galois Sub-hierarchy-building Algorithms", *Proceedings of ICFA: pp 166-180*, 2007.
 - [21] K. Misue, P. Eades, W. Lai and K. Sugiyama, "Layout Adjustment and the Mental Map", *Journal of Visual Languages and Computing*, Vol. 6, pp. 183-210, 1995.
 - [22] G. Battista, P. Eadeaus, R. Tamassia and I. Tollis, "Graph drawing. Algorithms for the visualisation of graphs", Prentice Hall, 1999.
 - [23] P. Eades, "A heuristic for graph drawing", *Proceedings of the 13th Manitoba Conference on Numerical Mathematics and Computing*. *Utilitas Mathematica*, vol. 24, pp. 149-160, 1984.
 - [24] R. Freese, "Automated Lattice Drawing", *Lecture Notes in Artificial Intelligence*, 2961, Springer, Berlin, 112-127, 2004.
 - [25] T. Hannan and A. Pogel, "Spring-Based Lattice Drawing Highlighting Conceptual Similarity", *Proceedings of ICFA'06*, pp. 264-279, 2006.
 - [26] S. A. Golder "Measuring Social Networks with Digital Photograph Collections", *Proceedings of ACM Conference on Hypertext and Hypermedia*. Pittsburgh, Pennsylvania, 2008.