ImageGrouper: Search, Annotate and Organize Images by Groups

Munehiro Nakazato¹, Ljubomir Manola² and Thomas S. Huang¹

¹ Beckman Institute, University of Illinois at Urbana-Champaign,
405 N. Mathews Ave. Urbana, IL 61801, USA
{nakazato, huang}@ifp.uiuc.edu

² School of Electrical Engineering, University of Belgrade
manola@yubc.net

Abstract. In Content-based Image Retrieval (CBIR), trial-and-error query is essential for successful retrieval. Unfortunately, the traditional user interfaces are not suitable for trying different combinations of query examples. This is because first, these systems assume query examples are added incrementally. Second, the query specification and result display are done on the same workspace. Once the user removes an image from the query examples, the image may disappear from the user interface. In addition, it is difficult to combine the result of different queries.

In this paper, we propose a new interface for Content-based image retrieval named *ImageGrouper*. In our system, the users can interactively compare different combinations of query examples by dragging and grouping images on the workspace (Query-by-Group.) Because the query results are displayed on another pane, the user can quickly review the results. Combining different queries is also easy. Furthermore, the concept of “image groups” is also applied to annotating and organizing a large number of images.

1 Introduction

Many researchers have proposed ways to find an image from large image databases. We can divide these approaches into two types of interactions: *Browsing* and *Searching*. In image browsing, the users look through the entire collections. In most systems, the images are clustered in hierarchical manner and the user can traverse the hierarchy by zooming and panning [3][4][10][16]. In [16], browsing and searching are integrated so that the user can switch back and forth between browsing and searching.

Meanwhile, enormous amount of research have been done for *Content-Based Image Retrieval* (CBIR) [7][18][24]. In CBIR systems, the user searches image by visual similarity, i.e. low-level image features such as color [25], texture [23] and structure [27]. They are automatically extracted from images and indexed in the database. Then, the system computes the similarity between the images based on these features.

The most popular method of CBIR interaction is *Query-by-Examples*. In this method, the users select example images (as positive or negative) and ask the system to retrieve visually similar images. In addition, in order to improve the retrieval further, CBIR systems often employ *Relevance Feedback* [18][19], in which the users

can refine the search incrementally by giving feedback to the result of the previous query.

In this paper, we propose a new user interface for digital image retrieval and organization, named *ImageGrouper*. In *ImageGrouper*, a new concept *Query-by-Groups* is introduced for Content-based Image Retrieval (CBIR.) The users construct queries by making groups of images. The groups are easily created by dragging images on the interface. Because the image groups can be easily reorganized, flexible retrieval is achieved. Moreover, with the similar operations, the user can effectively annotate and organize a large number of images.

In the next section, we discuss how groups are used for image retrieval. Then, the following sections describe the use of image groups for image annotation and organization.

2 User Interface Support for Content-based Image Retrieval

2.1 Current Approaches: Incremental Search

Not many researches have been done regarding user interface support for Content-based Image Retrieval (CBIR) systems [16][20]. Figure 1 shows a typical GUI for CBIR system that supports *Query-by-Examples*. Here, a number of images are aligned in grids. In the beginning, the system displays randomly selected images. The effective ways to align images are studied in [17]. In some cases, they are images found by browsing or keyword-based search.

Under each image, a slide bar is attached so that the user can tell the system which images are relevant. If the user thinks an image is relevant, s/he moves the slider to the right. If s/he thinks an image is not relevant and should be avoided, s/he moves the slider to the left. The amount of slider movement represents the degree of relevance

Fig 1. Typical GUI for CBIR Systems

Fig 2. Example of “More is not necessarily better”. The left is the case of one example, the right is the case of two examples.

(or irrelevance.) In some systems, the user selects example images by clicking check boxes or by clicking on the images [6]. In these cases, the degrees are not specified.

When the “Query” button is pressed, the system computes the similarity between selected images and the database images, then retrieves the N most similar images. The grid images are replaced with the retrieved images. These images are ordered based on the degree of similarity.

If the user finds additional relevant images in the result set, s/he selects them as new query examples. If a highly irrelevant image appears in the result set, the user can select it as a negative example. Then, the user press “Query” again. The user can repeat this process until s/he is satisfied. This process is called *relevance feedback* [18][19]. Moreover, in many systems, the users are allowed to directly weight the importance of image features such as color and texture.

In [22], Smeulders et al. classified *Query by Image Example* and *Query by Group Example* into two different categories. From user interface viewpoint, however, these two are very similar. The only difference is whether the user is allowed to select multiple images or not. In this paper, we classify both approaches as Query by Examples method. In stead, we use term “Query by Groups” to refer our new model of query specification method described later.

Query-by-Example approach has several drawbacks. First of all, these systems assume that *the more query examples are available, the better result we can get*. Therefore, the users are supposed to search images *incrementally* by adding new example images from the result of the previous query. However, this assumption is not always true. Additional examples may contain undesired features and degenerate the retrieval performance.

Figure 2 shows an example of situations when more query examples could lead to worse results. In this example, the user is trying to retrieve pictures of cars. The left column shows the query result when only one image of “car” is used as a query example. The right column shows the result of two query examples. The results are ordered based on the similarity ranks. In both cases, the same relevance feedback algorithm (Section 5.2 and [19]) was used and tested on Corel image set of 17,000 images. In this example, even if this additional example image looks visually good for human eyes, it introduces undesirable features into the query. Thus, no car image appears in the top 8 images. An image of car appears in the rank 13th for the first time.

This example is not a special case. It happens often in image retrieval and confuses the users. This problem happens because of *semantic gap* [20][22] between the high-level concept in the user’s mind and the extracted features of images. Furthermore, finding good combinations of query examples is very difficult because image features are numerical values that are impossible to be estimated by human. Only way to find the right combination is *trial and error*. Otherwise, the user can be trapped in a small part of image database [16].

Unfortunately, the traditional user interfaces were designed for *incremental search* and are not suitable for *trial and error query*, if not impossible. This is because in these systems, query specification and result display must be done on the same workspace. Once the user removes an image from the query examples during relevance

feedback loops, the image may disappear from the user interface. Thus, it is awkward to bring it back later for another query.

Second, the traditional interface does not allow the user to put aside the query results for later uses. This type of interaction is desired because the users are not necessarily looking for only one type of images. The users' interest may change during retrieval. This behavior is known as *berry picking* [2] and has been observed for text documents retrieval by O'Day and Jeffries [15].

Moreover, because of *semantic gap* [20][22] mentioned above, the users often need to make more than one query to satisfy his/her need [2]. For instance, a user may be looking for images of "beautiful flowers." The database may contain many different "flower" images. These images might be completely different in terms of low-level visual features. Thus, the user needs to retrieve "beautiful flowers" as a collection of different types of images.

Finally, in some case, the user had better start from a general concept of objects and narrow down to specific ones. For example, suppose the user is looking for images of "red cars." Because image retrieval systems use various image features [23][27] as well as colors [25], even cars with different colors may have many common features with "red cars." In this case, it is better to start by collecting images of "cars of any color." Once enough number of car images are collected, the user can specify "red cars" as positive examples, and other cars as negative examples. Current interfaces for CBIR systems, however, do not support these types of query behavior.

Another interesting approach for Query by Examples was proposed by Santini et.al [20]. In their *El Niño* system, the user specifies a query by mutual distance between example images. The user drags images on the workspace so that the more similar images (in the user's mind) are located closer to each other. The system then reorganizes the images' locations reflecting the user's intent. There are two drawbacks in *El Niño* system. First, it is unknown to the users how close similar images should be located and how far negative examples should be apart from good examples. It may take a while for the user to learn "the metric system" used in this interface.

The second problem is that like traditional interfaces, query specification and result display are done on the same workspace. Thus, the user's previous decision (in the form of the mutual distance between the images) is overridden by the system when it displays the results. This makes trial and error query difficult. Given the analogue nature of this interface, trial and error support might be essential. Even if the user gets an unsatisfactory result, there is no way to redo the query with a slightly different configuration. Any experimental result is not provided in the paper.

2.2 Query-by-Groups

We are developing a new user interface for CBIR systems named *ImageGrouper*. In this system, a new concept *Query-by-Groups* was introduced. Query-by-Groups mode is an extension to Query-by-Example mode described above. The major difference is that while Query-by-Example handles the images individually, in *Query-by-Group*, a "group of images" is considered as the basic unit of the query.

Figure 3 shows the display layout of *ImageGrouper*. The interface is divided into two panes. The left pane is the *ResultView* that displays the results of content-based

Fig 3. The ImageGrouper

retrieval, keyword-based retrieval, and random retrieval. This is similar to the traditional GUI except for there are no sliders or buttons under the images.

The right pane is the *GroupPalette*, where the user manages each image and image groups. In order to create an image group, the user first drags one or more images from the *ResultView* into *GroupPalette*, then encloses the images by drawing a rectangle (box) as we draw a rectangle in drawing applications. All the images within the *group box* become the member of this group. Any number of groups can be created in the palette. The user can move images from one group to another at any moment. In addition, groups can be overlapped to each other, i.e. each image can belong to multiple groups. To remove an image from a group, the user simply drags it out of the box.

When the right mouse button is pressed on a group box, a popup menu appears so that the user can give query properties (positive, negative, or neutral) to the group. The properties of groups can be changed at any moment. The colors of the corresponding boxes change accordingly. To retrieve images based on these groups, the user press the “Query” button placed at the top of the window (Figure 3.) Then, the system retrieves new images that are similar to images in positive groups while avoiding images similar to negative groups. The result images are displayed in the *ResultView*.

When a group is specified as *neutral* (displayed as a white box), this group does not contribute to the search at the moment. This group can be turned to a positive or negative group later for another retrieval. If a group is *positive* (displayed as a blue box), the system uses common features among the images in the group. On the other hand, if a group is given *negative* (red box) property, the common features in the group are used as negative feedbacks. The user can specify multiple groups as positive or negative. In this case, these groups are merged into one group, i.e. the union of the groups are taken. The detail of the algorithm is described in Section 5.2.

In the example shown in Figure 3, the user is retrieving images of “flowers.” In the *GroupPalette*, three flower images are grouped as a positive group. On the right of this group, a red box is representing a negative group that consists of only one image. Below the “flowers” group, there is a neutral group (white box), which is not used for retrieval at this moment. Images can be moved outside of any groups in order to temporarily remove images from the groups.

The gestural operations of ImageGrouper are similar to file operations of a Window-based OS. Furthermore, because the user’s mission is to collect images, the operation “*Dragging Images into a Box*” naturally matches the user’s cognitive state.

2.3 Flexible Image Retrieval

The main advantage of Query-by-Groups is flexibility.

Trial and Error Query by Mouse Dragging. In *ImageGrouper*, images can be easily moved between the groups by mouse drags. In addition, the neutral groups and space outside of any groups in the palette can be used as *storage area* [8] for images that are not used at the moment. They can be reused later for another query. It makes trial and error of relevance feedbacks easier. The user can quickly explore different combinations of query examples by dragging images into or out of the box. Moreover, the query specification that the user made is preserved and visible in the palette. Thus, it is easy to modify the previous decision when the query result is not satisfactory.

Groups in a Group. *ImageGrouper* allows the users to create a new group within a group (*Groups in a Group*.) With this method, the user begins with collecting relatively generic images first, then narrows down to more specific images.

Figure 4 shows an example of *Groups in a Group*. Here, the user is looking for “Red cars.” When s/he does not have enough number of examples, however, the best way to start is to retrieve images of “cars with any color.” This is because these images may have many common features with red car images, though their colors features are different. The large white box is a group for “Cars with any colors.”

Once the user found enough number of car images, s/he can narrow down the search only for red cars. In order to narrow down the search, the user divide the collected images into two sub-groups by creating two new boxes for red cars and other cars. Then the user specifies the red car group as positive and the other cars group as negative, respectively. In Figure 4, the left smaller (blue, i.e. positive) box is the group of red cars and the right box (red, i.e. negative) is the group of non-red cars. This *narrow down search* was not possible on the conventional CBIR systems.

2.4 Experiment on Trial and Error Query

In order to examine the effect of ImageGrouper’s trial-and-error query, we compared the query performance of our system with that of a traditional incremental approach (Figure 1). In this experiments, we used Corel photo stock that contains 17000 images as the data set. For both interfaces, the same image features and relevance feedback algorithms (described in Section 5.2) are used.

For the traditional interface, the top 30 images are displayed and examined by the user in each relevance feedback. For ImageGrouper, the top 20 images are displayed in the *ResultView*. Only one positive group and one neutral group are created for this

Fig 4. The number of hits until the 10th round or convergence

Table 1: The number of feedback loops until convergence. 10 means that it did not converge until 10th round

Object	Traditional	Grouper
Red Car	4	10
Tiger	3	10
Bird	4	8
Yellow Flower	3	10
Citrus	2	4
Polar Bear	5	10
Elephant	3	10
Swimmer	2	10

test. On both interfaces, no negative feedback is given. Feedback loops are repeated up to 10 rounds or until convergence.

We tested over eight classes of images (Table 1). For each class, a query starts from one image example. In case of the traditional interface, the query is repeated by giving additional example from the result of previous query. When no more good example appears, the search stops (convergence). Meanwhile, for ImageGrouper, the search is refined incrementally at first. When the incremental search converges, trial and error search is applied by moving some images out of the positive group into a neutral group (This means that the number of positive examples is temporarily decreased.) Then, the search is refined incrementally again until another convergence occurs. The user repeats this until trial and error query has no effect.

Figure 4 shows the number of correctly retrieved images after the convergence (or the 10th rounds.) This value is proportional to the recall rate. Thus, larger value means better retrieval performance. Table 1 shows the number of relevance feedback loops until convergence. The value 10 means the query did not converge before the 10th round.

Clearly, ImageGrouper can achieve better retrieval (higher recall) even if underlying technologies (relevance feedback algorithm and visual features) are identical. In addition, the search with ImageGrouper is less likely to converge prematurely even when the search with the traditional interface converges into small number of images after a few iterations. This result suggests the importance of support for *trial and error* query.

3 Text annotations on images

No wonder keyword-based search is a very powerful method for searching images. The problem is that it works well only when all the images are annotated with textual information. For commercial photo stocks, it may be feasible to enter keywords to all images manually. For home users, however, it is too tedious.

Fig 4. Groups in a group.

Fig 5. Overlap between groups. Two images in the overlapped region contain both mountain and cloud.

When keyword search is integrated with CBIR like our system and [16], keyword-based search can be used to find the initial query examples for content-based search. For this scheme, the user does not have to annotate all images. In any cases, it is very important to provide easy and quick ways to annotate text on a large number of images.

3.1 Current Approaches for Text Annotation

The most primitive way for annotation is to select an image, then type in keywords. Because this interaction requires the user to use mouse and keyboard repeatedly in turn, it is too frustrating for a large image database.

Several researchers have proposed smarter user interfaces for keyword annotation on images. In *bulk annotation* method of *FotoFile* [9], the user selects multiple images on the display, selects several attribute/value pairs from a menu, and then presses the “Annotate” button. Therefore, the user can add the same set of keywords on many images at the same time. To retrieve images, the user selects entries from the menu, and then presses the “Search” button. Because of *visual and gestural symmetry* [9], the user needs to learn only one tool for both annotation and retrieval.

PhotoFinder [21] introduced *drag-and-drop* method, where the user selects a label from a scrolling list, then drags it directly onto an image. Because the labels remain visible at the designated location on the images and these locations are stored in the database, these labels can be used as “captions” as well as for keyword-based search. For example, the user can annotate the name of a person directly on his/her portrait in the image, so that other users can associate the person with his/her name. When the user needs new words to annotate, s/he adds them to the scrolling list. Because the user drags keywords into individual images, bulk annotation is not supported in this system.

3.2 Annotation by Groups

Most home users do not want to annotate images one by one, especially when the number of images is large. In many cases, the same set of keywords is enough for several images. For example, a user may just want to annotate “My Roman Holiday, 1997” on all images taken in Rome. Annotating the same keywords repeatedly is painful enough to discourage him/her from using the system.

ImageGrouper introduces *Annotation-by-Groups* method where keywords are annotated not on individual images, but on groups. As in *Query-by-Groups*, the user first creates a group of images by dragging images from *ResultView* into *GroupPalette* and drawing a rectangle around them. In order to give keywords to the group, the user opens *Group Information Window* by selecting “About This Group” from the pop-up menu (Figure 3). In this window, arbitrary number of words can be added. Because the users can simultaneously annotate the same keywords on a number of images, annotation becomes much faster and less error prone. Although *Annotation-by-Groups* is similar to bulk annotation of *FotoFile* [9], there are several advantages described below.

Annotating New Images with the Same Keywords. In *bulk annotation* [9], once the user finished annotating keywords to some images, there is no fast way to give the same annotation to another image later. The user has to repeat the same steps (i.e. select images, select keywords from the list, then press “Annotate”). This is awkward when the user has to add a large number of keywords. Meanwhile, in *Annotation-by-Group*, the system attaches annotations not on each images, but on groups. Therefore, by dragging new images into an existing group, the same keywords are automatically given to it. The user does not have to type the same words again.

Hierarchical Annotation with Groups in a Group. In *ImageGrouper*, the user can annotate images hierarchically using *Groups in a Group* method described above (Figure 4). For example, the user may want to add new keyword “Trevi Fountain” to only a part of the image group that has been labeled “My Roman Holiday, 97.” This is easily done by creating a new sub-group within the group and annotating only on the sub-group.

In order to annotate hierarchically on *FotoFile* [9] with bulk annotation, the user has to select some of images that are already annotated, and then annotate them again with more keywords. On the other hand, *ImageGrouper* allows the user to visually construct a hierarchy on the *GroupPalette* first, then edit keywords on the *Group Information Window*. This method is more intuitive and less error prone.

Overlap between Images. An image often contains multiple objects or people. In such cases, the image can be referred in more than one context. *ImageGrouper* support this *multiple references* by allowing overlaps between image groups, i.e. an image can belong to multiple groups at the same time. For example, in Figure 5, there are two image groups: “Cloud” and “Mountains.” Because some images contain both cloud and mountain, these images belong to both groups. They are automatically referred as “Cloud and Mountain.” This concept is not supported in other systems.

4 Organizing Images by Groups

In the previous two sections, we described how *ImageGrouper* supports content-based query as well as keyword annotation. These features are closely related and complementary to each other. In order to annotate images, the user can collect visually similar images first, using content-based retrieval with *Query-by-Groups*. Then s/he can annotate textual information to the group of collected images. After this point, the user can quickly retrieve the same images using keyword-based search.

Conversely, the results of keyword-based search can be used as a starting point for content-based search. This method is useful especially when the image database is only partially annotated or when the user is searching images based on visual appearance only.

4.1 Photo Albums and Group Icons

As described above, *ImageGrouper* allows groups to be overlapped. In addition, the user can attach textual information on these groups. Therefore, groups in *ImageGrouper* can be used to organize pictures as “photo albums [9]” Similar concepts are proposed in *FotoFile* [9] and Ricoh’s *Storytelling* system [1]. In both systems, albums are used for “slide shows” to tell stories to the other users.

In *ImageGrouper*, the user can convert a group into a *group icon*. When the user selects “Iconify” from the popup menu (Figure 3,) images in the group disappear and a new icon for the group appears in *GroupPalette*. When the group has an overlap with another group, images in the overlapped region remain in the display.

Furthermore, the users can manipulate those group icons as they handle individual images. They can drag the group icons anywhere in the palette. The icons can be even moved into another group box realizing *groups in a group*.

Finally, group icons themselves can be used as examples for content-based query. A group icon can be used as an independent query example or combined with other images and groups. In order to use a group icon as a normal query group, the user right clicks the icon and opens a popup menu. Then, s/he can select “relevant”, “irrelevant” or “neutral.” On the other hand, in order to combine a group icon with other example images, the user simply draws a new rectangle and drags them into it.

Organize-by-Groups method described here is partially inspired by the *Digital Library Integrated Task Environment* (DLITE) [5]. In DLITE, each text documents as well as the search results are visually represented by icons. The user can directly manipulate those documents in a workcenter (*direct-manipulation*.) In [8], Jones proposed another graphical tool for query specification, named *VQuery*. In *VQuery*, the user specifies the query by creating Venn diagrams. The number of matched documents is displayed in the center of each circle.

While DLITE and *VQuery* were systems for text documents, the idea of *direct-manipulation* [5] is applicable more naturally to image databases. In text document database, it is difficult to determine the contents of text documents from the icons. Therefore, the user has to open another window to investigate the detail [5] (in case of DLITE, a web browser is opened.) On the other hand, in image databases, images themselves (or their thumbnails) can be used for direct-manipulations. Therefore, instant judgment by the user is possible [16][22].

5 Implementation

A prototype of *ImageGrouper* is implemented as a client-server system, which consists of *User Interface Clients* and *Query Server*. They are communicating via Hyper-Text Transfer Protocol (HTTP).

5.1 The User Interface Client

The user interface client of *ImageGrouper* is implemented as a Java2 Applet with Swing API (Figure 3). Thus, the users can use the system through Web browsers on various platforms such as Windows, Linux, Unix and Mac OS X.

The client interacts with the user and determines his/her interests from the group information or keywords input. When “Query” button is pressed, it sends the information to the server. Then, it receives the result from the server and displays it on the *ResultView*. Because the client is implemented in multi-thread manner, it remains reactive while it is downloading images. Thus, the user can drag a new image into the palette as soon as it appears in the *ResultView*.

Note that the user interface of *ImageGrouper* is independent of relevance feedback algorithms [18][19] and the extracted image features (described below.) Thus, as long as the communication protocols are compatible, the user interface clients can access to any image database servers with various algorithms and image features. Although the retrieval performance depends on the underlying algorithms and image features used, the usability of *ImageGrouper* is not affected by those factors.

5.2 The Query Server

The *Query Server* stores all the image files and their low-level visual features. These visual features are extracted and indexed in advance. When the server receives a request from a client, it computes the weights of features and compares user-selected images with images in the database. Then, the server sends back IDs of the k most similar images.

The server is implemented as a Java Servlet that runs on the Apache Web Server and Jakarta Tomcat Servlet container. It is written in Java and C++. In addition, the server is implemented as a *stateless server*, i.e. the server does not hold any information about the clients. This design allows different types of clients such as the traditional user interface [13] (Figure 1) and 3D Virtual Reality interface [14] can access to the same server simultaneously.

For home users who wish to organize and retrieve images locally on their PCs’ hard disks, *ImageGrouper* can be configured as a standalone application, in which the user interface and the query server are resident on the same machine and communicate directly without a Web server.

Image Features. As the visual features for content-base image retrieval, we use three types of features: *Color*, *Texture*, and *Edge Structure*.

For color features, HSV color space is used. We extract the first two moments (mean, and standard deviation) from each of HSV channels [25]. Therefore, the total number of color features is six. For texture, each image is applied into *wavelet filter bank* [23] where the images are decomposed into 10 de-correlated sub-bands. For each sub-band, the standard deviation of the wavelet coefficients is extracted. There-

fore, the total number of this feature is 10. For edge structures, we used *Water-Fill edge detector* [27] to extract image structures. We first pass the original images through the edge detector to generate their corresponding edge maps. From the edge map, eighteen (18) elements are extracted from the edge maps.

Relevance Feedback Algorithm. The similarity ranking is computed as follows. First, the system computes the similarity of each image with respect to only one of the features. For each feature i ($i = \{\text{color, texture, structure}\}$), the system computes a query vector q_i based on the positive and negative examples specified by the user. Then, it calculates the feature distance g_{ni} between each image n and the query vector,

$$g_{ni} = W_i^T (\vec{p}_{ni} - \vec{q}_i) W_i \quad (1)$$

where $\vec{p}_{ni}$ is the feature vector of image n regarding the feature i . For the computation of the distance matrix W_i , we used *Biased Discriminant Analysis* (BDA.) The detail of BDA is described in [26]. After the feature distances are computed, the system combines each feature distance g_{ni} into the *total distance* d_n . The total distance of image n is a weighted sum of each g_{ni} ,

$$d_n = \vec{u}^T \vec{g}_n \quad (2)$$

where $\vec{g}_n = [g_{n1}, \dots, g_{nI}]$. I is the total number of features. In our case, I is 3. The optimal solution of the feature weighting vector $\vec{u} = [u_1, \dots, u_I]$ is solved by Rui et al. [19] as follows,

$$u_i = \sum_{j=1}^I \sqrt{f_j / f_i} \quad (3)$$

where $f_i = \sum_{n=1}^N g_{ni}$, and N is the number of positive examples. This gives higher weight to that feature whose total distance is small. This means that if the positive examples are similar with respect to a certain feature, this feature gets higher weight. Finally, the images in the database are ranked by the total distance. The system returns the k most similar images.

6 Future Work

We plan to evaluate our system further with respect to both usability and query performance. Especially, we will investigate the effect of *Groups in a group* query described in Section 2.3. As mentioned in [11], traditional precision/recall measure is not very suitable for evaluation for interactive retrieval systems. Therefore, we may need to consider appropriate evaluation methods for the system [12][22].

Next, in the current system, when more than one group is selected as positive, they are merged into one group, i.e. all images in those groups are considered as positive examples. We are investigating a scheme where different positive groups are considered as different classes of examples [28].

In addition, for the advanced users, we are going to add support for group-wise feature selection. Although our system automatically determines the feature weights, the advance users might know which features are important for their query. Thus, we will allow the users to specify which features are supposed to be considered for each group. Some groups might be important in terms of color features only, while others might be important in terms of structures. Finally, because the implementation of

ImageGrouper does not depend on underlying retrieval technologies, it can be used as a benchmarking tool [12] for various image retrieval systems.

7 Conclusion

In this paper, we presented *ImageGrouper*, a new user interface for digital image retrieval and organization. In this system, the users search, annotate, and organize digital images by groups. *ImageGrouper* has several advantages regarding image retrieval, text annotation, and image organization.

First, in content-based image retrieval (CBIR), predicting a good combination of query examples is very difficult. Thus, *trial-and-error* is essential for successful retrieval. However, the previous systems are assuming *incremental* search and do not support *trial-and-error* search. On the other hand, *Query-by-Groups* concept in *ImageGrouper* allows the user to try different combinations of query examples quickly and easily. We showed this lightweight operation helps the users to achieve higher recall rate. Second, with *Groups in a Group* configuration, *narrowing down search* was made possible. This method helps the user find both positive and negative examples, and provides him/her with more choices.

Next, typing text information to a large number of images is very tedious and time consuming. *Annotate-by-Groups* method eases the users of this task by allowing them to annotate multiple images at the same time. *Groups in a group* method realizes hierarchical annotation, which was difficult in the previous systems. Moreover, by allowing groups to overlap to each other, *ImageGrouper* further reduces typing.

In addition, our concept of image groups is also applied for organizing image collections. A group in *GroupPalette* can be shrunk into a small icon. These *group icons* can be used as “photo albums” which can be directly manipulated and organized by the users.

Finally, these three concepts: *Query-by-Groups*, *Annotation-by-Groups* and *Organize-by-Groups* share the similar gestural operations, i.e. *dragging images and drawing a rectangle around them*. Thus, once the user learned one task, s/he can easily adapt herself/himself to the other tasks. Operations in *ImageGrouper* are also similar to file operations used in Windows and Macintosh computers as well as most drawing programs. Therefore, the user can easily learn to use our system.

Acknowledgement

This work was supported in part by National Science Foundation Grant CDA 96-24396.

References

1. Balabanovic, M., Chu, L.L. and Wolff, G.J. Storytelling with Digital Photographs. In *CHI'00*, 2000.
2. Bates, M.J. The design of browsing and berrypicking techniques for the on-line search interface. *Online Review*, 13(5), pp. 407-431, 1989.
3. Bederson, B.B. Quantum Treemaps and Bubblemaps for a Zoomable Image Browser. *HCIL Tech Report #2001-10*, University of Maryland, College Park, MD 20742.
4. Chen, J-Y., Bouman, C.A., and Dalton, J.C. Heretical Browsing and Search of Large Image Database. In *IEEE Trans. on Image Processing*, Vol. 9, No. 3, pp. 442-455, March 2000.

5. Cousins, S.B., et al. The Digital Library Integrated Task Environment (DLITE). In *2nd ACM International Conference on Digital Libraries*, 1997.
6. Cox, I.J., Miller, M.L., Minka, T.P., Papathomas, T.V. and Yianilos, P.N. The Bayesian Image Retrieval System, *PicHunter*: Theory, Implementation, and Psychophysical Experiments. In *IEEE Transactions on Image Processing*, Vol. 9, No. 1, January 2000.
7. Flickner, M., Sawhney, H. and et al. Query by Image and Video content: The QBIC system. In *IEEE Computer*, Vol. 28, No.9, pp. 23-32, September 1995.
8. Jones, S. Graphical Query Specification and Dynamic Result Previews for a Digital Library. In *UIST'98*, 1998.
9. Kuchinsky, A., Pering, C., Creech, M.L., Freeze, D., Serra, B. and Gwizdka, J. FotoFile: A Consumer Multimedia Organization and Retrieval System. In *CHI'99*, 1999.
10. Laaksonen, J., Koskela, M. and Oja, E. Content-based image retrieval using self-organization maps. In *Proc. of 3rd Intl. Conf. in Visual Information and Information Systems*, 1999.
11. Lagergren, E. and Over, P. Comparing interactive information retrieval systems across sites: The TREC-6 interactive track matrix experiment. In *ACM SIGIR'98*, 1998.
12. Müller, H et al. Automated Benchmarking in Content-based Image Retrieval. In *Proc. of IEEE International Conference on Multimedia and Expo 2001*, August, 2001.
13. Nakazato, M. et al., UIUC Image Retrieval System for JAVA, available at <http://chopin.ifp.uiuc.edu:8080>.
14. Nakazato, M. and Huang, T.S. 3D MARS: Immersive Virtual Reality for Content-based Image Retrieval. In *Proc. of IEEE International Conference on Multimedia and Expo 2001*.
15. O'Day V. L. and Jeffries, R. Orienteering in an information landscape: how informationseekers get from here to there. In *INTERCHI '93*, 1993.
16. Pecenovic, Z., Do, M-N., Vetterli, M. and Pu, P. Integrated Browsing and Searching of Large Image Collections. In *Proc. of Fourth Intl Conf on Visual Information Systems*, Nov, 2000.
17. Rodden, K., Basalaj, W., Sinclair, D. and Wood, K. Does Organization by Similarity Assist Image Browsing? In *CHI'01*. 2001.
18. Rui, Y., Huang, T. S., Ortega, M. and Mehrotra, M. Relevance Feedback: A Power Tool for Interactive Content-Based Image Retrieval. In *IEEE Transaction on Circuits and Video Technology*, Vol. 8, No. 5, Sept. 1998.
19. Rui, Y. and Huang, T. S., Optimizing Learning in Image Retrieval. In *IEEE CVPR '00*, 2000.
20. Santini, S. and Jain, R. Integrated Browsing and Querying for Image Database. *IEEE Multimedia*, Vol. 7, No. 3, 2000, pp. 26-39.
21. Shneiderman, B. and Kang, H. Direct Annotation: A Drag-and-Drop Strategy for Labeling Photos. In *Proc. of the IEEE Intl Conf on Information Visualization (IV'00)*, 2000.
22. Smeulders, A.W.M., Worring, M., Santini, S., Gupta, A. and Jain, R. Content-based Image Retrieval at the End of the Early Years. In *IEEE PAMI* Vol. 22, No. 12, December, 2000.
23. Smith, J.R. and Chang S-F. Transform features for texture classification and discrimination in large image databases. In *Proc. of IEEE Intl. Conf. on Image Processing*, 1994.
24. Smith, J.R. and Chang S-F. VisualSEEK: a fully automated content-based image query system. In *ACM Multimedia'96*, 1996.
25. Sticker, M. and Orengo, M., Similarity of Color Images. In *Proc. of SPIE*, Vol. 2420 (Storage and Retrieval of Image and Video Databases III), SPIE Press, Feb. 1995.
26. Zhou, X. and Huang, T. S. A Generalized Relevance Feedback Scheme for Image Retrieval. In *Proc. of SPIE* Vol. 4210: Internet Multimedia Management Systems, 6-7 November 2000.
27. Zhou, X. S. and Huang, T. S. Edge-based structural feature for content-base image retrieval. *Pattern Recognition Letters, Special issue on Image and Video Indexing*, 2000.
28. Zhou, X. S., Petrovic, N. and Huang, T. S. Comparing Discriminating Transformations and SVM for Learning during Multimedia Retrieval. In *ACM Multimedia '01*, 2001.