

Mining Social Networks and Their Visual Semantics from Social Photos

Michel Crampes
EMA-LGI2P

Parc Scientifique Georges Besse
30035 Nîmes Cedex, France

michel.crampes@mines-ales.fr

Michel Plantié
EMA-LGI2P

Parc Scientifique Georges Besse
30035 Nîmes Cedex, France

michel.plantie@mines-ales.fr

ABSTRACT

With the new possibilities in communication and information management, social networks and photos have received plenty of attention in the digital age. But there has been little research about the possibility that social photos carry some implicit social information and about the capacity to use this information for mining social data. In this paper, we show how social photos, captured during family or friends' events, representing individuals or groups, can be used to build social networks and express part of their semantics. Our contribution takes a wedding as an example. We present a method with different facets which give different results. The resulting social networks can be used by people as a mirror of the event and also as a means for sharing photos with personalized albums.

Categories and Subject Descriptors

D.2.2 [Design Tools and techniques]: User interface G.2.2 [Graph Theory]: Graph algorithms

General Terms

Graph theory, Social Networks, Hypergraphs, Galois Lattice, Formal concept analysis

Keywords

Social photos, Social Tribe.

1. INTRODUCTION

Considering the important role that digital photos seem to play during social events, it is interesting to explore their capacity to act as a new social medium that witnesses social relations. It is also pertinent to explore new applications that can be derived from them to strengthen socialization. Applications like Facebook or Flickr already manipulate photos within social networks. However, in these cases the networks are built using people and photos only act as passive objects that are voluntarily shared. They do not play an active role in building social networks. This paper explores a new paradigm whereby social photos are used to extract social networks and part of their semantics. The objective is to provide people with a mirror of social events and a means for sharing photos and other social relations after the events.

We then consider how to build personalized photo albums, which is known to be a hard problem. On the last ACM 2010 MultiMedia International Conference, one of the keynote invited

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.

WIMS'11, May 25–27, 2011, Sogndal, Norway.

Copyright 2011 ACM 978-1-4503-0148-0/11/05 ...\$10.00.

speakers qualified this problem as a major consideration which is still to be solved [10].

To consider photos as social actors it is necessary to interact with them as active objects. This requirement must be applied at several levels of the photo and social network life cycle. Each level is developed in the sections that follow. Section 3 describes techniques to organize social photos in a way that reveals social relations between people. In order to fulfill this objective, we use Formal Concept Analysis methods, and we extract social networks using three different forces. We introduce hypergraphs as a means to model social networks and we give a method to identify social tribes from social networks. We finally show how photo albums can be derived from social tribes. Sections 4 and 5 present a practical application from social network extraction to personalized album diffusion. The results are evaluated in section 6 and the whole contribution is discussed in section 7.

Note: Both co-authors have equally contributed to the preparation and the writing of this paper.

2. PHOTOS AND SOCIAL NETWORKS, STATE OF THE ART

Organizing social photos and eliciting social networks from their content has recently been a very active research topic, although we do not know of papers that integrate photo organization, social network extraction, and personalized album creation as we do in this paper. One close approach to our research is presented in [8]. The author justifies the fact that photos representing people can be used to elicit social networks. The evaluation method is well grounded and the results agree with ours, although the referenced paper does not extract and display social networks with semantic foundations. Another approach in [18] is very close to our work where the authors build social networks from photo collections by examining the co-appearance of individuals in photos. In this paper they use algorithms to maximize the modularity of the graph partition to detect social clusters. Our work uses new models of social links that better express the social relations between persons. In this context, traditional clustering and graph network measures like [6] are useful but do not express fully the semantics of the social relations between persons. In paper [2] authors use hypergraphs for music applications purposes to represent multiform datas from persons. Hypergraph as we will show are very useful to represent social data and especially social networks derived from photos.

The social network theory as depicted in [17] deals with network morphology and try to find "super-nodes" or referent elements in a network allowing frictionless information circulation along the network. Our goal is quite different since we want to form networks derived from the knowledge we may extract from photos as expressed by the people simultaneous

presence and their respective frequencies in a particular event. This knowledge is enough to infer social links between persons. Well known techniques of clustering are grounded on distances between elements. Our approach is different and is to be placed upstream compared to these techniques since we first want to define distances related to social links and then infer rules and properties from these links that keep the model's semantics. In [16] we have already sketched the extraction of tribes and the description of a photo diffusion policy. These issues are developed in this paper with more contribution on graph reduction, semantic justification, and an important evaluation methodology with results.

3. DEFINITIONS AND METHODS

In order to build social networks from a collection of social photos we must first ignore aesthetic considerations or other possible biases such as the fact that many photos are taken in a row with the same subject. To achieve this goal we organize all photos in groups showing unique subsets of people. All photos must be indexed beforehand with the names of the people. Regrouping the photos can be considered from the point of view of Formal Concept Analysis as constructing object concepts [7]. In a nutshell, let $\{P, X, F\}$ be a set with P being the set of photos, X the set of people who are seen in all photos and F a Boolean matrix with $f_{ij} = 1$ if $x_i \in X$ belongs to photo j and $f_{ij} = 0$ otherwise. $\{P, X, F\}$ is called a context where P are the objects and X the attributes. Formally it can be seen as a bipartite graph where F represents the edges between elements of P and elements of X . A maximal group of photos which contains a subgroup of people is a 'concept': a group of objects whose composition would change through adding or subtracting an attribute. The set of objects in a concept are called the extension of the concept and the set of attributes its intent. A partial order is applied to the context through an inclusion operator to build a hierarchy. An object may appear in several concepts when it shares subsets of attributes with other objects. Those concepts which contain objects at the lowest level of the hierarchy are object concepts [4]. The photo collection reorganized into groups of photos which contain unique groups of people is a set of object concepts.

Each concept shows a particular group of people (the intent of the concept) and contains a collection of photos (the extent of the concept) with this same group of people. From this organization we can derive social networks through the computation of distances between any pair of people. Different distance formulas are possible and we shall only introduce a few of them within the limits of this paper.

We define the following variables:

- $[.]$ is a set (also symbolized as a capital letter); X is the set $[x]$.
- $\text{Car}()$ is a cardinality of a set.
- x_i is an individual among all people in a set X ($x_j \in X$)
- C is the set of concepts c ($c \in C$)
- c is a concept, i.e. a set of photos, ($c \in C$) whose intention is the set of people it contains
- $c[x_i, x_j]$ is the set of concepts containing the couple $[x_i, x_j]$.
- $c[x_i \vee x_j]$ is the set of concepts containing at least x_i or x_j .
- $\text{Car}(c[x_i, x_j])$ is the cardinality of the concept containing $[x_i, x_j]$.

3.1 Extracting social graphs

In the following paragraphs as in the remaining of the paper we will use the term social force or more simply "force" to mean the strength of social relations between persons.

3.1.1 The simple force of a couple

We define the simple force of a couple (two individuals in some form of relationship) as the 'frequency' of this couple, i.e. the number of occurrences of this couple among the concepts divided by the number of concepts:

$$\text{Simple force}(x_i, x_j) = \frac{\text{Car}(c[x_i, x_j])}{\text{Car}(C)} \quad (1)$$

This interesting metrics represents how often a couple has been captured by photographers in different social situations. If the couple is seen in many situations then the relationship between the individuals is rather stable. However this distance does not express the strength of the liaison between two people as we do with the two following forces.

3.1.2 Proximity of a couple

In agreement with [9] we consider that if there are many people in a photo, the links between them are weak. Consequently to assess the strength of a couple, the number of persons present in photos with this pair should be taken into account. With this in mind, the author in [9] defines the strength of a pair of individuals. We see three pitfalls in this definition. First it is introduced only through common sense and it is not justified. Second, the value is not bounded (for instance with 1). Third, our main concern is that it is applied to all photos. As a consequence, if a large number of photos of the same people have been taken in a row, they will incorrectly strengthen the links between these people. A more evolved formula says that if a couple of persons often appear in small groups, logics says that this couple is closer than if it often appears in large groups of persons. We define the proximity for a couple as the sum of the inverse of the number of persons being members of concepts containing this couple, divided by the total number of concepts. Then, the more a couple is diluted in a big number of persons, the weaker the contribution of this concept to the proximity value.

$$\text{Proximity}(x_i, x_j) = \frac{\sum_{c \in c[x_i, x_j]} \frac{2}{\text{Car}(c[x_i, x_j])}}{\text{Car}(C)} \quad (2)$$

3.1.3 Cohesion of a couple

Intuitively we may understand that the above social model does not express some observations automatically extracted from the set of photos. Among them, one is of great interest. The more we see two persons separately, the more they are independent from each other. Conversely, when two persons are always seen together the hypothesis of particular links between them is strengthened. We call this type of force "cohesion". The couple $[x_i, x_j]$ has a strong cohesion when x_i and x_j are often seen together in the same concept with regard to the times they appear in a concept.

$$\text{Cohesion}(x_i, x_j) = \frac{\text{Car}(c[x_i, x_j])}{\text{Car}(c[x_i \vee x_j])} \quad (3)$$

3.2 From graphs to tribes and hypergraphs

3.2.1 A social network is a hypergraph

The previous step gives us three different matrices $B_{n,n}$ where n is the number of people that are present in all the photos. These matrices show three different relations among persons who were present at the particular event where the photos were taken. They can be considered as weighted graphs where each node is a person. Commonly authors consider that a social network is a graph where vertices are the persons and edges are the links between these persons. In our case we have three different social networks which are modeled as graphs.

However, if we analyze more thoroughly the situation, a social network is a hypergraph as [15] suggests. Hypergraphs are mathematical objects generalizing the graph concept, considering that edges may link more than two vertices up to any number of vertices. Formally a hypergraph H is a couple (V, E) where $V = v_1, v_2, \dots, v_n$ is a non-empty set (usually limited) and $E = E_1, E_2, \dots, E_m$ is a family of the non-empty parties of V [2]. The elements of V are the vertices of H . The elements of E are the hyperedges of H . Semantically hyperedges can be considered as groups of elements that share the same features.

Going back to the set of people which can be identified in the photos, we define a social tribe (or a tribe) as a subset of people that share common interest in some issue. In our case this issue is the different contents of the photos. Formally social tribes can be considered as hyperedges of a hypergraph which is to be built onto the set of people and which represents the social network as the grouping of people according to different centers of interest.

From the matrices above, we now must extract tribes. We must compute the hypergraph incidence matrix $H = T(G)$ where line i represents a tribe, column j represents a person and, $h_{ij} = 1$ if j is in tribe number i .

3.2.2 Introductory computations: density and weighted density of a subgraph

In a non valued graph (i.e. edges have no weight), graph theory defines graph density as the ratio between the number of edges and the number of edges of the corresponding complete graph (clique). The edge number of a complete graph is the maximum

number of edges in a graph; it is equal to $\frac{p \times (p-1)}{2}$, p being the number of vertices. So density is $\frac{2 \times nb_{edges}}{p \times (p-1)}$.

In a valued graph (i.e. edges are weighted), we define the weighted density of a graph as the ratio of the sum of weights of edges and the maximum number of edges of the corresponding complete graph (clique)

$$D_w = \frac{2 \times \sum_{i=1}^{nb_{edges}} edge_weight(i)}{(nb_{vertices} \times (nb_{vertices} - 1))} \quad (4)$$

3.2.3 Tribe mining

There may be many particular methods for extracting hyperedges (tribes) from a weighted graph. The choice of the method depends upon the goal of the construction and the underlying semantics of the weights. Since we want to build tribes with the aim of distributing photos to the people according to their interest into the persons that can be seen on the photos, the hyperedges must represent the common interest of people into other people. To do so for each social graph, we look for the densest subgraphs of the

corresponding weighted matrix. For a graph with no weights the problem is known to be NP-hard [8], and it is still complex with weights [14]. However in our particular case we can reduce the complexity through reducing the original graphs to a limited number of edges without losing much information. This is due to the fact that in social events if a person is loosely linked to anybody else, she is bound to be linked only to few people. Consequently, we reduce the graphs to a limited number of edges using a threshold for the weights according to the context as it will be explained in the experimentation below. Then all subgraphs are computed on the reduced graphs as well as their weighted density. We apply another threshold onto the weighted density to keep the subgraphs that are the most connected. A heuristics for choosing this threshold will be presented in section 5. A hyperedge is finally defined for each remaining subgraph.

3.3 Personalized albums strategies

3.3.1 Distribution principles

Our main goal is to distribute personalized albums to all the people who were present in the social event. The simplest extreme strategies are either send the photos only to the people who are visible on the photos (strategy 1), or send all photos to everybody (strategy 3). These two strategies are the most common strategies. They are implicit in web services like FaceBook according to the privacy choices.

However they can be criticized with respect to personalization. Strategy 1 limits the photos which are sent and someone may feel frustrated because she doesn't receive photos of her close friends or family members where she does not appear. For instance a grandfather would not receive photos of a married grandson and the bride in a wedding. Strategy 3 is too permissive and everybody must filter the photos according to their preferences. In between there are several possible strategies and one of them (strategy 2) is explored below using the tribes that were built with the above method.

3.3.2 Weighted photo-tribe distance

The core of the personalized diffusion strategy is to consider that all persons in a tribe are interested by the same photos. Which leads us to the following question: what are the photos that are interesting for a particular tribe? Several decision criteria are again possible. We propose hereafter to consider a particular distance between the photos and the tribes which will then be discussed. Conversely we can consider the proximity of the photos to the tribes.

We use the Jaccard distance (see [12]), that measures the dissimilarity between two sets (or conversely the Jaccard index which measures the proximity). For the sets A and B the Jaccard distance is:

$$J_d(A, B) = \frac{|A \cup B| - |A \cap B|}{|A \cup B|}$$

In our case, i.e. the distance between a photo and a tribe, the two sets are the persons in the photo and the persons in a tribe. The Jaccard distance is used here as a first possibility. We have conducted experiences with it. It is an interesting approach because it is well known, it is symmetric and its meaning is clear: the more the two sets are different, the more the distance. However it has some important biases which give us limited results as we shall see below. For instance the bigger is a tribe, the fewer photos it will receive, which is counter-intuitive. Conversely, the more people there are on a photo, the fewer

people will receive it. We will propose other distances with better results in future papers.

We then select the photos for a particular tribe which have a distance below a certain threshold (or conversely a proximity above a threshold). Tuning the distribution threshold opens up new questions such as: will it be the same value for all tribes or are there particular values for each tribe? What is (are) this (these) value(s)? It is obvious that there are no unique answer and that the answers depend upon the context. In the experimentation which will be presented hereafter we use a unique threshold of 0.5 which will be justified.

4. EXPERIMENTATION: MINING A WEDDING SOCIAL NETWORK AND ITS SEMANTICS

In this section we apply the three types of forces onto a corpus of social photos and analyze the results to open up semantic issues. Social photos are most of the time taken during family events, such as weddings, or parties between friends. A wedding ceremony is interesting because on one hand many photos are taken during the event and on the other hand the civil relations are known beforehand; they may be used as a semantic referent which describes the social network. In the following we describe how the forces can mine these social relations and how they can unveil other social relations representing alternative semantics.

4.1 The corpus and the resulting graph




Figure 1. The civil graph of the wedding.

Our experimental corpus contains 144 photos which have been taken during a particular wedding ceremony. In fact there were more photos but a previous selection was made to avoid too much redundancy because many photos had been taken in a row with the same subjects. Since the next step will be to extract social networks from these photos we need a referent graph which represents the known civil relations to validate the expected results on a topologic and on a semantic point of view.

To avoid a possible bias we delegated the task of building the referent graph to two post graduate students; they were asked to independently draw a graph of the civil relations in the wedding using information given by the married couple. We call these two hand-made graphs the 'civil graphs'. The goal of the two civil graphs is to identify a civil precise role for each person. The resulting two civil graphs were very similar with a central role given to the groom and bride couple. The nodes were obviously the same because both experimenters started with the same 28

people, the roles were defined nearly with the same words, and most of the hand-made links were similar in both graphs. We concluded that any of the two graphs could be used as a referent for the following steps and one of them was chosen. Figure 1 shows the chosen referent.

4.2 Extracting social networks

From the photo corpus we extracted 127 concepts.

Taking all the concepts as input, we computed the similarity values between all couples of people according to the three different social forces. Since there were 28 people involved in the wedding ceremony the application of social forces gives three 28 x 28 square matrices, one for each type of force.




Figure 2. Cohesion original graph with 149 liaisons.

We then computed a weighted graph for each social force. The graph nodes are the persons and the liaisons are their observed social distances which are computed from the matrices after the following normalization:

The strongest force (i.e. the smallest distance) is considered as the norm. For instance in the simple force the groom-bride liaison is the strongest with a weight equal to 0.50294. We take it as the reference value 1. All other forces are then transformed into:

$$\text{force}_{ij} = 1000(1 - \text{simple force}([x_i, x_j]) / \text{fmax})$$

where fmax is the reference force.

The dissimilarity between individual i and j is equal to 1 - force_{ij}.

These dissimilarities are not yet distances on a mathematical sense because the transitive law of mathematical distance may not be true at all times. A person X may not be closer to a person Y than the sum of closeness between X and Z, added to the closeness between Y to Z. However the symmetric law is verified, and implies that a social network graph derived from photos is undirected. Dissimilarity as they are stated here are able to correctly express closeness between nodes in a social graph and hypergraph.

The three resulting graphs show up 146 liaisons for the simple force, 138 liaisons for the proximity force and 149 liaisons for the cohesion force. All these computations were carried out by a team of three post graduates with precise requirements and the results were checked by the authors of this paper. Figure 2 shows a display of the result for the cohesion force. The figure was produced by our graph drawing software 'Molage' with which it is possible to interactively define display heuristics of different types to get the best visual result according to aesthetic criteria defined in the graph drawing community [1]. When looking at this figure it is obvious that the graph is complex and its topology hidden for interpretation: it is far from resembling to the one in Figure 1. Moreover when considering the three graphs built with

the three forces (the two other graphs are not shown in the paper) it is impossible to visually compare them. We need a good display that somehow looks like Figure 1 and at the same time shows other semantic features. The first thing to do is obviously to reduce the number of edges.

4.3 Network reduction

A better display may be obtained if we get rid of many edges and keep only those edges that represent the semantics of the network at its best. With this respect we need to find the ‘right number of edges’ and select the ‘good’ ones. The most obvious approach is to keep the smallest edges corresponding to the smallest distances since they represent the social forces of the couples. The number of edges to keep must meet the following constraints: not too many edges to avoid visual complexity, enough edges to still have a significant display, and the same number of edges between the three reduced graphs corresponding to the three social forces to be able to compare their semantics.




Figure 3. number of edges vs distance for the 3 social forces.

To find this value, we can compare the three graphs according to their number of edges versus threshold distance values shown in Figure 3. In this figure for each social force the number of accepted edges is given on the x-coordinates and the value of the distance normalized to 1 on the y-coordinates. When looking at this figure, the proximity force and the simple force behave very much the same as far as the number of edges is concerned. Few couples are captured before the value 900 of the distance. The cohesion force is much flatter and it captures more couples than the two other forces at smaller distances.




Figure 4. Reduced self displayed simple force graph (33 edges)

The choice of a significant number of edges is also partly given by the referent graph in Figure 1 because this is how the civil semantics was manually defined with 41 edges. Taking into

account these two points of view we chose 33 as the right number of edges for the three different networks and took the corresponding distance threshold value for each graph. Figures 4, 5 and 6 give the final displays for the three forces reduced to 33 edges.

4.4 Semantic analysis of the networks

The three original graphs with all liaisons could not be visually analyzed and compared because of the number of edges. The reduction process based on the number of liaisons in the referent network and an empirical choice of a distance threshold produces three graphs with smaller density that can be more easily analyzed and compared.

There are two first evidences: i) the simple force and the proximity graphs look very much the same and ii) they have a different but very similar topology (Figure 5 and 6). Roughly 40% of the people are disconnected and the remaining people, mostly the close family, gather around the bride and the groom like a starfish. Semantically, these two social forces extract networks that bring to light the deep meaning of the marriage: the union of two families through the union of two people.




Figure 5. Reduced self displayed proximity graph (33 edges)




Figure 6. Reduced self displayed cohesion graph (33 edges)

More precisely, the simple force considers the number of times a couple appears on different photos. This social force takes into account the number of object-concepts (groups of photos with the same people). Obviously the bride and the groom are the heroes: their nodes are very close and most of the liaisons start from them. Members of the close family are most of the time captured with them (links only with the bride and the groom) and sometimes together either with or without the bride and the groom. The other people are passive participants that are less often on photos and consequently show no liaisons (at least on the reduced graph which is not the case on the original graph).

With the proximity force the more a couple is isolated on a photo, the stronger it is. A couple lost among many other people is less prone to social relations. It is surprising that the resulting network looks very much the same as the simple force network. This is a result which was not expected. It suggests the idea that a strong couple is bound to be present on a photo on its own or with few people. It can be noticed that both forces take into account the number of concepts as their denominator.

The cohesion network is very different. This force does not take into account the number of object-concepts. A couple is strong if the two persons are present together on the photos and rarely separated. In the resulting reduced graph in Figure 6 only three people are disconnected. The graph is flat with long chains of nodes. The heroes of the wedding, the bride and the groom, are still there with the most important node degrees. They now clearly belong to two small cliques which are their respective close families. The other participants have fewer relations with other people and as a result become side actors of the ceremony. This network is the easiest to read. The big difference with the referent network is that in this social network the human relations are more important than the civil relations. It captures individual encounters and the participants may remember whom she or he talked to or was seated by. There are only three disconnected people and they look like they were lost. Actually they were the less implicated people in the wedding. It would be good to go back and reintroduce some edges from the original network to reintegrate them if the network had to be shown to the participants. The cohesion graph is a partial mirror of an important moment of social life for all the participants.

5. PHOTO DIFFUSION POLICY FOR THE WEDDING EVENT

In Section 4 a photo dissemination policy was presented. In this section this policy is applied to the ground truth example of the wedding. The method consists of analyzing the original graphs and extracting tribes. Photos that are the closest to each tribe after application of the Jaccard distance are prone to be sent to all the members of the corresponding tribe.

5.1 Tribe mining

To compute tribes according to the method presented in the above section we must compute all possible subgraphs and their density from the three original graphs, then choose a density threshold and select all subgraphs above this threshold. Applying this brute force method is very time consuming because there are $E = n(n-1)/2$ potential edges with n individuals and there are 2^E potential sub-graphs. The algorithm is NP-complete. However since density is directly linked to the weights on the edges, if we reduce the number of edges according to their weight we lose little information. Consequently we can compute all sub-graphs from

the reduced graphs presented above with 33 edges. The computation of all subgraphs is still time consuming with the possibility of generating 2^{33} sub-graphs in the most general case (although trivial sets such as singletons and the empty set are discarded). For our experimentation we tried a second reduction with only 13 connected people in the reduced graph on the basis of the following consideration: all the people who will be disconnected in the final graph (i.e. all final singletons) will be sent the photos on which they are visible according to the diffusion strategy 1. Consequently if they are loosely linked with other people they will not be part of a tribe but they are still bound to appear on the photos they will receive with the persons they are linked to. The 13 individuals who are selected for extracting tribes are the remaining connected nodes when decreasing the weights on the graphs. From these 13 people we identified the remaining subgraphs and computed their weighted density. To select the tribes we consider those subgraphs that have a weighted density above a particular threshold. This parameter is defined as follows. For each force we take tribes up to the point where all nodes are at least in one tribe: the threshold is consequently different for each force.

N°	tribe	density	weighted density	description
1	15-16-11-8	1	0,16	newlyweds, bride-parents
2	15-16-22-23	1	0,14	newlyweds, bride's witnesses
3	15-16-11-8-18	0,9	0,135	newlyweds, bride-parents, groom's father
4	15-16-11-8-17	0,9	0,132	newlyweds, bride-parents, groom's mother
5	15-16-11-8-18-17	0,93	0,095	newlyweds, bride-parents
6	15-16-11-8-18-17-13	0,714	0,094	newlyweds, bride-parents, bridesmaid
7	15-16-22-23-24	0,7	0,08	newlyweds, bride's witnesses
8	2-8-15-16-11-13-18-17-21-22-23-24	0,423	0,01	connected component

Table 1. Reduced : tribe densities from social simple force

Table 1 shows the generated tribes and their weighed densities for the simple force social network. The semantics of each tribe is given on the right. There are many more tribes than those presented in this table, but for illustration we have chosen to show in the table only subgraphs that have a high density. We also do not show the subgraphs of 2 and 3 persons.

Table 2. Reduced : tribe densities from cohesion social force

The last subgraph contains the whole connected component which means that we had to go to the whole reduced graph to find all people in a tribe. We may notice that the weighted density value is less and less important as the number of persons gets higher.

Table 2 shows the tribes for the cohesion force. A very interesting point is that the three subgraphs that are presented are not connected and all the people are affected to a tribe. The tribe organization witnesses the fact that the cohesion graph is flat.

Table 3 shows some of the tribes extracted from the proximity force. The tribe with the couple 9-7 is an isolated subgraph. One can see that although the simple force graph look very much like the proximity graph, these two forces give different tribes.

N°	tribe	density	weighted density	description
1	15-16-22-23	0,66	0,0473	newlyweds, bride's witnesses
2	15-16-11- 8-18-17	0,6	0,0368	newlyweds, newlyweds' parents,
3	15-16-11-8- 18-17-13	0,571	0,0346	newlyweds, newlyweds' parents, bridesmaid
4	09 07	1	0,03	cousin, bride's cousin
5	8-11-15-16 13-18-17 2-5-23-22	0,381	0,0176	newlyweds, newlyweds' parents, clergyman bride-witnesses bridesmaid

Table 3. Reduced : tribe densities from proximity social force

5.2 Photo diffusion policy and constitution of personalized photo albums

Three strategies of photo diffusion have been introduced in section 3. Strategy 1 consists of sending photos only to the people who are present on the photos. In the case of the wedding, a total number of 415 copies would be sent. Strategy 3 consists of sending all photos to everybody. 3429 copies would be sent, i.e. nearly 9 times more than strategy 1.

In between we are looking for strategy 2 who takes into account some personalization policy for building photo albums. To define which albums to build we compute for each force the distance between each photo and each tribe. We use for this matter the Jaccard distance which measures dissimilarity between two sets as stated in section 3. The two sets here are persons on one photo and persons of a tribe. We then sort in increasing order the photo-tribe couples according to this distance.

To decide which photos must be sent to the tribes we need threshold values as stated in section 3. There are not many criteria to objectively choose these values. In fact they should depend totally upon the context (the event, the people, the number, the quality of the photos and certainly other parameters). We think that the best idea would be to let everyone choose a kind of radius

N°	tribe	density	weighted density	description
1	23-24	1	0,444	bride-witness, girlfriend, groom's father
2	9-7-2-12	1	0,154	cousin, bride-cousine, groom's sister, bride's sister
3	15-16-11- 8 13-17-18	0,9	0,139	newlyweds, bride-parents, newlyweds' parents, bridesmaid

of interest. However this solution would not satisfy most of the people because they had rather receive photos without being bothered by any cumbersome decision factor. We must decide on the values for this experiment. We chose a unique Jaccard distance value of 0.5 for all the tribes and the forces. It means that a person is interested in a photo as far as there are at least a majority of 51% of common faces between a person's tribe and the people on the photo.

For the simple force social network the tribe policy yields a number of 324 distributed photos. For the cohesion social network the photo number is 177. Finally with the proximity social network the photo number is 301.

This strategy 2 applied on any social network gives less photo copies than strategy 1 (and obviously than strategy 3). This is a normal situation since tribes have been built onto only 13 people; they do not take into account all people of the wedding event. The personalized strategy generally allows to send a smaller photo quantity while increasing the relevance and interest of photos for each event participant as we show in the next section.

6. EVALUATION

6.1 Methodology

It is particularly difficult to define a methodology to assess the results of the whole process which takes a collection of social photos and yields different social networks and a personalized strategy for distributing social photos. The main difficulty is the absence of objective reference models. However we present in this section a first methodology which will be improved in future works. The process is assessed using the results of the experimentation as inputs. They are not statistically significative because there is only one experiment, but the number of photos and of the people gives some indication of soundness.

There are three results to assess: i) the three generated social graphs that result from the application of the three forces, ii) the tribes 'contents and iii) the photos which are finally distributed to the people. The unique referent we have is the civil graph of Figure 1. It will be used for the three assessments with different points of view. Finally we will apply different types of measurement according to the type of data to be evaluated.

6.2 Evaluation of the weighted social graphs

The first part of the process consists of building the three different weighted graphs representing social networks, one of which being shown in Figure 3. The question is: do these graphs still represent the original real civil network. With this respect we need to compare the topology (if not the semantics) of these graphs to the

referent graph in Figure 1. The problem is that there are 41 edges in the referent graph and nearly 150 edges in the other graphs. We applied a reduction down to 33 edges of the three graphs for the sake of aesthetic presentation. We need a closer structure to the referent graph. We looked for a better weight threshold for each graph to get the closer to 41 edges. We found 37 edges for simple force, 35 for proximity and 41 for cohesion. From this new reduction, to compare the three graphs to the referent graph we used two types of measures.

6.2.1 Degree assessment

In a graph the degree of a node is the number of other nodes that are linked to it. If we take the degree of each node in the graphs being assessed we can compare it with the degree of the corresponding node in the referent graph. For each test we have two series of values: the degrees of the nodes in the referent and the degrees of corresponding nodes in the social graphs. The Pearson's correlation coefficient indicates the correlation intensity between two variables. It is a measure between +1 and -1, with +1 representing perfect correlation and 0 representing independence. Applied to the three social graphs with respect to the civil graph, we get the three values shown in Table 4.

civil/simple	civil/proximity	civil/cohesion
0,75049635	0,75305661	0,67650283

Table 4. The Pearson test

The overall structures of the simple and the proximity social network do look like the civil graph although they are not similar. They also look alike. The cohesion structure is somehow different still keeping the original overall structure. These results which were visible in Figures 4 to 6 are confirmed here.

However this result does not say much about to which points the edges are preserved by the first the forces and second by the reduction process. The edges are the topology of the semantics.

6.2.2 Topology assessment

To compare the topology of the social graphs to the topology of the civil graph we use another method which makes use of precision and recall. These are two metrics that measure the performance of an information retrieval process. They will be used in several places during our evaluation. Precision is the percentage of true positives with respect to the number of all retrieved objects. Recall is the percentage of true positives with respect to the number of all expected objects. The F-score is the harmonic mean of these two values.

In our case we consider that in a social graph there are edges that are true positives if they were present in the civil graph and false positives if they were not. The precision is the percentage of civil edges among all the edges in the social graph. The recall is the percentage of civil edges among all the edges of the civil graph.

Applied to the three social graphs we get the values shown in Table 4. These results need to be corrected by the fact that there were less edges in the simple force and the proximity social graphs. However, even after correction, it appears that the F-score is rather low, particularly for the cohesion graph. It means that the forces create new liaisons between people and forget some civil liaisons. We conclude that far from constituting a drawback this deviation show that the forces see the people in the event with a different angle than the civil angle. It needs a more qualitative assessment which is presented in the following section.

	Simple force	proximity	cohesion
--	--------------	-----------	----------

recall	34,1%	36,5%	40,0%
precision	37,8%	38,6%	34,1%
f-measure	35,9%	37,5%	36,8%

Table 5. Social graphs versus civil graph

6.3 Evaluation of the tribes

The tribe extraction process that we presented is rather complex. The resulting tribes are derived from the constructed social graphs and another computation which extracts subgraphs and their weighted density. Do these tribes represent any sound semantics? When we look at the semantics which is indicated in the right columns in Tables 1 to 3, it appears that the tribes represent close family and inter-family relations at least for the selected 13 people selected from the social graphs. The inter-family relations are not visible in the civil graph with the exception of the groom-bride couple. We temporarily conclude that the tribe construction conserves close relations and unveils new social relations that one would expect from a wedding:

- 1) groom and the groom's close family
- 2) bride and the bride's close family
- 3) groom's friends
- 4) bride's friends
- 5) groom's and bride's enlarged families
- 6) groom's and bride's brothers' and sister's friends

The simple force and the proximity tribes reproduce family relations. The cohesion tribes are more representative of friends relations. These conclusions will be subject to deeper quantitative and statistic investigation in the future taking into account several biases such as the small number of people and the choice of the weighted density threshold.

6.4 Evaluation of photo diffusion

6.4.1 Diffusion methods relevance, and objective criteria

It is even more difficult to establish a referent to assess the different diffusion strategies. However when asking people what sort of personalized diffusion policy they would favor for family photos, two main preferences were given: i) each person should receive all photos in which she is; ii) a person should receive the photos in which she is (first policy) and, among the remaining photos which do not contain the person, she should receive the photos in which other closest persons are present. Applied to the particular case of the wedding photo collection, we have to specify who the closest persons are. We can find them in the civil graph: parent-child relation, brother-sister relation, close relation. Other relation types are considered more distant such as: family (other than previous ones), relations of family members, friends of close friend relations, etc.

These two referent policies can be applied to the wedding photo collection to give the photos which should be distributed to each person. It gives us the two following assessment criteria:

C1: according to the first policy 415 photos should be distributed.

C2: according to the second policy, 901 more photos should be distributed.

These two numbers will be used to assess our distribution strategies

6.4.2 Evaluation of three photo diffusion strategies

- 1) Strategy 1 (each photo is sent to the persons present on photos): precision, recall and f-measure are 100% for the C1 criteria. This is an expected result because this strategy 1 is the same as the policy that produced criteria C1. For the C2 criteria precision is 100%, recall is 46% and f-measure is 63,1%. Precision is perfect because according to strategy 1 only the photos in which the persons are present are sent to them, but recall is not so good because according to criteria C2 the people would expect more photos.
- 2) Strategy 2 (our solution with photo-tribe Jaccard distance) has been evaluated for each force against criteria C1 and C2. Criteria By computing photo-tribe distances we obtain different photo diffusions only because tribes are defined differently. Table 4 shows the results of recall, precision and f-measure for each force. The proximity force gives the best results and will increase the satisfaction degree of each person. However the cohesion force sends different photos to persons according to cohesion tribes. These tribes promote friend relationships more than family relationships. Finally the simple force is the simple cooccurrence solution and it gives quite good results.

	simple force	cohesion	proximity
recall	57,05%	50,28%	55,49%
précision	95,19%	97,21%	95,79%
fmeasure	71,34%	66,28%	70,27%

Table 4. statistics for 3 forces diffusion strategies

- 3) Strategy 3 (all photos to everybody): recall is 100% of course, precision is very low 12,1% and f-measure is bad 21,6% against the C1 criteria. For the C2 criteria recall is obviously 100%, precision is 26,28% and f-measure is bad 41,6%.

Looking at the results above it clearly appears that there is not a perfect diffusion strategy. However the tribe strategy is an interesting compromise between the two extreme strategies delivering at least

Diffusion evaluated strategies based on semantics like solution 2 sends out a limited number of photos compared to the global solution 3 but increases the degree of satisfaction for users. Photo diffusion usually is a real headache for persons in charge of diffusion, and not wishing to send all photos to everyone but rather the most interesting ones. This task finds here a real elegant and efficient solution.

7. DISCUSSION AND FUTURE WORK

Three main reduced social network graphs were produced during this work and they were interpreted form observing their topology and their display with respect to a referent graph. Their validity and their interpretation can be discussed since several biases may have resulted from the methodology. The referent graph was one of the two graphs manually produced by the experimenters after information from participants of the wedding and observation of the photos. The original information was factual: civil roles and people that were present on the photos. The topology and the display of the two graphs are very similar and little information is lost during this phase. Indexing was more challenging since there may be many people on some photos among which some of them may not be recognizable or may be far in the background. Should

they be identified and tagged? For simplification the experimenters indexed the persons that seemed to be the most engaged in the scene on this kind of photos. A more precise way would be to give a weight to the people on the photos, or even better to give semantics using RDF to describe the scene. A future work will be to find a methodology to exploit this semantics using also time to build the networks. For instance we expect that it would be possible to produce a narrative of the ceremony and see the social networks evolve according to what is happening. The use of FOAF (<http://xmlns.com/foaf/0.1/>) is considered for this purpose.

The production of the original networks from the photo corpus does not present any biases since we only apply algorithms to produce link matrices. The reduction process is more questionable. It was used to both analyze the semantics of the social networks and provide the participants with a mirror of the ceremony and of their roles. A lot of information is lost since more than 75% of the links are dropped. To avoid this loss a first solution is to let the user use a slider to ask for the right number of edges till she or he is satisfied with the display. This is not possible because as we have shown above the original display is not any more optimal when dropping edges (see Figure 2 and 4). A better solution is to start from a final reduced graph such as the one in Figure 5 and only locally show all edges when hovering with the mouse over a vertex. We are presently experimenting with this mixed strategy and the results are encouraging. It is briefly presented in section 4.4 because it is used for inviting participants to choose whom they would like to see in the pictures they will receive.

The original networks may be used to produce photo albums after extracting tribes. Since there are three original networks, which one is the best? From the analysis of the reduced networks, it seems that the cohesion network is more social than ceremonial. The reduced cohesion network pleads for the corresponding original parent to be chosen for tribe extraction. This may be true if the reduced network's social quality is also present in the original graph. This is an issue that needs to be investigated in the near future.

The display algorithms are questionable because it is difficult to balance aesthetic criteria with semantic criteria. The solutions we have explore till now give good results as can be seen in the different figures. But they require a combination of forces which needs to be controlled and some manual tuning. For a generalized usage a solution which does not require human intervention must be found.

The most important challenges are scalability on the one hand and semantics on the other hand. Scalability is required when there are many people in the ceremony, and when several ceremonies will be accessible on the web in the context of cross family social networks. The problem is not linked to technical complexity since there will never be thousands of persons, unless the principles we are developing here are used in other domains than human physical relations (gene interaction in biology for instance). The difficulty is more linked to usage and interpretation of results by users. The techniques developed in this paper propose solutions to handle this difficulty: graph reduction, expressive displays and good interaction for controlling complexity. The semantic challenge is certainly the most difficult but also the most exciting. It is discussed in the conclusion.

8. CONCLUSION

In this paper we presented a hypothesis (social photos can be used to extract social networks), different strategies for extracting social networks (simple force, proximity force and cohesion force), and methods and tools for graph reduction and display. An experimental process was presented to validate the method and to introduce semantic analysis of the extracted result. To conclude, two key questions can be asked. What are the extracted social networks and their graph representation useful for and for whom? Are they good representatives of the real social networks? The extracted social networks can be used for preparing personalized photo albums which do not contain too many photos but contain photos that are interesting for the receiver. It would be interesting to use the 'precision' and 'recall' criteria like in the Information Retrieval community to assess the relevance of personalized photo albums. However it is not easy since these criteria need the categorization of 'relevant' photos which is not a rigorous task in such subjects.

The answer to the first question brings us back to the photos from which the networks are derived. Photos are taken by people to keep in memory the event, the people and what they experienced during the event. If the extracted networks represent some semantics which is embedded in the photos, they also capture the persons' memory and experience. We limited our exploration of photo content only to the presence of people and the extraction of social relations on the presence or absence of couples of people. From this simple information two kinds of very different reduced networks were extracted: the wedding can be seen either as a focused civil ceremony with two young heroes, or as a flatter network of people engaged in physical relation. The methodology we proposed extracted two different views that meet what anybody experiences in this sort of event and as such we can pretend that it extracts part of the social semantics from social photos.

Many improvements are possible, and our models may be extended to take into account semantics when extracting the social network. Like the bias of the photographer, who may not be impartial or the photo manufacturer or the photo sender. Our future works will explore these models, with more profound testing. Particularly the scaling factor will be studied when using hundreds of photos or to reveal the social networks evolution in time and melting several social events. We will also use the principles exposed here to probe another field: building social networks from semantic analysis of shared documents between persons.

9. REFERENCES

- [1] G. Battista, P. Eadeaus, R. Tamassia, and I. Tollis. Graph drawing. algorithms for the visualisation of graphs. Prentice Hall, 1999.
- [2] Berge, C. Graphs and Hypergraphs. New York: Elsevier. 1973.
- [3] J. Bu, S. Tan, C. Chen, C. Wang, H. Wu, L. Zhang, and X. He. Music recommendation by unified hypergraph: Combining social media information and music content. Proceedings of the ACM international conference on Multimedia , Firenze, Italy, October 2010.
- [4] M. Crampes, S. Ranwez, J. Villerd, F. Velickovski, C. Mooney, A. Emery, and N. Mille. Concept maps for designing adaptive knowledge maps. *Information Visualization* , Palgrave, pages 5:211–224, 2006
- [5] P. Eades. A heuristic for graph drawing. *Congressus Numerantium*, pages 42:149–160, 1984.
- [6] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, pages 40/1:35–41, 2009.
- [7] B. Ganter and R. Wille. Formal concept analysis. *Mathematical Foundations*, Springer, 1999
- [8] Goldberg, A. V. Finding a maximum density subgraph. Tech. rep. 1984.
- [9] S. A. Golder. Measuring social networks with digital photograph collections. *Hypertext and Hypermedia ACM Conference*. June 19-21. Pittsburgh, Pennsylvania, 2008.
- [10] R. Jain. Life = experiences(events) and vision. *Proceedings of the ACM international conference on Multimedia* , Firenze, Italy, October 2010.
- [11] T. Kamada and S. Kawai. An algorithm for drawing general undirected graphs. *Information Processing Letters*, page 31(1):7(15), 1989.
- [12] J. L. Kaufman, P. Rousseeuw, and J. B. P. Finding groups in data : An introduction to cluster analysis. WILEYInterscience, 1990.
- [13] D. Koschützki, K. A. Lehmann, L. Peeters, S. Richter, D. Tenfelde-Podehl, and O. Zlotowski. Centrality indices. Brandes, U. and Erlebach, T. (Eds.) *Network Analysis: Methodological Foundations*, LNCS 3418, Springer-Verlag, pages 16–61, 2005.
- [14] Samir Khuller , Barna Saha, On Finding Dense Subgraphs, *Proceedings of the 36th International Colloquium on Automata, Languages and Programming: Part I*, July 05-12, 2009, Rhodes, Greece
- [15] P. Mika. Ontologies are us: A unified model of social networks and semantics. In *International Semantic Web Conference*, pages 522–536, 2005.
- [16] M. Plantié and M. Crampes. From photo networks to social networks, creation and use of a social network derived with photos. *ACM Multimedia International Conference*, Firenze, Italy, October, 2010.
- [17] L. J. Strahilevitz. A social networks theory of privacy. *American Law & Economics, Association Annual Meetings*, 2005.
- [18] P. Wu and D. Tretter. Close & closer: Social cluster and closeness from photo collections. *Proceedings of the seventeen ACM international conference on Multimedia*, 2009.