

# An Exploration of Space-Time Constraints on Contextual Information in Image-based Testing Interfaces

Unmil Karadkar<sup>1</sup>, Marlo Nordt<sup>1</sup>, Richard Furuta<sup>1</sup>, Cody Lee<sup>2</sup>, Christopher Quick<sup>2</sup>

<sup>1</sup> Center for the Study of Digital Libraries, College Station, TX 77843-3112, USA  
{unmil, mnordt, furuta}@csdl.tamu.edu

<sup>2</sup> Michael E. DeBakey Institute, College Station, TX 77843-4466, USA  
{codylee, cquick}@tamu.edu

**Abstract.** Digital image collection interface layouts vary in the nature and degree of contextual information they provide to their users, thus enabling or impeding specific tasks. We are exploring image presentation techniques to support image-centric cognitive tasks in the context of cardiovascular systems research and education. To investigate the effect of image layout on user performance, we conducted an experimental evaluation of three image layouts for three representative tasks in this domain. The layouts varied the spatial and temporal presentation of images, thus providing different contextual information to the test subjects. Our results indicate that the degree of contextual information provided by the image layouts affected user performance, as did their research expertise. These results will inform the design of user interfaces for performing image-focused cognitive tasks as well as the development of interfaces for training novice researchers.

## 1 Introduction

Government agencies, educational institutions, scientific communities, commercial entities, and individuals all develop and deploy Web-based digital image collections to further their goals. The two-dimensional, grid-based, thumbnail layout has emerged as the de-facto standard for browsing Web-based image collections and search results [1, 6, 12, 19]. While the Frappr! interface employs a novel map-based image overlay, it retains the familiar 2-D thumbnail layout for displaying the images associated with each location [10].

A significant body of research has focused on designing targeted interfaces for presenting digital images in the context of personal photograph collections and digital libraries. These interfaces employ a variety of techniques, such as zooming [4], radial quantum layouts [18], collaging [13, 15], temporal cues [11, 14], 3-D immersive environments [5], and mechanisms for creating topical sub-collections [17, 20] to support users in recognition tasks such as searching and browsing large image collections for finding images that match specific criteria.

In contrast, to the recognition-oriented tasks, researchers and technicians in medical or atmospheric science areas evaluate images for identifying patterns and for answering specific questions. For example, doctors use X-ray images to determine the

possibility and extent of bone fractures and weather forecasters predict the arrival of storm fronts. In these domains, trained professionals analyze images and apply their knowledge from the real world to reach critical conclusions.

We are exploring techniques for testing specific research skills and for training novice researchers in methods of cardiovascular research using a video and image corpus. In particular, we are investigating the role of image layouts and the contextual information embodied by these layouts in image-based question-answering tasks. Towards this purpose, we conducted an experimental evaluation of three image layouts that differed in the spatial and temporal organization of images. The layouts we chose varied the nature of the contextual information available to test subjects as well as the amount of this information. Our subjects performed three tasks with each of these layouts to test for specific, representative skills that cardiovascular systems researchers acquire through experience.

Our results indicate that the degree of contextual information provided by the interfaces affected user performance, as did their research expertise. The results highlight significant differences in the strategies used by experts and novices and will inform the design of software for training cardiovascular systems researchers. These results have broader applications in areas that involve image-based cognitive tasks.

The remainder of this paper is organized as follows: the next section situates our research in the context of the cardiovascular systems research domain. We then describe the experimental design of the evaluation: the nuances of the three image layouts and the three representative tasks that evaluate cardiovascular systems research skills, followed by the findings of our study. We discuss the implications of these experimental results and conclude by presenting continuing research directions.

## **2 Cardiovascular Research**

Researchers in the Cardiovascular Systems Dynamics Laboratory, also known as the “bat lab”, conduct cardiovascular research through non-invasive, *in-vivo* studies on Pallid bats. The bat lab houses a colony of bats and extensive equipment for the study of their blood and blood vessels including high-resolution microscopes connected to computers for recording the experiments in a high-quality video format. The Pallid bat’s thin and transparent wing enables researchers to inspect blood cells, vessel walls, and much of the cardiovascular network visually through a microscope. While bats do not, by nature, stick their wings underneath microscope objective lenses, they have been trained to sleep in a special container designed for this purpose with one wing extended. Researchers in the bat lab observe the wings of napping bats and often save video of the microscope feed of these experiments for later analysis.

### **2.1 Research Skills**

Typically, researchers learn basic concepts by viewing microscope video feeds of the experiments where they can observe the bat’s cardiovascular system. They view a bat wing in its natural condition, as well as by modifying these conditions, such as by

applying pressure to occlude blood-flow to a specific part of the wing. Researchers can identify features such as major and minor arteries and veins, capillaries, and lymphatic vessels. Some features, such as blood vessels, are easy to identify as well as classify. Experienced individuals can quickly gauge whether they are looking at a vein, a second-level venule, or a capillary. Other features, such as lymphatic vessels, are more difficult to recognize. Researchers often use strategies that employ external knowledge in order to locate these features. For example, following a vein longitudinally is likely to help locate physical proximate lymphatic vessels. Acquisition of such skills helps them minimize the time required for basic tasks such as identification of features and focus on the research questions regarding these features.

## **2.2 Image-based Research Skills Assessment**

Our observation, that while conducting experiments, researchers identify basic features on the bat wing almost instantaneously, led us to explore the possibility of using still images instead of videos for testing these skills. Unlike videos, several images can be displayed simultaneously, and in different layouts, thus supporting differences in users' mental models and practices. We tested the use of images for specific tasks with the aid of a pilot user—an advanced graduate student—who effortlessly identified various features on a bat wing using printed photographs. Encouraged by the performance of this user, we designed an experimental evaluation to study the use of still image interfaces for testing basic cardiovascular research skills.

## **3 Evaluation of Image Layouts for Cardiovascular Research**

The evaluation targeted two goals: identification of image layouts for testing critical skills and exploration of digital image collections for training novice researchers. We designed a set of experimental tasks around basic skills that cardiovascular systems researchers apply routinely in the course of their experiments. Performing these tasks involved the use of information gained from neighboring images in a layout. To investigate the role of available contextual information, we chose three image presentation layouts that provide different contextual information to their users.

### **3.1 Subject Characteristics**

Our subject pool comprised of 15 graduate and undergraduate student researchers in the bat lab with varying degree of experience with analyzing bat wing videos. These students, 11 male and 4 female, were between 18 and 34 years old. As active bat cardiovascular researchers, the subjects possessed adequate knowledge of the area for participating in this evaluation without additional training. We classified the subjects who had worked with bat videos for one complete semester in the lab as novices (6 subjects) and those who had completed two or more semesters as experts (9 subjects).

### 3.2 Image Layouts

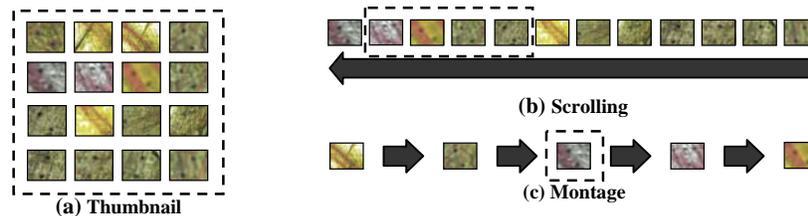
The subjects worked with the three different image layouts diagrammed in figure 1. These layouts provided users with diverse contextual information due to differences in spatial and temporal organization of images illustrated in table 1. The images enclosed by a dashed line in figure 1 correspond to the last column in table 1 and were displayed simultaneously to the users. All layouts displayed individual images of the same size in order to ensure that the size of the images, and hence, the level of visible detail did not affect subjects' perception of the interfaces. We created these layouts by placing the images in a PowerPoint presentation.

**Table 1.** Contextual differences in image layouts.

Layout	Spatial Dimensions	Temporal Dimension	Concurrent Images
Thumbnail	2	No	16
Scrolling	1	Yes	4
Montage	0	Yes	1

**Thumbnail.** The thumbnail layout displayed the images in a 4x4 grid, illustrated in figure 1(a), which allowed us to fit 16 images on one screen while keeping the individual images at a workable size. This layout is widely used for presenting image search results [1, 6, 12] and as a browsing interface for Web-based digital image collections [19, 24]. Simultaneous display of all images enables users to contextualize each image in terms of all others. Users can compare images and detect patterns within the sets. The thumbnail layout presents images in two spatial dimensions and does not employ the temporal dimension, as images in the set are static.

**Scrolling Filmstrip.** Images in the scrolling filmstrip layout, shown in figure 1(b), scrolled smoothly from the right to the left at a constant rate. Our subjects viewed these images from the left to the right and could not control the scrolling rate. The images were displayed at the same size as the thumbnail layout but only four images were visible at any one time. Thus, this layout restricted our subjects' contextual space. When working with an image, subjects could compare this image to its predecessors as well as successors thus providing a forward and backward context for the images. This layout used one spatial dimension, as the images scrolled horizontally, and the temporal dimension, as the displayed images changed over time.



**Fig. 1.** Testing interfaces – images enclosed by the dashed line were displayed simultaneously

**Montage.** The montage layout resembles a slide show where successive images replace earlier ones. This layout, shown in figure 1(c), displayed all the images for an equal amount of time and did not allow users to control this time. The direct context of each image in this layout consisted of a preceding single image, although subjects could have worked with a larger context by recalling additional preceding images. The montage layout did not support a preview or look-ahead mechanism to include the succeeding images in constructing image contexts. This layout used the temporal dimension but no spatial dimensions; all images were displayed at the same location.

### 3.3 Tasks

The subjects answered image-based questions regarding three of the research skills that they acquire via experiments on bat cardiovascular systems. We showed them sets of sixteen images via PowerPoint slides on a computer display. The set size was selected to use the computer's available display space optimally while retaining the integrity of the features presented within the images. The pilot subject verified that he could work with these images and layouts effectively. Our subjects performed the following three tasks, each of which emphasizes a different research skill and embodies a different thought process.

**Artery/Vein Recognition.** The subjects viewed images of bat blood vessels and responded whether the blood vessel indicated by an arrow pointing to it was an artery or a vein. For this task, the subjects used their prior experience and knowledge regarding the visual appearance as well as behavior of arteries and veins to recognize the nature of the indicated blood vessel. For example, the subjects were aware that arteries and veins usually run in pairs and veins are wider than arteries. In addition, while the blood in the arteries flows at a constant rate, the blood flow in veins pulses in a stop-and-go fashion. Thus, a picture in which the cells in a blood vessel are clearly visible is more likely to be that of a vein. We recorded the subjects' verbal responses for later analysis.

**Size Estimation.** The subjects estimated the diameter of the indicated blood vessel at the location of the arrow. This task required them to couple their expertise in identifying the class of a blood vessel (artery, second-level arteriole, or a capillary) with their knowledge of the size for this class of vessels to estimate the diameter of the indicated vessel as accurately as possible. Like the Artery/Vein recognition task, we recorded the subjects' verbal responses for later analysis.

**Lymphatic Vessel Wall Identification.** Lymphatic vessels blend with the surrounding tissue and, hence, are difficult to identify. To make this task tractable, we showed the subjects different images as the vessel went through a complete expansion-contraction cycle. The subjects then marked the location of the wall using PowerPoint's on-screen marker. This task required the subjects to compare the images within each set and analyze the differences between these images. This analysis, cou-

pled with their knowledge of lymphatic vessels enabled them to identify the vessel boundaries, a feature that is generally unidentifiable in a single still image.

### 3.4 Image Data Set

We obtained about 300 digital images for this study from video feeds of past experiments. The Artery/Vein recognition and Size estimation tasks employed about half of these pictures. The images used for these tasks were taken at two magnifications: 10X and 40X. Some images were reused, but each image was used only once for each of the tasks. For the Lymphatic vessel identification task, we identified video sections that displayed a complete expansion-contraction cycle for the vessel and captured a set of 16 images equally spaced over this section.

We obtained definitive answers for each question that our subjects would answer. A panel of three researchers provided the answers for the Artery/Vein recognition task. In order to ensure that we had conclusive answers for the size estimation task, we calculated the exact diameter of the arteries and veins at the locations where our subjects would estimate these values using calipers. We noted the exact boundaries of the lymphatic vessels from pre-recorded videos of observations of these vessels. Using this information, we prepared a comprehensive key that served as a standard for evaluating the test subjects' answers.

### 3.5 Experimental Design

The subjects performed the three tasks sequentially, completing one task before they started another. We reordered the tasks across subjects in order to balance the experiments. A third of the subjects performed the Artery/Vein recognition task (A), followed by the Size estimation task (E), and finally, the Lymphatic vessel identification task (L). The other two sets of subjects performed the tasks in the order E, L, A, and L, A, E.

For each task, subjects viewed the different image layouts in sequential cycles, balanced across the tasks. For task A, subjects viewed the Thumbnail (T), followed by the Scrolling filmstrip (S), and finally, the Montage (M). For task E, they viewed the layouts in the order M, S, T, and for task L, in the order S, T, M. The subjects viewed two cycles of these layouts for the Artery/Vein recognition and Size estimation tasks and answered one question per image (16 questions per layout). For the Lymphatic vessel identification task, they only answered one question per layout as they were looking at a sequence of images. For this task, we had the subject view three layout cycles in order to compensate for the reduction in answers. In order to enable the subjects to work on the images out of order if desired, each image within a view was labeled by a letter.

The subjects had a fixed amount of time for each task regardless of image layout. We timed our tasks based upon the performance of our pilot user. For the Artery/Vein recognition task, the subjects had three seconds for each image (48 seconds for each view), for the Size estimation task, they had four seconds for each image (64 seconds for each view), and for the Lymphatic vessel identification task, they had two seconds

for each image (32 seconds for each view). The evaluation time for each subject was about 30 minutes including the time required for orientation and debriefing. At the outset, we introduced the subjects to each of the layouts. At the beginning of each task, we illustrated the task via an example. Users also could perform a short practice task to acquaint themselves with the task more thoroughly, prior to performing the actual task.

## 4 Results

We collected data from the subjects through a demographic questionnaire, a task questionnaire, responses to the task-specific questions and observation of subjects and their strategies as they performed the tasks. All but two of our subjects (an expert and a novice) had worked with bat videos more than once a week over the course of their cardiovascular research experience. This section presents the results of ANOVA analyses of subject responses.

### 4.1 Artery/Vein Recognition Task

Our subjects reinforced the conventional wisdom that Artery/Vein recognition is one of the first skills that new cardiovascular researchers acquire by returning the best scores for this task. In terms of the different image layouts the subjects, experts as well as novices, consistently performed the best with the familiar thumbnail layout. The experts correctly answered 13.59 questions out of 16 on average (84.94%) to outperform the novices ( $p=0.005$ ), who answered 9.75 questions correctly (61%). Figure 2 illustrates the differences in results of the two populations.

Overall, the experts' performance differed significantly across the individual layouts ( $p=0.005$ ). They performed consistently on the two rounds of the familiar thumbnail layout (89.6%). On the other two layouts, they performed slightly worse than the thumbnail layout on the first round of both the montage (79.2%) and the scrolling layouts (75%), but their performance improved significantly ( $p=0.0004$ ) on the second round as they quickly learnt the nuances of the new layouts. In fact, the

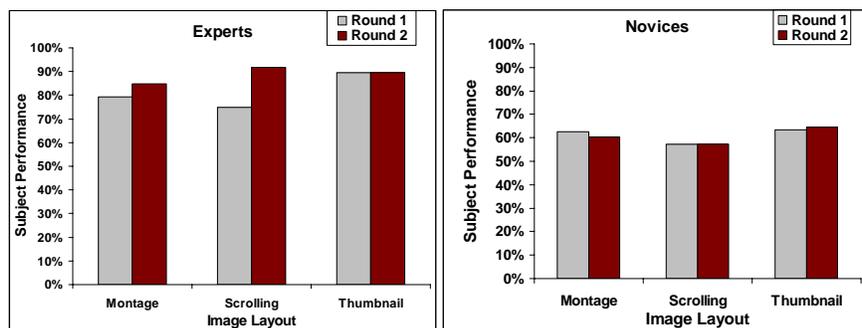


Fig. 2. Performance of subjects on the Artery/Vein recognition task

experts performed better on the second round of the scrolling layout (91.4%) than they did on either of the thumbnail rounds. This suggests that once they adjusted to the “sense of impending doom”—as a user characterized the scrolling view—the scrolling and montage views may be better than the thumbnail view for supporting this task. Although the experts did well with the scrolling layout, most subjects (11 out of 15) thought that they were most effective on the thumbnail layout.

The performance of our novice subjects was consistent for each layout over the two rounds. While the experts improved with practice on the unfamiliar image layouts, we did not see any improvement with the novices. We believe that this is partly due to the difference in expertise of the two groups. While the experts were constrained by the layout and got better with practice, the novices were constrained by their experience rather than the presentation of the images or the contextual information provided by these layouts. Overall, the scrolling layout was the least favored, with only two subjects vouching for this layout.

#### 4.2 Size Estimation Task

The subjects found it more difficult to estimate the diameter of arteries and veins than we had expected. While they regularly intuit the size of such features in their daily research in order to find features of interest, this general sense of proportion did not translate into accurate estimates. The subjects responded with a single number that was their best estimate for the diameter of the vessel in microns. As shown in figure 3, we graded their performance on this task at several levels. At the  $\pm 10\%$  level, answers that fell within a 10% range of the exact answers were considered correct. Thus, for a vessel 40 microns in diameter, subject who guessed the vessel to be between 36 and 44 microns received credit for this question. At the  $\pm 40\%$  approximation range answers between 24 and 56 microns were considered correct.

While the novices found it particularly difficult to estimate the sizes, even the experts could estimate only about 50% of the sizes within a  $\pm 40\%$  range of the actual size. We did not find a significant difference in performance across the various image layouts. However, while the experts’ responses on the scrolling and montage layouts were consistent between the two rounds,

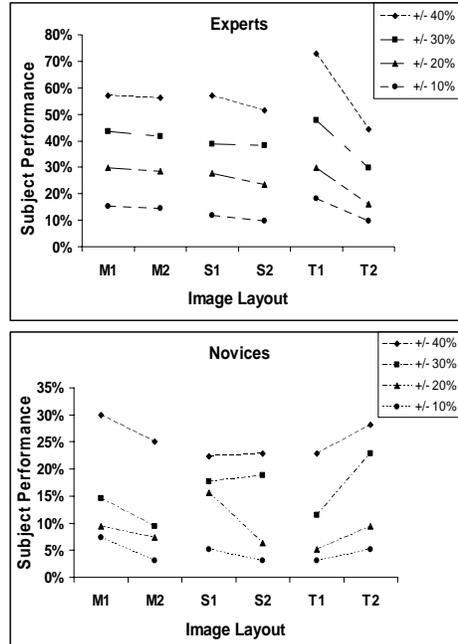


Fig. 3. Result of the Size estimation task

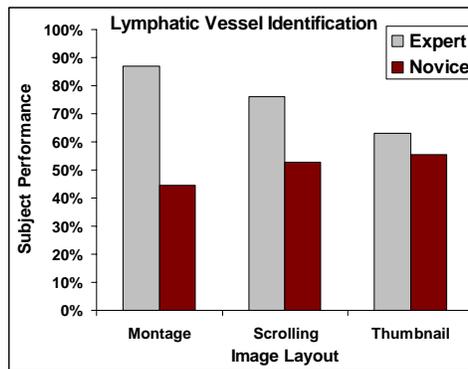
their performance on the thumbnail layout dramatically dropped for the second round. Overall, the experts significantly outperformed the novices on this task as well at all ranges of approximation. 56.5% of the experts' responses were acceptable at the  $\pm 40\%$  range, compared to 25.3% of the novices' responses ( $p < 0.002$ ). At the  $\pm 10\%$  range the acceptable responses were 13.2% and 4.5% respectively ( $p < 0.002$ ).

Some subjects liked the thumbnail view because they could compare the diameters of vessels across images. Thus, they estimated the diameter for a few images and compared the diameters across images to guess the others. We believe that this strategy may have backfired as the subjects may have missed the fact that the images were taken at different magnification factors, a fact that was stressed during their introduction to the task. Some novices had trouble managing the available time when working with the Thumbnail layout and did not have the time to attempt several questions. The subjects used the visual cues provided by the other layouts to respond to the images within the available time. 11 of the 15 subjects liked the thumbnail layout and although we did not find that the layout affected their performance, 12 believed that they worked best with the thumbnail layout.

### 4.3 Lymphatic Vessel Identification Task

While we only discovered the contextual use of information in the Size estimation task during the experiment, this task was designed to encourage use of information across images. For each set, the subjects drew the walls of the lymphatic vessel. Since we are interested in the boundaries of the wall, this approach yielded two responses for each set. We treated the two walls of the vessel independently and accepted answers where the right wall of the vessel was identified as the left wall, as the subject had believed that the vessel was parallel to its actual location.

As illustrated in figure 4, the experts performed better than the novices in correctly identifying walls of the lymphatic vessels. While the experts identified 75.3% of the vessel boundaries across all the layouts, the novices could only identify 51% ( $p = 0.03$ ). When comparing success on this task by the type of view, the experts' performance across the layouts varied significantly ( $p = 0.035$ ). They returned the best results for Montage (87%) and the worst for Thumbnail (63%). This was somewhat expected as the montage layout acted as a slowed down video interface and facilitated the comparison of differences in a set of somewhat similar images. However, the novices surprised us by scoring the lowest with this layout (44%) and the highest with the thumbnail view (55%) which we expected to be the most ill suited layout for this



**Fig. 4.** Performance of experts and novices on the Lymphatic vessel identification task

task. This result contradicted our expectation that the novices and experts alike would benefit the most from the video-like montage layout. In spite of their performance, 4 of the 6 novices (11 of 15 subjects overall) preferred the montage layout and 3 of these believed that they worked better with this layout.

## 5 Discussion

Familiarity with the image layouts may have influenced the subjects' choices for the most-liked and most-effective interfaces. While the subjects had not worked with the scrolling layout before, they adapted to this view quickly and their performance on this view was not significantly worse than the other layouts.

Time management also played a role in user performance with the different layouts. The thumbnail layout allowed users to manage their time freely and devote it to those images that they found challenging. However, this did not turn out to be the most effective test-taking strategy as some subjects lost track of the available time and consequently had to forego answering some questions. In contrast, the other views naturally constrained the time available per image and aided the subjects in managing their time.

Furthermore, while most interfaces for personal photograph presentations employ external contextual information, such as the timestamp [14] or the geographic location [10] the contextual information for the tasks of our interest is derived directly from other images that are displayed simultaneously. Our tasks required subjects to use this contextual information in two ways: in the Size estimation task, subjects used a few images as a baseline for estimating the size of objects in other images, much like unordered reference points; for the Lymphatic Vessel Wall identification, they used the images in sequence and focused on the differences between sequential images. Our tasks also require specialized training and the subjects are typically more interested in the characteristics of images and their content, unlike personal collections where users typically search for familiar images.

Our strategy of using still images to test for video-acquired skills seems somewhat analogous to the extraction of key frames from video [23]. The key frames are expected to convey a sense of the larger collection, in this case, the video from where they are extracted. This is somewhat similar to the concept of collection understanding [7]. Christel, et al., have explored the use of various video abstractions for key frame representation [8]. However, our images do not represent (sections of) videos; we are interested in the cognitive decisions of our users based upon the details within these images. Our technique has more in common with the researchers who slowed down videos of honeybees to analyze the still frames and discover the nature of their flight [2].

Our research contributes to the areas of Digital Libraries and Medical Imaging Systems. Prior research at the intersection of these areas has focused on digital collections for augmenting the virtual laboratory [3] and support for querying over medical information [16]. Other research for imparting scientific training has explored graphics intensive techniques to generate high-quality virtual reality training simulations

[9, 22]. Artificial Intelligence-enhanced tutoring environments have also been designed for training professionals in medical areas [21].

## 6 Future Work

Currently we are exploring mechanisms to shorten the training period for new researchers in order to make their semester-long research experience richer and more productive. This is critical since each semester the bat lab trains a large number of new undergraduate and graduate students in cardiovascular research methods. Typically, new students work with experienced researchers and acquire the necessary skills over the course of the semester. Lab equipment is a critical resource and is only sparingly available to trainees. The semester-long training presents a significant barrier when attracting young researchers and only a few continue their research into the following semesters. An understanding of the interfaces that novice and experienced researchers can work with gained from our study will aid us in designing these training systems for developing novice researchers into experts. Furthermore, since the three layouts that we selected for the purpose of this study are by no means the only candidates for information displays for cognitive tasks. We are currently exploring the contextual properties of other layouts.

In addition, while our subjects can intuit the nature of the blood vessels that they are dealing with, this did not translate into ability to accurately estimate the diameter of these vessels. Clearly, researchers can benefit from additional cues to aid them in accurately applying known real world dimensions to scaled photographs. The nature of the additional information that is necessary for such translations and the mechanisms for delivering this information unambiguously, yet unobtrusively, are significant challenges that merit investigation.

**Acknowledgements.** Jay Widmer, Arun Venugopal, and Missy Young were instrumental in the development of the image corpus.

## References

1. AltaVista - Image Search. 2006. <http://images.altavista.com/> viewed Feb. 2006.
2. Altshuler, D., Dickson, W., Vance, J., Roberts, S. and Dickinson, M. 2005. Short-amplitude high-frequency wing strokes determine the aerodynamics of honeybee flight. *PNAS*, Dec 2005; 102: pp. 18213 - 18218.
3. Bartolo, L., Lowe, C., Sadoway, D., and Trapa, P. 2005. Large Introductory Science Courses & Digital Libraries. In *Proc. of JCDL 2005 (Denver, CO)* ACM Press, pp. 366.
4. Bederson, B. 2001. PhotoMesa: A Zoomable Image Browser Using Quantum Treemaps and Bubblemaps. In *Proc. of UIST '01 (Orlando, FL)* ACM Press, pp. 71-80.
5. Börner, K., Dillon, A., Dolinsky, M. 2000. LVis—Digital Library Visualizer. In *Proc. of IEEE InfoViz 2000 (London UK)* IEEE Press, pp. 77-81.
6. Corbis: stock photography and digital pictures. 2006. <http://pro.corbis.com/> viewed Feb. 2006.

7. Chang, M., Leggett, J., Furuta, R., Kerne, A., Williams, J., Burns, S., and Bias, R. 2004. Collection understanding. In Proc. of JCDL 2004 (Tuscon, AZ) ACM Press, pp. 334-342.
8. Christel, M., Winkler, D., and Taylor, C. 1997. Multimedia abstractions for a digital video library. In Proc. of DL '97 (Philadelphia, PA). ACM Press, pp. 21-29.
9. de Lima, L., Nunes, F., Takashi, R., Rodello, I., Brega, J., and Sementille, A. 2004. Virtual Reality for Medical Training: A Prototype to Simulate Breast Aspiration Exam. In Proc. of VRCAI '04 (Singapore) ACM Press, pp. 328-331.
10. Frappr! – Group Maps for Online Communities. 2006. <http://www.frappr.com/> viewed Mar. 2006.
11. Graham, A., Garcia-Molina, H., Paepcke, A., and Winograd, T. 2002. Time as Essence for Photo Browsing Through Personal Digital Libraries. In Proc. of JCDL 2002 (Portland, OR) ACM Press, pp. 326-335.
12. Google Image Search. 2006. <http://images.google.com/> viewed February 2006.
13. Greenberg, S. and Rounding, M. 2001. The Notification Collage: Posting Information to Public and Personal Displays. In Proc. Of CHI '01 (Seattle, WA) ACM Press, pp. 514-521.
14. Harada, S., Naaman, M., Song, Y., Wang, Q., and Paepcke, A. 2004. Lost in Memories: Interacting with Photo Collections on PDAs. In Proc. of JCDL 2004 (Tuscon, AZ) ACM Press, pp. 325-333.
15. Kerne, A. 2001. CollageMachine: Interest-Driven Browsing Through Streaming Collage. Proc Cast01: Living in Mixed Realities. Bonn, Germany, pp. 241-244.
16. Kholief, M., Maly, K., and Shen, S. 2003. Event-based Retrieval from a Digital Library Containing Medical Streams. In Proc. of JCDL 2003 (Houston, TX) ACM Press, pp. 231-233.
17. Kuchinsky, A., Pering, C., Creech, M., Freeze, D., Serra, B., and Gwizdka, J. 1999. Foto-File: A Consumer Multimedia Organization and Retrieval System. In Proceedings of CHI '99 (Pittsburgh PA), ACM Press, pp. 496-503.
18. Kustanowitz, J. and Shneiderman, B. 2005. Meaningful Presentations of Photo Libraries: Rationale and Applications of Bi-level Radial Quantum Layouts. In Proc. of JCDL 2005 (Denver, CO) ACM Press, pp. 188-196.
19. NYPL Digital Gallery. 2006. <http://digitalgallery.nypl.org/nypldigital/> viewed Feb. 2006.
20. Rodden, K., and Wood, K. 2003. How Do People Manage Their Digital Photographs? In Proc. of CHI 2003 (Ft. Lauderdale, FL) ACM Press, pp. 409-416.
21. Suebnukarn, S. and Haddawy, P. 2004. A Collaborative Intelligent Tutoring System for Medical Problem-based Learning. In Proc. of IUI '04 (Funchal, Madeira) ACM Press, pp. 14-21.
22. Wang, Z., Chui, C., Cai, Y., and Ang, C. 2004. Multidimensional Volume Visualization for PC-based Microsurgical Simulation System. In Proc. of VRCAI '04 (Singapore) ACM Press, pp. 309-316.
23. Wildemuth, B. M., Marchionini, G., Yang, M., Geisler, G., Wilkens, T., Hughes, A., and Gruss, R. 2003. How Fast is Too Fast?: Evaluating Fast Forward Surrogates for Digital Video. In Proc. of JCDL 2003 (Houston, TX) ACM Press, pp. 221-230.
24. Yahoo! Photos. 2006. <http://photos.yahoo.com/> viewed Feb. 2006.