Interacting with Image Sequences: Detail-in-Context and Thumbnails

Oliver Kuederle, Kori M. Inkpen, M. Stella Atkins
{okuederl,inkpen,stella}@cs.sfu.ca
School of Computing Science
Simon Fraser University
Burnaby, BC, V5A 1S6, CANADA

M. Sheelagh T. Carpendale
sheelagh@cpsc.ucalgary.ca
Department of Computer Science
University of Calgary
Calgary, AB, T2N 1N4, CANADA

Abstract

An image sequence is a series of interrelated images. To enable navigation of large image sequences, many current software packages display small versions of the images, called thumbnails. We observed radiologists during typical diagnosis sessions, where image sequences are examined using photographic films and sophisticated light screens. Based on these observations and on previous research, we have developed a new alternative to the presentation of image sequences on a desktop monitor, a variation of a detail-in-context technique. This paper describes a controlled experiment in which we examined the way users interact with detail-in-context and thumbnail techniques. Our results show that our detail-in-context technique accommodates many individual strategies whereas the thumbnail technique strongly encourages sequential examination of the images. Our findings can assist in the design and development of interactive systems that involve the navigation of large image sequences.

Keywords: Image sequences, detail-in-context, thumbnails, medical imaging, information visualization

1 Introduction

Areas such as Magnetic Resonance Imaging (MRI), meteorology, or video editing typically involve viewing a large number of interrelated images. In MRI, an image sequence consists of successively scanned image slices of a volume such as the human brain, a knee, or a shoulder. In order to diagnose a patient's condition, radiologists traditionally examine MR image sequences on silver-based films that are mounted onto a large and sophisticated light screen. A typical examination often involves up to eight such films with a total number of more than a hundred images. Due to the high costs associated with the production and archival of these films, the maintenance of the light screen, and the occasional loss of patient data, many hospitals are now implementing solutions based on computer hardware and software. The use of desktop monitors, however, has an inherent disadvantage. The display area, in which the images must be viewed, is severely limited in terms of space. This is often referred to as the screen real-estate problem.

Based on previous literature and on our observations of radiologists in their workplace, we designed a new technique to display image sequences on a desktop monitor. This technique is a variation of a *detail-in-context* technique. Detail-in-context techniques visualize information using multiple magnification factors. High magnification factors are assigned to user-selected areas to provide detail. In order to provide contextual information, the magnification factors of unselected areas are typically adjusted to fit the remaining screen space. There has been significant research on variations of the detail-in-context technique, some of which have reported evaluations based on user studies. While some of these studies provide statistical results for users' performance with the various visualization techniques, very little is known about the way users interact with them. As a result, many questions concerning how such a technique can be adapted to a particular application remain unanswered.

Towards this end, we have run a controlled user study with the goal of gaining a better understanding of how users interact with two image presentation techniques: the *thumbnail* technique, which is used in many commercially available medical imaging systems, and our detail-in-context technique. In the next section, we describe some of the underlying ideas for these two display techniques. This is followed by a description of our user study. We then conclude with a discussion of the results, the impact on radiology, and pointers to future work.

2 Background

2.1 Detail-in-Context

Detail-in-context techniques, which date back to applications such as Furnas's "Fish-eye Views" [5] in 1986, have been evaluated in a number of studies.

While Björk and Redström [1], Fisher *et al.* [4], and Furnas [5] ran studies with inconclusive results, Hollands *et al.* [6], Leung *et al.* [9], and Schaffer *et al.* [10] were able to provide statistically significant differences between detail-in-context techniques scrolling views, a technique which displays information at a single magnification level and allows navigation with scrollbars. The results of the study of Hollands *et al.* did not provide evidence that the use of detail-in-context improved user performance. However, the studies of Leung *et al.* and Schaffer *et al.* reported superior performance of detail-in-context over the scrolling view. None of the mentioned studies provide an accurate description of how users interacted with each technique despite the fact that such information may help in the design of detail-in-context techniques for new applications.

2.2 Medical Imaging

Picture Archival and Communication Systems (PACS) offer functions to view medical images on a desktop monitor. The user interface of some of these systems is described in studies by Dayhoff and Kuzmak [2] and Erickson *et al.* [3]. While some systems can only display a fixed number of images at a time, others provide some context with thumbnail bars that contain small versions of the images that can be selected for further magnification in a separate window. Honea *et al.* [7] present an evaluation of five commercial software products developed for the PC. It was determined that none of the tested systems offered an adequate set of tools required during diagnosis. The authors state that this "seems to be the result of incomplete requirement definition, inadequate software development, or deliberate decisions to limit product development." [7]

2.3 Introducing Detail-in-Context to Medical Imaging

Van der Heyden *et al.* [11] observed radiologists during MRI examination and performed a requirements analysis based on these observations. The identified requirements suggested the use of a detail-in-context technique to display a large number of MR images on a desktop monitor. In an informal study involving three radiologists and screen shots of various detail-in-context layouts, van der Heyden *et al.* showed that detail-in-context was feasible since lower magnification factors are often sufficient to distinguish images. Although this research encourages the use of detail-in-context techniques in a radiology context, a follow-up controlled study would provide additional insight into users' interactions with this detail-in-context technique.

2.4 Field Observations

Extending the work of van der Heyden *et al.* [11], we conducted additional informal field observations of radiologists at work at a local hospital. While interacting with the photographic MRI films, the radiologists made gestures that suggested that the images on the films were viewed as part of a one-dimensional sequence, rather than part of the grid in which they were arranged. This observation led to an extension of the detail-in-context technique presented by van der Heyden *et al.* We describe this technique in the following section.

2.5 Our Implementation

Based on previous work and on our field observations, a number of constraints were identified for our detail-in-context technique, including:

1. All images in an image sequence are visible on the screen.
2. User-selected images have a fixed magnification factor. When running out of screen space, this factor is reduced for all selected images.
3. Images are aligned along rows.
4. Images do not move between rows.
5. Space between images remains black.
6. Images are no smaller than 30×30 pixels.
7. Unselected images are equally distributed to reduce the number of different magnification factors.
8. Consecutive layouts are interpolated in ten intermediate steps. These smooth transitions provide visual feedback to the user when the layout changes.

Figure 1 shows a screen shot of our detail-in-context implementation. The images of one image sequence are displayed in the main display area according to our layout algorithm. Space between images remains black. A mouse click selects an image and causes it to be magnified. A second mouse click de-selects the image and returns it to its minimized state. Additional functions include a menu to specify the magnification factor for the selected images and a button labeled

Figure 1. The detail-in-context implementation.

"Done/Next" to bring up the next image sequence.

A thumbnail technique was also implemented for the presentation of image sequences. This software is similar to the medical imaging package that we observed in use at a local hospital[1]. Figure 2 shows a screen shot of our implementation. On the right hand side of the screen, a thumbnail bar shows small versions of the images in the sequence; for square images the thumbnails are each 80x80 pixels. Clicking on an image causes it to be displayed in the top left corner of the large display area. The large display area, which occupies most of the screen, shows a subset of the image sequence at high magnification factors. Only consecutive images are displayed in the large display area and the layout can be changed by pressing one of the buttons on the left hand side. Similar to the detail-in-context implementation, the "Done/Next" button loads the next image sequence.

Both programs were written, compiled, and run with Sun Microsystem's Java 1.2.2 to allow execution on other platforms. In our research, the software was run on a Pentium III 500MHz PC with a 21" monitor at a resolution of 1024×768 pixels.

3 Empirical Study

3.1 Overview and Setting

We conducted an experiment that involved two conditions: detail-in-context and thumbnails. The experiment took place at Simon Fraser University, Canada, in April 2000. To gain a better understanding of how users interact with image sequences using the two display techniques, a combination of qualitative and quantitative analyses was performed on the data collected during the study. A more extensive description of the study and the results is given in Kuederle's M.Sc. thesis [8].

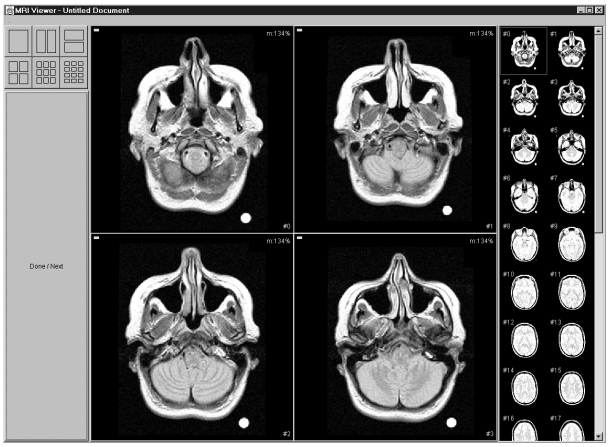

Figure 2. The thumbnail implementation

[1] The software is called "Advantage Windows" by General Electric Medical Systems

3.2 Participants

Thirty-two university students from various disciplines participated in the study. It was decided not to involve radiologists for three main reasons. Only a limited number of MRI radiologists were available in the Greater Vancouver area. Due to their heavy workload, they were not able to spend sufficient time to take part in the study. Furthermore, the logistics of obtaining real patient data would have delayed our study significantly.

3.3 Experimental Task

Background

The task in our study was modeled in part after the radiologists' work. The following aspects of their work were maintained:

- The presented image sequence showed a familiar object.
- Some image sequences contained an anomaly.
- The participants were asked to find and describe the anomaly.
- Images were displayed in order.
- Only grayscale images were shown.

Description

In both the detail-in-context and the thumbnail condition, five image sequences were presented to the participants. Each sequence showed a familiar object. The object was rotated around its vertical axis in fixed angle steps so that in each image of the sequence, it was shown from a different perspective. Refer to Figure 3 for an example.

An artificial anomaly was placed on the object. This anomaly was either black or white and its shape was one of the four suits: clubs, spades, hearts, or diamonds. The anomaly remained in the same spot on the object but was randomly removed in some images so its

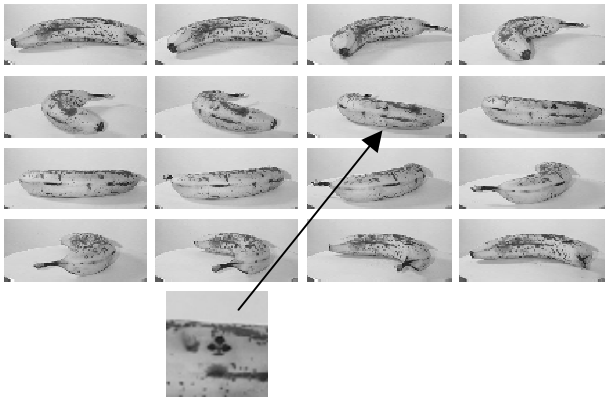

Figure 3. An example image sequence.

occurrence was unpredictable. Two image sequences did not contain any anomaly.

For each of the presented image sequences, the participants were asked to report the shape of the anomaly as well as which images it appeared in. We provided an answer sheet on which shape and image numbers could be circled.

Concerns

In our attempt to model a task that was similar to the radiologists' work, there were several aspects we did not account for:

- Radiologists spend far more time diagnosing patients than it took our participants to complete the study.
- Only one image sequence was shown at a time, while radiologists typically examine multiple sequences simultaneously (e.g. a proton density sequence in combination with a T2 sequence).
- Our participants did not have any prior training in the examination of image sequences.
- The presented images showed an object from different perspectives whereas in MRI, images show consecutive slices of a volume. We were concerned that extensive training was required for the ability to build a three-dimensional mental model, given two-dimensional slices.
- The participants were not required to report the intensity level of the anomaly.

Although our results may not be directly applicable to the area of radiology due to these tradeoffs, the selection of a more general population sample and the nature of our experimental task may allow for generalization to a wider variety of areas.

3.4 Procedure and Data Collection

After an introduction to our research, the participants filled out a background questionnaire that assessed their experience with computers and digital images. We then explained the experimental task. The participants completed the task in the two conditions. Before each condition, they were given the opportunity to practice with the software. After the second condition, they filled out a post-session questionnaire where they indicated their preference for one of the two display techniques. Additional space was provided for comments.

During the participants' interaction with the software, a computer log was recorded with events such as mouse clicks, magnification changes, and the end of a trial. This information was later analyzed to identify trends and patterns in the participants' interaction with each display technique. In addition, we performed statistical analyses on the participants' performance and preference.

3.5 Independent Variables

We identified three independent variables:

- *Display Condition:* There were two conditions: detail-in-context and thumbnails.
- *Image Sequence Set:* We created two image sequence sets of similar difficulty level. Each set contained five image sequences whose order within the set was randomized.
- *Gender:* An equal number of men and women participated in the study.

The participants were presented with two sets of image sequences, each set assigned to one condition. The experiment was a 2×2×2 (condition×set×gender) mixed design, with gender as the between subjects factor and condition and set as within subjects factors. All independent variables were counterbalanced, resulting in four condition×set groups with four females and four males in each group.

3.6 Dependent Variables

Nature of Interaction and Comments

The focus of our study was to investigate the way users interact with the two display techniques. Recording the participants' actions in a computer log allowed us to examine their behaviour in order to identify patterns, trends, and differences for each display technique. The information provided in the post-session questionnaire was used to collect feedback from the participants about each display technique.

Performance and Preference

We also performed statistical analyses on the participants' performance and preference for a display technique to determine if any significant differences existed between the two techniques.

4 Results and Discussion

4.1 Nature of Interaction and Comments

Trial Charts

In order to examine the way the participants interacted with the detail-in-context and the thumbnail technique, we visualized the computer log of each trial in a *trial chart*. Figure 4 shows a trial chart in the detail-in-context condition. Time (in milliseconds) is displayed on the horizontal axis and image numbers are displayed on the vertical axis. The left end of the trial chart denotes the beginning of a trial and the right end denotes the end of a trial. Gray bars indicate images that contain an anomaly. Solid dots represent mouse clicks on an image at a specific time. A horizontal line between two dots indicates the time during which an image was selected. Therefore, the dot to the left of that

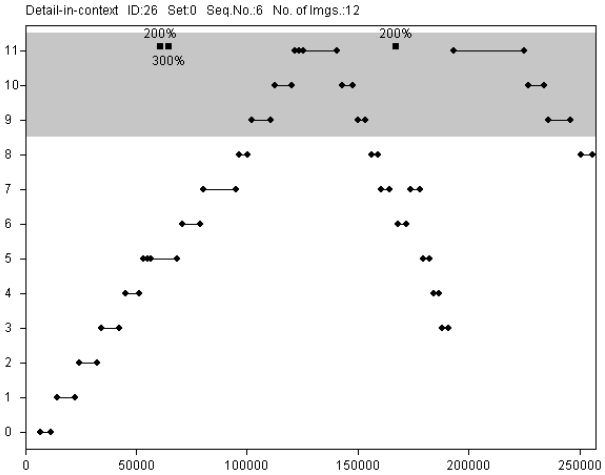

Figure 4. A trial chart in the detail-in-context condition (time in milliseconds on the horizontal axis, image numbers on the vertical axis).

line represents a magnification event and the dot to the right represents a minimization event. Additional events such as magnification factor changes (100%, 150%, 200%, or 300%) are displayed in the top row of the trial chart.

The trial charts in the thumbnail condition are similar, as shown in Figure 5. Solid dots represent mouse clicks on thumbnails. The lines between the dots indicate the order of the events. Layout changes in the large display area (1×1, 1×2, 2×1, 2×2, 3×3, and 3×4) are displayed in the top row of the trial chart. Additionally, empty dots represent clicks on images in the large display area (this did not have any effect on the software but was nonetheless recorded).

Examination Strategies

All trial charts in the thumbnail condition reveal an interaction pattern similar to the one displayed in Figure 5. During the majority of time for a trial, images were examined in sequential order with a few iterations on images with an anomaly. Note that, as shown in Figure 5, the first few mouse clicks were made in steps of four because, by default, the large display area was set to a 2×2 layout.

Sequential strategies were also observed in the detail-in-context condition. However, some participants selected images in a random order, as shown in Figure 6. In some instances, participants applied a mix of both strategies, typically consisting of an initial exploration phase and a final sequential check (see trial chart in Figure 7).

In the post-session questionnaire, eight participants stated that it was easier to spot anomalies in the detail-in-context condition because it provided a global overview of the image sequence. Five participants

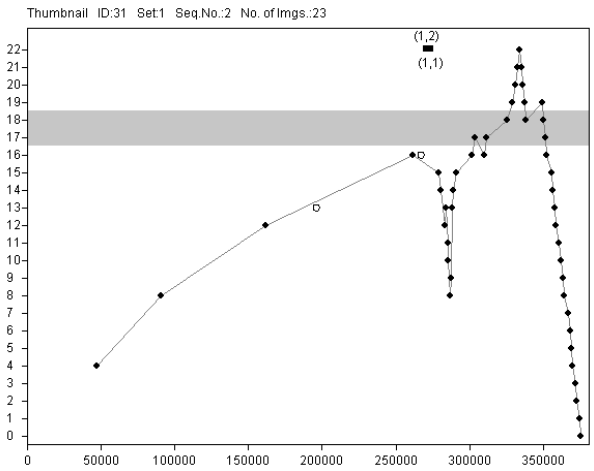

Figure 5. A trial chart in the thumbnail condition (time in milliseconds on the horizontal axis, image numbers on the vertical axis).

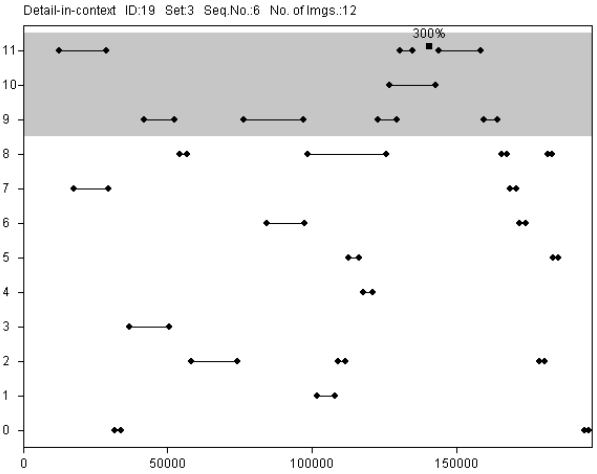

Figure 6. Random selections in the detail-in-context condition.

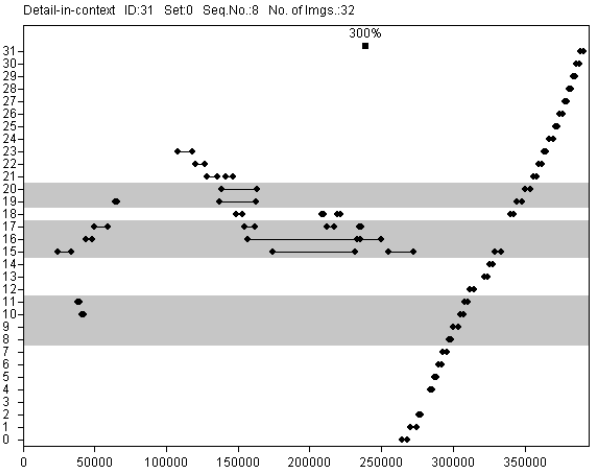

Figure 7. Final checks in the detail-in-context condition

mentioned that it was difficult to keep track of image numbers in the thumbnail condition because the thumbnail bar and the large display area were in two separate windows.

The detail-in-context technique displayed all images on screen at all times during a trial. Thus, for an image sequence with 15 square images, each image occupies 130x130 pixels. As we observed, this facilitated the selection of images that immediately attracted the users' attention. Some users examined images according to their current focus of attention while some applied a sequential strategy. The detail-in-context technique supported both. On the other hand, in the thumbnail technique, the thumbnails had a resolution of 80×80 pixels and, therefore, did not display sufficient detail to detect anomalies. Thus the thumbnail technique strongly encouraged the users to examine (by magnifying) small subsets of images at a time, selecting them in sequential order.

Image Comparisons

Some participants in the detail-in-context condition did not make any comparisons between multiple magnified images. Figure 4 shows such a trial. Other participants frequently made comparisons, magnifying typically three, sometimes up to six images at a time. In some cases, images were kept magnified as a reference while the remaining images were compared to that reference (e.g. images #15 and #16 in Figure 7). The comment most frequently made in the post-session questionnaire (by thirteen participants) was appreciation of the fact that with the detail-in-context technique, random combinations of images could be picked for comparison. Images that were not adjacent, e.g. image #2 and #14, could be magnified, whereas with the thumbnail technique, only images in close proximity could be magnified at the same time.

Skipping of Images

In the detail-in-context condition, participants frequently skipped images, not magnifying them at all. Figure 8 shows a trial in which, for example, images #1 through #7 were not magnified. We observed that a large number of images were skipped when there were fewer images in the sequence. For one of the image sequences, most of the computer logs did not contain any events because the anomaly could be clearly seen at a magnification of less than 100%.

In the thumbnail condition, images were skipped in only four out of 160 trials. The low resolution of the thumbnails made it more difficult to examine an image sequence without magnifying all images whereas in the detail-in-context condition, images were sometimes

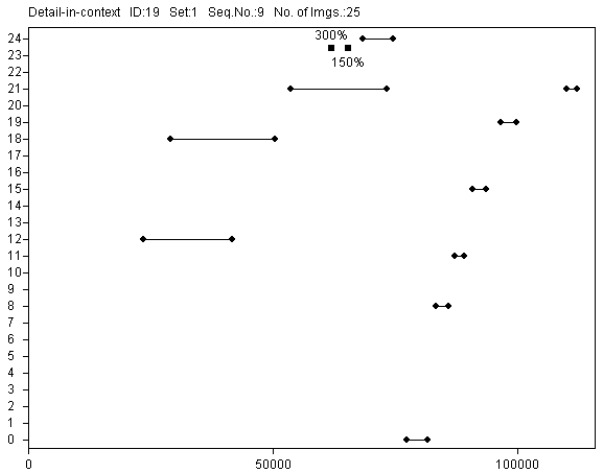

Figure 8. Skipping of images in the detail-in-context condition.

large enough to recognize the anomaly without further magnification.

The Space Tradeoff

In many trials in the detail-in-context condition, the magnification factor was set to the highest setting of 300%, as can be seen in Figures 4, 6, and 7. Correspondingly, in the thumbnail condition, the participants frequently selected a 1×1 layout (see Figure 5).

Six participants said that in the detail-in-context condition, the selected images often did not attain the desired size because the surrounding images always remained on the screen. Even when the magnification was higher than 100%, participants frequently requested larger images. The thumbnail technique, on the other hand, provided high magnification factors since only a few images were shown in the large display area.

Figure 5 shows empty circles that represent mouse clicks on images in the large display area. These clicks were observed for participants in all experimental groups. The clicks did not have any effect in the software. Frequently, they were followed by layout changes that resulted in higher magnifications for these images. This suggests that the images in the large display area afforded interaction, perhaps with the intent of an additional increase in magnification.

Motion Sickness

Eight participants reported motion sickness due to the animation in the detail-in-context condition. The low animation frame rate (approximately 10 frames per second) or the constant layout change may have posed problems for these participants.

4.2 Performance and Preference

We identified four performance measures: time to complete all trials in a condition, number of false negatives in a condition, number of false positives in a condition, and number of wrong symbols in a condition. False negatives were anomalies that were not reported. False positives were anomalies that were reported although non-existent. Wrong symbols were misinterpretations of the shape of the anomalies.

We performed repeated measures analyses of variance on all four measures (α=.05). A significant Time × First Condition interaction was found ($F(1,24)=9.004$, $p=.006$). Further analysis revealed that participants starting with the detail-in-context condition improved significantly in trial completion time ($F(1,12)=6.846$, $p=.023$). This effect was not found for participants starting with the thumbnail condition ($F(1,12)=3.162$, ns). Although the participants were given time to practice on the interface, the presentation of an unfamiliar task with the detail-in-context technique may have required some adjustment time. Thumbnails are used in a number of applications as well as on the World Wide Web. Some participants may have been familiar with this kind of representation. For the three accuracy measures false negatives, false positives, and wrong symbols, no significant effect was found.

A χ^2-analysis was performed to determine if the participants' preference for one of the two display techniques was significantly higher. No such difference was detected.

In the post-session questionnaire, we asked participants why they preferred a technique. Participants who preferred detail-in-context mentioned the good comparison capabilities and the global overview as the main reason for their choice. Most participants who preferred thumbnails stressed negative characteristics of the detail-in-context technique, such as the animation causing motion sickness and the fact that images had to be clicked on twice, i.e. for magnification and minimization.

5 Summary

In the design of interactive systems, questions may arise in the adaptation of an interaction technique to a specific application. Quantitative analyses often depend on a specific implementation and may not convey information about the way users interact with a system. In the area of detail-in-context viewing, the results of user studies have been inconclusive and may not be helpful in the design of real applications. We have run an experiment to study the way users interact with a detail-in-context technique and a thumbnail technique, both used to present image sequences on a desktop monitor. Our main results for the detail-in-context technique were:

- It accommodates a wide variety of individual strategies.
- It provides a global overview, facilitating the search for images of interest.
- It allows for comparisons between any pair of images.
- High magnification factors are rarely attained because space is required to display contextual images.
- Users may experience motion sickness if the animation frame rate is low and/or the layout changes frequently.
- Time may be required to familiarize users with detail-in-context techniques.

The main results for the thumbnail technique were:

- It strongly encourages examination of the images in sequential order.
- It discourages skipping of images.
- More space is available for selected images and, therefore, higher magnification factors can be attained.
- Multiple magnifications are restricted to consecutive images.

Our statistical analyses did not reveal significant differences in performance between the detail-in-context and the thumbnail technique. This suggests that both techniques are equally valid approaches to the presentation of image sequences for tasks similar to the one in our experiment.

6 Impact on Radiology

We presented our results to a radiologist at the University of British Columbia Hospital in Vancouver, Canada. Given our presentation of the differences between the two display techniques, the radiologist was able to provide valuable information about the way radiologists examine MR images, including the following:

- Radiologists are, by law, required to examine all images of a patient.
- During the radiologists' extensive training, they are taught to examine images in order, even if an obvious anomaly distracts their attention.
- Image sequences are examined multiple times. In each pass, the focus is on a different anatomic region.
- Only consecutive images are compared.

This suggests that the thumbnail technique may indeed be appropriate in the examination of MR images. It was noted, however, that specialized physicians are often only interested in a specific region of the images and

would therefore like to work with the detail-in-context technique because it allows them to focus quickly on critical images.

7 Conclusion

In this paper, we have presented an implementation and an evaluation of two techniques for the presentation of image sequences: detail-in-context and thumbnails. We have shown that both techniques are equally valid approaches to the presentation and navigation of image sequences. The way users interact with the implementation, however, differs for both techniques. Our main findings are:

- The detail-in-context technique accommodates a wide variety of individual strategies and provides good comparison capabilities.
- The thumbnail technique strongly encourages sequential examination of the images and allows for high magnification factors.

These findings serve to improve our understanding of how users interact with detail-in-context and thumbnail techniques. Designers can make use of this information to choose the image presentation technique which is appropriate for their specific task or user-base.

8 Future Work

We present two directions indicated by our research. In a study involving radiologists and real patient data, one could study how specialists interact with all these techniques in a diagnosis setting. We are interested in the way our results apply to a study that involves trained specialists.

We would also like to extend our work in other fields that involve the examination of image sequences, such as meteorology and video editing. In future projects, one could investigate how our findings can be applied to these areas.

9 Acknowledgments

We would like to thank Kellogg Booth and Yvette Cheung for their valuable input. Many thanks also to all members of the EDGE, GrUVi Lab, and the Medical Imaging Lab for their feedback. Thanks also to the Canadian NSERC and the German DAAD for funding this work.

References

[1] Björk, S. and Redström, J. An alternative to scrollbars on small screens. In *CHI'99 Extended Abstracts of the Human Factors in Computing Systems Conference*, pages 316-317. ACM SIGCHI, May 1999.

[2] Dayhoff, R.E. and Kuzmak, P.M. Providing complete multimedia patient data to consulting radiologists, other specialists, and the referring clinician. *Journal of Digital Imaging*, 11(3):134-136, August 1998.

[3] Erickson, B.J., Ryan, W.J., Gehring, D.G., and Beebe, C. Clinician usage patterns of a desktop radiology information display application. *Journal of Digital Imaging*, 11(3):137-141, August 1998.

[4] Fisher, B., Agelidis, M., Dill, J., Tan, P., Collaud, G., and Jones, C. CZWeb: Fish-eye views for visualizing the world-wide web. *Design of Computing Systems: Social and Ergonomic Considerations*, 2(21B):719-722, 1997.

[5] Furnas, G.W. Generalized fisheye views. In *Proceedings of the ACM Conference on Human Factors in Computer Systems*, SIGCHI Bulletin, pages 16-23. ACM, New York, USA, 1986.

[6] Hollands, J.G., Carey, T. T., Matthews, M. L., and McCann, C. A. Presenting a graphical network: A comparison of performance using fisheye and scrolling views. In *Proceedings of the Third International Conference on Human-Computer Interaction*, volume 2, pages 313-320, 1989.

[7] Honea, R., McCluggage, C., Parker, B., O'Neall, D., and Shook, K.A. Evaluation of commercial PC-based DICOM image viewer. *Journal of Digital Imaging*, 11(3):151-155, August 1998.

[8] Kuederle, O., Presenting Image Sequences – A Detail-in-Context Approach. *M.Sc. Thesis, School of Computing Science, Simon Fraser University*, August 2000.

[9] Leung, Y.K., Spence, R., and Apperley, M.D. Applying bifocal displays to topological maps. *International Journal of Human-Computer Interaction*, 7(1):79-98, 1995.

[10] Schaffer, D., Zuo, Z., Greenberg, S., Bartram, L., Dill, J., Dubs, S., and Roseman, M. Navigating hierarchically clustered networks through fisheye and full-zoom methods. *ACM Transactions on Computer-Human Interaction*, 3(2):162-188, 1996.

[11] Van der Heyden, J.E., Carpendale, M.S.T., Inkpen, K.M., and Atkins, M.S. Visual presentation of magnetic resonance images. In *Proceedings of the 9th Annual IEEE Conference on Visualization*, pages 423-426, New York, October 18-23 1998. ACM Press.