

An Evaluation of Techniques for Image Searching and Browsing on Mobile Devices

Dynal Patel, Gary Marsden

Department of Computer Science
University of Cape Town
gaz@cs.uct.ac.za

Matt Jones

Department of Computer Science
Swansea University
always@acm.org

Steve Jones

Department of Computer
Science
Waikato University
stevej@cs.waikato.ac.nz

ABSTRACT

This paper reports on a system for browsing and searching image collections on small-screen devices. The system design was informed by our studies of how people organize and access image collections on desktop computers. The final system was evaluated in a user study where users had to search for images with varying degrees of precision about what they were searching for. We discovered that individual users adopt a wide variety of search strategies and that future image management tools must support users through a wide variety of interaction techniques.

Categories and Subject Descriptors

H.5.2 [User Interfaces]: User Centered Design

General Terms

User Centered Design

Keywords

Digital photography, SDAZ, searching, photo collections

INTRODUCTION

In our previous work, we reported on the development of photograph viewing techniques for small screen devices [1]. These techniques, which we called ManualZoom and AutoZoom, were evaluated in a large user experiment. The results showed that the new systems performed better than standard thumbnail browsers across a variety of tasks. To improve and refine the techniques, we embarked on an ethnographic study of how people stored their digital photographs and the information they used when searching for them on their PC [2].

In summary we discovered that locating events is central to all search behaviours. Moreover, the following factors are important

to people when searching for events:

- Being able to exploit existing folder structures (structures were well remembered)
- Searching for folder labels in that structure (even if they did not realise it, individual users are remarkably consistent in their naming conventions)
- Viewing images in terms of absolute dates (often users could map an event to a date; typically birthdays or anniversaries).

Once events have been located, then the AutoZoom and ManualZoom techniques discussed in [1] can be used to good effect to refine the search to an individual image.

What is unknown at this point, however, is how these various facilities might be integrated into a single tool and, if the facilities are all available, how they might be combined in typical searches? If users are able to use dates, events, names etc. to find an image, what is their relative importance and in what combinations are they typically used? Or, are the search strategies required for mobile devices so different that only a subset, or single, search facility is required.

To answer these questions, we developed a mobile application which allows users to combine these methods in any way they chose when searching for target images. The implementation of this tool is documented in the next section. We then evaluated the tool using naturalistic observation and ultimately concluded that users have vastly different search strategies and require a great deal of flexibility from any image search tool.

1. EXISTING SYSTEMS

Most mobile search systems support one, or more, of the facilities listed above, but does not do so in a seamless way. On most cellular handsets, the user simply has a thumbnail browser with images arrayed in a chronologically ordered grid. Despite the fact that many current handsets can hold upward of 1000 images, there is no mechanism for clustering photos into events or folders – the assumption is that the images will ultimately be uploaded to the Web or a PC. Other devices which are designed to work more closely with desktop systems (e.g. iPod and iPaq) allow folder structures from the desktop devices to be replicated on the mobile device. Taking the iPod as an example, users can search for photos by folder and then, once in a folder, can search photos via thumbnail browsing (either on a grid or by RSVP). However, users cannot search for a named folder nor can they enter an absolute date. Furthermore, once the user has navigated to a

folder's contents, they are not then free to navigate beyond the contents of a folder. For example, if a user wanted to find a holiday image taken on the bus between Paris and Lyon, they could start their search by navigating to the "Paris" folder and scroll to the end; if, however, the image was in the "Lyon" folder, they would need to navigate back up to the folder level and select "Lyon" rather than continuing to scroll chronological thumbnails. Whilst this notion of discrete folder contents may make sense for albums in iTunes (or files in the folders of an iPaq) it does not best represent how users think about their photograph collections.

Our work then is about understanding how users construct their searches for images (e.g. is it always folder first?) and how to best support those searches in a mobile application.

2. SYSTEM DESIGN

2.1 Temporal Navigation

As noted in the introduction, PC users wanted to use absolute dates to search for images. Rather than use a calendar pop-up, which would obscure the screen, we adopted a time-line with range delimiters at either end, to allow users to dynamically express their query, as in Figure 1. The timeline is automatically generated using date information from each image. This date information is extracted from image EXIF headers when images are imported into the library. (Import occurs on a per-image basis as new images are detected in the collection.) The import process takes place at start-up. Any years or months that are not visible on the screen are automatically filtered out. For example, Figure 1a shows the application in its initial state. The entire timeline is visible on the screen and no filtering is applied by default. In Figure 1b and Figure 1c, the timeline filter has been used to shorten the timeline, showing only a few months (May to September 2004 for Figure 1b and March to June 2003 for Figure 1c). The user can shorten the timeline by tapping and dragging the arrows shown at its top and bottom (on the left edge of the screen). Only these time periods are accessible; the rest of the timeline is filtered out. Of course, this action is reversible by re-adjusting the timeline filter. We refer to this selection technique as a Timeline Filter (TF). The timeline also provides a way of directly accessing a year or month. For each year, it displays a year tab and the months beneath it, in reverse order from December to January. The first letter of each month is used as a label to distinguish months (see Figure 1b and Figure 1c). Months are also highlighted using alternating grey bands. The darker grey bands are used to show months that have pictures associated with them. The current year is highlighted by an orange border. The current month is highlighted in orange. Years and months can be selected by tapping on them. We refer to this feature as Timeline access (T).

2.2 Folder Browsing

Figure 2 illustrates the hierarchical folder-based search technique. The folder labels and structure are extracted from each photo collection. The functionality of this technique is similar to Microsoft Windows Explorer. The major difference is that the folders are always displayed in a single list and the hierarchical structure is encoded by indenting the connecting lines. The "expand" and "contract" icons are used to expand and collapse folders. The buttons on either side of the list of folders are used to navigate up and down the list of folders. A folder can be selected by tapping on it. Figure 2a shows that a folder called "Portugal" has been selected. This is visible from the highlighting and the

blue border. The expand/contract icons serve two purposes. Firstly, they show which folders contain subfolders. Secondly, they can be used to expand and contract folders. From Figure 2a we can see that the "Portugal" and "New Zealand" folders contain sub-folders. In Figure 2b, the "Portugal" folder has been expanded to show its sub-directories, "Lisbon" and "Santar", with "Santar" being the currently selected folder. In Figure 2c the navigation buttons have been used to scroll down the list of folders. The arrows on the buttons indicate that you can scroll down further to access more folders. The square icon (see Figure 2a) indicates that you cannot scroll any further.

The pictures that correspond to the currently selected folder are highlighted using a dark grey band and shown at the center of the screen. A few images from the next event are used to show what happened next. Selecting an image from the next event automatically selects the event folder containing that image.

2.3 Dynamic letter and keyword search

A dynamic letter and keyword search is implemented using a Patricia-trie data structure [3]. This structure provides an efficient way of storing and retrieving alpha-numeric strings. The Patricia-trie data structure is populated using keywords from folder names. The keywords are extracted using a rule-based approach to speed up the extraction process. For example, keywords are often delineated by using special characters (e.g. "Wildlife Park" or "Wildlife.Park" or "WildLife_Park") or by capitalization (e.g. "SeaWorld" or "WarMuseum"). By applying various annotation templates, we can quickly extract keywords from the folder labels. For each folder label, the entire label is also entered as a keyword. This is necessary to ensure a match is always found when the entire label is inputted as a search string.

Figure 3a shows the components that make up the dynamic search. These are the keyboard, search box, clear search icon, minimize/maximize search icon, close search icon and the toggle keyboard icon. The keyboard icon on the far right is used to show or hide the keyboard as shown in Figure 3b. The search interface can also be minimized as shown in Figure 3c. This enables other navigation techniques to be used while in search mode.

In Figure 3, when the letter "w" is entered, three matching folders are found. These folders have keywords beginning in "w". If the letter "o" was entered next, two of the folders would be filtered out and only the "Sea World" folder would be shown. The search is dynamic, so the search results are displayed as the search term is entered. The search is also fully reversible. It is possible to return to a previous state by deleting letters or to the original state by clearing the search.

2.4 Visual search techniques

The AutoZoom and ManualZoom techniques [1] are accessible only in the detail view shown in Figure 4. For those not familiar with our earlier work, the techniques work as follows:

- *AutoZoom*: The user scrolls by dragging the stylus on the screen. As the drag distance increases, so does the rate of scroll. However, the zoom level is automatically adjusted so that at high scrolling speeds, the images become smaller.
- *ManualZoom*: Similar to the above technique, except that this time the zoom level is controlled manually by dragging in the horizontal axis.

The detail view can be selected by tapping on the thumbnail view shown in Figure 1, Figure 2 and Figure 3. The thumbnail view is

then smoothly expanded to fill the screen. Alternatively, you can select this view by using the expand/contract view icon next to the keyboard icon. The two icons to the left of that are used to select the AutoZoom (see Figure 4b) or ManualZoom (see Figure 4c) techniques. The interface defaults to the most recently selected technique.

The current image is always shown in the center of the screen. The information bar at the top of the screen is used to show its position in a folder and also its path. For example, Figure 4a shows that the current image is the first image out of a total of 44 images in the 'Knysna' folder, making it the most recent image in the folder. It also shows that its original path is 'Knysna\image1455.jpg.' The light and dark grey bands visible in Figure 4b and Figure 4c are used to show folder demarcations.

2.5 Component Integration

The first step was to find a way to integrate the timeline with the folder structure. The solution was based on two findings from our earlier study, namely: users expect their folder structures to be preserved and they expect events folders to be placed in the order in which they temporally occur. Within each folder, the photographs are sorted according to the date on which they were captured. At each level of the hierarchy, the folders are sorted according to the most recent photograph they contain. This is necessary as some folders (e.g. flower pictures) contain images not restricted to a given event and can span the time period of several other folders which are tied to specific consecutive events. To be consistent, when a keyword search is performed, the search results are ranked and displayed according to this scheme. The same ordering is also applied when visually searching through the images.

When the timeline filter is applied, the resulting action is reflected throughout the system. The folder structure only shows folders within this visible time period. The scroll and zoom techniques can only be used to locate images within this range. If a folder has photographs from a number of months, only the images from the valid time period are shown. The search is also limited to this time period.

Similarly, when a search is conducted, restrictions imposed by the search are also reflected throughout the system. The timeline only highlights the dates that are associated with the search results. The scroll and zoom techniques can only navigate through photographs in the results set.

In order to bring out correlations or disparities between the multiple views we had to ensure that when an action occurs along one dimension, the changes to other dimensions are immediately effected and made apparent. We used perceptual cues to make these relationships more apparent, focusing the user's attention on the relevant parts of the screen. For example, in Figure 1c, the image in the center of the screen has just been selected. We can see the month and year in which it was taken in and also the folder in which it is stored.

Similarly, when a search is conducted, restrictions imposed by the search are also reflected throughout the system. The timeline only highlights the dates that are associated with the search results. The scroll and zoom techniques can only navigate through photographs in the results set.

3. Experimental Evaluation

The main goal of the experiment was to gain a deeper understanding of how the search strategies differ depending on how much is known about an event. The search tool described in the previous section was used as a vehicle to explore this research question. The goal then is not to evaluate the system as a final artifact, but rather use it as a technology probe [4] to investigate the types of behavior found when these various techniques are made available to users.

3.1 Subjects

Twelve subjects took part in the experiment, five were female and seven were male. The subjects were all university students. Five were undergraduates and seven were postgraduates. They were drawn from a variety of technical and non-technical backgrounds and selected as they were all enthusiastic users of digital cameras who had large photograph collections. While this sample may not be representative of the population as a whole, the range of differences that were observed helped us gain an understanding of the factors that affect multiple search strategies. Again, we are not evaluating the system as a finished article, but exploring the uses to which such features might be put. We reasoned that this group of avid photograph users would provide a 'lead user' [5] group to help us understand future usage patterns.

The subjects were all early adopters of digital photography with fairly large photo collections. The mean number of photographs was 5134 (s.d. 2612), with a minimum of 2379 and a maximum of 8673.

3.2 Method

An observational study was used to gather data on multiple search strategies. The 'Think Aloud' protocol was used to get users to describe their actions and their reasoning while performing various searching tasks. To aid this process, the searching tasks were designed to encourage social interaction with the evaluator. The entire experiment was captured using video. The video data was used for data analysis. Post-testing interviews were used to gain further insights into user behavior and decision processes.

3.3 Equipment

The software was running on an HP iPAQ 4100 with a 1GB SD card to contain the images. It was developed using a cross-platform toolkit called Gapi-Draw, allowing us to port to the Symbian platform, should that be required. The interaction was captured on video via a Noldus mobile camera. A typical screenshot is shown below in Figure 5.



Figure 5: Video feed from PDA screen

3.4 Procedure

Prior to the experiment, each participant submitted their digital photo collections. Each collection was processed by resizing the pictures to the same resolution as the screen. This was necessary to ensure that thousands of photographs could fit on the 1GB SD card used in the device. Other than the resizing, the photo collections were not altered in any other way.

On arrival, the subjects were told that they would be trying out a new photo browsing application for small display devices. A conceptual model extraction was conducted to ensure the subjects were familiar with all the functionality before beginning the main experiment. The subjects were then given a 10 minute break.

For the main experiment, the subjects were required to complete 12 tasks which involved searching for events, singles and properties. These tasks were chosen based on the work of Rodden et al [6], who found that searching for events (e.g. a birthday, holiday or wedding), single specific images (e.g. a photo containing a hot air balloon or a particular building) and images that share some common property (e.g. all photos containing uncle Joe) constitute the bulk of image search tasks. For each of the three task types, the questions were framed in four separate ways, namely: precise, vague, verificative and exploratory. Of course these categories are not mutually exclusive. For example, a task can be both precise and verificative or vague and exploratory. Instead of using them as strict categories, they were used more as guidelines for probing different information needs and to prompt a wide range of responses from the users. The tasks were also tailored for each subject to ensure that they were as relevant as possible. An example of a sample question set for event searching is shown below.

Events

- Precise: Can you tell me about your trip to Seattle?
- Vague: What did you get up to on your trip overseas?
- Verificative: Were you at this year's Float party?
- Exploratory: I'm thinking about going away for the long weekend, can you recommend any trips that you have been on?

The tasks were designed to initiate discussions and encourage the subjects to show and tell. This was essential as we did not want the subjects to know that we were interested in their search strategy, as they might 'stage' their strategy. The discussion was

also used to gain some insight into thought processes. For example, subjects might say; 'Oh let me show you this' or 'Where was that again?' or 'I remember seeing that somewhere.' All these statements were invaluable in understanding search strategies.

An informal interview was conducted after the main experiment. The interview was used to clarify some observations and to distinguish elements of the searching behavior that were due to interface artifacts, as opposed to being part of wider searching behavior.

3.5 Data captured

Each video was analyzed to extract quantitative and qualitative data for each task. In order to distinguish search strategies we collected the order in which the techniques were used (i.e. ManualZoom, AutoZoom, Keyword Search, Folder Navigation, Timeline, and Timeline Filter); the time spent using each technique and the number of times each technique was used.

For each task, the qualitative data consisted of a short description explaining why the subject chose to search in the way that they did. We noted down how much information was known about the target event. There were three knowledge categories: precisely-known, something-known and unknown [7]. In the precisely-known category, users have a specific event in mind. They are able to articulate what they are looking for and have a good idea of where to begin searching. In the something-known category, users have some idea of the event they are looking for. They might not be able to articulate it or even know where to begin searching, but they are usually able to recognize it when they see it. In the unknown category, users do not have any specific event in mind. They also have no idea where to begin looking. This data was obtained by collating user comments with our own observations.

Please note that these three user-response categories (*precisely-known, something-known and unknown*) are separate from the four ways in which we asked questions (*precise, vague, verificative and exploratory*). So a user may have a precise response to a vague question: Vague Question: "What did you get up to on your trip overseas?"; Precise Response: "Let me show you pictures of the balloon festival I attended." Our goal is to explore how the subject's use of the software is altered by their surety in what they are searching for; are different techniques used when they know precisely what they are looking for as opposed when they only know something or know nothing about what they are looking for?

3.6 Results

For the purposes of analysis, the list of search techniques was expanded by treating the AutoZoom, ManualZoom, Timeline and Timeline filter as separate techniques. This was to develop a better understanding of search strategies.

3.6.1 Search technique usage

Overall, the number of times each search technique was used across the 12 subjects is as follows: ManualZoom (MZ) = 125, AutoZoom (AZ) = 56, Keyword search (S) = 44, Folder navigation (F) = 121, Timeline (T) = 50 and Timeline filter (TF) = 2.

The ManualZoom technique was used over two times more than the AutoZoom technique. The Folder navigation was also used over two times more than the Keyword search or Timeline. The Timeline filter was only ever used twice.

Multiple search techniques were often used to complete tasks. Search techniques were sometimes used in different orders. Table 1 shows the search sequences that were used by the subjects. For example 'F->MZ' specifies that the Folder navigation technique was used first and the ManualZoom technique was used second. For each search sequence its states the number of occurrences and the number of subjects that used the same search sequence. The shortest sequence is one link long, while the longest is eleven links long. On average, each subject used 4.25 (s.d. 1.42) techniques. Each subject used 5.6 (s.d. 2.3) different search sequences. Ten out of the twelve subjects used one or more unique sequences.

Search Sequence	Count	# of subjects
AZ	4	1
S->MZ	19	7
F->MZ	33	7
F->AZ	17	5
T->MZ	6	4
S->AZ	4	3
T->AZ	1	1
T->F->MZ	16	5
T->F->AZ	11	5
S->F->MZ	7	1
TF->T->AZ	2	2
T->MZ->AZ	1	1
S->F->AZ	1	1
S->AZ->MZ	1	1
F->MZ->AZ	1	1
T->MZ->F->MZ	2	2
F->AZ->F->AZ	2	2
F->MZ->F->MZ	3	2
F->MZ->S->AZ	1	1
T->F->AZ->MZ	1	1
S->F->AZ->MZ	1	1
F->T->F->MZ	1	1
T->F->MZ->F->MZ	2	2
S->AZ->MZ->S->MZ	1	1
S->F->AZ->S->AZ	1	1
F->MZ->F->MZ->S->MZ	1	1
S->MZ->S->T->F->MZ	1	1
S->T->F->MZ->AZ->T->AZ	1	1
T->F->T->F->MZ->F->T->MZ	1	1
T->MZ->T->F->AZ->T->F->AZ->T->AZ->MZ	1	1

Table 1: All feature combinations, from shortest to longest

3.6.2 Search technique usage across task types

The purpose of this experiment is to inform the design of the image search application. Therefore, our primary area of concern is the software features employed by users in completing the three task types. Figure 6 shows the distribution of search techniques across the three tasks types. For each task type it also shows how often it was used. The similarity across tasks types indicates that the choice of a search technique is not dependent on the task type. We therefore wanted to see if the user's confidence in what they were searching for (precisely-knew, something-known and unknown) would affect the search strategies.

3.6.3 Search strategies based on information needs

Below we analyze the search sequences employed by users according to the level of confidence they had in their search target. For each of the three confidence levels we cluster common patterns into 'Groups' to give an overview of the broad categories the subjects employed.

3.6.3.1 Precisely-known knowledge category

To understand the significance of all the results presented above, it is necessary to have a closer look at the search strategies that are used in each knowledge category. Table 2 delineates the search sequences in the precisely-known category into four main search strategies. It is important to remember that in this category events are well-known. Subjects can articulate what they need and are able to formulate queries. Bearing this in mind, one would expect the Keyword search to be used most frequently. However, this was not the case (see Table 2).

In Group 1, subjects search for events by scanning through event folders for a particular event folder name. On finding the folder, they use the AutoZoom and/or ManualZoom techniques to locate target photographs. In Group 2, subjects locate events by performing a keyword search. Again, the visual search techniques are used to locate target photographs. In Group 3, subjects locate event folders via dates. They know exactly when it occurred, so they enter the date and then scroll to the event folder. In Group 4, subjects start out by using the strategy in Group 1 and then switch to the strategy in Group 3. They realize that the first strategy will take too long and switch to what they perceive to be a quicker strategy.

Group	Search Sequence	Percentage	Group Percentage
1	F->MZ	24	46
	F->AZ	20	
	F->MZ->AZ	2	
2	S->MZ	29	38
	S->AZ	7	
	S->AZ->MZ	2	
3	T->F->MZ	5	12
	T->F->AZ	5	
	T->AZ	2	
4	F->T->F->AZ	2	4
	F->T->F->MZ	2	

Table 2: Four feature combinations in the Precisely-Known knowledge category.

Each strategy represents what subjects believe to be the most cost effective in terms of time and effort (i.e. mental and physical demand). Often subject would prefer to scan through their folders than change into the ‘keyboard’ mode.

There were other factors that affected the choice of a strategy such as, user expertise, cognitive styles (i.e. the habitual and preferred way of doing cognitive tasks [8]) and abilities (i.e. factors that contribute towards intelligence and influence the performance of specific tasks [9]).

The ManualZoom technique was used at least two times more than the AutoZoom technique. The reason being that most users found that, once they got used to it, ManualZoom would give them a greater degree of control.

3.6.3.2 “Something known” knowledge category

Table 3 delineates the search sequences in the something-known category into seven main search strategies.

The most popular search strategy in this category (see Group 1) uses the TimeLine to narrow down the search and minimize the amount of scrolling. Once a suitable temporal position is found, subjects scroll through the folders until they recognize the event folder. The AutoZoom and ManualZoom techniques are then used to locate target images. The more precisely a temporal location can be specified, the less scrolling is required.

From our previous interviews with users of PC software, we expected the Timeline filter to be more prominent in this category. The reason given by many users was that their photo collection did not span a large enough period of time (typically 2-3 years) to make it worth limiting the timeline in this way.

The search strategy in Group 3 is more exploratory in nature. This is because subjects were not able to narrow the search down to a single event. The search strategy involved investigating each possible target event folder. If all the candidate folders were known, subjects would work through them from the most likely to the least likely. Otherwise, subjects would work through them in the order in which they were remembered.

Group	Search Sequence	Percentage	Group Percentage
1	T->F->MZ	18.2	34.8
	T->F->AZ	10.6	
	T->MZ	3.0	
	TF->T->AZ	3.0	
2	F->MZ	24.2	31.8
	F->AZ	7.6	
3	T->F->MZ->F->MZ	3.0	13.5
	T->MZ->F->MZ	1.5	
	S->AZ->MZ->S->MZ	1.5	
	S->F->AZ->S->AZ	1.5	
	F->MZ->F->MZ	1.5	
	F->AZ->F->AZ	1.5	
	F->MZ->F->MZ->S->MZ	1.5	
F->MZ->S->AZ	1.5		
4	S->F->MZ	7.6	9.1
	S->F->AZ	1.5	
5	AZ	6.1	6.1
6	S->MZ	3.0	3.0
7	S->F->AZ->MZ	1.5	1.5

Table 3: Seven feature combinations in the Something-Known knowledge category.

The search sequences in this group are characterized by inspections of more than one event folder. The differences in this group are due to how much is known about an event at the start of the task and also the knowledge that is accumulated during the task. For example, in search sequences ‘F->MZ->F->MZ->S->MZ’ and ‘F->MZ->S->AZ’ subjects were able to learn and discover sufficient information to enable them to remember the name of the event they were looking for, allowing them to search for it.

In Group 4, the keyword search was used to aggregate similar events, making it easier to compare them and search through them.

In Group 5, target events and photographs were found by visually searching through images. This search strategy was only used by one subject. The subject felt that less effort was required to scroll through the images than navigating in and out of folders.

In Group 6, repeated searches were conducted to try and guess the event folder name.

The search strategy in Group 7 takes advantage of the way people organize their photo collection. The keyword search is used to locate a pre-organized category. This essentially restricts the search space to the category. Subjects can then search through the category for a target events and photographs.

3.6.3.3 Unknown knowledge category

Table 4 delineates the search sequences in the unknown category into four main search strategies. We found that the majority of subjects would conduct an exhaustive search (see Group 1). Typically a date was selected first. This was a quick way of selecting the most recent event folder. Subjects would navigate through the list of folders, checking each in turn. At the folder level, the image preview was used to make decisions on whether or not to investigate further. The AutoZoom and ManualZoom techniques were used to carry out more thorough inspections.

Group	Search Sequence	Percentage	Group Percentage
1	T->MZ	22.2	60.2
	T->F->MZ	5.6	
	T->MZ->AZ	5.6	
	T->F->AZ->MZ	5.6	
	T->MZ->F->MZ	5.6	
	T->F->MZ->F->MZ	5.6	
2	F->MZ	16.7	27.9
	F->MZ->F->MZ	5.6	
	F->MZ->F->MZ->S->MZ	5.6	
3	T->F->T->F->MZ->F->T->MZ	5.6	11.2
	T->MZ->T->F->AZ->T->AZ->T->AZ->MZ	5.6	
	S->MZ->S->T->F->MZ	5.6	
4	S->T->F->MZ->AZ->T->AZ	5.6	11.2
	S->T->F->MZ->AZ->T->AZ	5.6	

Table 4: Four feature combinations in the Unknown knowledge category.

In Group 2, subjects began by rapidly scanning through the entire photo collection to get a feel for what is in their collection. If the target event folder was not found in the initial phase, subjects

would then revisit event folders in order of the most likely to the least likely.

In Group 3, temporal navigation was used as a mechanism for jumping quickly to various locations in the photo collection. At each of these locations, subjects would inspect any promising folders.

In Group 4, subjects began by trying keyword-based searches. They felt that it was worth a try as it would save them from having to conduct an exhaustive search. However, the search terms were too broad and vague to return any relevant event folders. In both search sequences, subjects had to resort to using one of the other search methods.

4. DISCUSSION

We found evidence confirming that if you provide users with multiple searching methods, they will indeed use them. In fact, users were quite creative in combining search methods to achieve their goals. This is highlighted by the fact that on average 5.6 (s.d. 2.3) search strategies were used by each subject. The variation between individuals shows that a “one fits all” solution is not ideal. This is because users need to be able to tailor their search strategy according to their information needs. These needs change as users learn and discover more about the information space.

One interesting result was that across all subjects the target type of the search (event, single image or property) did not impact the choice of feature used in completing the search (Table 6). Our interpretation would be that the information needs were quite similar across task types.

The blend of different search methods in our system was able to cater for varying information needs. When events were well-known, the keyword search was able to provide rapid access to the event. When events were less well-known, the temporal-navigation, folder-navigation, and visual search techniques offered an alternative way of navigating through a photo collection. The temporal navigation was used to easily and efficiently skip to various locations in the photo collection. The folder navigation was used to skim through a photo collection without having to see every photograph. The visual searching techniques were used when more thorough inspections were needed. The search techniques were sometimes used in unexpected ways. For example, the keyword search was used to filter and aggregate content. Through supporting multiple search techniques, users had more opportunity to be innovative. The seamless integration allowed users to specify information along multiple dimensions, enabling them to make use of any information that was available to them. The integration was also useful for investigative activities, as users were able to filter and aggregate data on multiple contextual dimensions. Through the multiple views, users were able to see correlations and disparities in the data. By supporting familiar folder navigation, users were able to use any expertise they might have acquired when navigating through the folders on their desktop computer.

However, there are compromises in using an integrated approach. A single configuration (or bundle) of search techniques is unlikely to suit all users. For example, in our study we found that when subjects were locating events, some would think of their events in terms of time, while other would think of them in terms of their pre-created categories. Temporally-oriented users prefer events to be ordered chronologically, while categorically-oriented subjects prefer an alphabetical listing. Choosing one over the other will

obviously impact negatively on some subjects. One solution is to allow users to choose the techniques that should be integrated together to cater for their individual needs. The key to supporting customization is to ensure that each technique is built around locating events. This makes it easier add or remove a search technique. To help us better understand these differing search strategies, we can look at the problem in terms of domain knowledge.

4.1 Impact of domain knowledge

One way of thinking of the three knowledge categories (precisely-known, something-known and unknown) is in terms of the knowledge gap. That is, the distance between the user’s current knowledge of an event and the target knowledge needed to locate it. When an event is precisely-known the gap between the current and target knowledge is minimal. When an event is unknown the gap between the current and target knowledge is large. When something is known about an event, the gap may vary between these two extremes depending on the current knowledge. The challenge is how to assist users in bridging the gap.

4.1.1 *Precisely known*

When a user knows sufficient information to directly access a target event, a designer’s challenge is to assist the user in making use of this information in the most effective way. In our study, we observed that some users were willing to compromise time for reduced effort. The goal should be to minimize both. The best way to support a user in this category is by designing a set of highly specialized and tuned search techniques. As a minimum, a searching toolkit should contain the following techniques:

- Keyword search: When users can articulate what they are looking for a keyword search provides almost instant access to an event.
- Categories that support direct target acquisition and exploit spatial memory: An example of this as an A-Z listing or Date listing that is directly accessible without any scrolling.
- Active and passive links: Provide mechanisms that allow users to ‘actively’ select things that they wish to remember. For example using tags, shortcuts, or favorites. Provide mechanisms that ‘passively’ remember without needing any user input. For example, dynamic lists of frequently accessed content or recently accessed content.

4.1.2 *Unknown*

Usually when a user does not know what event they are looking for, it is because they forgotten about the event or they can’t think of any suitable event. In both cases users need to be exposed to the correct information to trigger the recollection of the appropriate event. In our study we found that a common strategy was to conduct an exhaustive search of all material in the hopes of finding the trigger. However, this type of search can be time consuming and tedious. Some subjects were more strategic by first getting an overview of the photo collection and then using this information to conduct a more focused search. The goal for designers is to assist users in meeting their initial needs, such as finding a starting point or learning about the information space. The design challenge is to maximize the amount of knowledge discovered while minimizing the amount of interaction.

4.1.3 *Something known*

This category is challenging to design for because the requirements for a searching tool differs according to how much is

known. When the knowledge gap is large, users will require solutions that enable them to explore, discover and learn about their photo collection. As the knowledge gap narrows, user will require solutions that enable them to investigate. When the knowledge gap is small, users need solutions that provide immediate access to the information that they are seeking.

One must also consider the range of investigative search activities that users are engaged in. For example, we found that users were engaged in search activities such as; recall and recognition (folder search, visual search); filtering (narrowing down search via temporal navigation) and aggregation (grouping folders by similar keywords); and verification and negation (inspecting multiple candidate folders).

Furthermore, information about an event can be split along a number of contextual dimensions (who, when, where, what and why). At the start of a task users may only be able to specify a limited amount of information along each dimension. As users learn and discover more information, searching tools must ensure that this information can be used in a way that complements what they already know. The searching tool must allow users to input any information that is available to them, allowing users to search simultaneously along multiple dimensions.

The versatility of the integrated approach described in this paper makes it suitable for supporting the something-known category. However, designers must take great care in choosing a blend of search techniques to cater for the full range of users informational needs. The seamless integration is key to supporting investigative activities. The multiple inputs allow users to specify as much as they know, allowing them to search in a more natural way.

5. CONCLUSIONS

The system we have presented here is, we believe the first of its kind for small-screen devices. It has allowed us to explore how people think about their image collections when asked to undertake a series of typical image searching tasks. The goal was not to evaluate the system's usability so much as to create a system that helped us understand how users wish to search their personal image collections.

We discovered that photo searching strategies are dependent on how much information is known about an event. When events are well-known, search strategies are based on minimizing the time and effort required. When events are less well-known, the search strategies become more investigative and incorporate activities such as recognition, inspection, filtering, aggregation and verification. The goal is to narrow down the search. When events are unknown, search strategies are more explorative to support learning and discovery. For this category, the goal of the search strategy is to maximize the knowledge gained, while minimizing the amount of interaction.

Our system represents just one approach to this problem and we look forward to other researchers and designers building image management systems for small screen devices that support multiple forms of searching and browsing support. This is especially important as many new handsets can store thousands of

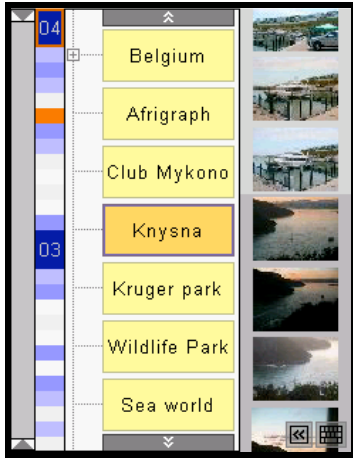
photographs yet have image browsers that were design to cope with much smaller collections.

6. ACKNOWLEDGMENTS


The handsets for this development work were provided by a donation from HP and bridges.org. This work was funded by Microsoft Research Cambridge. Thanks also to Marc Pelteret who managed to trick Word into re-sizing the tables.

7. REFERENCES


- [1] Patel, D., Marsden, G., Jones, S. & Jones, M. An Evaluation of Techniques for Browsing Photograph Collections on Small Displays *Proceedings Mobile HCI 2004*, (2004) 132-143.
- [2] Jones, S., Jones, M., Marsden, G., Patel, D. and Cockburn, A. An Evaluation of Integrated Zooming and Scrolling on Small Screens. *International Journal of Human Computer Studies*. 63(3) (2005) 271-303
- [3] Morrison, D. R. PATRICIA - Practical Algorithm to Retrieve Information Coded in Alphanumeric. *Journal of the ACM*, 15(4) (Oct 1968). 514-534.
- [4] Hutchinson, H., Mackay, W., Westerlund, B., Bederson, B. B., Druin, A., Plaisant, C., Beaudouin-Lafon, M., Conversy, S., Evans, H., Hansen, H., Roussel, N., and Eiderbäck, B. 2003. Technology probes: inspiring design for and with families. *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA, April 05 - 10, 2003). CHI '03. ACM, New York, NY, 17-24.
- [5] Von Hippel, E. Lead Users: A Source of Novel Product Concepts. *Management Science*, 32(7), (Jul., 1986), 791-805
- [6] Rodden, K., and Wood, K. How do people manage their digital photographs? *In Proc. of Human Factors in Computing Systems, CHI* (2003), 409-416, ACM Press.
- [7] Mäkelä, A., Giller, V., Tscheligi, M., and Sefelin, R. 2000. Joking, storytelling, artsharing, expressing affection: a field trial of how children and their social network communicate with digital images in leisure time. *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (The Hague, The Netherlands, April 01 - 06, 2000). CHI '00. ACM, New York, NY, 548-555.
- [8] Wang, P., Hawk, W.B, and Tenopir, C. Users' interaction with World Wide web resources: an exploratory study using a holistic approach. *Information Processing & Management*, 36(2), 229-251. 2000.
- [9] Kang, H. and Shneiderman, B. Direct Annotation: A Drag-and-Drop Strategy for Labeling Photos, *Proceedings of the International Conference on Information Visualisation* (July 19-21, 2000).



(a)




(b)




(c)

Figure 1: Temporal Navigation



(a)




(b)



(c)

Figure 2: Folder Browsing



(a)



(b)



(c)

Figure 3: Keyword search




Figure 4: AutoZoom and ManualZoom




Figure 6: Distribution of feature usage by task type