Investigating the use of Photo Collection Structures for Photo Searching

Dynal Patel
Department of Computer Science,
University of Cape Town,
Private Bag Rondebosch 7701,
Cape Town,
South Africa
dpatel@cs.uct.ac.za

Gary Marsden
Department of Computer Science,
University of Cape Town,
Private Bag Rondebosch 7701,
Cape Town,
South Africa
gaz@cs.uct.ac.za

ABSTRACT
In this paper, we investigate the use of photo collection structures in aiding the rapid retrieval of *events, single* photographs and sets of photographs sharing a common *property*. As these structures encode a great deal of useful contextual information, we advocate for the need for photo searching interfaces to exploit and expose this information. In our study of people's photo collection structures, we found that people organise their photographs into event folders. Providing rapid access to *events*, *singles* and *properties* equates to locating event folders quickly. When event names are well known, we advocate for the use keyword based searches. Temporal based navigation becomes increasing important when event folders are less well known. We found a significant amount of data showing that people do organise and structure their photo collections more than previous literature suggests. The number of different *property* folders we found, illustrates the range of different tasks people perform when structuring and organizing their photo collections. In concluding this paper, we also make a number of recommendations for photo searching interfaces.

Keywords
Photographs, Searching, Photo Collection Structures

INTRODUCTION
As the proliferation of mobile devices with cameras continues, digital photo searching and browsing is becoming an increasingly important research topic.

One of the major reasons why this problem has proved difficult to solve is because of the lack of meaningful and consistent textual annotations for each photograph. This has rendered conventional text-based information retrieval methods useless in photographic collections. Although some scholars might argue that a simple solution is to manually annotate each photograph, research has shown that people are put off by the huge effort that is required to explicitly annotate each photograph [10].

To overcome these issues, researchers have proposed a variety of techniques that use alternative meta-data, such as pixel or EXIF data, which is readily available for each image. Rodden [11] delineates this image content into four categories. These are: primitive level (the colours, textures, composition and simple shapes present); generic level (the types of objects and activities depicted); specific level (the names of people, places, landmarks and events); abstract level (involves interpreting the overall meaning of the image, such as the mood or feeling it evokes).

Unfortunately, the majority of research in this area does not consider the fact that a user might specify a request, or have a requirement, at any one of these levels. In fact, a study by Rodden [12] found that that people are most likely to look for: an event (set of photos relating to a particular well-defined event); a single (individual photograph); and a property (set of photographs that have a common property) in decreasing order of frequency. Also, these search requests are likely to require image content at each of the previous four levels.

Previous research [1][2][3][9][11] in this area has concentrated on one or two of these levels. However, it is clear that to cater for these photo searching requests, a more all-encompassing approach is required. Such an approach should utilize as much content and context as possible to allow users to formulate and conduct these divergent queries.

One important source of semantic data that has been over-looked when searching for photographs, is the photo collection structure. Within the photo collection's structure on a hard disk, people encode a vast amount of data that often crosscuts all four categories in Roddens' taxonomy. For example: primitive level (people may create a grouping that contains photographs with a similar artistic property); generic level (people may distinguish between different arbitrary pictures by placing them in generic categories such as: edibles; objects; people; and places); specific level (people often give events meaningful names such "21st b-day party"); and abstract level (The mood or feelings can usually be inferred from the event labels. The mood or feelings can sometimes be made more explicit by the use of adjectives or adverbs in the event label, for example "romantic dinner").

In the context of our multi-level approach, this paper is a first step that investigates the value of information encoded in photo collection structures in the context of the three searching tasks. It also discusses how these findings can inform the design of the photo browsing interface.

We begin by considering the background and related literature in this field. We then present our research question which we then test by experimental procedure and report our findings. Lastly, we analyse our results and draw conclusions.

BACKGROUND

In an initial study, Rodden and Wood looked at how people manage their print and digital photographs [10][11][12]. They found the organization of print photographs to be a low priority task. Prints were normally kept in their original packet for long periods of time. When users eventually get round to organizing them, they would usually place the most appealing prints into an album. Non-appealing prints would either placed in the original packet or thrown away. Normally, albums were classified by events, often with one album per event. Within the albums, print photographs were usually placed in chronological order, with only a few changes in the order for aesthetic effects. Albums were normally kept in chronological order. However, this ordering would usually be lost when photographs were browsed.

With digital photographs, Rodden and Wood found that users would normally organize their photographs into event folders. Each event folder was given a meaningful name. Within folders, photographs were ordered chronologically using the time data from the EXIF header the camera automatically stored with each image. They found that photographs were rarely annotated as users found the task too time consuming.

Although the work by Rodden and Wood provides some initial insights, no researchers have investigated how the semantic data encoded in a photo collection's structure, organization and composition can be used to find events, singles and properties quickly. We therefore set out to investigate this topic. We focus on identifying patterns, similarities and differences between participant photo collection structures. We distance ourselves from the two extremes of photo searching interfaces that are prevalent in current literature: annotation-based interfaces [4][5] and feature extraction-based interfaces [2][6][9][11][14]. While annotation-based interfaces can be quite powerful for non-personal photo collections [7], research has shown that they are not ideal for personal collections as the majority of people do not annotate their photographs. Although feature-based interfaces work well where a user wants to extracts low level feature such as color it becomes increasing difficult to specify and conduct high level searches. As Rodden [11] points out, "meaningful feature identification and interpretation at the generic level is an AI-complete problem. Finding a solution requires solving strong artificial intelligence in general." In our research, we make use of any semantic data the users are willing to add. We also extract any image content that is readily available. However, we only use image content that does not rely on strong AI to be able to extract, identify and interpret.

Lastly, research in photography is not limited to photo searching. A number of researchers have investigated other stages in the photo management cycle. For brevity, we only mention classification schemes that encompass the bulk of this research. A classification scheme proposed by Schneiderman [13] delineates research in digital photography according to an activities relationship table. Another by Frohlich et al [1] uses a time-space matrix to characterize various photoware technologies.

RESEARCH QUESTION

For sake of clarity, we begin by defining the following terms:

- *Event*: A burst of photographs that encompasses one activity (i.e. it cannot contain any sub-events)
- *Event folder*: A folder that encompasses one or more events. For example in Figure 1, "Knysna", "Balloon Festival" and "Raglan" are event folders.
- *Event grouping folders*: A folder that encompasses one or more event folders. For example, in Figure 1, "2004 Pics" and "New Zealand Trip" are event grouping folders.
- *Event directory count*: For a photo collection, the event directory count is made up of the event folder count and the event grouping folder count.
- *Folder level*: The current nesting depth in a folder hierarchy, where folder level 1 refers to content in the root directory and folder level 2 refers to content in a level 1 folder. For example, in Figure 1, "New Zealand Trip" is in folder level 1, "Raglan" is in folder level 2 and "James Camera" is in folder level 3.

- *Folder name component*: A basic element used in a folder name. In our sample data we identified 5 elements: year (e.g. 2005), month (e.g. 07, July), day (e.g. 21st), a non-date specific number (e.g. 1, num-1, part1) and a name or descriptive non-numeric or non-date character sequence (e.g. Waterfront). For example, in Figure 1, "21st Bday Party" and "2004 pics" each consist of two folder name components: day-name and year-name.
- *Folder naming scheme*: A unique combination of one or more folder name components. For example CHI2005, 2005CHI, and 2005_CHI are three different naming schemes. In Figure 1, there are three naming schemes: day-name (e.g. "21st Bday Party"); year-name (e.g. "2004 pics") and name (e.g. "Hells Gate").
- *Imported event folder*: A special case event folder that contains one or more photographs from one or more members of an individual's social network. For example, in Figure 1, "21st Bday Party" and "Raglan" are imported event folders as they contain photographs from other people.
- *Property folder*: A folder that stores a set of photographs that share a common property or theme.
- *Roll*: An unorganized batch of photographs that is imported from a picture taking device.

Figure 1: An example photo collection structure

Our research question is:

Are there any patterns, similarities or differences in the structure, organization or composition of photo collections that can be used to aid the rapid retrieval of events, singles and properties?

METHOD

Twelve participants (5 females and 7 males) volunteered to take part in the study. Nine participants were postgraduate university students, while 3 were university lecturers. All participants had an undergraduate degree. While this sample is unlikely to represent the population as a whole, it allows us to gain some initial insights into our research questions.

The study was carried out over a one-month period. One the first day, participants were asked to fill out an initial questionnaire. This was used to collect demographic data. They were also requested to burn their entire photo collection on to a DVD. Then over the course of the month they were then instructed to log how photographs were imported into their photo collection. After the one month period we conducted informal, semi-structured, interviews.

In order to investigate our question we extracted data from each participant's photo collection. We extracted: the photo count, event folder count, event grouping folder count, folder name component counts, event folder naming scheme count, imported event folder count, property folder count, minimum event folder size, maximum event folder size, mean event folder size and the photo collection range in months. Where applicable, we extracted the data for each of these variables over 5 folder levels. This data was used to drive the interviews and to clarify some of the observations we had made.

In order to investigate our research question, analysis of variance (ANOVA) tests were used extensively. Before conducting any ANOVA's we tested the assumptions of normality and homogeneity of variances to ensure they were not violated. Where the ANOVA assumptions were violated, a non-parametric alternative test was used. The Wilcoxons Matched Pairs Test was used to test significant difference between two samples, whereas the Friedman ANOVA & Kendall's concordance was used for multiple samples. While these tests are less powerful than the normal ANOVA, they do not rely on any very serious restrictive assumptions concerning the shape of the distributions. Furthermore, when any assumptions are significantly violated the Wilcoxons Matched Pairs Test or the Friedman ANOVA is likely to yield more accurate results than an ANOVA.

RESULTS

In our data set, participants had varying numbers of photographs (mean 1489.08; s.d. 1224.15; min 419; and max 4317). As a result, to make meaningful comparisons between participants we always normalized the data.

In this section we present each of our major findings. For each finding we describe how it was investigated using the statistical methods above. Where possible we also include data from our semi-structured interviews.

Event folder distribution
For each participant, we counted the number of event folders at each folder level. We then normalized these values by dividing the event folder count at each level by the total event folder count. For example, for one participant, the distribution of their event folders expressed as a proportion of the total event folder count was: 0.0 at level 1, 0.65 at level 2, 0.30 at level 3, 0.05 at level 4 and 0.0 at level 5. Figure 2 below shows this distribution for 12 photo collections. For each level it shows the mean proportion as well as the variation between 12 participants.

There was a significant difference in the event folder count between the 5 levels (CS = 13.08696 p = .00144). On closer inspection, we found that this because there were significantly more event folders at levels 2 (p=0.01) and 3 (p=0.5) than level 1. There was no significant difference between levels 2 and 3. In other words, event folders were just as likely to be found at level 2 or 3.

Figure 2. The distribution of event folders across 5 folders for 12 photo collections.

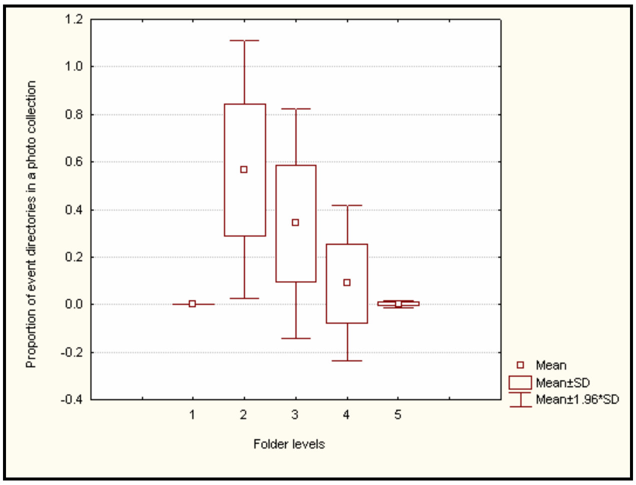

Event grouping folder distribution
For each participant we looked for evidence that shows some organization took place to help them locate events faster. In particular we looked for evidence of event grouping folders (see definition above). Eleven out of twelve photo collections contained event grouping folders. Across all participants we found that, on average, 17.29 % (s.d. 14.53) of the event directory count consists of event grouping folders.

There was significant difference in the event grouping folder count between level 2 to 5 (CS = 12.50000 p = .00193). Post-hoc tests found that there are significantly more event grouping folders at levels 2 and 3 than at levels 4 and 5 at the 0.05 level. There was no significant difference between levels 2 and 3.

Participants stated that organizing photographs in an event hierarchy using event grouping folders had four major benefits. Firstly, it enabled them to retrieve events more quickly. Secondly, it reduced the cognitive effort required to remember where an event folder was positioned relative to others. Thirdly, it provided a nice way to encode relationships between events (main events and sub-events). Fourthly, it provided a way of encapsulating a group of events so that related photographs were stored in one place.

Relationships between events and event folders
We processed each participant's photo collection and counted the number of events (see definition above) and event folders. Events were detected by looking for changes in activities in the event folders. For example, when looking through a "birthday" event folder, if there were dining photographs, followed by dancing photographs, this would be considered as two separate provisional activities. To label these activities as distinct events, the experimenter would check the photo capture times to look for a distinct jump between the activities. To reduce the bias only one experimenter was used to process the entire data set. The justification for this was to minimize the number of errors would have arisen from the automatic detection of activities.

The scatterplot in Figure 3 illustrates the relationship between the event folder count and the event count. There was a high correlation between events and event folders (r = 0.93 at the 95% confidence interval). Across all participants the ratio of events to event folders was 2.41 (s.d. 1.12). In other words, event folders are likely to contain more than one event. Furthermore, there were significantly more events than event folders. Our definition of an event is not the same as what participants classify as events. If the two definitions were identical, we would expect no significant difference between the two.

Figure 3.Scatterplot illustrating the correlation between the event folder count and the event count (12 participants).

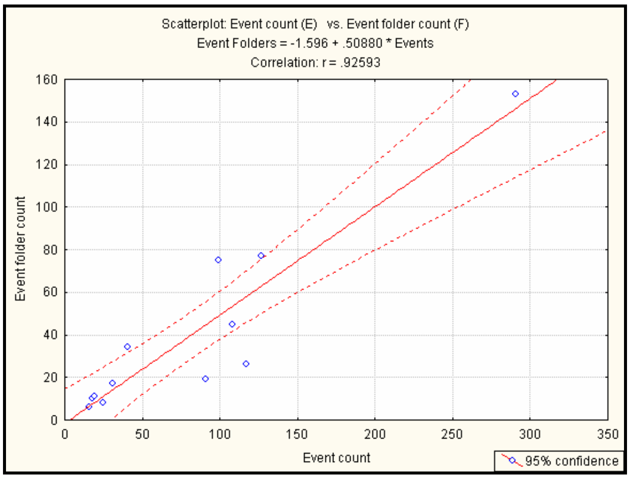

Photographs from a camera would normally be downloaded into a temporary folder. At some point the photographs would then be sorted into event folders. Sometimes photographs were placed into pre-existing event folders. For example, this would happen if an event got split between memory cards. Sorting photographs into folders was normally done visually and each event folder would be given a unique name.

Event folder date component distribution

For each participant, we counted how many event folders at each level had one or more date components in the label. We then normalized the values by dividing the count at each folder level by the total event folder count.

Between folder levels, there was a significant difference in the proportion of the total event folder count that had a date components in the label (CS = 12.62069 p = .00182). We found that the proportion at level 2 was significantly higher than at levels 3 (p=0.05), 4 (p=0.1) and 5 (p=0.1).

We decided to investigate further by looking at the frequency of event folder name components. For each participant, we counted the frequency of each event folder name component at folder levels 2 to 5. We normalized these values by dividing them by the event folder count at each level. Using the Wilcoxons test, we found that the proportion of the total year and month components count at level 2 was significantly greater than level 3 and 4 at the 0.05 level. There were no other significant differences.

We also found that within each folder level, the proportion of the total name component count used in the event folder labels was significantly greater than the proportion of each of the other components (year, month, day, number) at the 0.05 level. Also, at level 2, the proportion of the total year component count used in event folder labels was significantly greater than the proportion of each of the other date components, month (p=0.05) and day (p=0.05). There were no other significant differences.

From the interviews we found that dates were used in event labels for two main reasons. Firstly dates were used when they were relevant to the event, for example “2005 Trip to New Zealand” or “4 July Celebrations.” Secondly dates were used for temporal navigation. Within a folder level they would be used to create a sequential timeline. In folder level 2 dates were mostly used to provide an overview.

Event folder naming schemes

We counted the number of times each naming scheme was used in a photo collection. We then normalized these values by dividing the count by the total event directory count. For each participant, we extracted the most frequently used naming scheme. Across all participants we found that the most frequently used naming scheme is used 78.18% (s.d. 16.59) of the time.

We then looked at how consistently naming schemes were used across folder levels. To investigate this we looked at how consistently the most frequently used naming schemes were used between folder levels 2 to 5. Using Friedman ANOVA, we found that there was no significant difference in the proportion of the event folder count for each level that used the most frequently used naming scheme. In other words, the most frequently used naming scheme was used consistently between these levels.

To get some insight into which naming schemes were used most often, we counted how the number of times each naming schemes found in our data set was used by the participant. We then normalized the values by dividing the naming scheme count by the event directory count.

Table 1 shows the basic statistics we obtained for the 12 naming schemes found in our data set, where n = name, d = day, m = month, y = year, num = number). For example, if an event folder was named using the following “y-m-d-n” scheme, then its label would consist of the year, followed by the month, then the day and name.

Table 1: Basic statistics for the 12 naming schemes obtained in our data set. (num = number, n = name, d = day, m = month, y = year)

Naming Schemes	Mean	Min	Max	Std.Dev.
n	0.74	0.06	1.00	0.26
n-y	0.06	0.00	0.33	0.09
num-n	0.05	0.00	0.34	0.10
y-m-n	0.05	0.00	0.60	0.17
n-num	0.04	0.00	0.27	0.08
y-m-d-n	0.03	0.00	0.32	0.09
y-n	0.01	0.00	0.05	0.02
y-m-d	0.01	0.00	0.05	0.02
d-m-y	0.01	0.00	0.09	0.03
n-y-num	0.01	0.00	0.16	0.05
y	0.00	0.00	0.03	0.01
m-d	0.00	0.00	0.01	0.00

Using Wilcoxons Matched Pairs Tests we found that a simple “name” naming scheme was used significantly more than any other naming scheme at the p = 0.01 level. Other interesting results showed that a “name-year” naming scheme was used significantly more than the “year-name”, “day-month-year”, “year” and ”month-day” schemes at the p=0.05 level. Similarly, the “name-number” scheme was used significantly more than “d-m-y”, "y", "m-d" and the "d-m-y" scheme at the p=0.05 level.

Event folder sizes

For each participant, we ordered event folders according to the number of photographs contained in them, with the largest event folder being at the top. We then obtained the largest and smallest event folder sizes by obtaining values from the top and bottom of this listing. We also calculated the average event folder size for each participant. Table 2 shows the values we obtained.

Table 2. Descriptive statistics for the smallest, average and largest event folder sizes across 12 participants

	Mean	Min	Max	Std.Dev.
Smallest	3.08	1.00	7.00	1.93
Average	29.96	17.57	47.83	8.22
Largest	142.50	53.00	379.00	96.65

The small value of standard deviation relative to the mean tells us that average event folder size is fairly consistent. This result is surprising given the fact that users organize events into event folders, so one might expect events sizes to be more variant. The sizes for the smallest and largest event folder sizes are also not as divergent as one might think.

Imported event folders

For each participant, we counted the number of imported event folders (see definition above). We normalized this value by dividing the count by the total event folder count. Surprisingly, we found that across all participants, approximately 32% (s.d. 22.96) of all event folders were imported event folders.

When participants obtained photographs from other people, what they did depended on whether or not they attended the event. If they attended the event, then they would normally sort through the photographs and only keep the ones they found the most interesting, different or did not have. Some participants would then place these photographs in a sub-folder labeled with the donor's name. These participants were usually proud of their photographs for various reasons, such as their artistic quality, or simply they didn't want their sequence of photographs to be broken. Placing foreign photographs in a separate folder provided a way of keeping their photographs apart. Other participants simply merged the photographs. As one participant puts it, "merging photos means I don't have the hassle of going through lots of subfolders to find photographs; also I don't really feel the need to differentiate my photos." When participants received photographs from an event they didn't attend, the photographs became a lot more disposable. The rules for keeping these photographs were much more complex and needed a lot of context in the person's life. These pictures were usually stored in a property folder (see definition above). As one participant put it, "Generally speaking, these photographs do not make it into my collection. This is because I don't like to share other people's photographs. I feel that there needs to be a long term interest. The ones that I keep normally have some sentimental value. These selected few will normally be placed in a folder like 'Arb' as they are hard to classify."

Event folder creation over time

We correlated the photo collection range (most recent photographs minus the oldest photograph) against the total event directory count. Figure 4 illustrates the relationship between the two factors.

We found that there was a positive correlation (r= 0.66 at the 95% confidence interval) between the factors. So, the greater the photo collection range, the greater the total event directory count.

Similarly, we correlated the imported event count against the photo collection range. Again, we found that there is there is a positive correlation (r= 0.77 at the 95% confidence interval) between these two factors. The greater the photo collection's range, the more imported events there are.

Figure 4. Scatterplot illustrating the relationship between the total event directory count and the photo collection range (in months)

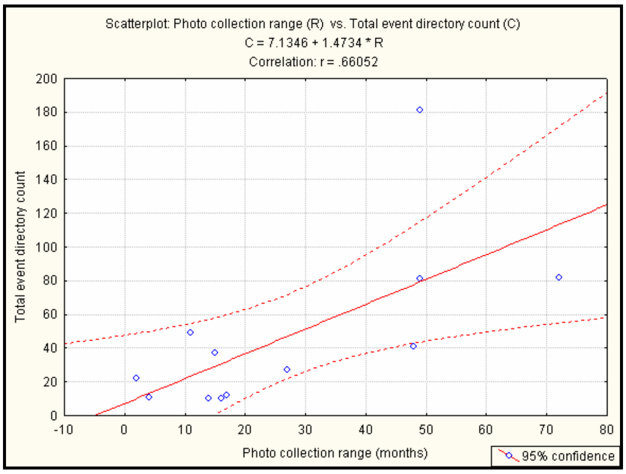

Distribution of single photographs

For each participant, we counted the number of photographs at each folder level. We then normalized these values by dividing the count by the total number of photographs. We found that there were significantly more photographs at folder levels 2 and 3 than levels 1, 4 and 5 at the 0.05 level. Also, the number of photographs at level 4 was significantly greater than the numbers at levels 1 and 5 at the 0.05 level. There were no other significant differences. Therefore, the majority of photographs were just as likely to be found at levels 2 or 3.

Shortcuts for finding single photographs

We searched through each participant's photographs looking for evidence of some organization that would help retrieve images quickly. We found that 2 participants created shortcuts in the root directory to individual photographs that were stored much deeper in the hierarchy. When asked why they didn't copy the image to the root directory, both participants responded by saying that didn't want any duplicate copies of images in their collection. Unfortunately, we were not able to collect any data on any shortcuts that were placed outside a photo collection. Through the interviews we were able to find that it is actually quite a common strategy to copy images that are accessed frequently to a destination that is easy to locate. In some cases the images were placed in the root directory. In other cases a copy of the image would be placed outside the photo collection, for example on the desktop. The rationale is that this strategy was much quicker than navigating through the folder structures. We found that 6 of the 12 participants placed photographs in their root directory. However, it is important to note that this was not the only reason why images were placed in the root directory. Sometimes images that were hard to classify into any folder were simply placed in the root directory.

Photo annotation

We found that 4 out of 12 of the participants annotated some of their photographs. However, for each participant this number was less than 1% of total collection.

Participants stated that they usually name photographs when emailing or posting them on a website or just to help them remember unique names of places or people. The general consensus was that too much effort was required to annotate photographs. Participants felt that a better way to distinguish images was to simply look at them.

Property folders

We searched through each participant's photographs looking for evidence of some organization that would help retrieve properties quickly. We found that users create property folders (see definition above). In fact, across all participants we found that 9.41 % (5.78) of all folders are property folders. By creating property folders, participants were able to process future requests for the same property much more rapidly.

We found five different types of property folders. The first type (A) of folder was for arbitrary photographs that were hard to classify into any event folder or well defined property folder. The second group (B) of property folders would be a specially organized group of photographs that were to be exported from the collection. In some cases these folders would only contain an event summary where only the "best" pictures from an event folder would make it into the property folder. Other times the pictures would be taken from different events. The third type (C) of property folders were working directories or temporary stores. The common property in this case was the fact that all these photographs were in a transient state, waiting to be organized or structured in some way into either events or other types of property folders. The fourth type (D) of property folders were picture groupings. Typical examples from this category are pictures of family or scenic pictures taken from various holidays or pictures of pets. These property folders were created in one of three ways. Firstly, they could be created following a search request for some property. Secondly, they could be created by giving the second type of property folder a more descriptive property label, instead of one that details to whom and why the photographs are being exported. Thirdly, they could be created when a common theme emerges in the first type of property folder. The fifth type (E) of property folders were property event grouping folders. For example, all birthday event folders were organized into a single folder titled "Birthdays." The important distinction with event grouping folders is that the encompassing folder is the main event and the sub folders are related sub-events.

Table 3: The proportion of each property group in a photo collection using a sample of 12 participants

	Mean	Std.Dev.
D	0.50	0.32
E	0.27	0.29
C	0.12	0.19
B	0.07	0.11
A	0.04	0.07

For each participant we counted how the number of times each property group appeared. We then normalized these values for each participant by dividing the count for each property group by the total count for all property groups. Table 3 shows the mean and standard deviation across the 12 participants.

Using Wilcoxons, we found that there were significantly more type D property folders than types A, B and C at the 0.05 level. We also found that there were significantly more type E property folders than types A and B at the 0.05 level. There were no other significant differences.

In our data set, we also found that property folders only have a label component in the folder names. The absence of date components may suggest that the date ordering of properties is less useful. As one participant suggests, "I have never thought of my folders in terms of properties and events. Maybe I make this distinction sub-consciously … when searching an event I usually order my folders by date, whereas with properties I would probably do an alphabetical listing and then search for the label. Within the folder I look through the pictures in date order." We correlated the total property folder count against the photo collection range and found that there is a positive correlation (r= 0.73 at the 95% confidence interval) between the total property folder count and the photo collection period. So the larger the photo collection range the more property folders there are.

DISCUSSION

From our results, it is clear that a substantial amount of information is encoded in the structure of photo collections and that this information can be exploited to aid the fast retrieval of events, photographs and properties. In this section, we discuss the implications of these findings for future photo searching interfaces.

As users organize photographs into event folders, providing rapid access to events equates to providing rapid access to these event folders. However, with singles or properties this is not as straight forward. When pre-organized structures (shortcuts or property folders) are in place, rapid access to photographs or properties involves providing quick access to these structures. However, when no pre-organized structures are available, participants are required to rapidly traverse event folders looking for the matching photographs. In this case, the rapid retrieval of singles and properties equates to providing rapid access to event folders. Therefore, we deduce that the corner-stone of any photo searching interface should be to provide rapid access to event folders. Rapid access to shortcuts or property folders is essential when similar searches are likely to be repeated.

Event grouping folders were one of the mechanisms used to allow fast access to event folders. Within the event hierarchy, events were given more generic names that served as an overview for sub-events. Sub-events were given more descriptive names to allow users to distinguish them. The use of a hierarchical scheme resulted in few folders at the top. It was also used to encode relationships between main events and sub-events. Both these factors reduced the mental workload required to remember event names and also where events were positioned in relation to each other. Selecting one of these event grouping folders allowed you to filter off irrelevant information. In essence, what we observed and have described here is a hard-coded, less fluid version of Schneidermans' mantra (Overview first, filter and details on demand) [8]. Photo browsing interfaces should seek to exploit these hard-coded structures and should make the transitions as fluid as possible.

Participants had a preconceived idea that organizing photographs in a hierarchy, as opposed to a single list, would allow them to retrieve events quickly. Research has shown [3][9] that this is not necessarily true. In these examples, a hierarchical photo browser was out-performed by a thumbnail photo browser where the pictures were organized in a flat scrollable list, ordered by the creation time. One hypothesis is that hierarchical organization in this case as opposed to a single listing, might reduce the mental workload required to locate photographs. Participants might then link this reduced mental effort with speed. Regardless of what the outcome of this hypothesis might be, as users have a good mental image of the structure of their photo collection, one should not destroy this structure, but rather look at ways of exploiting it and making it more accessible as an additional aid for navigation.

The results clearly show that at any folder level, the name component was the most used component. Therefore, photo browsing interfaces should support some form of keyword or letter searches for events. For example, one could implement a scheme similar to the one used on many mobile phones to locate contacts, where the contacts that appear are matched against each letter in the keyword. The list of potential contacts is filtered dynamically as the keyword is typed. One could also sort events alphabetically and provide an interface similar to the Apple iPod where the number of letters that appear is linked to the scroll speed. The assumption with these suggestions is that the user actually knows the name of the event they are looking for. If possible, these searching techniques should not break the structure of a photo collection, but should rather highlight appropriate sections creating a seamless browsing and searching interface.

We also observed that dates and time were important in naming. More date components were found in event folders at higher levels. We also found that the year component was used significantly more at level 2 than any other event folder name component. Dates were used for temporal navigation and to ensure the ordering of events was maintained. When the exact name or location of an event cannot be remembered, it is important to allow users to browse temporally.

Therefore, photo browsing interfaces should be both event and calendar driven. Both should seamlessly compliment each other to allow user to locate an event rapidly when its name is known or unknown. In the absence of any event information the calendar interface is the primary interface. For such a solution to work there are a number of challenges that must overcome. Temporal organization is likely to break event structures where an event folder contains photographs from a different camera, where the clocks are not synchronized. This is a major problem because of the number of imported events is very high. One simple fix is to ensure that the capture time for each photograph is adjusted so that it is not older than oldest photograph in the event folder or newer than the most recent photograph in the event folder. Another issue is that property folders may be broken up as often the pictures are taken from different event folders. A simple fix involves adjusting capture times for each photograph in a property folder so that it is not newer than the oldest photograph in the preceding folder or older than the newest photograph in the following event (assuming newest folders are placed at the beginning of the arrangement), where folders are ordered by creation date.

Event naming schemes are inconsistent. However, it is evident that people are willing to at least give each folder a meaningful name. Importing tools should be created to help users create more consistent and "richer" naming schemes that can further aid the rapid retrieval of events. These importing tools should allow users to visually sort photographs into event folders.

Within an event, people need to be able to sample the photographs to help them when performing searching task. From our results, users rarely annotate photographs and prefer to distinguish images visually. A photo searching interface should allow people to see as many pictures as is required to quickly complete the task. This number is unlikely to be static, so some provision should exist to dynamically adjust the number of pictures. Users should also be able to view the images at multiple zoom levels so that they can compare the images at the desired zoom level.

Almost a third of all event folders are imported event folders. This indicates that a substantial portion of events are co-experienced. This could be a further dimension that could be explored to differentiate between event folders. It is important to note that some individuals do not like to merge foreign photographs, and keep them separate by placing them in an imported event folder. In this case it is essential that the photo searching interface maintains this separation.

The number of imported event folders and property folders is positively correlated with time. This tells us that people do not give up organizing their photo collections. Any photo searching interface should be scalable so that it can accommodate any number of folders that are created. We also found that the event folder size is fairly consistent. Interestingly, the minimum and maximum event sizes are also fairly consistent. So, a photo searching interface must be able to cater for these extremities.

The majority of event folders are just as likely to be distributed at levels 2 and 3. This is important for caching algorithms as they need to cache horizontally and vertically. This will ensure that all relevant information is made available in the shortest amount of time.

Perhaps our most surprising discovery was the number of property folders in a photo collection. This evidence suggests that people do actually perform more organization with digital photographs than previous literature suggests. The diversity of property folders is a further testimony to this. Our two most important findings concerning property folders are that the names only have a label component and that the date ordering of these folders is not particularly important. Both of these finding have implications on finding properties rapidly. When searching for property folders, users rely on the name, hence a photo searching interface must allow searches and other organizational schemes such as alphabetical ordering. Some property folders that are created are the result of a search request. It is therefore essential that a photo searching interface continues to support this. Searches should be tracked so that dynamic lists can be created such as the showing the most frequently browsed properties, events or photographs. Again the photo structure should not be broken but rather highlighted. For example, event folders could be assigned a color that shows how frequently the event has been visited. This color could fade over time, providing users with an extra navigational aid. However, what is important is the notion of preserving the structure of a photo collection and providing seamless searching and browsing that highlights important areas of the photo collection.

Lastly we found that some users created shortcuts with the goal being to access the data as rapidly as possible. This was essential for repetitive tasks. This functionality should also be available in a photo browsing interface.

CONCLUSION

We found that a significant amount of data is encoded in a photo collection structure that can be used to aid the fast retrieval of events, singles and properties. We also found that by investigating the structure, composition and organization of photo collections we were able to get some insight into design principles for photo searching interfaces.

One of our findings was that providing fast access to events equates to providing fast access to event folders. This should be done through a seamless event and calendar based photo searching interface. As label components were the most prevalent in folder names, we suggest that text searching mechanisms be added to the interface to aid the fast retrieval of events. These searching mechanisms should not break the photo collection structure, but should be used to highlight the important areas or regions. Searching and browsing should be integrated seamlessly into the interface. The calendar part of the interface is particularly important when the location of events can not be remembered.

Providing fast access to properties or singles equates to providing fast access to events. Hence, a photo searching interface should primarily be designed around the rapid retrieval of event folders. When pre-defined structures, such as shortcuts, already exist to aid the fast retrieval of events, singles and properties, these should be made easily accessible.

We found evidence that suggests that people do organize their digital photo collections more than has been previously reported. This is supported by the fact that property folders make up a substantial portion of folders in a collection. The diversity of these property folders suggests that people perform a number of different tasks when structuring and organizing their photo collections.

ACKNOWLEDGMENTS

We would like to thank the Harry Crossley foundation for funding this research.

REFERENCES

1. Frohlich, D., Kuchinsky, A., Pering, C., Don, A and Ariss, S. Requirements for Photoware. *In Proc CSCW 02*, ACM Press, 2002.
2. Gargi, U., Deng, Y. and Tretter, D. R. Managing and Searching Personal Photo Collections. Technical report, Hewlett Packard. 2002.
3. Graham, A., Garcia-Molina, H., Paepcke, A., and Winograd, T. Time as Essence for Photo Browsing Through Personal Digital Collections. *In Proc JCDL*, ACM Press, 2002
4. Kang, H. and Schneiderman. Visualisation methods for personal photo collection: Browsing and searching in the PhotoFinder. *In Proc. IEEE Intl. Conf. on Multimedia and Expo*, 2000.
5. Kuchinsky, A., Pering, C., Creech, M. L., Freeze, D., Serra B., and Gwizdka, J. Fotofile: a consumer multimedia organisation and retrieval system. *In Proc. ACM CHI Conf., pages 496-503*, 1999.
6. Oh, H., Chang, D., and Oh, G. Content-based retrieval system for image using human face information. *In Proc of SPIE*, Storage and Retrieval for Media Databases, vol 3972, p. 12-20.
7. Markkula, M, and Sorumen, E. (1998) Searching for photos: Journalist's practices in pictorial IR. Proc. of CIR1998, First UK conference on Image Retrieval.
8. Plaisant, C., Carr, D. A., and Shneiderman. Image browsers: Taxonomy and design guidelines. *IEEE Software*, 12(2):21-32, 1995
9. Platt, J. C. and Czerwinski, M. PhotoTOC: Automatic Clustering for Browsing Personal Photographs, Microsoft Research Technical Report MSR-TR-2002-17.
10. Rodden, K. and Wood, K. (2003). How do people manage their digital photographs? *In Proc. of Human Factors in Computing Systems*, CHI 2003, 409-416, ACM Press.
11. Rodden, K. Evaluating similarity-based visualisations as interfaces for image browsing, Technical report, UCAM-CL-TR-543, ISSN 1476-2986, Univeristy of Cambridge, 2002.
12. Rodden, K. How do people organise their photographs. *In Proc BCS IRSG Colloquium*, Electronic workshops in computing, http://www.eic.org.uk, 1999.
13. Shneiderman, B. Leonardo's Laptop: Human Needs and the New Computing Technologies, MIT Press, 2002
14. Sun, Y., Zhang, H., Zang, L. and Li, M. MyPhotos- A System for Home Photo Management and Processing, Multimedia, ACM Press, 2002