This is a preprint of an article published in Library & Information Science Research: Stvilia, B., Jörgensen, C. (2009). User-generated collection level metadata in an online photo-sharing system. *Library & Information Science Research. 31*(1), 54-65.

User-Generated Collection Level Metadata in an Online Photo-sharing System

Besiki Stvilia and Corinne Jörgensen
College of Information, Florida State University
101 Louis Shores Building, Tallahassee, FL 32306, USA
{bstvilia, cjorgensen}@fsu.edu

Abstract

Photoset and Group descriptions in Flickr, a large-scale online photo-sharing system, offer insight into the collection description and collection building practices of Flickr users. *Photosets*, assembled by individual users, appear to evolve from bottom-up, derived from the components of the individual users' context to evolve from the bottom up, and are based on selected attributes which a particular user's photos share. *Group* collections, on the other hand, seem to be organized more around general concepts or discussions relevant to the Group members' work and constructed top down by matching specific photo attributes with the purpose of the Group. This article identifies 10 categories of characteristics that Flickr users might use for forming these digital photo collections and discusses differences observed between Photoset and Group collection describing and building behavior. The categories are then compared with the classes and elements of some current metadata schemas and an ontology and the results of earlier research on individual behavior in describing individual items. The study shows that systematic investigation of user-generated collection level metadata in Flickr and other similar open-tagging sites is needed to help inform better design of collection metadata schemas and other information organization tools.

Introduction

Collections have long been an important approach to information organization in libraries and other information institutions such as archives and museums. However, until recently collections were designed and developed mostly by professionals (librarians, archivists, or museum curators) for use by both professionals and end-users. Descriptions of these collections, in the form of collection-level metadata, have been used to provide access and were created within a set of constraints derived from, in most cases, the local organizational context. With social networking sites now providing information organization capabilities to end-users as well as professionals, the impact of dual – or multiple – contexts on collection building, information seeking and use needs to be further examined.

Different techniques (e.g., questionnaires, interviews, focus groups, sorting tasks) have been used to collect information about user needs in information organization for individual items and collections. Still, few empirical data are available on the needs for collection-level metadata for different kinds of information resources, particularly for visual media. New large-scale photo-sharing systems such as Flickr (http://www.flickr.com), in which the users themselves author both collections and metadata, can provide unprecedented access to empirical data on types of user-generated photo collections and metadata, the activities in which the photo collections are used, and, by implication, end-user needs for subject and other types of access.

Flickr, a large-scale, popular photo-sharing and archiving system, is owned by Yahoo! (http://www.yahoo.com). The following figures demonstrate its phenomenal growth: the size of Flickr's registered user base in June 2005 was 775,000, and it hosted 19.5 million photos, with a monthly growth rate of approximately 30% (Kuchinskas, 2005). Furthermore, more than 80% of these photos were public (Gyford, 2005). Two years later the numbers of registered users had risen to an estimated 5 million and photos at 250 million (Ames & Naaman, 2007). The Flickr website itself does not give any information on the current number of photos or the number of users. However, Flickr does provide open access to the photo metadata and discussion logs of its public collections, along with a limited application programming interface (API), which enables researchers to access various kinds of data.

Many different facets and relationship can be used for organizing collections. To make information resources from traditional databases more accessible and visible as well as to reuse user created collection metadata one needs to learn *what* concepts and relationships users organize their collection by. Similarly, to influence the quality of user-provided collection metadata, one needs to understand *how* and *why* users create collections. Studying the photo collection-building practices of users of Flickr can help both LIS researchers and professionals better understand users' needs in using these collections and better align traditional information services and tools with those needs. It may also provide ways of encouraging these users to contribute better quality, standardized metadata and enable effective and efficient integration and access to these collections by the public.

Literature Review

Traditional collection definitions found in the literature usually include an identification of the basis and motivation for organizing items into collections, collection metadata functionalities – tacit or explicit – resulting from organization, or both. According to Hill, Janée, Dolin, Frew, and Larsgaard (1999), a digital library collection can be defined as a bag of objects put together based on some common characteristic, such as topic coverage, format, geographic coverage, temporal coverage, pertinence to a particular study, source of origin, physical location, or source of funding support. The authors list the following functionalities for collection-level metadata: (1) registering the collection with search and retrieval software, (2) providing information about the collection to network search agents for network discovery, and (3) managing collections. Buckland, when comparing paper-based collections with electronic ones, suggested four roles or functionalities for digital library collections: (1) preservation, (2) dispensing, (3) bibliographic, and (4) symbolic (Buckland, 1992).

Digital collections are often equated with digital libraries. Lesk (1997) defined digital libraries as "organized collections of digital information" which combine the traditional information organization practices and structures with the new forms of information representation and organization enabled by computer technologies. Indeed, with the increasing trend of moving ownership of collections away from libraries to digital content providers and aggregators, an increasing number of collection definitions now reference the technologies with which the collections are bundled (Covi & Cragin, 2004; Lee, 2000).

Digital collections are not restricted in space, and often their content, technologies, policies, and costs are determined externally by content providers or aggregators. This makes digital collections

more transient than traditional library collections, and they may require more complex and frequently updated metadata. Hill et al. (1999) identified two classes of metadata for modeling digital library collections: (1) inherent metadata, obtained through content analysis of the collection; and (2) content metadata, which cannot be obtained from the content of the collection and which are assigned by the collection provider. This classification of collection metadata is in agreement with traditional library cataloging practices, in which metadata are either extracted from or assigned to a resource (Taylor, 2006). In order to provide functional collection-level metadata, some sort of standardization has long been assumed to be an essential requirement.

The transitivity of distributed digital content and the confluence of different cultures (libraries, museums, archives, information technology specialists) in the digital environment may lead to uncertainty and different interpretations of the term "collection" (Palmer, Knutson, Twidale, & Zavalina, 2006); adding commercial content providers and aggregators to the mix increases the likelihood of these problems. Furthermore, practices of what and how much metadata need to be included in a collection record may vary for different types of information objects and metadata providers (Shreeves et al., 2005; Stvilia et al., 2004). Indeed, one may expect the requirements and sources of metadata for various forms of media to differ from those of textual resources based on their typified uses and attributes. For instance, in the realm of visual media, images may not need metadata indicating language unless they have textual annotations.

Visual media, because of their unique characteristics, have stimulated several different approaches. For instance, Rasmussen (1997) defines two approaches to indexing images based on the level of indexing and the technologies used: content-based and concept-based (Rasmussen, 1997). Content based indexing, a computer method for algorithmically parsing images, can be done automatically and can be an inexpensive method for assigning certain kinds of metadata to large numbers of images. However, content-based indexing and retrieval systems can successfully 'recognize' only pixel level attributes (e.g. color, shape, texture). They are not very successful in translating low-level pixel-derived data into higher-level content metadata users typically describe and search images with. In concept-based indexing, primarily a manual approach, human indexers assign terms to images representing higher level concepts and semantic relationships which are unable to be parsed in the machine environment. Museums and libraries have traditionally practiced concept-based indexing and have invested heavily in the tools needed to do this well, sophisticated controlled vocabularies, ontologies and schemas such as the Library of Congress Subject Headings (LCSH), the Art & Architecture Thesaurus (ATT), and the International Council of Museums' Committee on Documentation-Conceptual Reference Model (CIDOC-CRM). However, concept-based indexing, as a manual process, is expensive and is still both 'information lossy' (i.e. assigned keywords may not capture the full semantics of the image's content) and context sensitive, as different indexers may use different sets of keywords for describing the same semantic meaning (Smeulders et al., 2000).

While providers may optimize the content, amount, and quality of metadata based on the available resources and their intended use, studies have indicated that the metadata provided may not necessarily be the same as that needed by end-users (Jörgensen & Jörgensen, 2005; Palmer, Knutson, Twidale, & Zavalina, 2006). For instance, the Jörgensens' research suggests end-users describe many more attributes than are commonly provided for in standard controlled vocabularies and metadata structures.

The expanding grid of social content creation and sharing systems, which includes both textual and visual media (e.g. Wikipedia, Flickr, LibraryThing), has become an increasingly important, and inexpensive source of information that can include collection and item-level description and

metadata. Indeed, the final report of the Library of Congress (LC) Working Group on the Future of Bibliographic Control recommends that libraries enable their systems to connect to "appropriate user-added data available via the Internet" (The Library of Congress, 2008, p. 32). Most recently the Library of Congress deployed more than 3,000 photos from its most popular collection on Flickr (http://www.loc.gov/blog/?p=233). According to the LC, the goal is not only to increase the availability and awareness of its rich visual collections to the public, but to engage the Flickr community in improving the quality of the metadata for these photos. Many of the photos lack key metadata such as when and where a particular photo was taken, and who is shown in the photo. The hope is that Flickr users, by accessing these photos, can contribute some of this missing information.

The new social content-creation systems such as Flickr blur the line between information providers and end-users by enabling end-users to create their own information content, including metadata describing that content. Hence, it is reasonable to expect that both object and collection metadata in these systems would be optimized to the individual contributor—the end-user's context. Combining individual collection metadata into an aggregate collection, however, may lead to quality problems (Stvilia, Gasser, Twidale, Shreeves, & Cole, 2004). Indeed, after analyzing the quality of metadata tags used in Flickr, Guy and Tonkin (2006) noted that the main problem with Flickr tags may lie precisely in their use in two different contexts at the same time—individual and collective. The above-mentioned report from the Library of Congress also recommends that libraries help improve the quality of user-created tags by suggesting an entry vocabulary (The Library of Congress, 2008). The impact and utility of such aggregated data remain to be fully understood, but nevertheless the use of systems allowing end-user tagging and emergent metadata structures continues to grow.

Traditionally, printed photographs have been used in both these individual and collective contexts, for creating and sharing stories of different events and family histories. Considerable research has been done by the Computer-Supported Cooperative Work and Human–Computer Interaction communities on how to design systems enabling effective storytelling through digital photographs (Balabanovic, Chu, & Wolf, 2000; Frohlich, Kuchinsky, Pering, Don, & Ariss, 2002). These studies have investigated how and why people store and share photographs, and their needs in using a photoware system. The following requirements for photoware were elicited from user interviews: (1) archiving and capturing metadata with the least cost to the user; and (2) easy sharing with friends and family members. Indeed, Frohlich et al. (2002) found "Joy from the feedback and subsequent conversation around the photos" as one of the main motivations for sharing photographs. In addition, two more activities using photographs were identified—storytelling and reminiscence. Storytelling occurred when one showed photos to others who were not present when the photo was taken. Reminiscence, on the other hand, occurred when the members of the conversation were present at the time the photo was taken, and they recollected some aspects of the scene depicted in a photograph while leaving out most of the shared contextual knowledge.

Rodden and Wood (2003) investigated end-users' practices of digital photograph organization. Their study identified the following search capabilities users liked to have access to: (1) finding photos from a particular event, (2) finding individual remembered photos, and (3) finding photos from different events but with a shared property. While the majority of the respondents did not consider annotating individual photos to be worth the time, they thought it was important to have descriptive names for collections. More recently Cunningham and Masoodian (2007) analyzed the self-reports of photo taking, tagging and organizing behavior collected from 18 students. The study found that the students used time, event and location as facets to organize their photos, but

attached little metadata to them. In earlier research the same authors looked at the image information needs and seeking actions of 31 undergraduate students (Cunningham & Masoodian, 2006). The analysis of the student reports suggested four categories of image information needs: (1) specific, factual; (2) general namable; (3) general abstract; and (4) subjective, emotional.

In another recent study Ames and Naaman (2007) interviewed 13 'heavy' users of the camera - phone tagging application ZoneTag. The interviewees used ZoneTag to tag and upload photos to their Flickr accounts. In addition to identifying different uses of ZoneTag's features and their effects on the users' tagging behavior, the research analyzed users' tagging motivations. The following motivation types were identified: Self-Organization, Self-Communication, Social-Organization, and Social-Communication. Interestingly, the study found that the main motivation for the users to tag was to make photos accessible to the general public, followed by the need to search and select one's own photos. However, due to the small size and the type of the sample (a focused sample), this finding might not be generalizable to the general user population. These research reports bring us back to the issue raised earlier, the concurrent individual and collective contexts of tags and what the impact of this dual nature is on collection building and thus, information seeking.

Research Questions

A collection, as suggested by the literature review, is a high level conceptual manifestation of information organization. A collection definition may specify shared characteristics by which photos are organized, as well as the motivations and management rules for the collections. Many different facets and relationships can be used for organizing collections. Understanding the concepts and relationships that users organize personal and shared collections by can both inform the process of making information resources from traditional databases more accessible and potentially enable reuse of user created collection metadata. Similarly, understanding *how* and *why* users create both personal and shared collections can facilitate assessing and improving the quality of user provided collection metadata. Therefore, this research aimed to address the following research questions:

1) What types of collections do Flickr users of create and how do these types compare to each other? What are the criteria by which users name these collections?
2) What are some of the motivations for building these collections and the activities in which these collections are used?
3) Are there similarities between the characteristics and topics used to organize collections in Flickr and the findings of earlier research on attributes used in describing and sorting tasks with images?
4) How well do the elements and classes of currently used metadata schemas and ontologies support user created collection metadata?

While the focus of this research is on discovering factors which could inform metadata practices in distributed collection building, it should be noted that Flickr users, for the most part, are not necessarily initially thinking of an abstract concept such as metadata but are concerned with assigning words or "tags" to individual photos and groups of photos that will make these discoverable to themselves and/or to others in the future. It may be that, for some individuals, assigned tags then become an organizing principle as further images are added, as it has been shown that tags themselves in a collection tend to "settle down" to usage of a smaller group as

time progresses. This in is accordance with power law principles that have been observed in language usage and that have been observed in Flickr as a whole (Golder and Huberman, 2006).

Procedures

The research was a collaboration between two researchers, one with expertise in collection level metadata and one with expertise in visual item-level metadata. The broad goal of the methodology was to utilize these two types of expertise to analyze a single collection that shares characteristics of both individual and collective metadata creation. While different parts of the analyses were initially done by the researcher with the most expertise in that area, the researchers followed these analyses with discussions and agreed upon the categories arrived at.

The data for the research were collected in September 2006. To gain insight into end-users' digital photo describing and collection-building practices, the research began with an analysis of the two major types of collections which are determined by the Flickr system itself: individual *Photoset* and public *Group* descriptions. Each Flickr user may collect his or her photos into one or more individual sets called Photosets. Each Photoset has a title and may also have a description identifying the motivation for creating the Photoset or the circumstances under which it was created. Creating Photosets is not required of users. Like Photosets, Groups are also initiated by individual users but are open to others under varying conditions. The creator of a Group defines the purpose of the Group and sets the initial organizing principles and rules. The creator is automatically assigned administrative rights to the Group by the system. A Group administrator may assign the roles and access rights to the rest of the Group members. If a Group is public, other Flickr users may join the Group with or without an invitation. Private Groups can be joined only by invitation, which can be self-initiated. Group members can contribute photos to the Group's photo collection or pool and can also participate in discussions about the Group or about individual photographs.

Three thousand photographs were sampled from the Flickr Recent Photos page by using scripts developed in the Java language (http://java.sun.com/) and a Flickr Java application programming interface (http://sourceforge.net/projects/flickrj/). The sampled photos were associated with 879 users who had created a total of 13,212 Photosets hosted on Flickr. The researchers also used the Flickr IDs of those users to collect 11,140 unique public Groups in which the users participated. Next, 300 randomly selected Photoset descriptions and the descriptions and discussions of 200 Groups were content analyzed by the researcher with expertise with digital collections metadata. The same researcher performed the coding, using Atlas.ti software. The coding schemes employed in the content analysis included the types of collections observed, as well as shared activities of groups building the collections, and their motivations. The researcher started by applying an open coding procedure to the samples. The resultant codes were iteratively clustered to develop classifications or typologies (Bailey, 1994a). The researcher then used these classifications to recode the samples. Finally, the study used SPSS (http://www.spss.com) software to generate descriptive statistics and graphs for the samples.

Findings

Types of Collections

As described above, the Flickr system itself defines the two major types of collections, individually created Photosets and collectively developed Group collections, which have a defined (although sometimes quite broad) purpose. Analysis of the Photoset and Group descriptions revealed that although they were mostly unstructured, some recurrence of themes and semantic components was observed in both types of descriptions, suggesting they may share a degree of similarity. In Information Theory the quantity called *entropy* is used a measure of disorganization. A related quantity called *relative* entropy is a measure of the distance between two statistical distributions, which can be distributions of text (Cover & Thomas, 1991). Information Retrieval research often uses relative entropy as a measure of the distance or similarity between an individual document and a collection (e.g. Gabrilovich, Dumais, & Horvitz, 2004). This measure is used in the current research to measure the degree of similarity within the Photoset descriptions and the Group descriptions, and by inference the degree of similarity between the two types of descriptions. The median value of the relative entropy[1] of the Group descriptions was compared to the median value of the relative entropy of the Photoset descriptions. The median value of the relative entropy of the Group descriptions against the pooled set of descriptions was lower than the median value of the relative entropy of the Photoset descriptions against the pooled set of photoset descriptions (0.15 versus 0.84), suggesting that the Group descriptions might be more homogeneous or similar to each other than the Photoset descriptions. Indeed, a non-parametric Kruskal-Wallis test on the relative entropies of the descriptions against the pooled set of all descriptions showed a significant dependence on the description type (i.e. Photoset or Group) (χ^2 = 5800.78, df = 1, $\alpha < 0.001$).

Photoset Collection Types

The content analysis of the Photoset descriptions suggested Flickr users use eight criteria to group photos into Photosets: (1) *Activity*; (2) *Place*; (3) *Person*; (4) *Artistic*/*Photographic Technique*; (5) *Thing*; (6) *Random*; (7) *Time*; and (8) *Quality*; (Table 1). Note that these types were derived from both the content and the attributes of the photos, as evidenced by the text describing the category. Not surprisingly, for individual Photosets the most important criteria were Activity, Place, and Person. Given that these were photographs, photographic technique was also an important criteria, as were things depicted in the photos.

The largest number of sample Photosets were about an *Activity* (Table 1). This category included Photosets describing general *activities* without specifying their time and place; *events* (activities specified in time and space); and *processes* (ordered sequences of activities):

"MacBook Pro Unboxing" or "How to Get Your Hair to Curl . . ."

[1] The relative entropy was calculated as

$$\left(\sum_{x \in V} d(x) \log \frac{d(x)}{c(x)}\right) / N_d$$

where x is a term in term space V; $d(x)$ is the probability of the occurrence of x in the model of the document or object; $c(x)$ is the probability of the occurrence of x in the collection model; and N_d is the length of d. Laplace smoothing was applied for unseen words.

The most popular *event* Photosets were birthdays and weddings. Clearly, any activity involves places, people, things, actions, operations, and general sociocultural contextual components, such as language. The study found that the users would often name the Photosets by the activity name, even if the photos focused more on individual persons and places than on documenting the actions or operations of the activity. In some instances, the Photosets were created not for reflecting and documenting some past activity or event, but for directly supporting a specific activity in which the user was engaged:

> "eBay: Stuff I'm selling on eBay."

A few Photosets had clearly defined educational goals, such as to help the public recognize different consumer goods and groceries or to show processes such as how to make things, e.g. assembling a computer or cooking a dish:

> "I recently started taking photos of the food I cook, since it apparently impresses people on The Internet. If you want to check out photos of what it takes to make some of this stuff, check out the rest of my food."

Not surprisingly, *Place*, *Person*, and *Thing* were also very typical Photoset descriptor types. Photosets may be organized by a specific place (e.g. San Francisco) or by a location type, such as urban, rural, or mountains:

> "City: different urban/city shots i have."

Person Photosets, may focus on an individual or group. Some, such as the "mirror" or "365-Day Challenge" were simply self-portraits of the user. Other *Person* Photosets might focus on another person or a group of persons having some shared characteristic or relationship, such as membership in a band or a family:

> "Family Feature: Selected family snapshots."

The content of a *Thing* Photoset could be anything that was not a person or human: a cat, dog, plant, toy, car, food, or an example of art. A Photoset could be photos of one thing or of several things combined together based on some shared characteristic, such as a type, brand, or maker. Because photos were taken by a particular person the Photoset organizations were expected to reflect that individual's context, implicitly or explicitly:

> "Stuff I've Made: Quilts and Other Fun Stuff." "Courthouses: I love courthouse photos . . ."

The *sentimental value* of photos to the user might serve as a basis for organization as well:

> "Nostalgia: pics from the past"

The perceived *Quality* of photos determined by the user himself or herself or voted on by the Flickr community can be another organizing factor:

> "Like the title says, the photos I've taken that Flickr deems interesting."

A number of Photosets were organized by the *Artistic* or *Photographic Technique* used:

> "photos that are more interpretive than realist; created with double exposures, or photoshopped, or both."

Time, like place, has characteristics that could serve as a selection factor on its own: night, morning, winter. The user might use more than one attribute when selecting photographs for a set. These could be combinations of *Thing* and *Time*, *Thing* and *Place*, *Time* and *Place*, *Person* and *Place*, *Person* and *Time*, or all of the above:

> "X: Week Five;" "C&N San Francisco 2002: In 2002, when my daughters were 4 and 6, I took them to San Francisco for the first time;" "Photos from May of 2006."

Users might also use Photosets simply as archives of *Random* photos assembled without any specific goal or shared characteristic, or again in combination with one or more other attributes:

> "Random Stills: random stuff also photoshop practice mostly nature shots."

A Photoset could be a reportage of some existing collection or exhibition, and hence might inherit the organization of that collection:

> "The Buddhist Sand Mandala art at the Canadian Clay; Glass Gallery."

Finally, although most of the Photosets were provenance based and might be triggered by some event or shared characteristic of the photos in the user's possession, subsequent photos might be matched against the already established sets based upon a particular attribute:

> ". . . So I can tag a photo 'Saint Paul' in my iPhoto plugin uploader and it will automatically go into my 'Saint Paul' set."

The data thus indicate that while *Activity*, *Place*, *Person*, and *Thing* tend to be very common Photoset attributes, these may be uniquely influenced by a user's personal context.

Group Collection Types

At the time of the analysis, Flickr provided ten predefined Group categories for its users: (1) Computers & Internet, (2) Culture & Society, (3) Flickr, (4) Fringe & Alternative, (5) Languages, (6) Life, (7) Nature, (8) Recreation, (9) Regional, and (10) School & Learning (http://www.flickr.com/Groups_browse.gne (note that as of May 23, 2007, the link did not work and the list of Group categories seems to have been removed). Group creators could declare membership in up to three Flickr Group categories. Flickr, however, did not give any definitions or guidelines for how to apply the classification scheme to the Groups, nor did it give information on how the typology was formed. Therefore, users created their own Group descriptions (lack of use of the predefined categories may have been the reason for their removal).

In general, Groups were focused more on the general components of activities than on the individual events or actors as it is the case with Photosets. A Group description would define the *topic* and *scope* of the Group—what photos would be appropriate for its collection and why. This was done by listing the types that should be included or excluded:

> "However, we don't want too many of the family portrait and friends shots, these generally don't interest other people, unless they know them."

The content analysis of the sampled Groups identified similar categories as in the Photosets, but somewhat different distributions or emphases. In contrast to the Photosets, the Groups were organized more around *Things* than around personal *Activities* or events (Table 1). There were fewer Group collections focusing on weddings or birthdays. Instead, more Groups were dedicated to such things as motorcycles, martinis, cars, traffic signs, items of clothing, or biological species:

> "This Group is dedicated to Animals and Plants that are on the Endangered Species List."

Furthermore, the Group collections that could be placed under the *Activity* category focused more on actions than on the agents or the general context of an activity:

> "Cigarette Bums: People who bum cigarettes. If they want my cigarette, I want their picture"
> "People Watching or With a TV: *NO PICTURES OF JUST TELEVISIONS. People can be mesmerized by television. Sometimes I like to watch them watching."

The Group content analysis suggested two categories of collection-building factors in addition to the those identified in the analysis of the sample Photosets: (1) *Community* and (2) *Concepts*.

The *Community* category included Groups that were created with the goal of providing "active" space and support (emotional, information or knowledge, technology) to Flickr users by using photos, games, or Group discussions:

> "This Group pool exists to support the members of the Flickr Community that have Cancer or have friends/family battling Cancer."
> "Flickr Central is a place for the newbies to get a taste of what Flickr is about, and a place for the more experienced users to keep a finger on the pulse of our favorite addiction."

The coder also placed groups based on the language and the origin of their members in this category:

> "For people who speak Spanish, whether at the level of the Quixote, Borges and Neruda or 'uno más cerveza por favor.'"
> "This is a Group dedicated to photography enthusiasts from around Australia, who want to meet other australians or see what other Australians are doing in photography!"

The *Concepts* category included the Groups that attempted to express or document some immaterial concept or idea through photography:

> "Recursion: Pictures of Pictures; Pictures of things that take Pictures; Pictures of things that display Pictures."
> "Simplicity. . . . Contemplative photography. The art of simplicity. Uncomplicated and thought provoking. Sharing of the simple things in life."

Table 1. Distributions of collection categories among Photosets and Groups

Category	Photoset descriptions (%)	Group descriptions (%)
Activity	40.4	5.6
Place	24.6	21.3
Thing	11	29.9
Person	9.6	9.6
Photographic Technique	11	12.2
Random	1.5	1.5
Time	1.1	0.5
Quality	0.4	3.0
Community	0	7.1
Concepts	0	8.6

Behaviors and Motivations

The second research question asks: What are some of the behaviors and motivations in building these types of collections and what are the activities in which these collections are used?

User participation in creating Photosets and describing individual photos varies widely. The analysis showed that approximately 18% of the users from the sample did not create any Photosets, and more than 50% of the Photosets did not have any descriptions attached. The median number of photos in the Photosets was 22 and the median number of Photosets per user was four. The sample also suggested that both the number of photos in the Photosets and the number of Photosets per user might follow power law distributions (see Figures 1, 2), as has been suggested by other researchers for the descriptive terms applied to photographs in Flickr.

Similar to the Photosets, user participation in Groups varied significantly. Thirty-seven percent of the sampled users did not belong to any Group. The median number of Group memberships was

two. The median number of images and members in the Groups were 874 and 188, respectively (Figures 3, 4). In clear contrast to the Photosets, only 2% (vs. >50% for Photosets) of the Groups had no description attached, and the median description length was 264 characters. Also, Kruskal - Wallis tests showed that the Group collections differed significantly ($\alpha<0.001$) from the Photosets on the number of images, the length of title, and the length of description (see Table 2). In general, the Group collections had higher numbers of images, and longer titles and descriptions than the Photosets. Furthermore, the analysis showed that more than 50% of the photos in the sample did not have any tag attached. Hence, the Group or Photoset descriptive metadata often was the only metadata these photos had.

Table 2. Kruskal-Wallis test on distributions of title length and description length between the Group and Photoset collections (the original samples were normalized by removing the cases with empty descriptions).

	# Photos	Title Length	Description Length
Chi-Square	6,299.17	68.87	3,021.47
Df	1	1	1
Asymp. Sig.	< 0.001	< 0.001	< 0.001

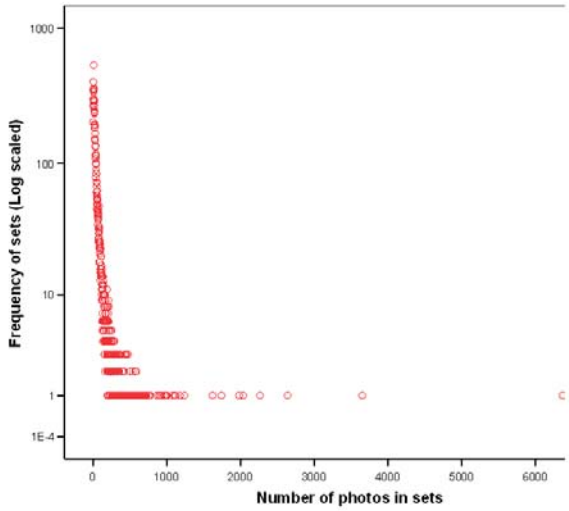

Figure 1: Distribution of the number of photos in the Photosets. The median number of photos was 22.

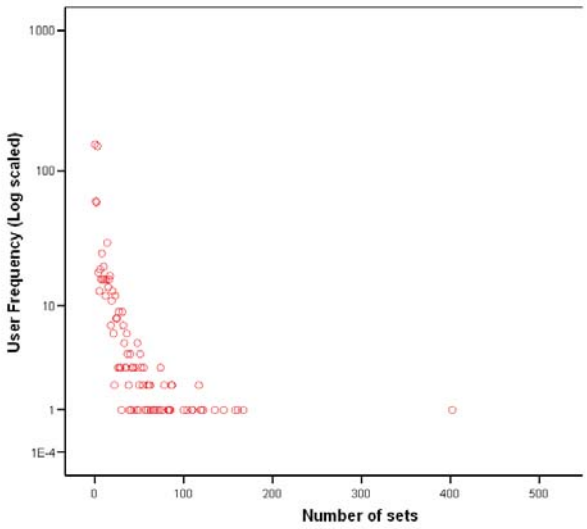

Figure 2: Distribution of the number of Photosets per user. More than 18% of the users from the sample had not yet created a Photoset; the median number of Photosets was 4.

In general, a Group description would define the topic and scope of the Group—what photos would be appropriate for its collection. This was done by listing the types that should be included or excluded. While Photoset creators tended to modify common attributes with their own individual user contexts, Group creators define the purpose of a Group and would often set rules or protocols for member participation or the behavior of the participants. For instance, a rule would set a limit on how many photos the members could add to the collection daily. This rule was clearly intended to control the quality of the collection—to make sure that all submitted photos were screened against the goals and organizing principles of the collection:

> "Maximum 1 per day. . . . your best!"
> "PLEASE limit your posts to 5 or less per day, and …read discussion on posting images linking to your webstore."

The descriptions could also include guidelines for constructing the title and metadata tags for the photographs:

> "In the tag you have to include 'window' and where was taken the picture."
> "Your pix's names must begin with the station's name in ALLCAPS. :-)"

Sharing often involves conversation. In Flickr, conversations are asynchronous and are carried out through discussion threads (archiving may not involve conversation). Groups with an explicit goal of sharing and criticizing the photos posted by their members would be expected to involve a higher amount of conversation than the other Groups or individual Photosets. Hence, the descriptions of such Groups would include the rules and guidelines for managing the member discussion threads:

> "For people that want the members to critique their work it is preferable to post the image to the discussion thread. To do that, go to the photo in your Flickr, select 'all sizes' on the left hand side. Select the small version and copy the link on the bottom . . .
> Take a look at other people's work and **leave 1 comment**. If you post 1 photo you should give a feedback to 1 person at least . . .
> Also, please do not post more than 3 photographs per day on the discussion thread, that way everyone gets a chance to post and have their work be seen and discussed. Thanks."

As noted earlier, the Groups were focused more on the general components of activities than on the individual events or actors. Some Group descriptions explicitly stated that:

> "However, we don't want too many of the family portrait and friends shots, these generally don't interest other people, unless they know them."

A Group purpose could change over time, and some of the Group descriptions reflected that:

> "The original goal was to make this the largest Group on Flickr, just to (literally) thank the 'creators of probably the best site on the internet. Right now our goal is to KEEP flickritis as the 'funnest' Group on flickr, if funnest is actually a word!"

Often Groups would attempt to provide community spaces for "active" viewing—playing games and voting for the best photos, learning photography, and getting feedback or critique on each other's work:

> "What would Simon Cowell say? (READ THE RULES!): . . . This Group is for 100% honest criticism—but it's not just an excuse to be nasty. Feel free to keep it witty, though!"
> "MatchPoint: photographical tournaments. It is intended to be fun to watch and fun to play. The tournament consists in a series of one-to-one matches. For each match, two players confront a picture of their choice. Winners of first round matches confront themselves in the second round, eventually reaching semi-finals and possibly the great final . . ."

Interestingly, photo games were also used for building context and telling a story of a place or a person

> "Atlanta Chain: each photo in the pool has at least two notes on them. One highlighting an element the NEXT picture should have and one highlighting the feature that was required in the PREVIOUS photo. Look at the last few photos in the pool and it should be easy to understand . . ."

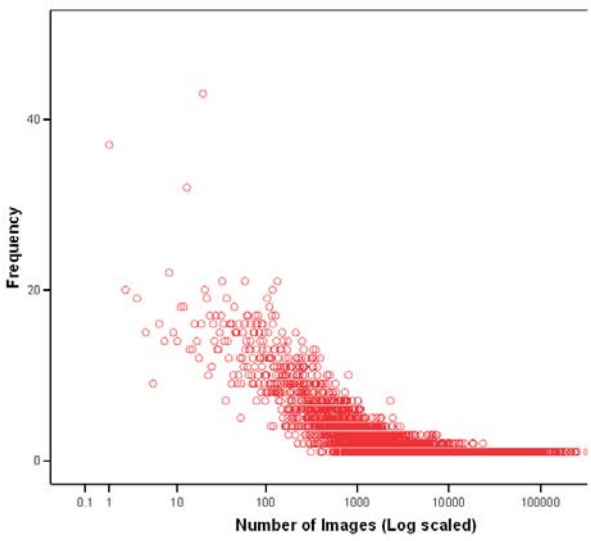

Figure 3: Distribution of the number of images in the Groups. The median number of images in the Group pools was 874.

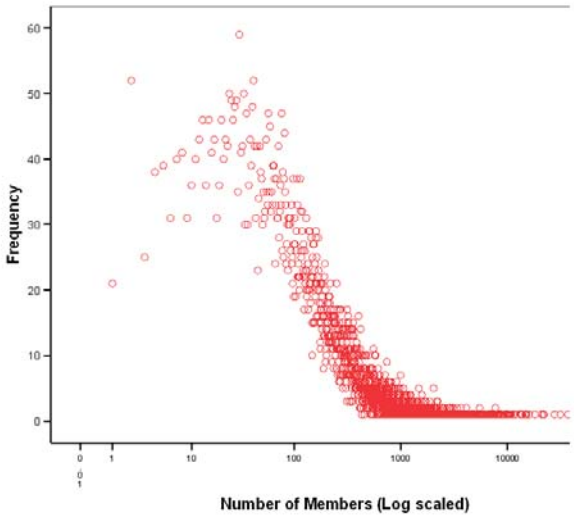

Figure 4: Distribution of the number of members in the Groups. The median number of members in the Groups was 188.

One of the Groups attempted to stratify the photographs into distinct but dynamic "quality" pools. The boundaries of the pools were determined by the number of times the photos had been viewed. Once a photograph accumulated the number of views outside the pool's range, it was moved to the next higher level of pool:

> "Your photo has 201 to 400 views? Put it here. After you add a photo, please view AND COMMENT ON the first THREE photos in the photo pool—at a minimum—in order to advance photos to the next Group. Move along to the next Group when your photo hits 401 views."

At the same time, a number of Groups seemed to have no particular goal but to amass the largest pool of random photographs, the largest number of members, or both. These collections were put under the *Random* category, although they could be called *vanity* collections as well:

> "There are Groups that want to have the most members, most viewed, most discussed. We want the most photos of any Group! At the time of this writing, the Group with the biggest pool has over 700,000 photos. Our goal is 10 Million!!"

One might expect the Group descriptions to include different amounts of descriptive information and guidelines, depending on the complexity of the Group's scope, goals, and member interaction. Indeed, the Groups organized for critiquing members' work and playing games had significantly longer descriptions than those sharing photographs of persons or things. The Friedman chi-square test for the Group types ranked by their description lengths was significant ($\chi^2 = 148.95$, df = 9, $\alpha < 0.001$). Interestingly, the same test on Photoset descriptions too showed a significant relationship between the photoset type and the length of description ($\chi^2 = 291.06$, df = 7, $\alpha < 0.001$).

Identifying user's motivations for building collections is essential for better aligning traditional collections to the user's needs, as well as for facilitating and improving the process of collection metadata creation by the user. The content analysis of the Photoset and Group descriptions suggested that Flickr users might have the following motivations when assembling collections: (1) to enable easy finding, (2) for easy sharing, (3) for archiving, (4) vanity, (5) "bibliographical," documenting a particular subject or concept (e.g., a sunrise), (6) supporting group or community activities (e.g., playing a game), (7) supporting an individual activity (e.g.,

documenting a process of setting up a computer for later use), and (8) no particular motivation - the collection was a product of the sum of many individual one-time activities.

Comparison with Earlier Research

The third question for the current research asks whether there are similarities between the characteristics and topics used to organize collections in Flickr and the findings of earlier research on describing and organizing images. This question can be answered by investigating to what extent attribute classes or item-level attributes from previous research would align with the data from the Flickr *Photoset* and *Group* descriptions.

As suggested by the literature review and subsequent data analysis, a collection definition may specify shared characteristics by which photos are organized, the motivations and/or management rules for the collection, or both. This study identified 10 categories of topics and characteristics that Flickr users used for organizing digital photo collections into Photosets and Groups. These included content entities and topics, as well as the properties and intended uses of the photos (Table 1). Earlier research by Jörgensen and others (Jörgensen, 1995; Jörgensen,1996; Jörgensen,1999; Brunskill and Jörgensen, 2002) investigated terms applied to images when different groups of participants describe and sort images under a variety of contexts. One of the motivations for the current study was to investigate the relationship of terms describing collections generated in a social networking system to terms generated by individuals describing images for a variety of purposes, including finding and sorting images.

We thus compared the Flickr data with quantitative data from two different task contexts, 1) describing images; and 2) sorting images into groups (Jörgensen, 1995). The describing task data were lists of terms from participants descriptions of images, including spontaneous single word descriptions under two conditions (no particular task and a task which was to enable finding the images at a later date) and descriptions of images from memory. Data from the sorting task consisted of participants' individual descriptions of groups into which images were sorted, as well as important attributes of the groups; these groups were created by participants to facilitate finding these images at a later time. This earlier research produced a large number of attributes that were grouped into higher level categories or classes for purposes of analysis. Classes of attributes that emerged from this earlier research were as follows: *Objects, People, Color, Visual Elements, Location, Description, People-Related Attributes, Art Historical Information, Abstract Concepts, Story, External Relation, and Viewer Response*.

As the earlier research was intended to establish the full range of image attributes characterized by participants in describing and sorting tasks, there were a larger number and a great deal more specificity of attributes in this research. Thus, these classes of attributes did not match the Flickr groups directly, but mapping of individual attributes to Flickr categories enabled comparison. Table 3 shows the mapping that was used for the comparison between the two sets of data. After mapping, these attributes and attribute classes were compared to the Flickr groups resulting from the recent content analysis to assess the extent of congruency among them. This mapping and analysis was done by the second researcher who had not performed the content analysis of the Flickr data.

Table 3. Flickr categories and attributes and attribute classes from previous research (Jörgensen 1995) included in the category.

Flickr category	Attributes and Classes from Jörgensen mapped to Flickr category
Activity	Activity, Event
Person	People*
Photographic Technique	Technique, Focal Point, Perspective, Orientation, Visual Component, Format, Motion, Composition
Place	Setting
Quality	Not directly mappable
Random	Unknown (from Sorting Task: "Not Sure," "By Itself," "Don't Understand," "No Category")
Thing	Object
Time	Time Reference, Time Aspect
Community	Not directly mappable
Concepts	Abstract, Symbolic, Theme, State

* This Class has a single attribute

Several interesting similarities and asymmetries were found in this data comparison (Table 4). Two categories from the Flickr data could not be directly mapped to data from the describing tasks: *Quality* and *Random*. Both occurred infrequently (3% or less) in the Flickr data and were eliminated from the comparison. *Time* also occurred less than 3% in both samples and was also eliminated.

Table 4. Distributions of Flickr collection categories and distributions of attributes across describing and sorting tasks in Jörgensen (1995).

	Flickr Data		Jörgensen data	
Category	Group descriptions (%)	Photoset descriptions (%)	Describing tasks (%)	Sorting task (%)
Activity	5.6	40.4	4.6	3.6
Person	9.6	9.6	10.0	8.6
Photo Tech.	12.2	11	11.3	6.6
Place	21.3	24.6	1.8	2.6
Thing	29.9	11	29.3	7.6
Community	7.1	0	NA	NA
Concepts	8.6	0	1.2	10.1

The major category with a similar distribution across the two sets of data was *Person*, the "who" of an image. Research in cognitive psychology has demonstrated that humans are highly sensitive to the presence of other humans in the visual field, so it is logical that the percentage of participants noting the presence of a person is consistent across these varying tasks and collection descriptions (Buswell, 1935; Yarbus, 1967). *Place*—the "where"—was high for the Flickr data and low for the image tasks; this makes sense because the users themselves provided data for Flickr, whereas participants in the image tasks were given no information about the image they were seeing and so could only guess at a location. *Activity*—the "what"—is low in the Group descriptions and in the image tasks, but high for Photoset descriptions. Again, this is sensible because the Photoset descriptions are a place for sharing images with friends and family, and these photos (labeled by the participants) are often of social events such and weddings and parties. *Photographic Technique* qualities are more frequently noted in Group and Photoset descriptions and in Describing Tasks than in the Sorting Task, in which techniques were less important and affective attributes such as emotion and mood played a more prominent part. One of the most

interesting comparisons was with the *Concepts* category; here the Group Description results were more similar to the Sorting Task, and the Photoset Descriptions were more similar to the Describing Tasks. As can be seen from the mapping, the Concepts category included such attributes as abstract concepts and themes, and symbolic meanings. The Sorting Task thus elicited a wider range of image attributes more similar to those in the photo Group collections, in which *Concepts* and *Community* were observed. One could theorize that the describing tasks were thus more similar to the act of tagging by individuals, whereas the sorting task drew upon more abstract conceptualizations, perhaps more typical of group goals and group participation.

Current Metadata Frameworks

Another goal of this study was to provide data to evaluate and inform the design of the existing visual resource metadata schemas and controlled vocabularies, that is, how well current metadata encoding tools could accommodate user supplied metadata. The fourth research question asks how well the elements and classes of currently used metadata schemas and ontologies match up with, and therefore support, user created collection metadata.

To answer this question the study compared the Flickr Group categories with the classes and properties of the CIDOC-CRM ontology (Crofts, et al., 2007), and the subject type values of the VRA Core 4.0 Schema (the VRA Core 4.0, 2007). The CIDOC-CRM is one of the most comprehensive ontologies for the material cultural heritage domain. It contains an extensible specification of conceptualizations of the domain — the definitions, classes, relations, and a vocabulary. The main purpose of the CIDOC-CRM is to provide a mediation/interoperability layer for different metadata schemas used in the cultural heritage domain (Crofts, et al., 2007). The VRA Core 4.0, on the other hand, is a specific descriptive metadata schema for the cultural heritage community, developed by the Visual Resources Association's Data Standards Committee. The current analysis found that the CIDOC-CRM ontology could support all of the Flickr collection types. The VRA Core 4.0 Schema, however, could not handle the *Activity*, *Photographic Technique*, *Quality*, and *Community* types (Table 5).

Table 5. Comparison of the Flickr collection types with the CIDOC-CRM, and VRA Core classes, properties, and subject types.

Flickr collection type	Flickr organizing characteristics	CIDOC-CRM classes and properties	VRA Core 4.0 restricted schema subject type values
Activity	Subject (content)	Activity, Event	NA
Person	Subject (content)	Person	personalName
Photographic	Attribute	Used general technique, Type	NA
Place	Subject (content)	Place	geographicPlace
Quality	Attribute	Type	NA
Random/Other	NA	NA	NA
Thing	Subject (content)	Stuff	scientificName, otherName
Time	Subject (content)	Date	NA
Community	Attribute	Activity	NA

Concepts	Subject (content)	Conceptual Object	conceptTopic

In addition to describing the subject of a collection, approximately 20% of the Group descriptions contained some elements of accrual policies, retention or deletion policies, or both. Most frequently (15%), the descriptions specified the rules for the number of photos the members could contribute over a period of time, as well as the rules for deleting and tagging photos. A comparison of the Group description components with the elements and vocabulary of the Dublin Core Collections Application Profile (DCCAP; http://dublincore.org/Groups/collections/collection-ap-summary/2007-03-09/) showed that the Flickr Group description rules were specified at the procedural level rather than at the policy level as in DCCAP. Also, at the time of writing, the DCCAP did not have elements and vocabulary for specifying a deletion policy and procedure(s), or for specifying annotation or tagging schemas for individual photos in a collection.

Discussion

To remain viable and useful, it is essential for traditional metadata schemas to be aligned well with actual user needs for and uses of metadata. This research outlined four exploratory research questions to investigate evolving user activities and the needs for information resources and related collection level metadata in Flickr.

The first research question sought to identify the types of image collections created by users and the criteria by which users name these types of collections. The analysis of the types of collections showed that individual user Photosets focused on different components and events of the user's context, whereas Group collections were organized more around general concepts, or discussion and specific activities (e.g., commentary on Group members' work). In addition, although individual Photosets evolved in a more bottom-up manner based on some shared attributes and subjects of the user's photos, Group collections were constructed top down by matching characteristics of photos with the predefined purposes and scope of the Group.

Interestingly, the quantitative tests found significant relationships between some of the collection attributes (e.g. description length) and the collection types both between the Photoset and Group samples and within each sample. This points not only to the stability of the collection types identified by this study, but also to possible latent differences in the contexts of creation of these collections.

The second research question aimed at identifying some of the user motivations for creating collections and the activities in which these collections were used. The literature identifies three generic types of metadata – descriptive, administrative and structural (Caplan, 2003). Descriptive metadata describes an information resource for the purposes of discovery, identification, selection and retrieval. Administrative metadata gives information relevant to managing a particular resource (e.g. accrual/selection rules, copyright, provenance, media, software & hardware used, etc,). Structural metadata provides information about how different information resources or the parts of the same information resources are related to each other. Most of the earlier research on end-user image metadata requirements identified the need for descriptive metadata (e.g. Cunningham and Masoodian, 2007). In Flickr, however, in addition to the descriptive information about a collection, some of the Group descriptions included elements of administrative metadata

such as rules for coordinating and managing contributions to the collection. This might be explained by the multiple roles users might have in Flickr – as a creator, a manager, and a consumer of images.

Furthermore, these 'collection policies' were explicitly linked to a specific activity (e.g., critiquing members' photos), and the needs of the activity shaped the accrual or use rules of the collection. In earlier research, Gordon (2001) suggested the need for a richer browsing experience of image collections by connecting image access points (subject terms) through their relationships with a shared general activity context. He proposed using a conceptual analysis of the *Library of Congress Thesaurus for Graphic Materials* (YR) to identify these shared activities. Some of the disadvantages of such conceptual classifications are relatively high cost and nonparsimonious categories (Bailey, 1994b). The majority of individual Photosets, and some of the Group collections in Flickr, are organized by the context of an *activity* they report on or are used in (see Table 1). Hence, Flickr can serve as a valuable source of inexpensive empirical data on the individual photo subject terms (tags) and activity relationships.

The set of user motivations found in this study included more elements than another study of Flickr users (Ames & Naaman, 2007). In particular, the findings of this study include motivations such as *archiving* and *vanity*. However, this could be explained by the focus of the current research which aimed to investigate users' motivations to create a collection, not their motivations to tag. Indeed, the vanity and preservation motivations found in Flickr were listed in the set of functionalities of digital collections proposed by Buckland (1992). However, in addition to all the functionalities from Buckland's set, the Flickr set also included motivations to support a particular group or individual activity. This result points to the diversity of Flickr uses, which ranges from personal information management to group work and social networking, and to providing services and information to the public.

The third research question was concerned with the degree of similarity between the characteristics and topics used to organize collections in Flickr and the findings of earlier research on attributes used in describing and sorting tasks with images. The comparison of the Flickr data with the quantitative data from two different task contexts analyzed earlier by Jörgensen (1995) showed that in the Flickr Photoset description environment, attributes such as *Person, Place,* and *Activity* emerged as typical Photoset types, similar to the more direct task of describing images. In the Flickr Group description environment, the *Concepts* category emerged as more important; in previous research, *Concepts* were more important in a sorting task, which gave participants more freedom to create their own categories. Given that Flickr image descriptions, as noted above, are created in both individual and collective contexts, a better understanding of the nature of the contexts and the attributes they produce may improve the use of these descriptions as access points in large collections. The results of Jörgensen's research, which aimed to establish the range of attributes typically described by users of images, also points to the need to consider multiple contexts and thus multiple access points.

The fourth research question focused on how well currently used metadata schemas and ontologies support user created collection metadata. Interestingly, while the CIDOC-CRM supported all the subjects of the Flickr Group descriptions, the VRA Core 4.0 did not support subjects such as quality. Also, the DCCAP did not support encoding of administrative metadata at the procedural level and did not support encoding of deletion policy information. A possible explanation can be the differences in the work organization models used by the Flickr Groups and by the traditional information providers (libraries, museums, and archives) to which schemas like the DCCAP are tailored. The Flickr Group work is more informal, and there is much greater

uncertainty and diversity in all of its components: information providers, metadata practices, types, and the quality and regularity of contributions. For traditional collections, most of these relationships are predefined, precoordinated, and relatively stable. They often are created by a single institution or a consortium of institutions, which may share work organization standards and best practices. These standards not only reduce some of the variances and uncertainties of collection building and metadata creation, but they often constitute a shared knowledge to the members of community and/or are maintained and specified by a designated institution (e.g. Library of Congress). Hence, collection level metadata records in traditional databases might not include the detailed specifications of the standards they conform to. With a growing number of information resources and collections created through less formal social content-creation models, it becomes important that collection metadata schemas such as the DCCAP are able to support encoding of more granular provider, procedural, and annotation information.

Conclusion

This study's exploration of user generated collection-level metadata in a social tagging system contributes to a better understanding of user-generated collections and the need in these collections for various types of collection level metadata. Its results inform research into and the design of metadata schemas and help to improve organization and access to images in both traditional databases and emerging social contexts of image description and use

The study has certain limitations. The study only looked at images within a specific photo collection and sharing site; therefore these results may apply only to similar collections of visual media with similar functionalities. The study only utilized English language samples and any implications of the research is limited to English speaking populations. Research for the future could include expanding the research to other languages as well. The types of collections and the user motivations for creating collections were identified by analyzing the content of collection descriptions. Since these descriptions were not originally created for such purposes, the identified types are only approximations of the true types and intentions, and potentially could be influenced by the coder's bias. Congruence between the current research results and results of prior research points to the validity of the work, however further study is needed, and multiple coders would increase the reliability of the results, as would study of multiple collections. Flickr has continued its rapid rate of growth, therefore the results of this study would be strengthened by a more recent, larger sample.

The study found that users would not usually tag individual photos and the Photoset or Group metadata were often the only metadata associated with those photos. Also, it was observed that the Groups built around 'active' spaces (e.g critiquing members' work or playing games) had significantly longer descriptions than those sharing photographs of persons or things. These invite a further investigation to determine if Group type can influence the amount of user contributions and the type and quality of item level metadata in Flickr.

The same image can be indexed by many different attributes, and concepts perceived in its content, and, thus can be relevant to many different subject or activity based collections. Having comprehensive and valid subject and activity concept trees or taxonomies is essential for both 'on demand' generation of a collection and better utilization of images. Future work includes an analysis of item level tags in the Photoset and Group collections, and their comparison with the collection types and metadata schemas that accommodate these tags, as well as with earlier

research results on item-level description. The goal is to build mappings between the collection types and item-level subject metadata that will be grounded in empirical data and that can be used in automatic classification and collection building.

References

Ames, M., Naaman, M. (2007). Why we tag: motivations for annotation in mobile and online media. In: B. Begole, S. Payne (Eds.), *Proceedings of the SIGCHI.* (pp. 971-980). New York, NY: ACM.

Bailey, K. (1994a). *Methods of social research* (4th ed.). New York: The Free Press.

Bailey, K. (1994b). Typologies and taxonomies: An introduction to classification techniques. *Sage University Papers Series. Quantitative applications in the social sciences.* Thousand Oaks, CA: Sage.

Balabanovic, M., Chu, L., & Wolf, G. (2000). Storytelling with digital photographs. In T. Turner (Ed.), *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 564–571). New York: ACM.

Brunskill, J. and Jörgensen, C. (2002). Image Attributes: A Study of Scientific Diagrams. Proceedings, Annual Meeting of the American Society For Information Science 39 (1), 365–375.

Buckland, M. (1992). Collections reconsidered. In *Redesigning library services: A manifesto* (pp. 54–61). Chicago: American Library Association.

Buswell, G. T. (1935). *How people look at pictures*. Chicago: University of Chicago Press.

Caplan, P. (2003). *Metadata fundamentals for all librarians.* Chicago, IL: American Library Association.

Cover, T. & Thomas, J. (1991). *Elements of Information Theory.* New York, NY: Wiley.

Covi, L., & Cragin, M. (2004). Reconfiguring control in library collection development: A conceptual framework for assessing the shift toward electronic collections. *Journal of the American Society for Information Science and Technology, 55*, 312–325.

Crofts, N., Doerr, M., Gill, T., Stead, S., & Stiff, M. (Eds.). (2007). *Definition of the CIDOC conceptual reference model.* Retrieved November 18, 2007, from http://cidoc.ics.forth.gr/docs/cidoc_crm_version_4.2.2.pdf

Cunningham, S., Masoodian, M. (2006). Looking for a picture: an analysis of everyday image information searching. In: G. Marchionini, M. Nelson (Eds.), *Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries,* (pp. 198-199). New York, NY: ACM.

Cunningham, S., Masoodian, M. (2007). Metadata and organizational structures in personal photograph digital libraries. In: D. Goh, T. Cao, I. Sølvberg, E. Rasmussen (Eds.), *Proceedings of ICADL 2007, The 10th International Conference on Asian Digital Libraries,* (pp. 458-467). Berlin, Germany: Springer.

Frohlich, D., Kuchinsky, A., Pering, C., Don, A., & Ariss, S. (2002). Requirements for photoware. In *Proceedings of the 2002 ACM Conference on Computer-Supported Cooperative Work* (pp. 166–175). New York: ACM.

Gabrilovich, E., Dumais, S., Horvitz, E. (2004). Newsjunkie: providing personalized newsfeeds via analysis of information novelty. In: S. Feldman, & M. Uretsky (Eds.), *Proceedings of the 13th international conference on World Wide Web*, (pp. 482-490). New York, NY: ACM.

Golder, S., and Huberman, B. A. (2006). Usage Patterns of Collaborative Tagging Systems. *Journal of Information Science, 32(2).* 198-208.

Gordon, A. (2001). Browsing image collections with representations of common-sense activities. *Journal of the American Society for Information Science and Technology, 52*, 925–929.

Guy, M. & Tonkin, E. (2006). Folksonomies: Tidying up tags. *D-Lib Magazine, 12*(1). Retrieved June 27, 2006, from http://www.dlib.org/dlib/january06/guy/01guy.html.

Gyford, P. (2005, 23 March). I'll show you mine. . . . *bbc.com*. Retrieved January 13, 2007, from http://news.bbc.co.uk/1/hi/magazine/4374971.stm

Hill, L., Janée, G., Dolin, R., Frew, J., & Larsgaard, M. (1999). Collection metadata solutions for digital library applications. *Journal of the American Society for Information Science, 50*, 1169–1181.

Jörgensen, C. (1995). *Image attributes: An investigation.* Ph.D. Thesis, Syracuse University, Syracuse, NY.

Jörgensen, C. (1996). Testing an Image Description Template. *Proceedings of the 59th Annual Meeting of the American Society for Information Science*, Baltimore MD, Oct. 21–24, 209–213

Jörgensen, C. (1999). Retrieving the Unretrievable: Art, Aesthetics, and Emotion in Image Retrieval Systems. *Proceedings SPIE* (International Society for Optical Engineering) 3644, 348–355, Human Vision and Electronic Imaging IV, Bernice E. Rogowitz, Thrasyvoulos N. Pappas; eds.

Jörgensen, C., & Jörgensen, P. (2005). Image querying by image professionals. *Journal of the American Society for Information Science and Technology, 56*, 1346–1359.

Kuchinskas, S. (2005, 15 June). Flickr to add print to photo service. *internetnews.com.* Retrieved January 13, 2007, from http://www.internetnews.com/bus-news/print.php/3512866.

Lee, H. (2000). What is a collection? *Journal of the American Society for Information Science, 51*, 1106–1113.

Lesk, M. (1997). *Practical digital libraries: Books, bytes and bucks.* San Francisco: Morgan Kaufman.

Palmer, K., Knutson, E., Twidale, M., & Zavalina, O. (2006). Collection definition in federated digital resource development. In A. Grove (Ed.), *Proceedings 69th Annual Meeting of the American Society for Information Science and Technology (ASIST) 43.* Retrieved December 9, 2007, from http://imlsdcc.grainger.uiuc.edu/docs/ASIST06FinalSubmittedVersion.pdf.

Rasmussen, E. M. (1997). Indexing images. In M. E. Williams (Ed.), *Annual review of information science and technology* (Vol. 32, pp. 169–196). Medford, NJ: Learned Information.

Rodden, K., & Wood, K. (2003). How do people manage their digital photographs? In V. Bellotti (Ed.), *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 409–416). New York: ACM.

Shreeves, S., Knutson, E., Stvilia, B., Palmer, C., Twidale, M., & Cole, T. (2005). Is "quality" metadata "shareable" metadata? The implications of local metadata practices for federated collections. In H. A. Thompson (Ed.), *Proceedings of the Twelfth National Conference of the*

Association of College and Research Libraries (pp. 223–237). Chicago: Association of College and Research Libraries.

Smeulders, A., Gupta, A. & Jain, R. (2000). Content based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *22*, 1349-1380.

Stvilia, B., Gasser, L., Twidale, M., Shreeves, S., & Cole, T. (2004). Metadata quality for federated collections. In S. Chengulur-Smith, L. Raschid, J. Long, C. Seko (Eds.), *Proceedings of the International Conference on Information Quality—ICIQ 2004* (pp. 111–125). Cambridge, MA: MIT.

Taylor, A. (2006). *Introduction to cataloging and classification*. Westport, CT: Libraries Unlimited.

The Library of Congress (2008). *On the record: report of the library of congress working group on the future of bibliographic control.* Retrieved April 21, 2008, from http://www.loc.gov/bibliographic-future/news/lcwg-ontherecord-jan08-final.pdf.

The VRA Core 4.0. (2007). The VRA Core 4.0—A data standard for the cultural heritage community. Retrieved 7 December, 2007, from http://www.vraweb.org/projects/vracore4/index.html

Yarbus, A. L. (1967). *Eye movements and vision*. New York: Plenum Press.