Measuring Social Networks with Digital Photograph Collections

Scott A. Golder
Hewlett-Packard Labs
1501 Page Mill Rd., ms 1139
Palo Alto, CA 94304
+1 (650) 857-2020
scott.golder@hp.com

ABSTRACT
The ease and lack of cost associated with taking digital photographs have allowed people to amass large personal photograph collections. These collections contain valuable information about their owners' social relationships. This paper is a preliminary investigation into how digital photo collections can provide useful data for the study of social networks. Results from an analysis of 23 subjects' photo collections demonstrate the feasibility of this approach. The relationship between perceived closeness and network position, as well as future questions, are also discussed.

Categories and Subject Descriptors
H.5.4 [**Information Interfaces and Presentation**]: Hypertext/Hypermedia – *user issues.*

General Terms
Design, Human Factors.

Keywords
Social Networks, Photographs.

1. INTRODUCTION
Digital cameras have nearly replaced film cameras for personal photography. Camera owners can now inexpensively create large digital photograph collections, which can be immensely rewarding as keepsakes recalling happy times with friends and loved ones. For many, the "shoebox" metaphor – the undifferentiated, disorganized collection of photographs – persists even once photography transitions from film and print to digital media. However, digital photography presents many opportunities for organizing and analyzing photographs via auto-generated and user-created metadata. For example, such metadata can support filtering, searching and sorting the photographs. This metadata also constitutes detailed information about the social lives of the owners of the photographs and the people depicted in them.

This paper focuses on the identity of the people in the photograph as one of the most valuable pieces of metadata a digital photograph can contain. Previous research on the social meaning of and behavior around digital photographs supports this, showing that digital photography is a part of daily life through which people fulfill a variety of needs, from immediately sharing and communicating via photos, to taking pictures as a way of preserving memories [10,2]. When interventions with novel interfaces allowed users to annotate their photographs, users used the photos and annotations as a way of enjoying and sharing memories through storytelling [12,3]. When photos are annotated, they are often annotated with the names of who is in them [12], and when telling stories about photos, stories often start with the names of who is involved, e.g., "This is [*person*] and [*person*] doing [*something*]" [3].

This name metadata can be useful to researchers, as well. Measuring social networks, or the set of relationships demonstrating how a group of people are socially connected, is a popular method of analysis in sociology and has applications across a variety of disciplines. For example, network measurement reveals how people make friends [4] or find jobs [8]. Unfortunately, data collection, while crucial, is difficult; this is discussed in more detail in the next section. On its face, however, two or more people taking a photograph together generally implies that the people know one another and share some social connection. Therefore, collecting this data from large photo collections may be an accurate, naturalistic way of measuring a social network.

The identity of the people captured in photographs is starting to receive a great deal of attention. The photo tagging service *Flickr.com* allows all photos to be "tagged," or given descriptive keywords, and some use this feature as a way of storing names. Popular online social network service *Facebook.com* supports linking photographs to the web profiles of the people shown in them. The method proposed here employs a network-centric view of these names, taking into account not only who is in which photo, but whom people are depicted in photos with.

The following section provides a brief overview of research in social network analysis. Next the experiment, in which social network data was collected from subjects' photo collections, and the software application used to collect the data will be presented. The characteristics of the resulting data are discussed, demonstrating the validity of using photographs as a network elicitation technique.

2. SOCIAL NETWORKS ELICITATION

A social network is the collection of relationships that demonstrates how a group of people are socially connected to one another. An individual's social network consists of friends, family, and other people one relies on for social and material support of many kinds.

As discussed in the introduction, measuring social networks is done to study a variety of social processes, like friend and job finding. A persistent problem in social networks research has been that collecting a reliable representation of a social network is challenging [5,6]. Clearly, it is impossible to ask subjects to simply list everyone they know [6].

Social networks can be measured in a variety of ways, from performing experiments, to asking individuals to self-report their relationships, to measuring the records of individuals' communication.

One early and quite famous experiment is Milgram's tracking of letters through the postal mail system as a way of identifying links between people [11]. It is the progenitor, in some ways, of many modern studies involving communication patterns. While groundbreaking, such a study necessarily measured only a relatively small sample of individuals' networks.

Self-report is a widely-used method for eliciting network data from subjects. Question prompts, such as "who would you go to, to discuss a personal problem?" allow subjects to generate lists of names [5] through a naturalistic thought process, but may be incomplete and subject to memory failures.

Diary studies – in which subjects are asked to keep a log of whom they spoke to that day, for example – produce quite useful data and, since they are relatively contemporaneous with the communication, lose less to memory, but are subject to biases, such as understatement, overstatement or systematic failures [9].

The use of electronic communication tools, from which usage logs can be obtained, has opened up several opportunities in network measurement, including email archive analysis [13,14], but also PBX telephone network records [9]. In both cases, pairs of individuals are identified as being linked to one another on the basis of some interaction. These include a sender and a recipient, in the case of email, or caller and callee, in the case of the telphone. These techniques are passive; the data collection is a by-product of existing social processes. This means that collecting the data is not burdensome to subjects, nor is incompleteness problematic (at least, not in the domain of inquiry).

In a digital photograph collection, depiction in the same photo is considered here to imply a link between the individuals depicted. As a photo collection grows, all the co-depictions form a network of people and links that is a reasonably accurate representation of the network of the collection owner. It is important to note that what is being created is a representation of a social network, not the network itself; the network itself manifests itself through the relationships of its constituent members, and is ever-changing and largely indeterminate. Indeed, a representation of a social network is valuable only to the extent that it closely reflects the composition of the actual network.

As discussed above, photography is an existing social practice. Therefore, "collecting the data" -- i.e. taking the photos -- is not additionally burdensome. It shares this property with email, PBXes, and so on. Additionally, while one can be an unintentional recipient of a phone call or email, with new contacts thrusting themselves into one's network, taking photos with one's camera is a deliberate process, and one can rather safely assume that, in most cases, whomever the photographer captured, the act of photographing them is intentional.

To some extent, annotating photographs can be enjoyable when it is part of an existing social process done in service of some personal benefit, like using those annotations for enjoying the annotated collection. Annotating photos with names is akin to a social network name generator [6] in which photographs are shown and subjects need only recognize the individuals, a comparatively easy task, compared to remembering lists of people as in other network elicitation tasks.

To be sure, using personal photo collections to measure social networks presents some complications, mostly related to the collection owner's place in the network. For example, if the owner is capturing the photos, then he or she is not in the photo, so co-depiction may be less appropriate for relationships between collection owners and alters (individuals to whom one is connected), than between pairs of alters. In contrast, the collection owner has nearly complete control over whom photos are taken of, so appearance in multiple photographs may be a good measure of a person's relationship with the collection owner.

Taking photographs requires geographic proximity. For distant friends and relatives, taking photographs on a regular basis is impossible. However, given the increased likelihood of taking photos on holidays and at other special events in which close friends and family are likely to come together, this consideration is somewhat tempered. Finally, a certain *kind* of social network data is being collected here. Because of the recreational nature of photo taking, the method proposed here would, for many people, capture friends and family more than business contacts. However, even this is contextual – photos from a corporate holiday party might contain information about the informal workplace connections that enable work to be accomplished. Regardless of the domain, however, posing in a photograph with others is a deliberate act, and is generally indicative of some social connection. Identifying the context of that connection remains a task for those studying the network in question.

3. METHOD

The intuition behind this network construction method is that, when two people are depicted in a photograph together, this implies that they share some social connection. A basic way of constructing a network in this way is shown in figure 1. As photos 1-5 are added to the collection, the network grows to include each relationship. For example, photo 1 adds persons A and B to the network, and shows they are connected. In photo 2, person C is shown with both person A and person B, so one link from C to each of those people is established. Next, in photo 3, D and C are shown together, and so their link appears in the network. This process continues until all co-depictions are represented in the network and is demonstrated by the graph – the collection of nodes and edges or links – shown in figure 1.

This representation becomes complicated by acknowledging that, when the same two individuals are in many photographs together, this implies an even stronger link between the two. As such, a weighted graph would be one in which each link between two nodes is characterized by a value, or weight. In figure 1, for

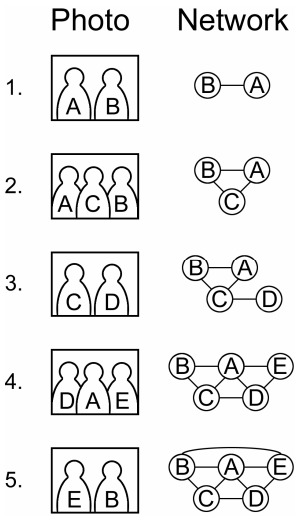

Figure 1. As photos 1-5 are added to a collection (left), the network representing the connections (right) grows.

example, the link between A and B would be 2, since those two individuals are in two photographs (1 and 2) together. In this way, the higher the link weight, the stronger the bond represented between the two individuals.

The network just described – a weighted network based on frequency of co-occurrence – is not a novel suggestion, to say the least. However, consider an extreme case: a classroom photo depicting 35 students. To say that this photo represents as strong a link between any pair of individuals as a photograph of only the two of them together would be absurd. Recognizing this fact, the link strength implied by any given photograph is decreased by the square root of the number of people in the photograph:

$$STR(P_a, P_b) = \sum_{i=1}^{m} \frac{1}{\sqrt{n_i - 1}}$$

where *STR()* is the strength of the social connection between two persons, m is the number of photos in the collection, and n_i is the number of people depicted in the ith photo.

Illustrating with an example, photo 1 in figure 1 would imply an increment in the link strength between A and B of 1 / √(2-1), or, simply, 1. In contrast, photo 2, which contains three people, would increase each of those links by 1 / √(3-1), or 0.707. This would give A and B an overall link strength of 1.707.

In summary, two principles dictate how network edge weights are computed. First, the more photos two individuals are depicted in, the stronger their social connection likely is. Second, the number of persons in a photograph necessarily affects the degree to which that photograph implies relationship strength.

3.1 Data Collection

23 subjects were recruited to identify all the people in a sample of their personal digital photo collections. These subjects are employed in the technology industry. While the subject pool is somewhat homogenous in their technical proficiency, they are geographically diverse. Nine subjects were in the United States, 5 in the United Kingdom, 6 in India, and 2 elsewhere.

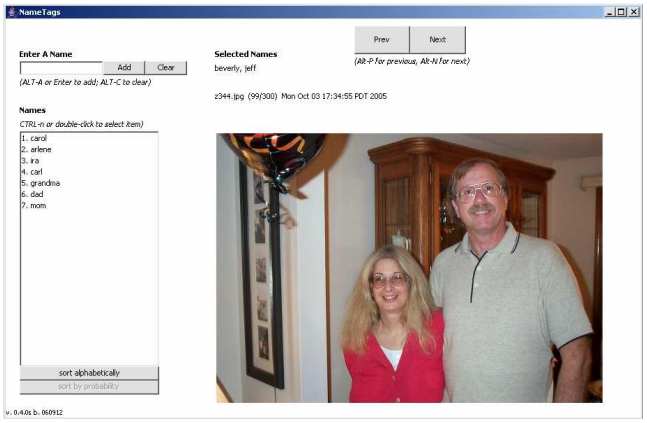

Figure 2. A screenshot of the NameTags photo browsing and tagging application, used for collecting data from the subjects.

The subjects were emailed a java application, which they were instructed to run on their home computer containing their personal photograph collection. The application, called NameTags (figure 2), selected a random sample of 300 images from the photo collection, and showed the photos to the subject one at a time. The subjects were instructed to identify by name the individuals depicted in each photo.

As they proceeded, a list of names appeared in the left column, so that they could choose names by mouse rather than typing them in each time. This list of names could be sorted either alphabetically or by the strength of the relationship between each person and the group of people already selected as being in the picture. This was intended to not only make the process easier and faster, but to also decrease the likelihood of multiple names being used by the same person. For example, "John Smith" might be called "john" at one point, and "johns" later.

After identifying the names of the people in each photograph, the subjects were asked to double-check that all people were identified uniquely, and allowed to use a name-merging interface to fix errors. The problem of multiple strings being used to identify the same item arises when users can freely index collections of items [7]. Other research on annotating photographs has recognized this problem [12], but allowing users to see their past choices and to merge duplicates has been shown to improve the quality of the resulting data [14].

Once the collection of names was finalized, users were asked to self-report how close they felt to each person named, using a 1-5 scale from "not at all close" to "extremely close" (figure 3). Closeness is a problematic concept; it is vague, and so people define it in a variety of ways, across many dimensions, depending on innumerable circumstances. Ferligoj and Hlebec [5] indirectly asked groups of students about closeness, asking which of their classmates they could turn to for academic resources and companionship as specific measures of more general social support. As vague as closeness is, it is nevertheless a rough measure akin to social support. Further, they found that response scale played a large role in data reliability, and that a five-point scale was the most reliable. Accordingly, our subjects were asked to categorize the names they listed from 1-5. They were not allowed to categorize themselves (n=23) but were allowed to not categorize any name they felt they couldn't decide on (n=37). In

Figure 3. Interface for categorizing closeness between the subject and the people in photographs by dragging names from the blue list at left to lists 1-5 (yellow).

the end, 972 of 1032 names were categorized (94%), and are the only ones used in any subsequent examination of the closeness variable.

4. NETWORK ANALYSIS

Across 23 users, 3852 photos out of 6848 photos were tagged with the names of people; recall that each subject had ~ 300 photos randomly selected, and note that only 56% of photos contained people. This is mostly consistent with Ames' [2] study of camera phones in which 49% of photographs contained people.

The total size of the individuals' collections ranged from 368 to over 15,000, with a mean of 3600 and a median of 2537. These numbers are likely minima; though subjects were asked to direct the application to their entire collection, some subjects informally reported having collections spread across multiple computers or CDs, and using the entire collection was therefore infeasible for them. In the set of photographs, a total of 1032 unique people were identified, including the 23 subjects (mean=44.87, median=33).

Contrary to expectations, most photographs contained a small number of people. Of the photos that did contain people, a surprising 1891 photos (49%) contained only one person, and 1263 (33%) contained two people. Very quickly, that count decreases; only 4.3% of photos contain 5 or more people. Similarly, only a few people appeared in a very large number of photographs; when subjects' collection size was normalized to 100, the distribution of pictures/person is a familiar power law distribution – the number of people decreases exponentially as the number of depictions increases [1].

Using photographs as a measure of position in the photographer's network is useful only to the extent that the photographer's feeling of connection to the individual is correlated with the measures the network model provides. First, the special case of being co-depicted in a photograph with the photo collection owner is considered. Next, the more general relationship between individuals' subjective closeness to the archive owners, and their network characteristics is examined.

4.1 Co-Depiction with the Archive Owner

As mentioned above, it takes special effort for a person to be in their own photo collection; being codepicted with an alter in one's own collection is a special marker of relationship strength. Indeed, co-depicted individuals score higher than non-codepicted individuals on several measures. Table 1 shows that alters co-depicted with photo archive owners were in more photos, were rated as subjectively closer, and were connected (i.e. co-depicted) with others who were themselves subjectively closer to the archive owner.

However, no significant difference was observed between co-depicted and non-co-depicted alters when considering network-centric characteristics such as degree (i.e. number of alters) or betweenness (a measure of "centrality" or importance) [15].

A possible explanation why this might be the case, is that the extra effort for camera owners to be in photos in their own collection is expended only for people closest to them. The fact that the co-depicted people are in many photos but are not necessarily connected to more people suggests that they constitute small, tightly-knit groups around the photographer. Indeed, a co-depicted person's alters are more likely to be co-depicted than not (57%), whereas this is not the case for a non-co-depicted person.

	Co-Depicted (n=235)	**Not Co-Dep.** (n=737)	**p**
Closeness to owner	3.68	2.77	< .001
# Photos overall	17.38	3.58	< .001
Alter Closeness	2.78	2.04	< .001
P(alter co-depict)	.57	.28	< .001
Degree	6.50	7.05	not significant
Betweenness	38.50	32.01	not significant

Table 1. Results of t-tests comparing co-depicted and non-co-depicted alters across a variety of dimensions.

4.2 Closeness and Network Position

Generally, individuals' closeness to archive owners is positively correlated with the characteristics discussed above, as shown in the first column of table 2. One might expect some of these things to be true; closer people are intuitively more likely to be in a larger number of photographs. Being close to someone and being in many photos affords the opportunity to be in photos with many different people; indeed, betweenness (i.e. centrality) is very strongly correlated with number of photos, and hence with degree (i.e. number of alters). If being close to the photographer is related to being in photos with a large diversity of others, then this may also explain why the relationship between closeness and alters' closeness is not particularly strong.

The final row in Table 2, clustering coefficient, is negatively correlated with the other factors, including closeness, and especially betweenness. Definitionally, a node that has a high clustering coefficient would have low betweenness, since there are few network paths that would rely on it. However, the clustering coefficient's negative correlation with closeness and number of photos supports the explanation that a close individual consitutes the basis for the inclusion of many others, and is in many cases

	Closeness	# Photos	Avg. Alter Closeness	Betweenness	Degree
Closeness					
# Photos	.467				
Avg. Alter Closeness	.259	.324			
Betweenness	.394	.705	.256		
Degree	.250	.458	.443	.600	
Clustering Coeff.	-.317	-.410	(not significant)	-.682	(not significant)

Table 2. Spearman correlations of several characteristics. Reported values are significant at p <= .01 (2 sided).

others' only link into the network. This is a good thing; it is exactly what we might expect in a social network.

4.3 Edge Weights

Until this point, the discussion has been limited to the egocentric aspects of the network; that is, only those factors that directly involved the photo archive owner. The edge weighting method discussed in section 3, by contrast, affected all nodes, or people in the network.

As mentioned above, fewer photos contained large numbers of people than initially expected. This condition limits the practical significance of dividing the incremental edge weight by the square root of the number of people. However, in specific scenarios where large group photographs are likely, such a step might prevent these photos from adding incorrect data to the network. Having more accurately weighted edges is of more significance when performing tasks like clustering (e.g. [16]). Graph clustering can be enlightening by identifying subgroups in the network, but would be most useful when the data sets are larger than the ones used in this analysis.

5. CONCLUSION

This paper represents a preliminary investigation in measuring social networks through the natural social activity of being depicted in photographs. That network-theoretic properties of the photo networks (e.g. betweenness centrality, degree) are correlated with subjects' self-reported relationship closeness suggests confidence that measuring a social network through photos computationally can be a useful, scalable proxy for the time-intensive process of collecting data from subjects.

The usefulness of the method is compounded when the photos are already tagged by users, or face recognition tags them automatically. This is increasingly possible as photo management and sharing applications such as Flickr and Facebook promote tagging sets of photos as they are uploaded, and as face recognition techniques improve.

Future work in this area includes examining larger sets of data, as well as the applicability of other network properties. This analysis demonstrated that higher closeness is correlated with higher betweenness, or centrality, with being in more photos, and therefore being connected to more people; future analysis will be extended to include specific relationship types shared with alters.

6. REFERENCES

1. Adamic, Lada A. and Bernardo A. Huberman. (2002). Zipf's Law and the Internet. *Glottometrics*. 3. 143-150.
2. Ames, Morgan. (2006). *The Social Life of Snapshots*. Thesis. UC Berkeley School of Information Mgt. and Systems.
3. Balabanovic, Marko, Lonny L. Chu and Gregory J. Wolff. (2000). Storytelling with Digital Photographs. *Proc. CHI 2000.*
4. Feld, Scott L. (1981). The Focused Organization of Social Ties. *American Journal of Sociology*. 86(5). 1015-1035.
5. Ferligoj, Anuska and Valentina Hlebec. (1999). Evaluation of Social Network Measurement Instruments. *Social Networks*. 21. 111-130.
6. Fu, Yang-chih. (2005). Measuring Personal Networks with Daily Contacts: A Single-Item Survey Question and the Contact Diary. *Social Networks*. 27. 169-186.
7. Golder, Scott A. and Bernardo A. Huberman. (2006). Usage Patterns of Collaborative Tagging Systems. *Journal of Information Science*. 32(2) 198-208.
8. Granovetter, Mark S. (1974). *Getting a Job: A Study of Contacts and Careers*. Cambridge, MA: Harvard U. P.
9. Higgins, Christopher A., Ronald J. McClean and David W. Conrath. (1985). The Accuracy and Biases of Diary Communication Data. *Social Networks*. 7. 173-187.
10. Kindberg, Tim, Mirijana Spasojevic, Rowanne Fleck and Abigail Sellen. (2005). An In-Depth Study of Camera Phone Use. *IEEE Pervasive Computing*. April-June 2005. 42-50.
11. Milgram, Stanley. (1967). The Small World Problem. *Psychology Today*. Vol. 1, issue 61.
12. Rodden, Kerry, and Kenneth R. Wood. (2003). How do People Manage their Digital Photographs? *Proc. CHI 2003.*
13. Tyler, Josh, Dennis Wilkinson & Bernardo A. Huberman. (2003). Email as Spectroscopy: Automated Discovery of Community Structure within Organizations. *Communities and Technologies 2003*.
14. Viégas, Fernanda B., Scott Golder & Judith S. Donath. (2006). Visualizing Email Content: Portraying Relationships from Conversational Histories. *Proc. CHI 2006*.
15. Wasserman, Stanley and Katherine Faust. (1994). *Social Network Analysis: Methods and Applications*. Cambridge.
16. Wu, Fang and Bernardo A. Huberman. (2004). "Finding Communities in Linear Time: A Physics Approach." *Eur. Phys. J.* 38(2). Springer.