A comparison of zoomable user interfaces and folders for grouping visual objects

LANCE GOOD*†, MARK J. STEFIK‡, AND BENJAMIN B. BEDERSON§

†Palo Alto Research Center, 3333 Coyote Hill Rd, Palo Alto, CA 94304, USA, phone: 1.650.812.4661, fax: 1.650.812.4258, email: good@parc.com
‡ Palo Alto Research Center, 3333 Coyote Hill Rd, Palo Alto, CA 94304, USA, phone: 1.650.812.4366, fax: 1.650.812.4241, email: stefik@parc.com
§Human Computer Interaction Lab, University of Maryland, College Park, MD 20742, USA, phone: 1.301.405.2764, fax: 1.301.405.6707, email: bederson@cs.umd.edu

*Corresponding Author Email: good@parc.com

Abstract

Making sense of information on computers is notoriously difficult because of limited size displays. Two interface techniques used to cope with these limitations are zooming and overview+detail. To evaluate these techniques, we ran a user study comparing a Zoomable User Interface (ZUI) and a folder-based overview+detail interface. The task used in the study required subjects to organize shapes into groups based on common visual properties. The study results showed 30% faster completion times with the ZUI than with folders. The ZUI's advantage most likely arose from its support for fast visual comparisons. Ultimately, we expect that these results can be extended to a range of problem solving tasks involving a variety of visual objects.

Keywords: Zooming; Folders; Overview+detail; Zoomable User Interfaces (ZUIs); Visual comparisons

AMS Subject Classification: 68U35; 91E45

1 Introduction

Modern user interfaces typically employ information visualization components to present immutable data but revert to more traditional interface controls to edit data. Yet it is often in these traditional controls where much of the sensemaking actually gets done, such as keeping track of discoveries made over time, incorporating existing knowledge about a domain, organizing information into groups, and synthesizing collections of information for a specific task. As a result, some recent systems have begun to explore techniques to blur the lines between fixed-data visualizations and user-editable controls.

One notable recent system in this space is VITE (Hsieh *et al.* 2002). VITE provides facilities for mapping multi-dimensional data to a 2D layout of multi-faceted visual objects. The visual objects in VITE use pure geometric properties such as color, position and size to represent discrete and continuous data attributes. Nominal data attributes are also represented in these objects as text labels. Unlike many previous visualization systems, VITE also provides a means for modifying the visualization which can optionally propagate back to the underlying data. The combination of geometric and textual properties along with a modifiable workspace leads to a range of system functionality from a pure visualization tool, such as a starfield display (Ahlberg *et al.* 1994), to a more authoring centered environment, like the VKB or VIKI spatial hypertext systems (Marshall *et al.* 1997, Shipman *et al.* 2001). This type of system offers the potential to integrate a user's visual exploration and sensemaking activities.

We have begun to explore these types of integrated visualization and sensemaking workspaces. One problem that we have observed in using these environments for both visualization and sensemaking is a lack of screen space. Users who are solving these types of tasks with more traditional physical media tend to spread their information out on tables, whiteboards, and other physical surfaces. Yet, computer displays have roughly 30 to 40 times less display space than even an ordinary table or whiteboard (Henderson *et al.* 1986). Research suggests that the limited view of these displays can make certain types of sensemaking tasks more difficult than with traditional media because of a lack of good global awareness (Monty 1990, Severinson-Eklundh 1992, Sharples 1996).

This previous research suggests that this problem of inadequate global awareness may be alleviated through improved overviews. However, research has not clearly pinpointed the reasons why good overviews are necessary for these kinds of tasks. We suggest that overviews are needed to make fast visual comparisons between chunks of information. In a large workspace, like a whiteboard or a table, a user can spread the information out and do comparisons between chunks of data with eye and head movements. On a computer screen this same task often requires the user to scroll, navigate, or rearrange to see the information chunks to be compared. Consequently, the task is much more time consuming and thus cognitively disruptive.

Naturally, the frequency of visual comparisons varies across the different activities involved in the task. One place where comparisons are especially common is in organizing ideas into groups and hierarchical structures. As a result, effective overviews are particularly important to support this activity.

Three main techniques have been introduced to provide improved overviews on computer displays. These include overview+detail, zooming, and fisheye. Our work focuses on the overview+detail and zooming techniques since these methods preserve the spatial arrangements used by authors to indicate semantic relationships.

This paper describes an experiment comparing the effects of a folder-based overview+detail interface and a Zoomable User Interface (ZUI) on user performance in completing an organization task. The task used in our study requires subjects to put shapes into groups based on common visual properties. This task is intended to be representative of a range of sensemaking tasks such as those involving more verbal information such as text.

2 Related Work

There have been a number of previous studies that have looked at users performing problem solving tasks with visual workspaces. One such study looked at the VITE system, which is perhaps the most closely related to the systems studied in the current paper (Hsieh *et al.* 2002). In this study, Hsieh et al. collected users' qualitative experiences with the system for two problem solving tasks. They found that users were able to construct appropriate mappings of data attributes to visual attributes and effectively manipulate information in the workspace to complete the tasks.

A longitudinal usage study was also performed for an earlier visual authoring system named Aquanet (Marshall *et al.* 1992). This study consisted of a small set of users making sense of a large collection of information in Aquanet over the course of 2 years. The study focused on the kinds of structures and relationships created using Aquanet in assessing machine translation systems and technologies.

Marshall et al. also compared users' qualitative experiences using paper and two versions of a visual sensemaking workspace (Marshall *et al.* 1997). The task in this study again required subjects to assess machine translation systems. However, in this study subjects were given strict time-constraints to make a decision. From the study, Marshall and Shipman derived several recommendations for future hypertext authoring tools.

There is also a large body of research related to concept mapping. However, this research has primarily focused on the effectiveness of concept mapping in fields such as education, counseling, planning, and group communication (see Abrams *et al.* (2004)). We are not aware of studies that focus on interface techniques to make users more effective at organizing concept maps.

A number of researchers have also quantitatively compared zooming to other techniques for visualizing information. Hornbæk et al. compared ZUIs with and without an overview for a series of map tasks (Hornbæk *et al.* 2002). The primary results from the study indicate that subjects are faster using the ZUI without an overview when the task involved a multilevel map or required navigation.

Plumlee et al. performed a comparison of zooming versus a multiple window strategy for tasks involving visual comparisons in large 2D workspaces (Plumlee *et al.* 2002). They present a model of a user's expected performance on these tasks based on working memory. A user study then confirmed the predictions of the model that zooming interfaces are more efficient for tasks with low visual memory demands and multi-window interfaces are more efficient for tasks with higher visual memory demands. One fundamental assumption in this study was that objects in the workspace are not recognizable when they are viewed zoomed out. In contrast, the assumption in our study is that system designers will take care to provide objects with good representations so that they retain as much value as possible when viewed zoomed out.

There has been little research at the intersection of these two areas. In other words, controlled studies have not compared zooming to other interface techniques for the kinds of sensemaking tasks described here. Below we present an initial study in this area focusing on a specific subcomponent of the larger sensemaking task.

3 Task Description

We are interested in combining visualizations and interactive workspaces to assist users in solving complex real world sensemaking problems. Completing these types of sensemaking tasks requires a combination of many different activities. Our work focuses on one particular activity that involves organizing information into groups of related objects. Choosing such a narrow focus allows us to isolate the effect of the two interface techniques on a common, well-defined subcomponent of the complete sensemaking task.

Additional motivation for this focus on organization comes from previous research in computer-based workspaces, such as Notecards (Monty 1990) and VIKI (Marshall *et al.* 1992), that suggests that many of the difficulties in sensemaking tasks arise in defining the structure and relationships between ideas. Further, the organization component seems to be one of the places where the problem of getting an overview of the workspace is particularly difficult on a computer display. In a longitudinal study of a user's experiences in Notecards, Monte reported that the subject wanted to print his information and 'spread [it] out on a table' in order to organize it (Monty 1990). Similar research cited by Severinson-Eklundh also found that users of computer-based authoring tools generated paper printouts to consider the document's organization (Severinson-Eklundh 1992).

[Insert figure 1 about here]

The particular instance of organization task used in our experiments involves dividing a large collection of shapes into smaller equal-sized groups of shapes with two common visual properties. For example, the completed task might have one group containing circles with green backgrounds, another group containing objects with blue backgrounds and pink borders, etc. An example of this task can be seen in figure 1. The objects have a number of properties on which they can be compared including shape, background color, background texture, border color, border texture, and size. Each of the six properties in our study had six possible values leading to 6^6 or 46,656 different possible shapes. The task starts with a random arrangement of shapes and is completed when the shapes had been sorted into a specified number of equal-size groups.

Table 1 presents the task from a more traditional categorization perspective where objects are described in terms of features. For shapes, these features are visual properties but for more complex objects these features could include topics, measures of complexity, level of detail, etc. In this sense, the shape task is representative of a range of organization tasks such as those that involve text, web pages, images, or pen strokes. Previous research in human cognition has also demonstrated similarity between categorization tasks involving different types of objects. For example, Anderson presented a model of human categorization that simulated human behavior across a variety of categorization tasks involving both discrete-valued features and continuous-valued features (Anderson 1991). Admittedly, there are also differences between the shape task and these other tasks. In particular, because it uses perceptual properties such as color, size, shape, and texture, this shape task introduces the issue of preattentive versus conscious processing. However, the role of preattentive processing in this task was likely to be limited since the task required the subject to match pairs of visual properties. Research has shown that combinations of preattentive features are typically not preattentively processed (Triesman 1985). The tasks may also differ in their use of verbal versus visual working memory. This means that subjects may be able to hold different numbers of objects in memory between the different tasks. However, as noted by Plumlee (2002), in practice the number of features that can be held in memory for visual versus verbal object types is likely to be comparable.

[Insert table 1 about here]

4 Overview Visualization Techniques

As mentioned, there are three common techniques for providing overviews of large computer-based workspaces: fisheye, overview+detail, and zooming. The fisheye technique provides overviews by introducing warping or distortion. For spatial arrangement tasks, distortion is often not desirable as it is more likely to interfere with users' perception of spatial relationships. As result, distortion-based approaches were not explored in the current study. Identifying the most appropriate distortion-based interfaces for these tasks and comparing them to zooming and overview+detail remains a topic for future research.

Overview+detail techniques make use of multiple windows to simultaneously display the workspace at multiple levels of detail. Typically, one window shows an overview of the entire space while a second window displays the workspace in greater detail.

Severinson-Eklundh noted that there are two kinds of possible overviews for a workspace, including physical and logical (Severinson-Eklundh 1992). Physical overviews portray a geometrically accurate rendering of the workspace. When used with an overview+detail approach, physical overviews must significantly scale a large workspace to fit the available window size. This limits the utility of the overview if objects in the space contain detailed information, such as textures or text, since this detail is often obscured at dramatically reduced magnifications. In contrast, logical overviews portray the conceptual structure of the space. These overviews preserve the high level organization of the workspace but omit the geometric layout. Instances of logical overviews such as folders or outlines are often used for organizing information since they can represent more appropriate abstracted information based on the overview window size.

Zooming interfaces combine a large workspace with navigation mechanisms for viewing the space at varying magnification. Views with alternate magnifications are obtained sequentially, typically through mouse or keyboard operations. In this paper, ZUIs are intended as a specific kind of zooming interface where the transitions between levels of magnification are animated.

5 Study

One of the primary reasons for using a workspace for a sensemaking task is because it functions as a large external memory. This external memory allows users to work with more information than can fit in their organic working memory at one time. However, the external memory also introduces additional time to access information. In a large workspace where all the information is visible at once, the user can retrieve information with eye and head movements which introduce a relatively short access delay. In smaller workspaces such as a computer displays, motor movements with more significant performance times may be necessary to manage and retrieve information in the workspace.

The cost of accessing information in a small computer workspace depends on the strategy used to manage the space. Since this space acts as a kind of memory, these strategies can be compared to virtual memory strategies on a computer. Henderson et al found a similar link between user behavior in a windowed computer environment and working sets from the virtual memory literature (Henderson *et al.* 1986).

The assumptions of most virtual memory strategies are that programs make use of limited size working sets and have locality of reference (Henderson *et al.* 1986). This means that only a limited amount of information is needed in working memory at once. Strategies that exploit this assumption make use of a technique called paging, which divides the data into page size chunks that are swapped in and out of working memory. When these assumptions are violated, memory management strategies lead to thrashing where a large percentage of time is spent swapping pages in and out rather than doing actual work.

The overview+detail technique facilitates a space management strategy that is comparable to paging in virtual memory. The overview acts as an access mechanism for selecting a ‘page’ of the workspace to display in the detail view. However, the assumptions of the virtual memory strategies described above do not necessarily hold for organization tasks. Indeed, rather than locality of reference, these tasks often require the user to access objects throughout the workspace. This type of random access is needed to make comparisons and find similarities between objects that are distributed throughout the workspace.

Zooming facilitates an alternative space management strategy that allows the author to zoom out to see the entire workspace in reduced detail. The amount of detail visible in these views depends on the display size, the size of the workspace, and the quality of representations for the zoomed-out objects. If enough detail is preserved in these views then they can support the kinds of non-local visual comparisons that are needed to efficiently complete the organization tasks.

Based on this analysis, our hypothesis is that zooming will provide better support than overview+detail for authors completing the organization tasks described above. The specific overview+detail interface in the study below uses a logical folder-based overview. A logical overview was chosen since they are more commonly used and better suited for a wide range of organization tasks. These types of logical overviews appear in applications such as file browsers, image libraries, slide authoring tools, email clients, etc. In keeping with the traditional behavior of these overviews, the contents of each folder are contained in their own disjoint 2D space. The specific zooming interface used in the study is implemented as a single large workspace with interactions for zooming in and out. It also uses animations when navigating between levels of magnification.

6 Method

6.1 Participants

Fourteen regular computer users participated in the study. Five of the subjects were female and nine were male. Subjects ranged in age from 18 to 50 and were given a gift certificate for their participation in the study.

6.2 Equipment

The study was run on Pentium 3 and Pentium 4 machines (ranging from 700 MHz to 2 GHz) with at least 256 MB of memory and 1600 x 1024 pixel flat panel displays.

6.3 Procedure

The experiment used a within-subjects design. There was one independent variable which was interface type (zooming or folders). The dependent variables were completion time and two subjective ratings of interface quality. The primary interface metaphor for moving the shapes in both interface conditions was simple drag and drop of shapes. Both interfaces were developed in Java with Swing and the Piccolo ZUI Toolkit (Bederson *et al.* 2004).

[Insert figure 2 about here]

The folder interface condition mimicked a file browser with a folder hierarchy in a resizable pane on the left side of the screen and a freeform 2D workspace in a pane on the right side of the screen. The folders came from a standard java tree widget where the folder labels had a minimum clickable area of 55 x 22 pixels (as compared to the minimum of 33 x 17 pixels in Windows Explorer). Each folder in the hierarchy provided access to its own (conceptually) infinite non-zoomable 2D space that was made visible in the right hand pane when the folder was selected. This 2D space allowed at least 10 shapes to be visible at once and provided scrollbars when shapes in the workspace extended outside the current view. The interface also provided facilities to create, name, rearrange, and delete folders. When objects were moved between folders, the tool automatically positioned objects as close to the centre of the folder's space as possible without overlapping other objects. The folder interface is shown on the left in figure 2.

The ZUI condition provided a freeform 2D workspace as in the previous condition, but instead of a folder hierarchy, it provided mouse interactions to zoom in and out in the workspace. A single left click on a shape zoomed the workspace in to one-to-one magnification centred on the shape and a single right click anywhere zoomed the workspace out to show all shapes. Here again, scrollbars were provided when shapes extended outside the current view, for example when the view was zoomed in. The ZUI is shown on the right in figure 2. The shapes in this zooming condition also implemented a type of semantic zooming (semantic zooming is described by Bederson et al. (1996)) that tried to visually preserve shape properties. So while the absolute size of the shapes was affected by zooming, the other properties, including a shape's size relative to other shapes, were largely unaffected. One primary exception was that background colour and pattern sometimes became difficult to see as the view was zoomed out and border widths increased relative to the size of the objects. Similarly, as the view was zoomed out the relative differences in size between shapes became less apparent. This type of problem required participants to zoom in to resolve any ambiguities.

In both interface conditions, the system implemented a simple form of 'bumping' which prevented objects from overlapping one another. Because this kind of bumping is difficult to implement for the range of shapes used in the study, the system implemented an approximation of the optimal occlusion-avoidance behaviour. Although there were cases where this algorithm defaulted to bumping with the bounding boxes, in practice it gave reasonable results for the most common cases.

It should be noted that although occlusion may not be particularly disruptive for shapes, it is likely to be more disruptive for more abstract object types such as text, web thumbnails, or pen input. In fact, users of a tool for organizing thumbnails of web bookmarks were observed to lose occluded thumbnails (Robertson *et al.* 1998). Bumping also effectively increases the demand for space, which is likely to highlight differences between the two interfaces. A similar effect could have been achieved by increasing the number of shapes in the tasks, but this would have also increased completion times.

The study task involved grouping shapes based on common visual properties. The shapes had six properties (described previously) that were described to the subjects prior to completing the tasks. For each task, subjects were asked to identify groups of 10 shapes that had two of the six properties in common. These two properties were not specified to the subjects ahead of time, as a primary component of the task was intended to be the discovery of the group membership criteria. A sample task is shown in figure 2. Subjects were informed that a unique set of property pairs divided the objects into groups of 10 but that individual objects may fit into more than one group.

One visual property that required special consideration in designing the study tasks was object size. In particular, it was necessary for subjects to be able to clearly differentiate between differently sized shapes when zoomed in. This was most important for comparing the sizes of different shape types, for instance diamonds and squares. Informal testing of different variations in shape sizes led to a difference of 40 pixels between sizes. A minimum size constraint of approximately 125 x 125 pixels was also enforced for the shapes so that subjects would be able to clearly identify the background textures for all the different (background texture)-(border type)-(shape) combinations when zoomed in. The combination of these different constraints led to an average bounding box of 205 x 205 pixels for the shapes in the study. However, it should be noted that several of the shapes (i.e. star, hexagon, triangle, circle, and diamond) covered only a fraction of the area in their bounding box which allowed for better packing of shapes than the bounding box would imply.

For each condition, subjects were first given a tutorial on how to use the interface. The task was then explained and the six shape properties were described. Subjects were told to indicate grouping through spatial proximity in the zooming condition and with folders in the folder condition. The experimenter demonstrated a sample task using the current interface, and the subjects were then given their own practice grouping task to complete. Finally, the subjects were given the actual timed task and were told to complete the task as fast as they were able. We allowed all subjects to work until they completed the tasks. This precluded having to make subjective judgments about the quality of incomplete groupings.

At the beginning of each task, the shapes were distributed in a random order. For the zooming condition, the shapes were distributed in a single large grid. In the folder condition, the shapes were distributed in folders such that each folder initially contained 10 objects. The practice tasks were composed of 30 objects requiring three groups of 10 and the actual tasks were composed of 50 objects or five groups of 10.

There were two instances of both the practice tasks and the actual tasks. The ordering of the interface condition, practice task, and actual task were all counterbalanced to reduce ordering effects.

There was a trade-off in the design of the starting condition for the folders task. Instead of distributing the shapes into the five folders, shapes could have been put into a single vertically-scrollable root folder to start with. This would have the advantage that all objects would be in a single conceptual space to start with. However, because the screen could show approximately 10 shapes at a time, the subject would have to scroll through approximately five screens worth of shapes at the beginning of the task. It was ultimately decided that the five folder icons provided faster and more consistent access to the shapes. Subjects could quickly see all 50 objects with five clicks on the folder icons. Moreover, they could directly jump to a consistent set of 10 objects with a single click.

Nevertheless, it should also be noted that folders did not have to be used with the folder-based interface. Although subjects were asked to put their final group selections into individual folders, they were free to move all the objects into a single folder during the task, using the scrollbars instead of folder icons to navigate. Because the tool automatically positioned objects as close to the centre as possible, subjects could quickly move the objects into a single folder and get a well-packed space in which to scroll around. The time needed to move all the objects into a single folder was relatively insignificant in comparison to the overall completion times. In the experiment, one subject was observed to use this strategy.

The software logged all relevant operations, though this information was ultimately not analyzed. Following the two tasks, subjects were given a questionnaire regarding their experiences with the two conditions. The questionnaire included background questions, two subjective ratings, and several free response questions.

7 Results

Paired two sample t-tests were used to analyze the results from the study. These were chosen over ANOVAs because only one independent variable and two samples were being compared. Standard deviations are presented for each of the measures. However, it should be noted that because the study design was within-subjects, statistical significance is dependent on the difference between measures for individual subjects rather than the individual measures themselves.

[Insert figure 3 about here]

One subject's results were discarded because of several significant cell phone interruptions during the timed tasks. For the remaining subjects, a paired two sample t-test was performed for task completion times. This data indicated a statistically significant effect of interface type $p<0.03$ with a 30% faster completion time for zooming. The mean completion times, shown in figure 3, were 17.5 (σ 8.0) minutes for folders and 12.2 (σ 7.0) minutes for zooming. The range of completion times were 5.3 to 30.0 minutes for folders and 4.6 to 25.8 minutes for zooming. All subjects were able to complete the tasks successfully and without errors.

[Insert figure 4 about here]

There was also a statistically significant effect of interface type on both questionnaire rating items. For the question ‘What is your overall reaction to using this tool for the task?’ (1=Frustrating,5=Satisfying) the means were 4.2 (σ 0.6) for zooming and 2.8 (σ 1.0) for folders with $p<0.001$. For the question ‘How often did you feel like the software interrupted your thinking?’ (1=Rarely,5=Often) the means were 1.8 (σ 0.9) for zooming and 3.4 (σ 1.3) for folders with $p<0.001$. These results are displayed in figure 4.

8 Discussion

The results indicate that the zooming interface provides a considerable time savings over the folder interface for the shape grouping task. A number of factors likely contribute to this advantage but the most notable seems to be the time involved in making comparisons between shapes. In the zooming condition, comparisons were usually made with rapid eye and head movements that took a few hundred milliseconds. In contrast, the folder condition often required subjects to look at objects in multiple folders. Moving between folders introduced a mouse pointing operation that increased the overhead of many comparisons to more than a second. Because of the large number of comparisons involved, this relatively small difference in time became significant over the complete task.

Other factors affecting user performance are suggested by the work of Shipman et al. They describe the harmful aspects of formal interfaces similar to folders (Shipman *et al.* 1999). These kinds of formalisms can add cognitive overhead to intellectual work like organization tasks. In particular, Shipman et al. found that formalisms introduce extra steps and decisions in chunking and labeling information for authoring tasks.

In the folder interface above, extra steps and decisions were often needed to manage ungrouped shapes. As the task progressed, users frequently needed to make decisions about how to partition the remaining ungrouped shapes among folders. In contrast, the zooming interface required fewer of these decisions since all the shapes were arranged in the same continuous space. Extra effort was also needed in the folder interface to produce descriptive names for the folders to help remember their contents. The zooming interface avoided this extra labeling effort since all group contents were readily visible in a single overview.

Shipman et al. also explain how formalisms can prematurely force the user to make decisions about creating groups and other structures. Once put into a formal structure, the information becomes more difficult to convert to alternate structures. This kind of structural inertia was observed with the folders interface above. Using the folder interface, subjects would often mark a folder with a descriptive name or unique position to indicate that the contained group was completed. Informal feedback from the study indicated that this behavior provided subjects with a sense of orderliness and progress. However, subjects were also observed to be hesitant to revisit groups that had been marked as complete. As a result, this behavior often delayed their completing the task when these groups contained an error. This problem was less frequently observed with spatial arrangement in the zooming condition, suggesting that the formal separation of groups in the folder condition may be harmful for maintaining awareness of groups once the user considers them complete.

9 Conclusion

This work was intended to find out why sensemaking tasks are often more difficult to complete with software tools than with physical tools. Previous research has indicated that organizing information is a particularly difficult component of these tasks on computers because of the limited size displays. To understand the mechanisms behind these observations, our work evaluated how well two interface techniques, overview+detail and zooming, support the organization activity on a space-constrained computer display.

The task in our study required subjects to organize shapes into groups based on similar visual properties. The study demonstrated a 30% improvement in completion time for a zoomable workspace over a folder-based overview+detail workspace. Subjects also indicated strong preferences for the zoomable workspace. These results are somewhat surprising given that most subjects have extensive experience with folders in tools such as file browsers and email clients and only minimal exposure to zoomable interfaces.

The observed differences between interface conditions seemed to be primarily driven by the shorter times needed to visually compare objects in the continuous zoomable workspace versus the segmented folder workspace. In addition, the formality of the folder interface was also a likely factor contributing to the poorer performance.

We also suggest that the results from this study will extend beyond shapes to a range of organization tasks including those with more abstract object types like text objects, web pages, images and pen strokes. Because these tasks all involve making comparisons among groups of objects, the facilities of a zoomable workspace are likely to have similar advantages to those seen for shapes.

Finally, it should be noted that although the visualization problems addressed here will be mitigated by the availability of larger computer displays, these problems will persist for large sensemaking tasks and for portable devices such as laptops and PDAs. Consequently, these results are intended to guide interface designers in applying visualization techniques in these settings for future software authoring and sensemaking tools.

10 Future Work

Zooming is only useful insofar as the objects in the space have good representations at multiple magnifications. In this respect, shapes are ideal for zooming since they retain most of their value at a wide range of scales. Future studies will look at how well these results extend to additional domains such as text, images, web pages, and pen input.

[Insert figure 5 about here]

Automated semantic zooming techniques are likely to be important for improving the usability of abstract data in zoomable environments. One particular technique we have implemented to create more meaningful representations of text objects under zooming is automatic text reduction (Good *et al.* 2002), shown in figure 5. Automatic text reduction combines content reductions and font size reductions to provide more meaningful views of text objects under zooming. Other automatic semantic zooming techniques include enhanced thumbnails (Woodruff *et al.* 2001) and automatic thumbnail cropping (Suh *et al.* 2003), which provide more meaningful reduced representations of web pages and photographs respectively. Additional work is needed to determine the extent to which these technologies support sensemaking tasks, such as those described previously, in zoomable environments.

Acknowledgements

This research has been funded in part by contract # MDA904-03-C-0404 to Stuart K. Card and Peter Pirolli from the Advanced Research and Development Activity, Novel Intelligence from Massive Data program.

References

ABRAMS, R., KOTHE, D. and IULI, R., 2004, A Collaborative Literature Review of Concept Mapping. Available online at: www2.ucsc.edu/mlrg/clr-conceptmapping.html (accessed 20 Oct. 2004).

AHLBERG, C. and SHNEIDERMAN, B., 1994, Visual Information Seeking: Tight coupling of dynamic query filters with starfield displays. In *Proceedings of CHI '94* (New York: ACM Press), pp. 313-317.

ANDERSON, J. R., 1991, The Adaptive Nature of Human Categorization, *Psychological Review*, **98**, 409-429.

BEDERSON, B. B., GROSJEAN, J. and MEYER, J., 2004, Toolkit Design for Interactive Structured Graphics. *Transactions on Software Engineering*, **30**, 535-546.

BEDERSON, B. B., HOLLAN, J. D., PERLIN, K., MEYER, J., BACON, D. and FURNAS, G. W., 1996, Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics. *Journal of Visual Languages and Computing*, **7**, 3-31.

GOOD, L., BEDERSON, B. B., STEFIK, M. and BAUDISCH, P., 2002, Automatic Text Reduction for Changing Size Constraints. In *Proceedings of CHI 2002 Extended Abstracts* (New York: ACM Press), pp. 798-799.

HENDERSON, D. A. and CARD. S., 1986, Rooms: the use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics*, **5**, 211-243.

HORNBÆK, K., BEDERSON, B. B. and PLAISANT, C., 2002, Navigation Patterns and Usability of Zoomable User Interfaces with and without an Overview. *ACM Transactions on Computer-Human Interaction*, 9, 362-389.

HSIEH, H. and SHIPMAN, F., 2002, Manipulating Structured Information in a Visual Workspace. *UIST 2002, ACM Symposium on User Interface Software and Technology, CHI Letters*, **4**, 217-226.

MARSHALL, C. C. and ROGERS, R. A., 1992, Two Years before the Mist: Experiences with Aquanet. In *Proceedings of Hypertext '92* (New York: ACM Press), pp. 53-62.

MARSHALL, C. C., and SHIPMAN, F. M., 1997, Spatial Hypertext and the Practice of Information Triage. In *Proceedings of Hypertext '97* (New York: ACM Press), pp. 124-133.

MONTY, M. L., 1990, Issues for supporting notetaking and note using in the computer environment. Ph.D. Dissertation, University of California, San Diego.

PLUMLEE, M. and WARE, C., 2002, Zooming, Multiple Windows, and Visual Working Memory. In *Proceedings of AVI 2002* (New York: ACM Press), pp. 59-68.

ROBERTSON, G., CZERWINSKI, M., LARSON, K., ROBBINS, D. C., THIEL, D. and VAN DANTZICH, M., 1998, Data Mountain: using spatial memory for document management. In *Proceedings of UIST '98* (New York: ACM Press), pp. 153-162.

SEVERINSON-EKLUNDH, K., 1992, Problems in achieving a global perspective of the text in computer-based writing. *Instructional Science*, **21**, 73-84.

SHARPLES, M., 1996, An account of writing as creative design. In *The Science of Writing: Theories, Methods, Individual Differences and Applications*, C. M. Levy and S. Ransdell (Eds.) (Hillsdale: Lawrence Erlbaum Associates), Ch. 8.

SHIPMAN, F. and MARSHALL, C. C., 1999, Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems. *Computer Supported Cooperative Work*, **8**, 333-352.

SHIPMAN, F., HSIEH, H., AIRHART, R., MALOOR, P. and MOORE, J.M., 2001, The Visual Knowledge Builder: A Second Generation Spatial Hypertext. In *Proceedings of Hypertext '01* (New York: ACM Press), pp. 113-122.

SUH, B., LING, H., BEDERSON, B. and JACOBS, D., 2003, Automatic Thumbnail Cropping and its Effectiveness. *UIST 2003, ACM Symposium on User Interface Software and Technology, CHI Letters*, **5**, 95-104.

TRIESMAN, A., 1985, Preattentive Processing in Vision. *Computer Vision, Graphics, and Image Processing*, **31**, 156-177.

WOODRUFF, A., FAULRING, A., ROSENHOLTZ, R., MORRISON, J. and PIROLLI, P., 2001, Using Thumbnails to Search the Web. *CHI 2001, ACM Conference on Human Factors in Computing Systems, CHI Letters*, **3**, 198-205.

Figure 1

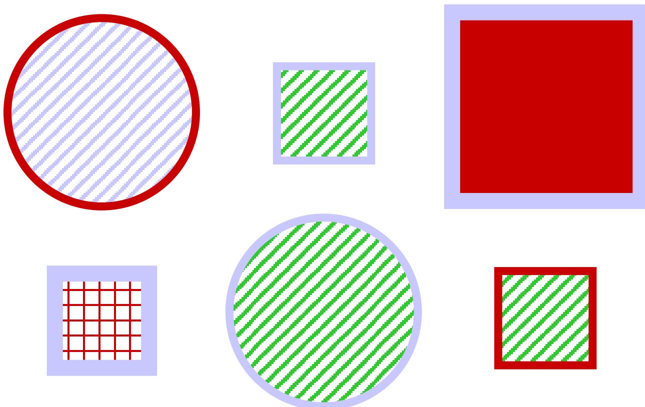

Figure 2

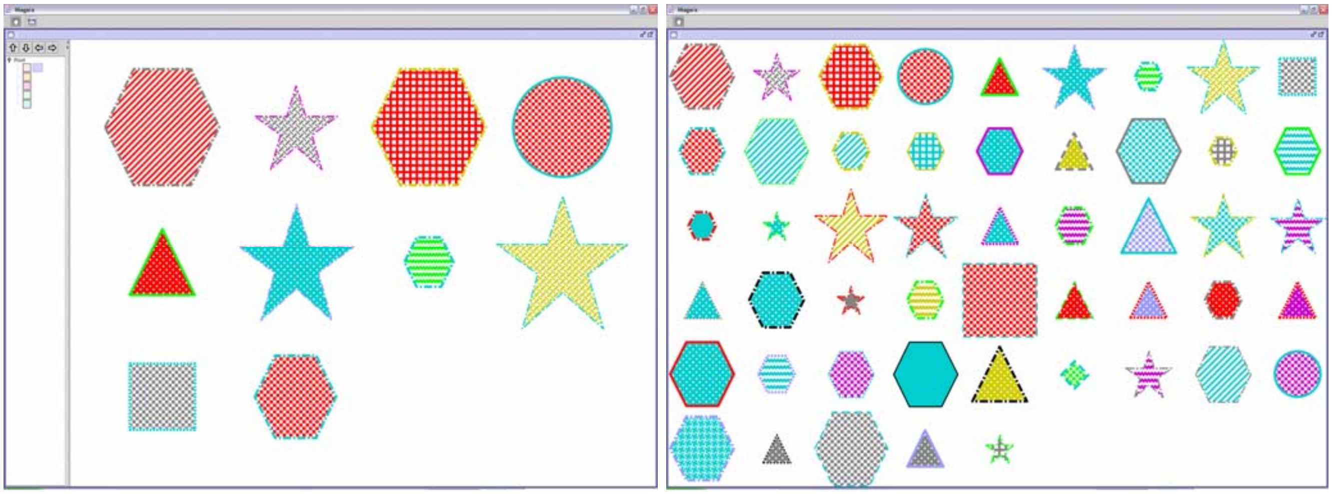

Figure 3

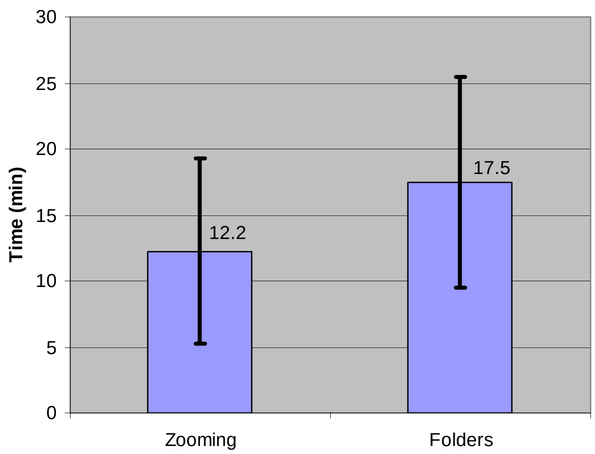

Figure 4

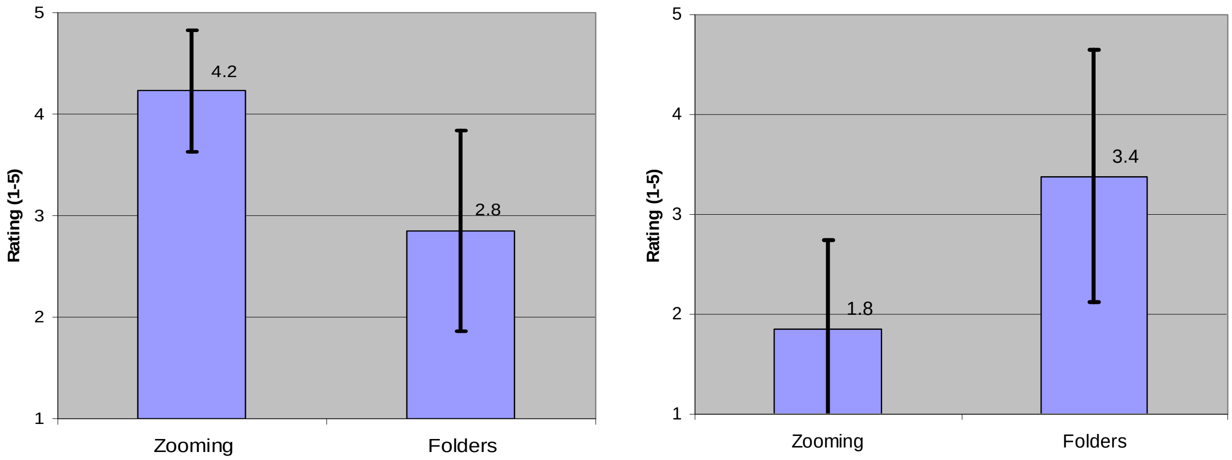

5
4
3
2
1
Rating (1-5)
4.2
2.8
Zooming
Folders
5
4
3
2
1
Rating (1-5)
1.8
3.4
Zooming
Folders

Figure 5

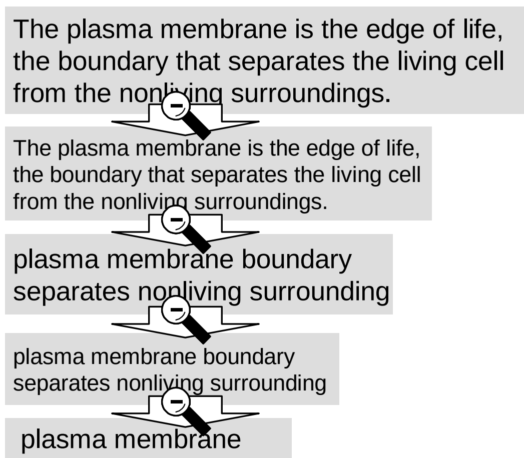
The plasma membrane is the edge of life, the boundary that separates the living cell from the nonliving surroundings.
The plasma membrane is the edge of life, the boundary that separates the living cell from the nonliving surroundings.
plasma membrane boundary separates nonliving surrounding
plasma membrane boundary separates nonliving surrounding
plasma membrane

Table 1

		Features			
Shape	**Border Color**	**Border Type**	**Background Type**	**Background Color**	**Size**
Square	Light	Thick	Grid	Dark	Small
Square	Light	Thin	Stripes	Medium	Small
Square	Light	Thick	Solid	Medium	Large
Square	Dark	Thin	Stripes	Medium	Small
Circle	Light	Thin	Stripes	Medium	Large
Circle	Dark	Thin	Stripes	Light	Large

Figure Captions

Figure 1 - An example shape grouping task. The shapes can be divided into two groups of three where each group has two visual properties in common. The first group is formed by the three squares with light borders. The remaining objects form the second group, and are defined by having diagonal striped textures and thin borders.

Figure 2 - The starting state for the two interface conditions in the shape study. The folder interface is on the left and the ZUI is on the right.

Figure 3 - A graph comparing the mean completion times for zooming versus folders. The error bars indicate standard deviation.

Figure 4 - a) On the left, a graph comparing subject ratings for the question 'What is your overall reaction to using this tool for the task?' (1=Frustrating, 5=Satisfying). b) On the right, a graph comparing subject ratings for the question 'How often did you feel like the software interrupted your thinking?' (1=Rarely, 5=Often). In both, the error bars indicate standard deviation.

Figure 5 - An automatic semantic zooming technique for text. Automatic text reduction alternates between reducing the font size and reducing the content.

Table Captions

Table 1 - The example task from Figure 1 shown in terms of features. The rows are sorted in descending order by shape and then border color. The lightly filled cells indicate the two common properties in the first group and darker filled cells indicate the two common properties in the second group.