

Comparison of Navigation Techniques for Large Digital Images

Bradley M. Hemminger,^{1,3} Anne Bauers,² and Jian Yang¹

Medical images are examined on computer screens in a variety of contexts. Frequently, these images are larger than computer screens, and computer applications support different paradigms for user navigation of large images. The paper reports on a systematic investigation of what interaction techniques are the most effective for navigating images larger than the screen size for the purpose of detecting small image features. An experiment compares five different types of geometrically zoomable interaction techniques, each at two speeds (fast and slow update rates) for the task of finding a known feature in the image. There were statistically significant performance differences between several groupings of the techniques. The fast versions of the ArrowKey, Pointer, and ScrollBar performed the best. In general, techniques that enable both intuitive and systematic searching performed the best at the fast speed, while techniques that minimize the number of interactions with the image were more effective at the slow speed. Additionally, based on a postexperiment questionnaire and qualitative comparison, users expressed a clear preference for the Pointer technique, which allowed them to more freely and naturally interact with the image.

KEY WORDS: User interfaces, human factors, medial image display, interaction techniques, pan, zoom, performance evaluation

INTRODUCTION

Viewing images larger than the user's display screen is now a common occurrence. It occurs both because the spatial resolution of digital images that people interact with continues to increase and because of the increasing variety of smaller resolution screens in use today (desktops, laptops, PDAs, cell phones, etc.). This leads to an increased need for interaction techniques that enable the user to successfully and quickly navigate images larger than their screen size.

People view large digital images on a computer screen in many different kinds of situations. This paper draws from work in many fields to address one of the most common tasks in medical imaging, finding a specific small-scale feature in a very large image. An example is mammographers looking for microcalcifications or masses in mammograms. For this study, large images are defined as images that have a spatial resolution significantly larger than their viewing device, i.e., at least several times larger in area. It may additionally be constrained by the user operating within a window on that screen that further constrains the available resolution. For instance, a user may wish to navigate a digital mammogram image that is $40,000 \times 50,000$ pixels on a personal computer screen that is $1,024 \times 768$ pixels in a window of size 800×600 pixels.

¹From the School of Information and Library Science, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3360, USA.

²From the 4909 Abbott Ave. S., Minneapolis, MN 55410, USA.

³From the Department of Radiology in School of Medicine, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3360, USA.

This research was supported in part by the National Institutes of Health (NIH) Grant # RO1 CA60193-05, US Army Medical Research and Materiel Command Grant # DAMD 17-94-J-4345, NIH RO1-CA 44060, and NIH PO1-CA 47982.

Correspondence to: Bradley M. Hemminger, Department of Radiology in School of Medicine, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3360, USA; tel: +1-919-9662998; fax: +1-919-9668071; e-mail: bmh@ils.unc.edu

Copyright © 2008 by Society for Imaging Informatics in Medicine

Online publication 27 June 2008
doi: 10.1007/s10278-008-9133-0

In the past, computer and network speeds limited the speed at which such large images could be manipulated by the display device, limiting the types of interaction techniques available and their effectiveness. As computer and network speeds have increased, it is now possible to interactively manipulate images by panning and zooming them in real time on most computer-based display systems, including the graphics cards found on standard personal computers. The availability of interactive techniques supporting real-time panning and zooming provides for the possibility of improved human-computer interactions. However, most interactions in existing commercial applications as well as freely available ones do not take advantage of improved interaction techniques or necessarily use the techniques best suited for capabilities of their particular display device. To test different interaction techniques, five different interaction techniques supported by imaging applications were selected.

In order to quantitatively compare the performance of different techniques, we must be able to measure their performance on a specific task. There are many types of tasks and contexts in which users view large images. In this study, we chose to examine the task of finding a particular small-scale feature within a large image. This task was chosen because it is a common task in medical imaging, as well as in other related fields such as satellite imaging.^{1,2} In addition to the interaction technique, the speed of updating the image view may affect the quality of the interaction. Several factors can affect the update rate, including processor speed and network connection speed. Increasingly, radiologists read from teleradiology systems, where images may be displayed on their local computer from a remote image server. To model this situation where images may be loaded over a slower internet connection, as compared to directly from the local computer memory, two display update rate conditions were tested. The slower update rate also corresponds to the typically slower computational speeds of small devices (PDAs, cell phones) and serves to model these situations as well. A change in the speed of image updates on the screen can dramatically affect the user experience resulting from the same interaction technique. To address this issue, we tested five different interaction techniques, with each technique evaluated with both a fast and a slow update rate.

BACKGROUND AND RELATED WORK

There has been interest in viewing large digital images since the start of digital computers and especially since the advent of raster image displays. Several decades ago, researchers began to consider digital image interpretation in the context of image display.³ Today, digital image viewing and interpretation plays a vital role in many fields, including much of medical practice. Digital images are now routinely used for much of medical practice including radiology.⁴⁻⁶

This paper is concerned with navigational and diagnostic uses (as defined by Plaisant et al.⁷) of digital images when displayed on screens of significantly smaller size. We limited our focus to techniques used on standard computing devices, i.e., not having special displays or input devices and used geometric zooming. Nongeometrical methods (like fisheye lens zooming) are not considered because the size and spatial distortions that occur to the images are not acceptable in medical imaging practice. Interfaces that provide the ability to zoom and pan an image have been termed “zoomable” interfaces in the human-computer interaction literature.⁸ Two well-developed environments that support development and testing of general zoomable interfaces are the Pad++⁹ and Jazz toolkits.¹⁰ To date, few studies have examined digital image viewing from the perspective of maximizing effective interface design for the task of navigating and searching out features within a single large image. There is, however, a significant body of literature in related areas.

Studies on Related Topics

Many researchers have examined the transition from analog to digital presentations, especially in medical imaging.¹¹⁻¹⁶ Substantial work has been done with nongeometrical zoomable interfaces including semantic zooming,^{8,17} distortion-based methods (fisheye),¹⁸⁻²⁰ and sweet spots on large screens.²¹ A summary of these different types of methods can be found in Schaffer et al.²² Additionally, much work has focused on searching through collections of objects. Examples include a single image from a collection of images,^{9,23-26} viewing large text documents or collections of documents,^{22,27} and viewing web pages.²⁸ Meth-

ods that involve changing the speed of panning depending on the zoom scale may have some relevance to our results. These methods have been developed to allow users to move slowly at small scales (fine detail) and more quickly over large scales (overviews). Cockburn et al.²⁹ found that two different speed-dependent automatic zooming interfaces performed better than fixed speed or scrollbar interfaces when searching for notable locations in a large one-dimensional textual document. Ware and Fleet³⁰ tested five different choices for automatically adjusting the panning speed, primarily based on zoom scale. They found that two of the adaptive automatic methods worked better than three other options, including fixed speed panning, for the task of finding small-scale boxes artificially added to a large map. Their task differs from our study in that their targets were easily identified at the fine-detail scale. Difficult-to-detect targets require slower, more careful panning at the fine-detail scale, which probably negates the advantage of automatic zooming methods for our task.

Closely Related Studies

One of the first articles addressing navigational techniques for large images was the article of Beard and Walker,³¹ which found that pointer-based pan and zoom techniques performed better than scrollbars for navigating large-image spaces to locate specific words located on tree nodes. They followed this work with a review of the requirements and design principles for radiological workstations^{32,33} and an evaluation of the relative effects of available screen space and system response time on the interpretation speed of radiologists.^{34,35} In general, faster response times for the user interface, larger screen space, and simpler interfaces (mental models) performed better.³³ This was followed by timing studies that established that computer workstations using navigational techniques to interact with images larger than the physical screen size could perform as well or better than their analog radiology film-based displays.^{11,16,34,35} Gutwin and Fedak²⁰ studied the effect of displaying standard workstation application interfaces on small screen devices like PDAs. They found that techniques that supported zooming (fisheye, standard zoom) were more effective than just panning and that determining which

technique was most effective depended on the task. Kaptelinin³⁶ studied scrollbars and pointer panning, the latter method evaluated with and without zooming and overviews. His test set was a large array of folder icons, with the overall image size nine times the screen size. Users were required to locate and open the folders to complete the task. He found the pointer panning technique performed faster than scrollbars and was qualitatively preferred, likely due to it not requiring panning movements to be broken down into separate horizontal and vertical scrollbar movements. Also, he found the addition of zooming to improve task speed. Hemminger³⁷ evaluated several different digital large-image interaction techniques as a preliminary step in choosing one technique (Pointer) to compare computer monitor versus analog film display for mammography readings¹⁶. However, the evaluation was based on the users' qualitative judgments and did not compare the techniques quantitatively.

Despite the relative lack of research in the specific area of digital-image-viewing techniques, many applications exist for viewing digital photographs, images, and maps. Online map providers such as Mapquest (available at <http://www.mapquest.com>, accessed September 2005) and Google Maps (available at <http://maps.google.com/>, accessed September 2005), as well as the National Imagery and Mapping Agency³⁸ and the United States Geological Survey³⁹ provide map viewing and navigating capabilities to site visitors. Specialized systems, such as the Senographe DMR (GE Medical Systems, Milwaukee, WI, USA), are used for detection tasks by radiologists; software packages such as ArcView GIS⁴⁰ support digital viewing of feature (raster) data or image data. Berinstein⁴¹ reviewed five image-viewing software packages with zooming capabilities, VuePrint, VidFun, Lens, GraphX, and E-Z Viewer, which were frequently used by libraries. The transition from film to digital cameras for the consumer market has resulted in a wide selection of photographic image manipulation applications.

These tools use a variety of different interaction techniques to give viewers access to images at different resolutions. There are two basic classes of interactions involved. The first is zooming, which refers to the magnification of the image. The spatial resolution of the image as it is originally

acquired is referred to as the “full resolution.” Different zoom levels that shrink the image in spatial resolution are provided so that the image can be shrunk down to fit the screen. The second operation is panning, which refers to the spatial movement through the image at its current zoom level. Most tools use some combination of these two techniques. Prominent paradigms for zooming in and out of images and some example applications that use them include: the use of onscreen buttons—toolbars,^{35–39} clicking within an image to magnify a small portion of that image (FFView available at <http://www.feedface.com/projects/ffview.html>, accessed September 2005), or clicking within the image to magnify the entire image with the clicked point at the center (ArcView GIS⁴⁰). Prominent image-panning paradigms and example applications include the use of scroll bars (Mapquest available at <http://www.mapquest.com>, accessed September 2005; Microsoft Office Picture Manager and Microsoft Office Paint available at <http://microsoft.com>, accessed September 2005; Adobe PhotoShop available at <http://adobe.com/>, 2005),⁴⁰ moving a “magnification area” over the image in the manner of a magnifying glass (FFView available at <http://www.feedface.com/projects/ffview.html>, accessed September 2005), clicking on arrows or using the keyboard arrows to move over an image (Mapquest available at <http://www.mapquest.com>, accessed September 2005), panning vertically only via the mouse scroll wheel (Adobe PhotoShop available at <http://adobe.com/>, 2005),⁴² and dragging the image via a pointer device movement (Google Maps available at <http://maps.google.com/>, accessed September 2005; Microsoft Office Picture Manager and Microsoft Office Paint available at <http://microsoft.com>, accessed September 2005).

Thus, while many systems exist to view digital images and digital image viewing is considered an important component of practice in many fields, there is no guidance from the literature regarding what geometric zoomable interaction techniques are best suited for navigating large images and, in particular, for the task of finding small features of interest within an image.

MATERIALS AND METHODS

The main hypothesis was to determine which of five different commonly used types of interaction

techniques were the most effective for helping observers detect small-scale features in large images and which of the techniques were qualitatively preferred by the users. Secondary aims include testing the main hypothesis when interaction techniques had slow update rates (such as might occur in teleradiology) and trying to identify major features of the interaction techniques that caused their success or failure. The study was comprised of both quantitative and qualitative parts. The quantitative part was the experiment to measure the users’ speed at finding features in large images when using different interaction techniques. There were three qualitative parts of the study: observations by the experimenter of the subjects during the experiment, a postexperiment questionnaire, and a qualitative comparison by the subject of all five interaction techniques on a single test image.

Pilot Experiment

To ensure we had developed the image-viewing techniques effectively and chosen appropriate targets within the images, we ran a pilot experiment. Three observers, who did not participate in the study, participated in the pilot. They each viewed 60 images using each of the five fast versions of the techniques to ensure that appropriate targets had been selected and to identify problems with the implementations of the techniques themselves. They then viewed ten images using each of the five slow versions of the techniques. Feedback from the pilot observers was used to refine the techniques and to eliminate target choices that, on average, were extremely simple or extremely difficult to locate. Measurements of the pilot observers completion times were also used to estimate the number of training trials needed to reach proficiency with the techniques. Once the experiment began, the techniques and targets were fixed.

Experimental Design

Quantitative

This study evaluated five different interaction techniques at two update rates (fast, slow) to determine which technique and update rate combinations were the most effective in terms of speed at finding a target within the image. Because the same interaction technique when used at a different update rate can have a substantially different user interac-

tion, each of the combinations is treated as a separate method. An analysis of variance study design using a linear model for the task completion time was chosen to compare the performance of the ten different methods. The images used in the study were large grayscale satellite images with very small features to be detected. These images were chosen because they are of a similar size to the largest digital medical images; they were representative of the general visual task as well as the medical imaging specific task, and they allowed the use of student observers. In a prior work of Puff et al.,⁴² it was established that the student's performance on such basic visual detection tasks served as a cost-effective surrogate for radiologist's performance.

The task of finding a small target within a large image is naturally variable, affected by the image contents and each observer's individual searching style. To minimize variance in each user's performance, users received a significant amount of training to become proficient with the interaction method on which they would be tested. The number of study trials was also chosen to be large enough to help control for this variability. This led to having each user only perform with a single interaction method because the alternative (a within subject design) would have been prohibitive due to the number of trials required if each participant was to test with all ten interaction methods.

A total of 40 participants were recruited by flyers and e-mail for the study. Participants had to be over 18 years of age and have good vision (corrected was acceptable). They were students, faculty, and staff from the University of North Carolina at Chapel Hill (primarily graduate students from the School of Information and Library Science). Thirty-one participants were women and nine were men.

Each participant completed five demonstration images, 40 training images, and 120 study images for the experiment. They were each randomly assigned one of the ten interaction methods, which they used for the entire study. At the beginning of the first session, the participant completed an Institutional Review Board consent form. Then, the experimenter explained the purpose and format of the study and demonstrated the image-viewing tool with the five-image demonstration set. Next, the participant completed the training set of 40 images, followed by the study set. The study set consisted of 120 images in a randomized order, partitioned into four sets. The presentation order of the four image

sets was counterbalanced across observers. Participants read images in multiple sessions. Most observers read in five separate sessions (training set and four study sets), although some completed it in fewer by doubling up sessions. Participants were required to take mandatory breaks (10 min/h) during the sessions to avoid fatigue. At the beginning of each new session, the participant was asked to complete a five-image retraining set to refamiliarize them with the interaction tool before beginning the next study image set. If time between sessions exceeded 1 week, participants were required to complete a ten-image retraining set.

Qualitative

During the experiment, the researcher took notes on the observer's performance, problems they encountered, and unsolicited comments they made during the test. When participants had completed all of the image sets, they completed the post-experiment questionnaire ("Appendix 1"). Last, they were asked to try all of the interaction techniques using an additional test image to compare the methods and then rank them.

Images, Targets, and Screen Size

To test the viewing mechanisms, participants were asked to find targets, or specific details, within a number of digital grayscale photographs of Orange County, NC, USA. These photographs are $5,000 \times 5,000$ pixels in size and were produced by the US Geological Survey. Since participants were asked to find small details within the images, knowledge of Orange County did not assist participants in task completion. The targets were subparts of the full digital photograph and are 170×170 pixels in size. They were parts of small image features such as landscapes, roads, and houses, which could be uniquely identified but only at high resolution. Target locations were evenly distributed across the images, so that results from participants who began each search in a particular location would not be biased. "Appendix 2" shows the distribution of targets within the images, for the 160 images in the training and test sets. The screen resolution of the computer display was $1,152 \times 864$ pixels, and the actual size of the display area for the image was $1,146 \times 760$ pixels. Thus, only about 3.5% of the full-resolution image could be shown on the screen

at one time. “Appendix 3” shows a full image and an example target from that image.

Presentation and Zoom Levels

We tested five types of image-viewing techniques in the study. Each technique supported the following capabilities:

- Ability to view both the image and the visual target at all times. The visual target was always onscreen at full resolution so that, if participants were viewing the image at full resolution, they would be able to see the target at an identical scale.
- The entire image could be seen at once (by shrinking the image to fit the screen).
- All parts of the image were able to be viewed at full resolution, although only a small portion of the full image could be seen at once when doing this.
- Ability to choose a portion of the image as the target and get feedback as to whether the selection was correct or not.

An example screenshot is shown in Fig. 1, showing the Pointer interaction method at zoom

level 3 (ZL3). The target can be seen in the upper-right corner.

Users would strike a key to begin the next trial. The application would time how long it took until they correctly identified the target. Identification of the target was done by the user hitting the spacebar while the cursor was over the target. Users would continue to search for and guess the target location until they found it correctly.

Four levels of zoom were defined to represent the image from a size where the whole image could be seen at once in ZL1 to the full-resolution image in ZL4. The choice of four zoom levels was determined by having the difference between adjacent zoom levels be a factor of 2 in each dimension based on previous work that found this to be an efficient ratio between zoom levels, performing faster than continuous zoom for similar tasks^{33,37}. The image sizes for the four zoom levels were 675×675 pixels (ZL1), 1,250×1,250 pixels (ZL2), 2,500×2,500 pixels (ZL3), and 5,000×5,000 pixels (ZL4). Thus, when viewing the image at ZL4, only about 1/28th of the image could be seen on the screen at any one time. The MagLens and Section techniques used only one intermediate zoom level, in both cases similar to ZL3 of the



Fig. 1. Sample screen from the Pointer interaction technique. The target is shown on the *top right*. The navigation overview is on the *upper left*, with *crosshairs* showing the current cursor location. The user is currently at Zoom Level 3 and positioned slightly above and left of the center of the full image.

other three techniques. The same terminology (ZL1, ZL2, ZL3, ZL4) is used to describe the zoom levels consistently between all the methods, with their specific differences described in the next section. “Appendix 4” contains an illustration of the four zoom levels. Resizing the image between zoom levels was done via a bilinear interpolation.

Interaction Techniques

Based on our review of the literature and techniques commonly available, we chose five different interaction techniques to evaluate.

ScrollBar

The ScrollBar technique allows the participant to pan around the picture by manipulating horizontal and vertical scroll bars at the right and bottom edges of the screen, similar to many current image and text viewing applications, in particular Microsoft Office applications. Zooming in and out of the image is accomplished using two onscreen buttons (ZoomIn and ZoomOut), located in the upper-left-hand corner of the screen. Four levels of zoom were supported. Image zooming is centered about the previous image center.

MagLens

The MagLens technique shows the entire image (ZL1) while providing a square area (512×512 pixels) that acts as a magnifying glass (showing a higher-resolution view underneath it). Using the left mouse button, the participant may pan the MagLens over the image to view all parts of the image at the current zoom level. Clicking the right mouse button dynamically changes the zoom level at which the area beneath the MagLens is viewed. Only three levels of zoom were supported (ZL1, ZL3, ZL4) because the incremental difference of using ZL2 for the MagLens area was not found to be effective in the pilot experiment and was eliminated. Thus, if the zoom level is set to ZL1 the participant is viewing the entire image at ZL1 with no part of the image zoomed in to see higher resolution. If the participant clicks once, the MagLens square would then show the image below it at ZL3 while the image outside of the MagLens stays at ZL1. Clicking again would increase the zoom of the MagLens area to ZL4,

and a further click cycles back to ZL1 (no zoomed area). This interface style is found on generic image-processing applications, especially in the sciences, engineering, and medicine.

Pointer

The Pointer technique allows the participant to zoom in and out of the image by clicking the right (magnify) and left (minify) mouse buttons. Zooming is centered on the location of the pointing device (cursor on screen). Thus, the user can point to and zoom in directly on an area of interest as opposed to centering it first and then zooming. The Pointer method supports all four zoom levels. Panning is accomplished by holding the left mouse button down and dragging the cursor. We found that many users strongly identified with one of two mental models for the panning motion: either they were grabbing a viewer above the map and moving it, or they were grabbing the map and moving it below a fixed viewer. This corresponded to the movement of the mouse drag matching the movement of the view (a right drag caused rightward movement of the map) or the inverse (right drag caused leftward map movement), respectively. A software setting controlled this. The experimenter observed their initial reaction during the demonstration trials and configured the technique to their preferred mental model. The individual components (panning by dragging) and pointer-based zooming are often implemented, although this particular combined interface was not commonly available until recently (for instance it is now available in GoogleMaps (available at <http://maps.google.com/>, accessed November 2007) using the scroll-wheel for continuous zoom). It is similar to the original Pad++ interface⁹ which used the center and right mouse buttons for zooming in and out. The Pointer interface used in this study is the same one qualitatively chosen as the best of these same five (fast) techniques in a medical imaging study by Hemminger.³⁷

ArrowKey

The ArrowKey technique works similarly to the Pointer technique but uses the keyboard for manipulation instead of the mouse. The arrow keys on the keypad are used to pan the image in

either a vertical or horizontal direction in small discrete steps. As with the Pointer interface, a software toggle controlled the correspondence between the key and the direction of movement and was configured to match the user's preference. The ArrowKey method supported all four levels of zoom. Zooming is accomplished by clicking on the keypad Ins key (zoom in) or Del key (zoom out). The technique always zooms into and out of the image at the point that is at the center of the screen. This interface sometimes serves as a secondary interface to a pointer device for personal computer applications; it is more common as a primary interface on mobile devices which have only small keypads for input.

Section

This technique conceptually divides each image into equal size sections and provides direct access to each section through the single push of a key. A section of keys on the computer keyboard were mapped to the image sections so as to maintain a spatial correspondence, i.e., pushing the key in the upper right causes the upper-right section of the image to be shown at a higher resolution. In our experiment, the screen area was divided into nine rectangles, which were mapped to the one to nine buttons on the keyboard's numeric keypad. The upper-left-hand section of the image would be selected and displayed at ZL3 by hitting key 7, the upper center by key 8, the upper right by key 9, and so forth. Once zoomed in to ZL3, the participant may zoom in further to ZL4 to see a portion of the ZL3 image at full resolution by striking another one of the one to nine keys. Thus, this technique allows the participant to view a total of 81 separate full-resolution sections, all accessible by two keystrokes. For instance, to see the upper rightmost of 81 sections, the participant would hit key 9 followed by key 9. To zoom out of any section, the participant presses the ZoomOut (insert) key on the numeric keypad. An overlap of the sections is intentionally built in at the section boundaries, as illustrated in "Appendix 5." This allows participants to access targets that may otherwise have been split across section boundaries. The

Section method supports three levels of zoom (ZL1, ZL3, and ZL4) similar to MagLens because the pilot experiment found the use of ZL2 to be a detriment for this technique. This interaction is sometimes implemented with fewer sections (for example quadrant-based zooming). It is less common than the other choices and probably more suited to mobile devices that have numeric keypads but not attached pointing devices.

Navigation Overview

Many systems provide a separate navigation window showing the user what portion of the entire image they are currently viewing^{7,43}. In our work evaluating several zoomable interfaces for medical image display³⁷, we found that, when the zooming interactions operated in real time and the full image could be accessed in less than 1 s (for instance via two mouse clicks or two keystrokes), users preferred to operate directly on the image instead of looking to a separate navigation view. Hornbaek et al.⁴⁴ reported similar findings for an interface with a larger number of incremental zoom levels (20). They found that users actually performed faster without the navigation view and switching between the navigation and the detail view used more time and added complexity to the task. Because some of the techniques tested in this study (particularly the slow update rate ones) might not perform as well without a navigation view, a navigation window (100×100 pixels in the upper-left corner) was included as part of all of the techniques. Based on the pilot study and guidelines^{7,31,44-46} established for navigation overview windows, the overview window was constructed so that it was tightly coupled to the detail window, showed the current location of the cursor, and kept small to leave as much of the screen real estate for the detail window as possible, which was crucial for this study's task.

We developed ten viewing tools corresponding to the ten methods and implemented them as Java 2.0 programs, running on a Dell 8200 computer with 1 GB of memory, and a 20-in. color Sony Trinitron cathode ray tube monitor. The viewing tools, an example image and instructions, are available at <http://ils.unc.edu/bmh/pubs/PanZoom/>.

RESULTS

Quantitative

We analyzed the training (first 40 images) and test images (numbered 41–160) to see if the observers reached asymptote performance with their interaction method by the end of their training, so that their test results would not be biased by observers continuing to significantly improve during the study trials. Time for each subject was modeled using least squares as a function of trial number with a modified Michaelis–Menten function which is nonlinear, monotonic, and decreasing to an asymptote. All observers reached asymptote performance by the end of training with most achieving it within the first 10–15 training cases. An example observer’s reading times with asymptote curve fit is seen in “Appendix 6.”

The primary quantitative analysis was to compare the ten different methods (five techniques each at two speeds) based on how quickly observers could complete the feature-finding task using that method. Table 1 summarizes the mean time and standard deviation for each method, calculated across all observers and all trials. To determine whether a particular method performed faster than another, the mean task completion times were compared using the SAS (Cary, NC, USA) GENMOD repeated-measures regression test (1 *df*, complete analysis in “Appendix 7”). A *P* value of 0.05 or less indicates the null hypothesis that the techniques have the same performance, which is rejected, and that the performance of the two techniques is statistically significantly different from each other. Using the results from this analysis, we grouped the methods into performance groups. Table 1 shows the mean task completion times in seconds averaged across all observers for each method and the performance groupings. Methods were placed in the same performance group if they had similar mean times and did not have statistically significant differences in mean times from all other members of the performance group (using SAS GLM Tukey’s Studentized Range Test). This segregated the methods into four groups (Table 1). Part of the reason for grouping the techniques is that the

group rankings are probably more informative than the individual rank ordering of methods, due to the large standard deviations in detection times due to image and observer effects, as seen in Table 1. A further regression analysis was conducted to compare these resulting groups. All of the groups were found to be statistically significantly different from one another (*p* value < 0.05), with the exception of group 1 versus group 2. A power analysis based on the existing data show that the study would have to increase from four to seven observers per method in order to reduce the variance sufficiently to demonstrate the difference between group 1 and group 2 at a statistically significant level.

A regression analysis was also performed to examine the significance of the other two factors (observer and image). The largest determining factor was the method, with the observer and image effect each approximately one third the magnitude. Table 2 shows how much each of the main effects contributes to determining the speed of detecting targets.

The last analysis determined whether the slow versions of techniques generally performed the same or differently than the fast versions of the techniques. A comparison of differences in mean task completion times between the fast and slow versions of each of the five techniques to zero (SAS GENMOD analysis, 5 *df*) determined that the fast techniques were statistically significantly different from the slow ones (*P* value of 0.047). It is evident that they are faster from Table 1, with the exception of the MagLens fast technique which

Table 1. The Mean Task Completion Times (seconds)

Performance group	Method name	Mean time	SD
1	ArrowKey Fast	76	76
1	Pointer Fast	79	100
1	ScrollBar Fast	84	108
2	Section Fast	97	127
2	Section Slow	97	131
2	ScrollBar Slow	98	94
3	MagLens Slow	117	165
3	ArrowKey Slow	119	134
3	Pointer Slow	128	146
4	MagLens Fast	155	176

Table 2. Main Effects

Factor	<i>df</i>	<i>F</i> value	<i>Pr>F</i>
Method	0	21.96	<0.0001
Observer	30	8.76	<0.0001
Image	115	9.42	<0.0001

observers had some difficulties with, resulting in it being the poorest performer.

Qualitative

A significant amount of valuable information resulted from observing the participants, from the survey, and from the postexperiment testing. We summarize only the highlights here but have included much of the rich qualitative details in “Appendix 8.”

Our observations of the observers closely matched both their comments and their rankings of the techniques. Table 3 shows the rankings of the interaction techniques by the observers, based on their trying each of the techniques at the conclusion of the study. Observers assigned the techniques rankings of 1–5 (1 being the best). The Pointer technique was listed by almost all observers as the best technique. The rest of the techniques all clustered at slightly below average.

Reasons the observers gave for favoring the Pointer method was the natural control it gave them in panning around the image, precise control of the zooming, maintaining context (location in the overall image), and speed of operations. The ArrowKey method was also favored for its speed and precise control of panning and zooming. Participants did not rank it as high because they found the panning motion to be “less smooth” and it was “harder to scan” than with the Pointer method. They did find the ArrowKey technique very effective for systematic searching. Some users found the MagLens interaction desirable because you always maintained the context of where you were in the overall full-resolution image. It was also considered to be a more familiar paradigm than some of the other techniques like the Section. However, many users felt it was difficult to use in

practice, saying it was “hard on the eyes” and “is a pain” and several observers who used it complained that it was disorienting to use, with one becoming dizzy as a result. The ScrollBar technique was considered “familiar” yet “old-fashioned.” Users felt it gave them good control but with too limited flexibility (i.e., only being able to pan in one dimension at a time versus two for most of the other techniques). Only two of the eight participants who had used the technique in the study ranked it in their top two choices. The Section technique was the least favored of all the techniques. Panning of the image is not directly supported by this technique, in that users have to step up a zoom level and then back down again in an adjacent section to effect a “pan” operation. Users felt this did not allow a natural panning exploration to occur, that too many button clicks were required to pan around, and that the constant zooming in and out frequently caused a loss of context.

The navigation view was very rarely used except for experimenting with it in training. The few instances where it was observed being used during the test cases were in the slow versions.

DISCUSSION

Our results indicate that some interaction techniques perform quantitatively better for feature detection types of tasks. Integrating the results from the quantitative and qualitative portions of the study did yield several consistent overall themes, and a clearer understanding of the benefits and shortcomings of the individual techniques is presented in this section. It is important to remember, though, that the performance of inter-

Table 3. Observers’ Rankings of Techniques on a 1–5 Scale

Technique	Average ranking
Pointer	1.77
MagLens	3.12
ArrowKey	3.28
ScrollBar	3.30
Section	3.51

action techniques will clearly depend on the task, and these results may not hold for other types of tasks. Additionally, the chosen surrogate visual detection task is not representative of all types of medical imaging tasks.

Overall Themes

Intuitive and Easy-to-Use Interface Favored From the qualitative feedback, users expressed clear preferences for intuitive, easy-to-use, and highly interactive user interface techniques. There were common elements to the techniques that performed well quantitatively and were preferred qualitatively. The top three performing techniques supported natural and easy ways to perform image panning. They supported both systematic and intuitive target searching. The most preferred method, Pointer, was favored in a large part because it had the most natural interaction for panning, with hand motion of the pointer corresponding to moving the image viewpoint. The most preferred methods (Pointer, ArrowKey) supported easy control of zooming, in that zoom levels could be selected without the observer moving their hand. Techniques that had more challenging mental models (Section) or difficult interactions (MagLens) were not favored and did not perform as well.

Simple Interface Favored Techniques that minimized interactions (keystrokes, mouse clicks, hand motions) tended to perform better, as might be predicted by *Goals, Operators, Methods, and Selection*⁴⁷ modeling of the techniques. The Pointer and ArrowKey had the most efficient interactions due to the hand remaining on input device (mouse or arrow keys, respectively), and only one interaction (click) is required for both pan and zoom operations. The Scrollbar method was perhaps the least efficient due to having to move the pointer between three areas and click on small controls (vertical and horizontal scrollbars and the zoom buttons). This was reflected in the user's comments and rankings which made it clear that they did not favor this technique because it did not support natural and quick panning and was too cumbersome for more generalized tasks. However, the Scrollbar method performed well quantitatively for the

feature detection task because all the users of this technique adopted a systemic way to scan the image (they scrolled across the image a "row" at a time using only one scrollbar control). Additionally, using multimodal interfaces may add mental distractions for the user. It is possible that the MagLens and Scrollbar interactions may have suffered from this because these two interactions utilized both the mouse and keyboard while the other techniques were primarily keyboard-based (Section, ArrowKey) or mouse-based (Pointer).

Faster and Real-Time Interactions Preferred Users clearly favored the faster update rate versions of techniques and also performed better with them in all cases except the MagLens technique, where the fast version had worse performance likely due to the users losing context and getting confused about what part of the image they had already viewed.

Individual Techniques

ArrowKeys This was one of the top performers and, while it was significantly behind the Pointer technique in user preference, it was generally favorably reviewed by observers. While this technique was not as natural as the mouse-panning interaction of the Pointer technique, the small discrete movements (left, right, up, down) were easily understood and utilized by the observers. As with the Pointer method, the slow version of this technique did not perform as well because of the reduced interactivity of the pan operation.

MagLens While this technique was familiar to most users, and favored by some, it was generally not preferred by those who used it in the experiment, and it performed the worst overall of all the techniques. Interestingly, the fast version was by far the slowest in performance. Users of the fast version tended to try to interactively pan more. When they did this, they lost their position (context) and often became disoriented with respect to what territory they had covered already. The users of the slower version tended to adopt a more methodical search pattern for covering the image at a high zoom level and ended up being more efficient.

Pointer This was one of the top performing techniques and the clear favorite of the observers. The interface lends itself well both to systemic tasks like the feature detection task of this experiment and more general tasks, such as manipulating large images or following map driving directions. As computer and graphics card speeds have increased, the panning part of the Pointer interaction (dragging the mouse) is becoming fairly common, and having interactive zoom is beginning to appear in tools. Currently, most tools have a separate interaction for zooming, as in MapQuest, which zooms by mouse clicks on a scale on the screen or keystrokes on the keyboard. This is less efficient than having both the zoom and the panning operations accomplished from the pointing device.³³ An easy way to do this is to zoom via the scrollwheel now commonly found on mouse devices, and this has been adopted by recent applications (for instance GoogleMaps⁴² now supports this). This technique is strongly dependent on a fast interaction. The natural connection between the panning motion of the mouse and the movement of the image on the screen was lost due to the update delay in the slow version of the pointer interaction. The result was that the slow version was not favored by users and was next to last in performance.

Scrollbar The Scrollbar method was familiar to users. They found it satisfactory for one-dimensional scrolling, as is commonly found in text viewers. However, it was generally viewed as cumbersome for navigating in two dimensions because of having to separately manipulate the vertical and horizontal scrollbar controls. In this experiment, users were able to adapt the task to a series of systematic searches along “rows” of the image, reducing their usage to manipulating a single scrollbar control to move across one “row” at a time. This allowed them to perform efficiently with both the fast and the slow versions of the technique.

Section The Section method was the least favored by the observers because most were not familiar with the technique, and the mental model was not as natural to them. However, users were able to become efficient with this technique, and both the fast and the slow version were in the top five in performance. It appeared that the slow version performed as well as the fast version because users tended to not rely on many quick panning motions but instead adopted a systematic section-by-section search pattern, which was not significantly affected by the difference in the slow and fast update rates.

This experiment dealt with a particular feature detection task, and given sufficient training users were, in most cases, able to adapt to the technique they utilized to efficiently perform the task. For most of the techniques, this resulted in the users scanning out the image in rows, with the height of the row being the size of the image seen at either ZL3 or ZL4 (depending on user preference). This type of serialized scanning interaction is formalized in several disciplines; for instance, it was popularized by Laszlo Tabar as a method of training radiology residents in detecting microcalcifications in mammography. The ArrowKey, Scrollbar, and Section techniques support this type of highly structured, linear movement in vertical or horizontal directions, especially well. They are less well suited to supporting navigation in two dimensions, such as following natural objects or anatomy. Observers commented that the Pointer method seemed much more effective for these types of interactions as well as for more general-purpose navigation.

Several factors affect the choice of the technique to utilize in a given situation. In addition to the task, the update rate of the display device and the types of interactions supported by the display device (keyboard only, cell phone–PDA keypad only) are key factors. For devices such as personal computers that commonly have pointing devices and fast update rates, the Pointer method

would likely be an effective choice across a wide range of applications. If the update rate is not fast, then a different technique than the Pointer method may be more optimal (e.g., Section or Scrollbar). The ArrowKey and Section interfaces do not require a pointing device and thus may be better suited for small mobile devices such as cell phones and PDAs.

Since the fast versions of the techniques performed significantly better than the slow versions, there is not a single technique that can be considered the best choice for working well under both update conditions. Thus, applications that may be used under both conditions should consider offering more than one interface technique to the user. For this particular task, if only a single technique could be supported, then the Section and Scrollbar techniques might be good candidates since both the slow and fast versions of these techniques were in the top two performance groups.

ACKNOWLEDGMENTS

Thanks to the Interaction Design Lab which hosted the space for the observer experiment and to Chris Weisen of the Odum Institute, who helped with the statistical analysis. Prior work studying different interaction techniques for radiology workstation designs helped lay the groundwork for this study. This included grants from Fischer Imaging and Hologic as well as federally funded grants NIH RO1 CA60193-05, DOD DAMD 17-94-J-4345, NIH RO1-CA 44060, and NIH PO1-CA 47982.

APPENDIX 1: POSTEXPERIMENT QUESTIONNAIRE

Observer# _____ Interaction Technique _____
 Speed Slow _____ Fast _____

1. In what ways was the interaction technique you tested successful (in helping you locate known targets on an image larger than the size of your electronic display)?

2. In what ways was the interaction technique you tested difficult to use, or made your task more difficult than necessary?
3. What do you think would be the ideal interaction technique for the task you were asked to do?
4. After trying all the techniques in the study, please rank them best to worst, and describe their comparative advantages for this task?


Method	Pros	Cons
#1		
#2		
#3		
#4		
#5		

5. Do you have any suggestions for improving this experiment?

APPENDIX 2: DISTRIBUTION OF TARGETS WITHIN IMAGES

Vertical and horizontal axes are the vertical and horizontal axes of the images used in the experiment (images are 5,000×5,000 pixels). The points depicted in the figure each correspond to the center of a target location used with a study image.

Distribution of Targets within Images



APPENDIX 3: TARGET AND ORIGINAL IMAGE

The bottom image is the full image (down-interpolated to fit on the screen). Highlighted on it in yellow is a target area. The target area is shown in the top image at original (full) resolution.



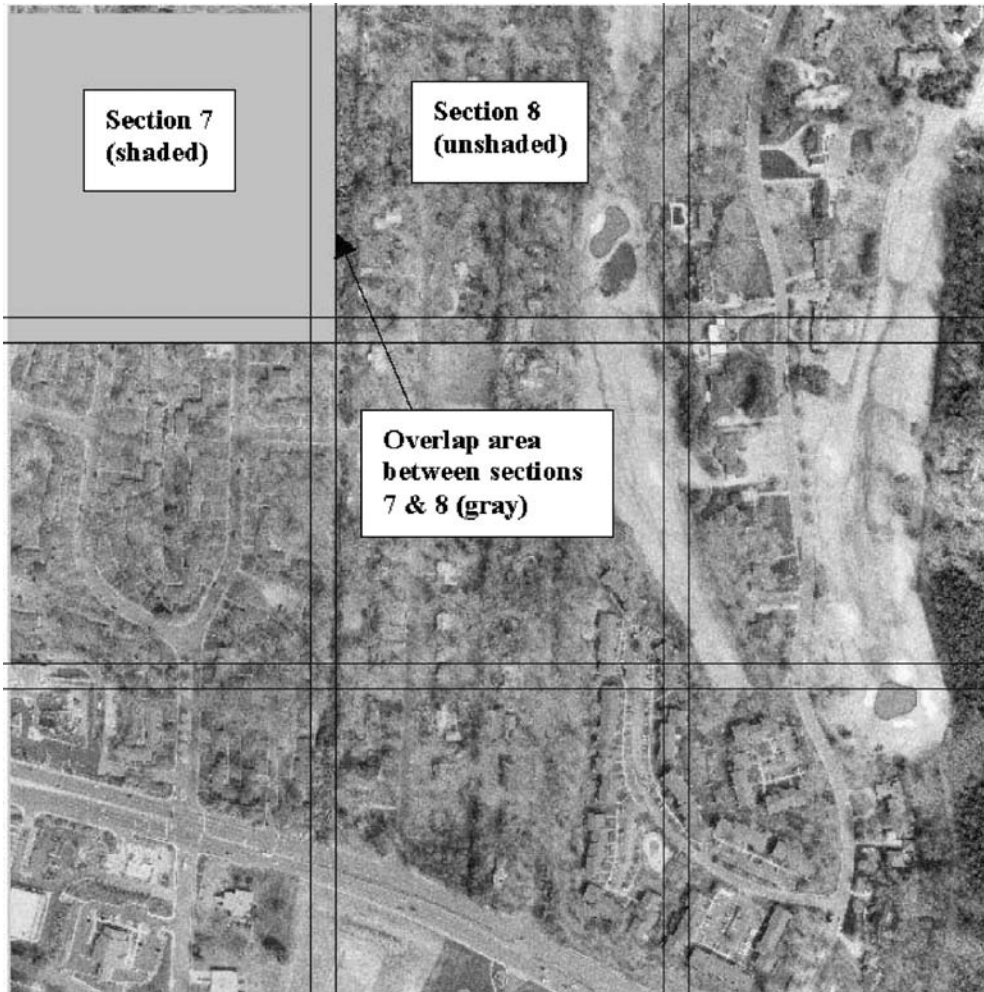
APPENDIX 4: ZOOM LEVELS

Below is an image from the study (down-interpolated to fit on the screen). This corresponds to Zoom Level 1 (i.e., you can view the complete image on the screen). Zoom Levels 2, 3, and 4 are highlighted to show the proportional area of the original image that would be seen when viewed at those zoom levels. Thus, under Zoom Level 4, the user would see only as much of the image as is seen in the pink highlighted section. The target size is seen as the small blue box within Zoom Level 4.




APPENDIX 5: SECTION ZOOM OVERLAP

The image below shows the first-level partitioning of the image for the sectional interaction technique. Adjacent sections intentionally overlap so that important information such as targets are not chopped off at the boundaries but can be approached and viewed from any adjacent section. Section 7 (upper left) can be seen to overlap into the adjacent sections on the right (Section 8) as well as below. Similarly, entering Section 8 allows overlap to the same “shared” area between Sections 7 and 8.



APPENDIX 6: EXAMPLE READING COMPLETION TIMES



APPENDIX 7: METHOD VERSUS COMPLETION TIME ANALYSIS: A COMPLETE LISTING OF PAIR-WISE COMPARISONS

The SAS System 13:46 Friday, July 1, 2005
 The GENMOD Procedure

Table 4. Contrast Estimate Results

Label	Estimate	Standard error	Alpha	Confidence limits		Chi-square	Pr>Chi sq
1 vs 2	5.4146	11.4895	0.05	-17.1044	27.9335	0.22	0.6375
1 vs 3	-70.8000	13.2604	0.05	-96.7900	-44.8100	28.51	<0.0001
1 vs 4	7.9250	17.9412	0.05	-27.2390	43.0890	0.20	0.6587
1 vs 5	-13.0771	10.3176	0.05	-33.2991	7.1450	1.61	0.2050
1 vs 6	-13.9833	12.9151	0.05	-39.2966	11.3299	1.17	0.2789
1 vs 7	-44.0583	16.3799	0.05	-76.1623	-11.9544	7.23	0.0071
1 vs 8	-33.1125	12.8421	0.05	-58.2826	-7.9424	6.65	0.0099
1 vs 9	-34.9500	15.9704	0.05	-66.2513	-3.6487	4.79	0.0286
1 vs 10	-12.3521	25.9969	0.05	-63.3052	38.6010	0.23	0.6347
2 vs 3	-76.2146	14.1289	0.05	-103.907	-48.5224	29.10	<0.0001
2 vs 4	2.5104	18.5923	0.05	-33.9298	38.9506	0.02	0.8926
2 vs 5	-18.4917	11.4123	0.05	-40.8593	3.8759	2.63	0.1052

Table 4. (Continued)

Label	Estimate	Standard error	Alpha	Confidence limits		Chi-square	Pr>Chi sq
2 vs 6	-19.3979	13.8054	0.05	-46.4560	7.6601	1.97	0.1600
2 vs 7	-49.4729	17.0906	0.05	-82.9698	-15.9760	8.38	0.0038
2 vs 8	-38.5271	13.7371	0.05	-65.4513	-11.6029	7.87	0.0050
2 vs 9	-40.3646	16.6985	0.05	-73.0930	-7.6361	5.84	0.0156
2 vs 10	-17.7667	26.4505	0.05	-69.6087	34.0753	0.45	0.5018
3 vs 4	78.7250	19.7358	0.05	40.0434	117.4066	15.91	<0.0001
3 vs 5	57.7229	13.1936	0.05	31.8640	83.5818	19.14	<0.0001
3 vs 6	56.8167	15.3108	0.05	26.8081	86.8252	13.77	0.0002
3 vs 7	26.7417	18.3281	0.05	-9.1807	62.6640	2.13	0.1446
3 vs 8	37.6875	15.2492	0.05	7.7996	67.5754	6.11	0.0135
3 vs 9	35.8500	17.9630	0.05	0.6431	71.0569	3.98	0.0460
3 vs 10	58.4479	27.2665	0.05	5.0067	111.8892	4.59	0.0321
4 vs 5	-21.0021	17.8918	0.05	-56.0694	14.0652	1.38	0.2405
4 vs 6	-21.9083	19.5055	0.05	-60.1385	16.3218	1.26	0.2614
4 vs 7	-51.9833	21.9537	0.05	-95.0118	-8.9549	5.61	0.0179
4 vs 8	-41.0375	19.4573	0.05	-79.1730	-2.9020	4.45	0.0349
4 vs 9	-42.8750	21.6499	0.05	-85.3080	-0.4420	3.92	0.0477
4 vs 10	-20.2771	29.8246	0.05	-78.7322	38.1780	0.46	0.4966
5 vs 6	-0.9063	12.8465	0.05	-26.0849	24.2724	0.00	0.9438
5 vs 7	-30.9813	16.3258	0.05	-62.9792	1.0167	3.60	0.0577
5 vs 8	-20.0354	12.7731	0.05	-45.0702	4.9993	2.46	0.1167
5 vs 9	-21.8729	15.9149	0.05	-53.0655	9.3197	1.89	0.1693
5 vs 10	0.7250	25.9629	0.05	-50.1614	51.6114	0.00	0.9777
6 vs 7	-30.0750	18.0798	0.05	-65.5108	5.3608	2.77	0.0962
6 vs 8	-19.1292	14.9499	0.05	-48.4305	10.1722	1.64	0.2007
6 vs 9	-20.9667	17.7097	0.05	-55.6770	13.7437	1.40	0.2364
6 vs 10	1.6313	27.1002	0.05	-51.4842	54.7467	0.00	0.9520
7 vs 8	10.9458	18.0277	0.05	-24.3879	46.2796	0.37	0.5437
7 vs 9	9.1083	20.3748	0.05	-30.8255	49.0422	0.20	0.6548
7 vs 10	31.7063	28.9123	0.05	-24.9608	88.3733	1.20	0.2728
8 vs 9	-1.8375	17.6565	0.05	-36.4436	32.7686	0.01	0.9171
8 vs 10	20.7604	27.0655	0.05	-32.2870	73.8078	0.59	0.4431
9 vs 10	22.5979	28.6823	0.05	-33.6183	78.8141	0.62	0.4308

The first column "label" in the statistical analysis above refers to pair-wise comparison of two methods. The methods are referred to by their numeric code; the coding scheme is shown below.

Method 1: ScrollBar Fast
 Method 2: Pointer Fast
 Method 3: MagLens Fast
 Method 4: ArrowKey Fast
 Method 5: Section Zoom Fast
 Method 6: ScrollBar Slow
 Method 7: Pointer Slow
 Method 8: MagLens Slow
 Method 9: ArrowKey Slow
 Method 10: Section Slow

APPENDIX 8: OBSERVER QUALITATIVE COMMENTS

This is a direct summary of the experimenter's notes from the experiment, including unsolicited observer comments, answers to the postexperiment questionnaire,

and the observer's final comments after trying all five interaction techniques at the end of the experiment.

Individual Techniques

This section examines the target-finding techniques used by participants for each method. We explore what techniques the method seems to encourage and how effective these techniques were, both in the slow version and in the fast version. We will look at data gathered by the experimenter and data that participants provided in the postexperiment survey.

Scrollbar Participants using the ScrollBar technique had the lowest average time per image for both the fast and slow tools, along with some of the smallest standard

deviations from average time. The experimenter notes help to explain what strategies the tool enabled participants to use to complete the task so successfully. Participants tended to use a combination target-finding strategy that allowed them to take advantage of the technique's utility in navigating to particular areas of the photo, as well as its facility in systematic searching. Many participants would examine the entire image at the ZL1 to choose a location in which to begin scanning. Then, they would zoom into a higher-resolution level (generally ZL3 or ZL4) and begin systematically scanning the picture for the target, beginning the search in the area they had chosen when looking at the entire image. Using this technique, they were able to closely examine the area of the photo where they suspected the target was located. They could freely pan around this area by clicking on the horizontal and vertical scrollbars and dragging them.

If they did not find the target in a particular area, they could use a more systematic approach to scan over the entire picture. To ensure that all areas of the picture were covered, several ScrollBar participants would scroll to a corner and scan for the target. If they did not find it, they would click in the empty section of the scrollbar track to move the scrollbar (and therefore the photo) a controlled amount. In this way, participants were able to ensure that they covered all areas of the photo while scanning. This combination of facilities that assist participants in both intuitive searching and brute-force scanning made the ScrollBar technique successful. As one participant commented in answer to the question of how the technique was successful, "It became easier to not search the same areas twice...I began searching in a pattern if the small image was not easily apparent."

As noted above, participants were more efficient completing the task with the slow version of the ScrollBar technique than all of the other slow methods, as well as two of the fast methods (Sectional and MagLens). The method seemed to help participants compensate very effectively for the delay. Participants using the slower technique tended to adopt an approach that maximized zooming and minimized panning. They would choose a section of the picture, zoom into it (to ZL2 or ZL3), and search for the target within it. If the target was not found in this area, they would zoom back out and choose another area to examine. This technique allowed participants to focus on clicking to pan as well as zoom, which is much faster than dragging to pan when a delay is present. Instead of scanning the entire picture, participants clicked to the areas that were most likely to contain the target first. If participants were not successful in finding the target using this technique, they could scan the entire photo by zooming into the ZL4 and using clicks, instead of the slower pans, to scan the entire photo.

Two of the three participants using the slow ScrollBar method further eliminated the amount of clicking and panning required by never zooming in ZL4. If they wanted to look at part of the image at full resolution, they would simply select an area for target confirmation and examine it closely. If the choice did not match the target, they would cancel the choice. One participant noted that this "allowed several modes of zooming in and let one easily scan in quadrants."

Both the fast and slow versions of the technique did garner some complaints from the participants. Four of the six participants commented that they did not like being placed in the center of the image when they zoomed in. One participant noted, "The zoom feature was fairly inaccurate in placement." Another described it as "disconcerting." Two people commented that they would have preferred to be zoomed into a corner instead of the center of the image. Participants also noted that they did not like holding down the scroll bar to see the parts of the image located beneath the target and crosshairs box. One commented, "It's a pain that you have to hold that thing [the scroll bar] down if you want to see everything too..." Two ScrollBar participants speculated on ideal search techniques. One commented that she would like "a zoom in/out controlled by cursor placement... and possibly a smooth, faster way of scrolling...What I think I would like best would be a keypad technique with general placement around the picture so parts could be jumped to quickly." She ranked the Sectional method one and the Pointer method two. Another explained that she would prefer "more precise controls—not limited to scroll bars. Bird's eye view—move cursor over picture, where it zooms for you." She ranked the MagLens method one.

ArrowKey The fast version of the ArrowKey technique performed very well; it was not statistically different from the fast ScrollBar technique, and standard deviations were quite low. However, the slow version of the ArrowKey was the second to last performer in average target identification time, with high standard deviations. An examination of the way participants used this technique to find targets may provide insight into why this was the case.

Like the ScrollBar, the ArrowKey enabled participants to employ a combination of systematic and intuitive searching techniques. Generally, they would choose an area from ZL1 where they felt the target was most likely to be located. They would zoom into ZL2 and examine the area for the target using a panning movement. Panning with the ArrowKey technique entails using the arrow keys to move around the image. Some participants began the task using slow, measured clicking of the arrow keys to pan around, examining the image after each click. This method is very systematic

but quite slow. One participant using the slow ArrowKey method chose to hold down the arrow keys to move the image more rapidly in an attempt to compensate for the delay. She lost control of the image several times and it scrolled completely off the screen. All participants, after experimenting with these different panning techniques, settled on a rapid-fire clicking of the arrow keys to pan the image. This seemed to be the most effective panning motion for both the fast and slow versions of the technique. Participants used this motion at ZL2, ZL3, and ZL4.

When some participants could not locate the target from ZL2, they chose to zoom to ZL3 to search for the target, using the same panning motion. If they did not find the target in the selected area, they would pan around the entire image at ZL3. Conversely, other participants would pan around the image at ZL2 if they did not find the target in the initially selected area. Participants who were able to identify targets ZL1 or ZL2 generally were faster than participants who routinely scanned the image at ZL3 or ZL4.

Participants were generally able to find targets the first time they panned over the image at a zoom level low enough for them to identify the target (as described above, usually ZL2 or ZL3), indicating that systematic searching with the ArrowKey is very effective. One participant noted that she liked that “movement [of the image] was easy to judge...when I pushed on an arrow I had a good idea of where I’d end up.” Another commented that the panning motion “feels pretty natural.” One commented, “Movement in blocks was bothersome, though I got used to it.” All of the participants with both versions of the technique avoided ZL4. One person noted, “Zooming in three times...I have to move the image little by little...it becomes very annoying.” Two of the six participants “browsed” at ZL4 by selecting targets to see if they were correct or not.

Participants using the slow version of the ArrowKey were significantly slower on average than their counterparts using the fast version, although they used many of the same techniques to identify targets. Since the movement of the image with each press of an arrow key is so defined, panning at higher zoom levels (ZL3 and ZL4) was extremely slow and penalized participants far more than panning at ZL2. The participant who was able to regularly select targets from ZL1 and ZL2 was quite a bit faster on average than those participants who selected targets ZL3 or ZL4. One participant commented, “A faster method might have prevented me from catching a glimpse of the target as I did periodically.” While these participants were able to take advantage of the ArrowKey’s utility for systematic searching, they were penalized with very slow average times per image.

By default, the ArrowKey technique moves the picture in the same direction of the clicked arrow (for example, if you click the UP arrow the image moves up) but participants can reverse the cursor direction, so clicking the UP arrow moves the image down, in the same manner as a ScrollBar. Four of the six participants chose to reverse the cursor direction; one participant did not reverse the cursor direction but commented, “The ways the arrows moved the picture felt counter-intuitive.”

One ArrowKey participant noted that she “would have liked to be able to choose an area to zoom in on without centering the area first.” Two ArrowKey participants provided their ideas about an ideal technique; both of them framed their ideas as improvements of the ArrowKey. One participant explained, “This [technique] was fine—could be improved by adding a smooth scroll.” The other participant expressed a related idea: “It would be nice to have a way—like in Photoshop—to make both short and long ‘nudges’ when moving/searching across an area.”

Pointer The fast version of the Pointer technique performed virtually the same as the fast ArrowKey technique, while the slow version was one of the worst performers. The Pointer technique enables many of the searching techniques used by participants with the ScrollBar and the ArrowKey, while providing several utilities that helped participants overcome the technique delay. As with the other methods, participants using the Pointer tried to avoid ZL4, finding panning at this level to be prohibitively slow. Participants using the fast version of the technique used a combination of zooming and panning that tended to focus on a panning technique. These participants panned by clicking on the picture and dragging it across the screen at a medium speed, sometimes speeding up or slowing down the panning motion, depending on how closely they wished to examine a particular area.

Two of the participants would begin searching with a more intuitive approach, choosing the most likely area for the target to zoom into first and then proceed to a full-image scan at the ZL1 or ZL2. One participant mentioned that she felt scanning for the target was faster than trying to deduce where it was and searching for it in a particular location. While both of these participants employed scanning heavily, they did avoid the parts of the images where they felt the target was less likely to be located. One participant said she liked the technique because “it helped me focus on the parts of the image that I thought were important and disregard the rest of the image.”

The third participant using the fast version of the Pointer employed more zooming than panning techniques to find targets. She would zoom in to ZL2 or ZL3 where she thought the target might be, and, if she did not

find it, she would zoom back out and choose a different location. This technique was not at all systematic; although she found a number of targets very quickly, she took such a long time on other targets that the standard deviation for her target-finding times was quite high. Her average time was also significantly lower than that of the first two participants.

Participants using the slow version of the Pointer relied on somewhat different strategies to locate targets that helped them to compensate for the delay in the technique. After experimenting with different combinations of panning and zooming to navigate around images, all three of these participants moved to a target-finding technique that concentrated more on zooming than on panning. The two most successful of the three participants minimized dragging to pan by carefully examining the entire picture at ZL1 before choosing an area to zoom into using ZL2. They performed the same actions when choosing to zoom in to ZL3. If they could not find the image using zooming techniques, they would pan around the image at ZL3 instead of ZL4. One of these participants commented that she liked the targeted panning that the technique allows: "I was able to drag and circulate around an area." The least successful of the three participants did spend a good deal of time panning at ZL3 and ZL4. He compensated for the slowness of the technique by clicking on the image at one edge of the screen and dragging the cursor to the other edge of the screen, thereby examining the image in chunks, instead of using the constant panning motion that participants with the faster technique employed. The Pointer technique's ability to accommodate direct zooming, enabling a focus on zooming instead of panning, as well as its flexibility in the ways participants could pan with it, helps to explain why the slow version of this technique helped participants compensate for the delay more than the slow version of the ArrowKey.

In general, the Pointer participants were very comfortable working with the technique. One commented, "I have an established comfort level with mousing and zooming." However, they did make several comments about how they would like to see the technique improved. Two participants mentioned that they sometimes had difficulty with left mouse clicks; when they would click to zoom in, nothing would happen. One commented, "If you were switching from drag to zoom and moved the mouse slightly the system often didn't read the switch." Two participants commented that they would like for the image to recenter itself if they zoomed all the way out to ZL1 (full-image view), so that they could restart the search process with the image already centered. Two participants would have liked to be able to select targets using the mouse instead of the keyboard; one of these suggested using a three-button mouse.

One participant mentioned that she would be interested in a technique that used the keyboard instead of the mouse to move the image because "my eyes are faster than my hand;" she thought a keyboard technique might enable faster scanning. However, when she saw the ArrowKey and Sectional techniques she commented that they had "too many buttons." No other participants speculated about techniques that may have helped them performed the task in a better way. This indicates that they all found the technique to be easy and intuitive to use.

Sectional The Sectional fast and slow methods performed about as well as one another; they were ranked as the fifth and sixth fastest methods, respectively. This technique was very good for systematic searching but had several major disadvantages that prohibited it from performing as well as the Pointer, ArrowKey, and ScrollBar fast methods. Participants using the fast and slow versions of the Sectional technique employed a systematic method for searching for targets. They would choose one of the nine sections to zoom into, from ZL1 to ZL3. While two of the fast participants tended to start in the same quadrant every time, the rest of the participants examined the picture to determine the section most likely to contain the target. They would zoom into the chosen section and then zoom into each section within it. Only one of the participants tended to find most of the targets at the first level of zoom. Unlike with the other techniques, participants did not tend to avoid the highest zoom level (ZL4). This may be because the quadrant zoom only uses three levels of zoom instead of four and because, since this method does not allow for panning, participants were not concerned with incurring the penalty for panning at the lowest level.

Virtually all of the participants commented on the technique's usefulness for systematic scanning; one commented, "It was quite easy to be methodical." Another participant explained it was "fairly easy to systematically zoom in on targets. Once I developed a kind of methodology for finding targets, I was able to zoom in and out quickly using the keyboard." Reliance on a scanning system could be a disadvantage at times. Participants using the fast version of the Sectional tended to scan through the picture very quickly. All three of them noted that at different times they would become so involved with the rapid systematic search that they would miss a target or forget where they had already looked. Interacting with the method placed a mental burden on the participants, causing them to lose focus on the detection task at times. One participant noted, "One problem I have is that I start with my system and then I get distracted and start somewhere else, and then I forget where I've gone and where I've

been.” Another said about the Sectional, “Although methodical, if you lost your train of thought you found yourself guessing as to whether or not you had been in that particular quadrant.” They all struggled to make sure they slowed themselves down when scanning at the lowest level of zoom, so that they could keep track of where they had been and be sure they had not missed the target.

Participants using the slow version of the technique tended to be more careful and methodical than their counterparts using the fast version. They carefully chose sections to zoom into from ZL1 and from ZL3. This helped them minimize the number of clicks it took them to find a target. Like the participants using the fast Sectional, they found they had better results finding targets when they approached the task more systematically and less intuitively. This was the second fastest of the slow methods; because it is a method that does not require any panning, it was not plagued by the penalty panning incurs in the slow Pointer and ArrowKey methods.

All six of the participants using the Sectional complained that sometimes at ZL4 targets were split between two quadrants or located in a corner of a quadrant instead of the center. They wanted to have finer control of where they were zooming. One participant expressed this when he said the technique gave him “not enough control over exactly where I would want to zoom.” Participants noted that finding targets such as a road or a utility pole in a string of power lines was very challenging because the quadrant zoom does not enable linear searching or tracking features in arbitrary directions; they were much more successful with discrete targets.

At the same time, participants found that ZL3, which overlaps the edges of the sections to a significant degree, could be confusing. One complained, “The computer keeps showing me the same two double-wides, no matter which section I go to!” One observer felt the overlaps at ZL3 were not consistent: “When I hit the 3 [key], I expect to get 50% more information, but I only get 10% new information.” Although we fine-tuned the tool to ensure that the overlaps were consistent, he did not feel that he got an equal amount of information in each new section. Finding an appropriate amount of overlap between sections, so that users were able to see all features completely in at least one section, was therefore problematic. While some overlap seemed to be necessary, it is difficult to determine how much is optimal.

Four participants also noted that they did not find the crosshair tool useful; since they were navigating around the image with the keyboard instead of the mouse, it did not provide them with any information and was sometimes in the way.

All of the Sectional participants had ideas about an ideal interaction technique; they all requested finer control over zooming. Two participants explicitly mentioned that they would like to use a scrolling technique; one said, “A combination of section and scroll techniques might work well so you could get to a high level of zoom quickly and then scroll to see those areas that were not fully captured in that particular section.” Two other participants requested finer zoom control with the mouse. One said he would like “using the mouse to either select or click and drag an area for zooming in.” One participant requested a “notation of where I had already searched” so she would not lose track of the quadrants she had visited.

MagLens The slow version of the MagLens technique performed fairly well; it was faster on average than the Pointer and ArrowKey slow techniques. However, the fast MagLens was the worst performing technique in the test set. While the MagLens technique can be particularly useful for spot checking for targets, its lack of support for systematic searching may have placed it at the bottom of the list of target-finding techniques.

Participants using both versions of the MagLens used similar strategies to search for targets. They would examine the full image to identify locations where the target was likely to be located. They would then zoom in one or two times in the likely locations and pan around those areas looking for the target. This selective magnification technique was fairly successful for most participants. The participants with faster average times per image, using both versions of the technique, were very adept at picking out targets using this method.

If selective magnification was not successful, participants would move to a full scan of the image. Full scanning involved moving the magnification lens, at either ZL3 or ZL4, over the entire image in a lawnmower motion. Five of the participants avoided scanning with ZL4 if possible, only moving to that zoom level after a full scan with ZL3 did not produce a result. As one participant explained, “If you use the highest level of zoom [ZL4], it is easier to see objects but harder to scan, because you lose the context of where you are looking.”

In comparison to participants using other techniques, MagLens participants spent a lot of time examining the full image. This is likely related to the fact that they had access to the full image even when they were utilizing the two zoom levels. Unlike users of the Sectional technique, participants seemed to struggle with the two levels of zoom. Although no participants explicitly requested an extra level of zoom, one participant explained, “Though two levels of zoom were necessary for locating the targets, scanning on the highest level [ZL4] was nearly impossible, but it was difficult to recognize the objects on the other level [ZL3].”

Participants using both versions of the MagLens struggled with knowing exactly where they had already scanned. This problem is exacerbated (in both versions of the technique) by the fact that the area that is magnified in the lens is much smaller than the area that is covered by the lens; in other words, when a small area is being magnified, a large area around it is neither visible under the lens or visible in the full-image view. See Fig. 2 for an illustration of this loss of context.

One participant described this when she complained about the “loss of accuracy” the technique causes. Participants with the fast MagLens technique found it extremely difficult to scan systematically; two of the participants mentioned that they sometimes went too fast and scanned over targets, while one participant mentioned that scanning the picture at ZL4 made her feel

motion sick. These comments indicate that the participants did not have a good sense of exactly what portions of the picture they had magnified and which portions they had not yet viewed in the lens and found the physical sensation caused by panning a small object over a large object was nauseating.

Participants using the slow version of the MagLens technique also had difficulties performing full-image scans, but they were on average more successful. Two of the participants complained about the delay in the slow version. One commented, “The motion of the image as the cursor moved was jerky and there seemed to be a delay so that it was hard to tell how quickly and how far to move the mouse in relation to where I wanted to zoom in on the image.” However, this slight slowness seemed to help participants to control the mouse better than participants with the fast MagLens. This may have




Fig. 2. MagLens magnification area vs area covered by the lens. ZLB is ZL3 and ZLC is ZL4 for the MagLens interaction method.

improved their ability to be systematic in the scanning process, leading to better average times overall.

Two participants mentioned that they never use the crosshair tool; since the full image is always visible, they do not find it helpful. One participant would have liked for the magnifying lens itself to have an outline so that it could be more easily distinguished from the rest of the image. Another participant wanted to move between the two magnification levels without having to turn off the magnification lens in between.

Four participants who used the MagLens technique provided ideas for an ideal technique. Two participants suggested augmenting the existing technique to include some notification, either in the form of a grid or a color overlay, indicating where they had already scanned. One participant said, "it would have been easier if the whole screen enlarged instead of one square superimposed on the screen." Another mentioned that she would have preferred a method that utilizes the keyboard instead of the mouse; she explained, "If I was able to use the keyboard to control the cursor, I [would] feel more comfortable and I also feel this technique [would be] more flexible."

Attitudes Toward Techniques

Since each study participant only had the opportunity to use a single technique throughout the study, we were unable to complete a comprehensive evaluation of participant attitudes towards the techniques. However, we did give them the opportunity to see all five fast versions of the techniques after they had completed the study. Participants were shown the five techniques and then asked to rank them from the technique they would most like to use (1) to the one they would least like to use (5). They were also asked to note the pros and cons of each method. Results are included in the main paper "Results" section.

In many cases, participants' rankings of the methods echoed their comments made during the experiment or on the other sections of the survey. They chose methods that they believed would give them the features they wished their method had, such as more precise zooming controls or a better understanding of the context in which they were searching. Specific participant comments on the methods are discussed below.

Pointer Participants ranked the Pointer technique as one they would most like to use by an overwhelming margin. Many participants commented that they believed this technique would give them "more control" over zooming and panning around the image. It was described as "fast" and several participants commented that they felt it would provide a good sense of their location within the image (so that they would not get lost within the

image). Several participants described the technique as "user-friendly" or "intuitive." One participant commented that the technique "mimic[ed] Net searching." However, two participants felt the technique would cause them to lose their orientation within the image, and one participant described searching with the Pointer as "hit or miss."

The four participants who used the fast version of the technique all ranked it first; two of them noted that it gave them "more control" than the other methods. Participants who used the slow version all ranked it between first and third.

ArrowKey Participants ranked the ArrowKey in the middle behind the Pointer. They commented that, like the Pointer, it gave them "more control" while completing the task but that it was "harder to scan" with it than with the Pointer. Participants found the ArrowKey useful for "systematic" and "controlled" searching. However, one participant described the panning motion as "less smooth" and several participants noted that they felt using the keyboard buttons was overly complicated. One participant explained that there were "too many buttons"; another commented that there was "too much to do" to make the technique work.

Two of the participants who used the fast version of the ArrowKey technique ranked the method first. Users of the slow version ranked it in the bottom three of the techniques.

MagLens Although this technique did not perform well in the efficiency portion of the study, many participants ranked the MagLens as well as the other non-Pointer methods in the survey. A number of participants commented that they found the motion of the lens to be "smooth" and liked that "you don't lose context in the image" because the entire image is always visible. They also described it as easy to use; two participants commented that the technique might feel "familiar" to many users. However, a number of participants noted that the MagLens might be "hard on the eyes." One participant noted that it might make her "dizzy," and another said, "This one just plain is a pain." There seemed to be a difference between just trying it once and using it more extensively. The observers who actually used MagLens during the study generally ranked it very poorly. It was only ranked first or second in two cases, both by MagLens slow technique users.

ScrollBar The ScrollBar was ranked in the middle group in the survey, although it was one of the top three fast performers and was the most efficient slow performer. Two participants felt the technique gave them "more control" and one noted that it is "fluid and fast." However, several others complained that it did not provide enough mobility or flexibility. Several partic-

ipants described it positively as “familiar,” indicating that users might be comfortable with it. However, two other participants described it as “old-fashioned.”

The participants who used the ScrollBar did not tend to rank it highly. One fast participant ranked it first; one slow participant ranked it second; and the rest ranked it in the bottom three techniques. Those that did not rank it high complained that it was “not flexible.” In general, users indicated that while they were able to adapt their search style to make efficient use of the technique during the study, they did not like using this technique as well as the others (except sectional).

Sectional The Sectional ranked as the technique participants would least like to use. Most participants felt the technique would cause them to lose the context of the surrounding image, and they had concerns that targets might get “cut off.” Several participants noted the technique did not provide them enough control. A few participants felt the keyboard system was too hard and contained “too many clicks.” However, a few participants was very enthusiastic about the technique. One participant liked the ability to “dissect the picture.” Another commented that it may be “hard to get lost” within the image. Several participants thought the technique was fast.

One participant with the fast version ranked the Sectional first, and one with the slow version ranked it second. All of the other Sectional participants ranked the technique in the bottom three.

REFERENCES

1. Ackerman R: Commercial imagery aids Afghanistan operations. *Signal* 56(4):16–19, 2000
2. Howard P: Coast Guard studies “digital imaging” systems. *Sea Power* 34(8):48, 1991
3. Mckeown D, Denlinger J: Graphical tools for interactive image interpretation. *Comput Graph* 16(3):189–198, 1982 Proceedings of the 9th Annual Conference on Computer Graphics and Interactive Techniques
4. Reiner BI, Siegel EL, Hooper FJ, Pomerantz S, Dahlke A, Rallis D: Radiologists’ productivity in the interpretation of CT scans: a comparison of PACS with conventional film. *AJR Am J Roentgenol* 176(4):861–4, 2001
5. Raman B, Raman R, Raman L, Beaulieu CF: Radiology on handheld devices: image display, manipulation, and PACS integration issues. *Radiographics* 24:299–310, 2004
6. Heyden JE, Carpendale MST, Inkpen K, Atkins MS: Visual presentation of magnetic resonance images. Proceedings of the conference on Visualization’98, October 1998, 423–426, 1998
7. Plaisant C, Carr D, Shneiderman B: Image-Browser Taxonomy and Guidelines for Designers. *IEEE Softw.* 12, 2 March, 21–32. doi:10.1109/52.368260, 1995
8. Perlin K, Fox D: Pad: an alternative approach to the computer interface, Proceedings of the 20th annual conference on Computer graphics and interactive techniques, September 1993, p.57–64, 1993
9. Bederson BB, Hollan J: Pad++: a zooming graphical interface for exploring alternate interface physics, Proceedings of the 7th annual ACM symposium on User interface software and technology (UIST), Marina del Rey, California, p.17–26, 1994
10. Bederson BB, Meyer J, Good L: Jazz: an extensible zoomable user interface graphics toolkit in Java. Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology (UIST), San Diego, California. p.171–180, 2000
11. Beard DV, Hemminger BM, Perry JR, Mauro MA, Muller KE, Warshauer DM, Smith MA, Zito AJ: Interpretation of CT studies: single-screen workstation versus film alternator. *Radiology* 187(2):565–569, 1993
12. Foley DW, Jacobson DR, Taylor AJ, et al: Display of CT Studies on a two screen electronic workstation versus a film panel alternator: sensitivity and efficiency among radiologists. *Radiology* 174:769–773, 1990
13. Andrichole KP: Productivity and cost assessment of computed radiography, digital radiography, and screen-film for outpatient chest examinations. *J Digit Imaging* 15(3):161–169, 2002
14. Fischer U, Baum F, Obenauer S, Luftner-Nagel S, Von Heyden D, Vosshenrich R, Grabbe E: Comparative study in patients with microcalcifications: full-field digital mammography vs screen-film mammography. *Eur Radiol* 12(11):2679–2683, 2002
15. Hermann KP, Obenauer S, Funke M, Grabbe EH: Magnification mammography: a comparison of full-field digital mammography and screen-film mammography for the detection of simulated small masses and microcalcifications. *Eur Radiol* 12(9):2188–2191, 2002
16. Pisano ED, Cole EB, Kistner EO, Muller KE, Hemminger BM, Brown ML, Johnston RE, Kuzmiak CM, Braeuning MP, Freimanis RI, Soo MS, Baker JA, Walsh R: Interpretation of digital mammograms: a comparison of speed and accuracy of soft-copy versus printed-film display. *Radiology* 223: 483–488, 2002
17. Frank AU, Timpf S: Multiple representations for cartographic objects in a multi-scale tree—an intelligent graphical zoom. *Comput Graph* 18(6):823–829, 1994
18. Furnas GW: Generalized fisheye views. *ACM SIGCHI Bulletin* 17(4):16–23, 1986 Proceedings of the SIGCHI conference on Human factors in computing systems
19. Hombæk K, Frøkjær E: Reading of electronic documents: the usability of linear, fisheye, and overview + detail interfaces. Proceedings of the SIGCHI conference on Human factors in computing systems 293–300, 2001
20. Gutwin C, Fedak C: Interacting with big interfaces on small screens: a comparison of fisheye, zoom, and panning techniques. Proceedings of the 2004 conference on Graphics interface GI ‘04, 2004
21. Baudisch P, Good N, Bellotti V, Schraedley P: Focus and Context: Keeping things in context: a comparative evaluation of focus plus context screens, overviews, and zooming. Proceedings of the SIGCHI 2002 conference on Human factors in computing systems, 259–266, 2002
22. Schaffer D, Zuo Z, Greenberg S, Bartram L, Dill J, Dubs S, Roseman R: Navigating hierarchically clustered networks

- through fisheye and full-zoom methods. *ACM Trans on Comput Hum Interact (TOCHI)* 3(2):162–188, 1996
23. Kenney AR, Rieger OY: Preserving digital assets: Cornell's digital image collection project. *First Monday*, 5(6), 2000
 24. Armitage LH, Enser PGB: Analysis of user need in image archives. *J Inf Sci* 23(4):287–99, 1997
 25. Gough V: Large Image Manipulation Project: designing a new library for processing of large images using a minimal amount of memory. *Linux Journal*, January 1999
 26. Combs TTA, Bederson BB: Does zooming improve image browsing? Proceedings of the fourth ACM conference on Digital libraries, 1999
 27. Hornbæk K, Frøkjær E: Reading patterns and usability in visualizations of electronic documents. *ACM Trans Comput Hum Interact* 10:119–149 doi:10.1145/772047.772050, 2003, (Jun. 2003)
 28. Hightower RR, Ring LT, Helfman JI, Bederson BB, Hollna JD: PadPrints: graphical multiscale Web histories, Symposium on User Interface Software and Technology, Proceedings of the 11th annual ACM symposium on User interface software and technology, 121–122, 1998
 29. Cockburn A, Savage J, Wallace A: Basic level interaction techniques: running and testing scrolling interfaces that automatically zoom. Proceedings of the SIGCHI conference on Human factors in computing systems. 71–80, 2005
 30. Ware C, Fleet D: Context sensitive flying interface. Proceedings of the 1997 symposium on Interactive 3D graphics, 127–130, 1997
 31. Beard DB, Walker JQ: Navigational techniques to improve the display of large two-dimensional spaces. *Behav Inform Techn* 9(6):451–466, 1990
 32. Beard DV, Brown P, Hemminger BM, Misra R: Scrolling radiologic images: requirements, designs, and analyses. *Computer Assisted Radiology (CAR)* Berlin July 3–6, 1991.
 33. Hemminger BM: Design of useful and inexpensive radiology workstations," S/CAR 1992 Proceedings, Symposia Foundation, 1992
 34. Beard DV, Hemminger BM, Denelsbeck KM, Brown PH: A time motion comparison of several radiology workstation designs and film. S/CAR 1992 Proceedings, Symposia Foundation, pp588–594, 1992
 35. Beard DV, Hemminger BM, Pisano ED, Denelsbeck KM, Warshauer DM, Mauro MA, Keefe B, McCartney WH, Wilcox CB: Computed tomography interpretations with a low-cost workstation: a timing study. *J Digit Imaging* 7(3):133–139, 1994
 36. Kaptelinin V: A comparison of four navigation techniques in a 2D browsing task. *ACM Special Interest Group on Computer-Human Interaction Conference Proceedings*, 282–283, 1995
 37. Hemminger BM: Softcopy display requirements for digital mammography. *J Digit Imaging* 16(3):292–305, 2003
 38. National Imagery and Mapping Agency (NIMA): NGA Raster Roam tool. Available at <http://www.nima.mil/>. Accessed September 2005
 39. United States Geological Survey (USGS): The National Map Viewer. Available at <http://www.usgs.gov/>. Accessed September 2005
 40. Environmental Systems Research Institute, Inc.: Arc-View GIS. Available at <http://www.esri.com/software/arcgis/arcview/>. Accessed September 2005
 41. Bernstein P: 1998. Now that I've got it, what do I do with it? No-nonsense image utilities for non-techies. *Online* 22 (1):86–90, 1998
 42. Puff DT, Cromartie R, Pisano ED, Muller KE, Johnston RE, Pizer SM: Evaluation and optimization of contrast enhancement methods for medical images. *Proc SPIE, Vis Biomed Comput* 1808:336–346, 1992
 43. North C, Shneiderman B, Plaisant C: User controlled overviews of an image library: a case study of the visible human. Proceedings of the first ACM international conference on Digital libraries, 74–82, 1996
 44. Hornbæk K, Bederson BB, Plaisant C: Navigation patterns and usability of zoomable user interfaces with and without an overview. *ACM Trans Comput-Hum Interact* 9 (4):362–389, 2002 (Dec. 2002), doi:10.1145/586081.586086
 45. Carr D, Plaisant C, Hasegawa H: Designing a real-time telepathology workstation to mitigate communication delays. *Interact Comput* 11(1):33–52, 1998
 46. Baldonado MQW, Woodruff A, Kuchinsky A: Guidelines for using multiple views in information visualization. In: Tarrantino L Ed. Proceedings of the 5th International Working Conference on Advanced Visual Interfaces (AVI'2000, Palermo, Italy, May 24–26). New York: ACM, 2000, pp. 110–119
 47. Card SK, Moran TP, Newell A: The psychology of human-computer interaction, Hillsdale: Lawrence Erlbaum, 1993