The Effect of Shading in Extracting Structure from Space-Filling Visualizations

Pourang Irani, Dean Slonowsky, Peer Shajahan
Department of Computer Science, Department of Statistics, Department of Computer Science
University of Manitoba, Canada
irani@cs.umanitoba.ca, dean_slonowsky@umanitoba.ca, peermoh@cs.umanitoba.ca

Abstract

Shading information is extracted by the human visual system during the earliest stages of the object recognition process. While shading can enhance the visibility of structures, it can have a negative impact on the judgment of sizes of elements in a structure. In certain visualization systems the underlying hierarchical structure is not noticeably explicit, such as in space-filling techniques. We hypothesize that in such cases, shading can make the structure more explicit. In this paper we report the results of a study comparing two space-filling visualizations: one with and the other without shading. Our results show that on structure-based tasks, users performed better with the tool when shading information was included than without. However we did not obtain statistically significant results to suggest that shading information degraded users' performance on tasks requiring comparison of local features such as file sizes. A subjective evaluation shows that users preferred interacting with the system when shading was available.

Keywords -- **2½-D representation, structure-from-shading, shading and perception, space-filling visualization, TreeMap, CushionMap.**

1. Introduction

Hierarchical data structures are abundant and interacted with on a regular basis. They describe the relationships among entities in organizations, in file systems, and in family genealogies. Information structured as a hierarchy is organized into levels. The hierarchy begins with a root at the top most level, and all other components are related through the root.

To adequately navigate or locate components within a hierarchy the structure needs to be evident to the user. The most common form of hierarchical representation is a node-link tree (Figure 1). The structure of the hierarchy is *explicit* and visually clear to the user (i.e. we can clearly see all the child-parent relationships in the hierarchy). However trees are not space efficient. A significant amount of space remains unused in the background as a result of creating an adequate layout for the nodes.

Space-filling visualizations have been developed to make more efficient use of the display area. These systems are characterized by their compactness and effectiveness at showing the sizes of elements in a hierarchy. The basic space-filling representation divides the display area into blocks of nested entities, with each block representing a node in the hierarchy. However, unlike node-link trees, some space-filling systems represent child-parent relationships by nesting the nodes and as a result do not reveal the hierarchical structure explicitly.

An example of a space-filling technique is the TreeMap [10]. In this system the display is partitioned into rectangular regions to map an entire hierarchy of nodes and their children (Figure 1). Each node occupies an amount of space relative to the weight of the item being represented in the hierarchy (such as the relative size of a file in a directory or the volume of shares sold on the stock market). TreeMaps are good for revealing global patterns in the data such as large pockets of empty space on a disk drive or nodes of a specific type (i.e. file type). However, the underlying hierarchical structure is not as visually explicit as that of a node-link tree (in Figure 1, parent nodes are not visible due to the nesting of components).

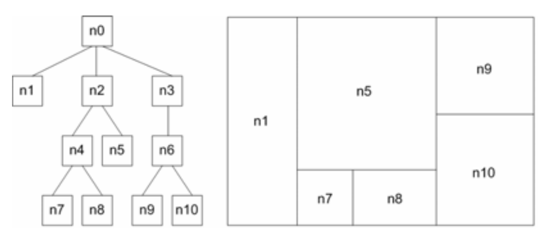

Figure 1 The node-link tree to the left clearly reveals the structure of the hierarchy. The TreeMap (an example of a space-filling visualization) to the right depicts the relative size of nodes but does not readily display the structure of the hierarchy.

To make the hierarchy visually more explicit, van Wijk et al. [14] developed a modified version of the TreeMap called the CushionMap. The CushionMap system uses shading to give the hierarchy a 2½-D impression, making

Proceedings of the Eighth International Conference on Information Visualisation (IV'04)
1093-9547/04 $ 20.00 IEEE

the structure more explicit. However this claim was not empirically validated. To test whether shading information makes non-explicit hierarchical structures more explicit we compared the TreeMap visualization to the CushionMap system. The results of our investigation are reported here.

2. Related Work

From the literature in information visualization, two particular evaluations of the TreeMap are relevant to our discussion. Turo and Johnson [13] compared TreeMap to UNIX shell command line syntax for simple directory and file browsing tasks. They found that subjects performed better with the UNIX shell commands on tasks of a local scope, i.e. comparing the size of two files. However subjects carried out their tasks quicker and more accurately with the TreeMap on tasks having a more global nature, such as finding the number of files in a directory.

In another study, Stasko et al. [11] compared TreeMap to a radial space-filling technique, called Sunburst. In the Sunburst technique the root of the hierarchy is placed in the center of the visual space and files and directories are laid out radially in wedges extending from the center. Each level of the hierarchy is a concentric circle and the deepest level is furthest away from the center. The size of the file or directory is represented by the angular sweep of each wedge. Their results show that for tasks involving file and directory size comparisons, subjects' performance were superior with TreeMap. On the other hand, for tasks that necessitated creating a mental map of the hierarchy, such as finding a file in the hierarchy, subjects performed better with Sunburst.

Overall, the results in [11] suggest that Sunburst conveys global hierarchical structural information better than TreeMap but at the cost of local feature information, such as file size. However, a drawback to a radial technique is that the area available for the display is not maximized. Ideally, a space-filling technique should make optimal use of the screen space, should be capable of displaying structure and at the same time facilitate size comparisons between elements in a hierarchy.

The CushionMap system was designed to exploit the use of shading to make the hierarchical structure more explicit in the TreeMap [14]. The focus of our evaluation has been to empirically test the effect of shading (such as that available with CuhionMaps) on tasks necessitating both structural information as well as local feature information. We first present results from the literature in perception research that suggest the importance of shading for perceiving structure in our environment. This underlying body of literature, referred to as structure-from-shading has motivated our investigation and the evaluation reported herein.

3. Structure-from-Shading

Research in the area of human perception shows that our visual system extracts shading information from a scene at a very early stage in the recognition process. In particular, there is evidence that such information is processed pre-attentively. A study by Enns and Rensink [4] investigated the influence of scene-based properties such as direction of lighting, surface locations and orientations, and surface reflectance on visual search. Their targets were composed of colored polygons with white, gray, and black pixels (Figure 2) some of which could be interpreted as three-dimensional objects. The task consisted of locating single target items among 1, 6 and 12 objects. They found that observers were significantly slower in finding the target when the items were two-dimensional. They concluded that rapid search is possible when the items consist of *spatial* and *intensity* relations that can be interpreted as *three-dimensional* objects.

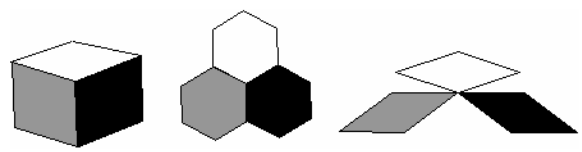

Figure 2 Targets used in Enns and Rensink's experiments [4]. The target to the left corresponded to projections of simple blocks under various lighting conditions. The pattern on the left was perceived faster than the 2D patterns on the right.

Sun and Perona [12] extended the work of Enns and Rensink [4] by investigating the pre-attentive perception of elementary three-dimensional shapes. To determine whether shading was more important than internal line crossings (which can contribute to determining the shape of a three-dimensional object) they compared the speed of processing single target patterns consisting of 3D shaded top-lit polyhedrals to their unshaded line drawing counterparts (Figure 3). Their results suggest that the shaded objects were processed faster and in parallel while the line drawings of the 3D shaded objects were processed serially. Their results are consistent with those of Braun [3] which showed that smoothly shaded circular targets, without any internal line edges, and which resemble 3D shaded bubbles, are processed in parallel and pre-attentively based on the perception of their 3D shape.

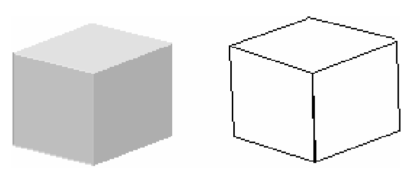

Figure 3 Sample targets used in the experiment by Sun and Perona [12]. A feature in common between the shaded item and the line drawing is the embedded Y-junction that assists in determining the shape of the object.

Beyond pre-attentive processing, shading information is also critical in structural object recognition. From a high-level view, object recognition is accomplished in a series of stages (Figure 4). At the first stage, the visual image is analyzed into primitives of edge elements, color and texture. This information is then used to segment the image so that the boundaries of objects can be extracted. Elementary shape from shading information is then used to extract object components [8] or "blobs" [1]. At the same time it also facilitates the creation of a structural skeleton, containing information about how the components are interconnected [7,2]. Ultimately, all of the information is combined which leads to object recognition.

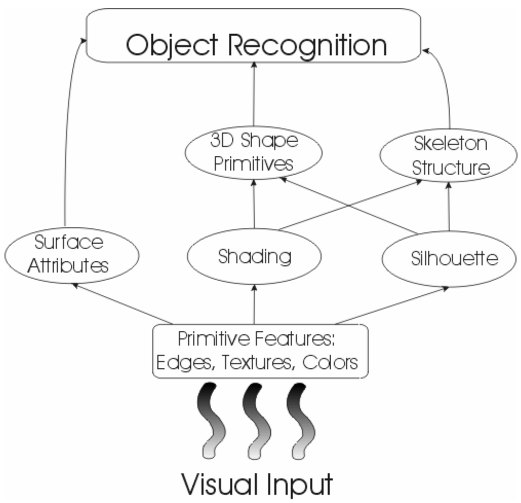

Figure 4 Theories of structural object perception propose a series of processing stages, culminating in object recognition. Shading plays a critical role in shape and structure extraction.

Irani and Ware [6] compared the effectiveness of visually parsing diagrams with and without shading. The non-shaded diagrams were constructed using solid nodes and links and only differed from the shaded diagrams in that the former were flat (Figure 5). Their results show that using 3D shaded nodes and links resulted in more accurate substructure identification (11.4% vs. 21% errors) and shorter times (4.1 seconds vs. 5.2 seconds). The subjects also accurately recognized more 3D diagrams than 2D non-shaded structures (20% vs. 34% error rate). Their results strongly suggest that shading facilitates visual parsing and recognition of diagrams.

The results from the above investigations support the idea that shading can enhance the visual parsing of structured representations of data. In particular, diagrams composed of nodes and links can benefit from shading information. Of course, in the case of node-link diagrams the structure of the hierarchy is already explicit. What has not been investigated is whether shading can enhance the representation of hierarchical data where the structure of the hierarchy is *not* explicit, such as in space-filling representations.

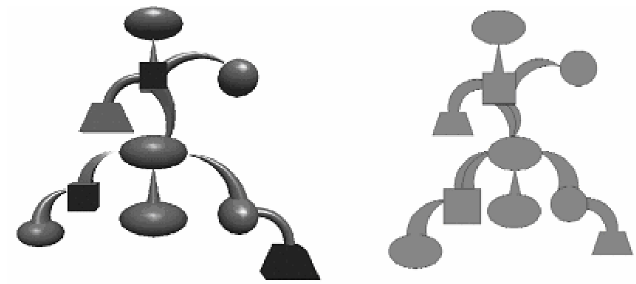

Figure 5 Diagrams used for comparing shading vs. no-shading in explicit structures. Adapted from [6].

While there is evidence that shading can facilitate structural identification, several studies have reported that simple shading information that gives a 3D impression can degrade users' accuracy in tasks relating to size comparisons. Carswell et al. [5] compared 2D bar and pie charts to their 3D counterparts. In their investigation, subjects were asked to make relative magnitude estimations based on looking at the two forms of graphs. Their results show that subjects performed better in magnitude judgment tasks with the 2D graphs.

Zacks et al. [15] conducted an investigation to find out whether graphs with a 3D impression influenced viewers' ability to extract information from it. In their experiments they varied the rendering characteristics and relative heights of the bars and asked participants to estimate the quantities portrayed. They found that the addition of 3D perspective depth cues lowered accuracy.

From the results of these studies we can conclude that shading information can be a detriment to tasks necessitating accurate judgment of size information. We have taken these results into consideration in the design of our experiment.

The remainder of this paper describes the specific visualizations we have evaluated and the results of our experiment.

4. Visualization Tool

To conduct the experiment, we used SequoiaView (version 1.3), an application developed by van Wijk et al. [14] at the Department of Mathematics and Computing Science of the Technische Universiteit Eindhoven [9]. SequoiaView presents file hierarchies in both the TreeMap and CushionMap views (Figure 6). The tool is equipped with multiple configuration options including the mapping of file types onto color. In our experiment we controlled all the options and only created two conditions: the presence and the absence of shading. We used the default mappings of file extensions onto color and only used the following file types: ".bmp", ".jpeg", ".mp3", ".exe" and ".dll". For

Proceedings of the Eighth International Conference on Information Visualisation (IV'04)
1093-9547/04 $ 20.00 IEEE

the experiment, the hierarchies (directories and files) were randomly sorted. We kept the filter settings untouched with the exception of not displaying hidden system files. The hierarchies used for our experiment are a variation of an actual hierarchy found on a machine in our lab. There were in total 120 subdirectories and 2300 files.

SequoiaView is built in such a way that by moving the mouse over each item, a hint displays the absolute path of the file. The user can navigate using an address bar that is located at the top of the tool, similar to that found in Windows File Explorer™. Double-clicking a subdirectory opens it so that the entire display is covered by the subdirectory. Browsing and navigating can be achieved by using the "up-one-level" button but this was not offered as an option to the users. Users were also asked to not modify any of the settings but instead to focus on the assigned tasks.

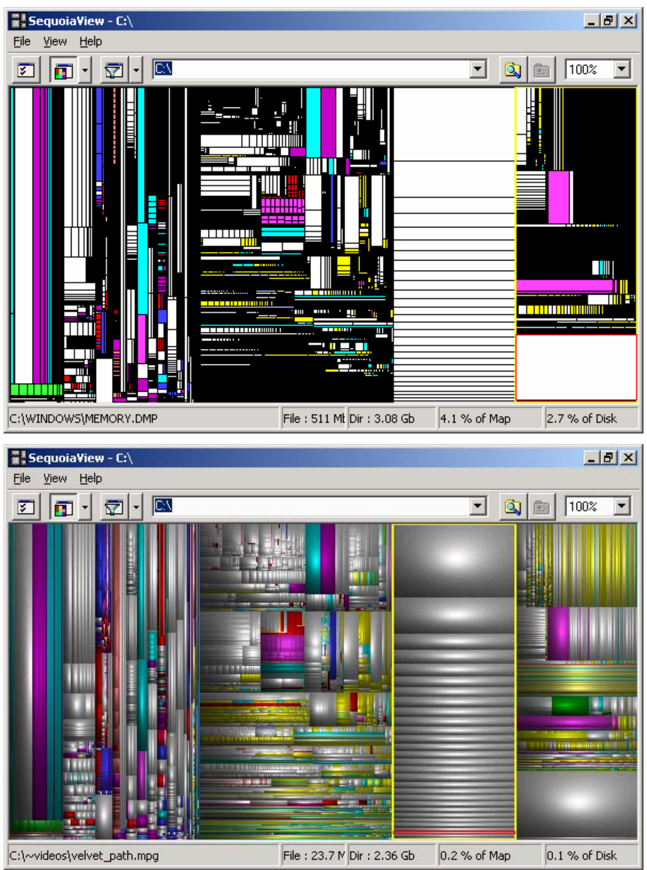

Figure 6 TreeMap and CushionMap representation of the local C: drive from SequoiaView.

5. Experiment

The experiment was designed to compare the TreeMap (hereafter referred to as TM) to the CushionMap (CM) on common file and directory management tasks. Each participant performed a series of tasks using both tools. To reduce learning effects we used two hierarchies (H1 and H2) which were different in terms of file names, order of files and directories. However both hierarchies were similar in structure, i.e. they each had the same number of subdirectories and files and with similar sizes. Half the participants started the experiment with the TreeMap and the other half started on the CushionMap. The set of trials consisted of {CM-H1, TM-H2}, {CM-H2, TM-H1}, {TM-H1, CM-H2}, and {TM-H2, CM-H1}. After completing the set of tasks on one tool they took a brief break and switched onto the other. All tasks required that subjects find files or subdirectories or perform size comparisons of files and of subdirectories.

Based on the earlier studies in perception discussed in section 4, we anticipated the following effects in our experimental data:

Hypothesis 1: The *shading* condition (CM) will result in higher performance on *structure* related tasks than the *no-shading* condition (TM).

Hypothesis 2: The *shading* condition (CM) will result in lower performance on tasks related to file and directory *size comparisons* than the *no-shading* condition (TM).

5.1 Method

5.1.1 Subjects. Twenty undergraduate students participated in the experiment and were randomly assigned to one of the two conditions: CushionMap first or TreeMap first. Subjects were primarily computer science and engineering majors. All were familiar with the concept of file and directory structures and had reasonable experience performing standard file management routines. None had any previous experience using SequoiaView and the TreeMap or CushionMap visualization tools.

5.1.2 Materials. Participants performed the experiment on a 17" monitor running SequoiaView over WinXP. Each task was read aloud to the participant and was placed on a sheet beside the computer for their referral during the experiment.

5.1.3 Procedure. Just prior to the experiment, subjects spent time getting familiar with both visualization systems. Then the experimenter read through a tutorial describing the various features of SequoiaView. The tutorial was given using a different hierarchy than those used in the experiment. It involved a series of tasks similar to the ones that would be given in the experiment. The experiment began only after the subjects indicated that they were comfortable using the tool and its interface. After the training session, each participant performed 10

Proceedings of the Eighth International Conference on Information Visualisation (IV'04)
1093-9547/04 $ 20.00 IEEE

tasks using the tools on the hierarchies. A short questionnaire was provided at the end of the experiment.

We measured participants' performance as a success (if the task was completed within 45 secs) or a failure (incorrect result or timeout). The time to execute the task was recorded in both cases. The experiment involved the following tasks:

- Count the number of directories in the hierarchy.
- Find the directory with the most number of files.
- Count the number of subdirectories in a given directory.
- Count the number of files in a given subdirectory.
- Find the directory with the most number of bit map files (.bmp).
- Count the number of sub-directories that contain bitmap (.bmp) files.
- Find the smallest directory in the hierarchy.
- Find the largest file in the hierarchy.
- Find the largest file in a given directory.
- Find the largest mp3 file in the hierarchy.

These particular tasks were chosen since they constitute representative tasks when working with file systems. Some examples include locating a particular type of file in a directory, finding a file occupying the largest amount of disk space, or comparing two directories by size when deciding which to delete. The hierarchies we tested our subjects on were developed such that the tasks would be relatively difficult to do.

The first six tasks were designed to test our first hypothesis. For instance, to count the number of directories or files in a subdirectory, the structure of the hierarchy would need to be relatively clear. The final four tasks were designed to test our second hypothesis. Note that to successfully complete certain tasks, such as finding a particular file based on its type, subjects would need to rely on color information. However, the basic nature of the task would still require that the subject compare sizes or use the hierarchy structure to successfully complete it.

At the start of each task the tool was refreshed to present the root view of the hierarchy. This ensured that all the subjects commenced the tasks from the same starting point. At the end of the experiment participants filled out a questionnaire stating their preference for either type of visualization across all the tasks.

5.2 Results and Discussion

To test the two hypotheses stated in the beginning of this section, we measured subjects' performance on each task with respect to two variables: time until completion (0 to 45 seconds) and successful/unsuccessful completion (0/1). For both hypotheses, we recorded the average response (over all tasks involved in the given hypothesis) for both of these variables. These measurements were taken for each of the 20 subjects, resulting in four data sets (structure-based, size-based $\times$ 2 response variables), each containing 20 pairs (CM,TM). For the dichotomous response variable, timeouts were classified as failures (0). Average completion times were consistent with the normality assumption in all datasets, whereas average success rates were far from normal.

Following the methodology employed in an earlier and related study [11], any task that was unsuccessfully completed or a timeout was excluded when calculating the average completion time. As a result, in the dataset of completion times for structure-based tasks, two individuals had an average response of 0 seconds using the TreeMap tool (they failed in all their tasks). Since this does not adequately measure their performance, these times were eliminated so that the corresponding dataset had only eighteen pairs of responses.

5.2.1 Effect of treatment and hierarchy order. Ten of the subjects were randomly allocated to "TM, then CM" and the remaining ten were allocated to "CM, then TM" (these two groups remained the same in all stages of the experiment). Likewise, independent of the allocation of subjects to treatment order, ten subjects were randomly allocated to "H1, then H2", the remaining ten to "H2, then H1". To justify our pooling of the subjects, we first verified that neither treatment order nor hierarchy order had a significant effect on our data.

A one-way ANOVA F-test was conducted to test if the order of treatments (CM or TM first) had a significant effect on average completion times. For both structure-based tasks and size-based tasks, we did not detect a significant difference between these two groups with respect to the variable CM–TM ($F_{1,16}$=2.5311, p=0.1312 and $F_{1,18}$=0.0112, p=0.9169, respectively).

Similar tests were conducted to test if the order in which subjects encountered hierarchies H1 and H2 had a significant effect on average completion times. Once again, a one-way ANOVA test detected no significant effect on CM–TM ($F_{1,16}$=0.2247, p=0.6419 for structure-based, $F_{1,18}$=0.2251, p=0.6409 for size-based).

A non-parametric Kruskal-Wallis test on the average success rate data detected no significant effects of treatment order (CM or TM first) on structure-based ($\chi^2(1)$=1.3401, p=0.2470) or size-based tasks ($\chi^2(1)$=0.4241, p=0.5149). Similarly, no significant effect was detected for order of hierarchy (H1 or H2 first) on either type of task ($\chi^2(1)$=0.6567, p=0.4177 for structure-based, $\chi^2(1)$=0.0149, p=0.9028 for size-based).

Proceedings of the Eighth International Conference on Information Visualisation (IV'04)
1093-9547/04 $ 20.00 IEEE

5.2.2 Analysis of structure-based tasks. Table 1 summarizes the results for average completion times and average number of tasks successfully completed. We recall that there were 6 structure-based tasks. A paired t-test confirms that there is a significant decrease in the mean completion times for CushionMap over TreeMap (t_{17}=-3.3, p=0.0021). We notice that subjects are approximately 33% faster with the CushionMap than with the TreeMap.

A non-parametric Wilcoxon rank-sum test shows a significant increase in average success rates for CushionMap over TreeMap on structure-based tasks ($p < 0.001$). By comparing the average number of tasks successufully completed with both systems, we see that the subjects were 45% more successful with the CushionMap than with the TreeMap.

These results provide very strong evidence in favor of hypothesis 1.

	Structure	Size
Average Completion Time (seconds)	TM=21.5 (6.1) CM=16.2 (3.7)	TM=17.9 (5.4) CM=20.2 (5.4)
Average # of tasks successfully completed	TM=2.7 (1.5) CM=4.9 (0.8)	TM=3.4 (0.7) CM=3.1 (0.9)

Table 1 Average completion times in seconds (for correct responses only) and average number of tasks successfully completed (standard deviations are in parentheses).

5.2.3 Analysis of size-based tasks. To test hypothesis 2, we used a paired t-test to compare the mean average completion times of the 4 size-based tasks. Whereas the observed mean average completion times for the 20 subjects was 2.3 seconds lower when using the TreeMap (13% faster with TM than with CM), this difference was not statistically significant (t_{19}=1.6707, p=0.0556).

Using a Wilcoxon rank-sum test with the one-sided alternative hypothesis, “average number of tasks successfully completed among CushionMap users is lower than those of TreeMap users for size-based tasks”, we failed to detect a significant advantage in using TreeMaps over CushionMaps (p=0.120). Indeed, the mean difference in average success rates for size-based tasks among the 20 subjects was merely 0.3 in favor of the TreeMap system.

Table 2 summarizes the results of the analysis described above. Statistically, our data strongly supports hypothesis 1, that shading will facilitate structure-based tasks. This confirms results from the literature in perception on the effect of shading for identification and extraction of structure. On the other hand, our results do not provide conclusive evidence that shading has adverse effects on size-based tasks and therefore do not support hypothesis 2. As a result, we cannot corroborate the work of others in suggesting that shading distorts the structure of the display, leading to misjudgments of local size features.

	Structure	Size
Completion Time	CM significantly faster that TM (p=0.0021)	No significant difference between CM and TM
Completion Success	Subjects significantly more accurate on CM over TM (p<0.001)	No significant difference between CM and TM

Table 2 Statistical significance of TM versus CM on structure-based and size-based tasks.

Based on our results, the degree to which each set of tasks (structure-based or size-based) is supported by either type of visualization is shown in Figure 7. While this chart may not accurately capture the entire picture, we can at least deduce that there is still a need for a space-filling tool that can adequately reveal global structure and at the same time allow users to compare local size features. We are currently investigating other forms of visualization methods that will satisfy these criteria.

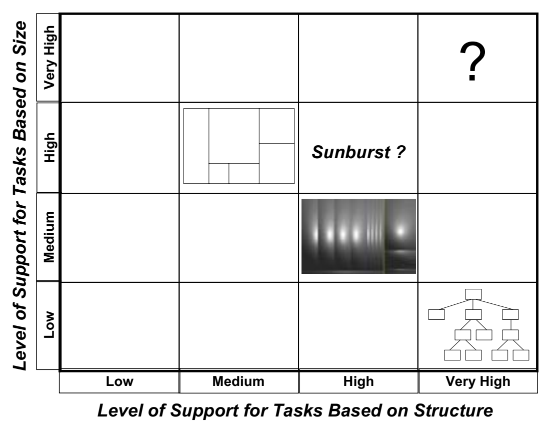

Figure 7 Space-filling techniques can be approximately positioned according to the degree to which they support structure-based and size-based tasks. Note that the top-right corner is empty – none of these systems fully support both types of tasks.

5.3 Subjective Evaluation

In addition to tracking performance measures, we also collected subjects' opinions of each tools utility. Participants replied to 10 statements on a Likert-scale with responses ranging from 1 (strongly disagree) to 5 (strongly agree). The average scores are summarized in Table 3.

The questions were based on the tasks that were completed earlier in the experiment. On average, subjects preferred the CushionMap visualization on all the assigned tasks. Their preference for CushionMap on size-related

Proceedings of the Eighth International Conference on Information Visualisation (IV'04)
1093-9547/04 $ 20.00 IEEE

tasks was not a result of superior performance as revealed by the experimental data. Anecdotally, several subjects reported that the 2½D effect from the shaded representation created a feeling of “wanting to click” onto the objects. This invitation to click on the nodes could have possibly contributed to the higher level of comfort with the CushionMap system.

Statement	TM	CM
1. I was able to count the number of directories using *toolname.*	3.65	4.40
2. I was able to find the bitmap (.bmp) files using *toolname.*	3.70	4.60
3. I was able to detect the type of files using *toolname.*	3.95	4.55
4. I was able to find subdirectories using *toolname.*	3.60	4.35
5. I was able to find the files inside a sub-directory using *toolname.*	3.05	3.95
6. I was able to find the largest file using *toolname.*	3.50	3.95
7. I was able to compare the sizes of files using *toolname.*	3.30	3.90
8. I was able to find the largest directory using *toolname.*	3.70	4.40
9. After the training session I knew how to use *toolname.*	4.00	4.35
10. I found *toolname* confusing to use.	3.05	2.05

Table 3 Subjective preferences averaged across subjects for both types of representations. 5 represents “strongly agree” and 1 “strongly disagree”.

6. Conclusions

In this paper we have presented the results of an experiment for testing the effect of shading on visually parsing non-explicit hierarchical structures. To conduct our evaluation, we used two previously developed space-filling visualizations, the CushionMap and the TreeMap. Supported by theories of structure-from-shading, we begun our investigation with the claim that shading will facilitate extraction of hierarchical structures. We also postulated that performance on size-based tasks will be impaired by the use of shading.

Our results confirmed the first hypothesis. Users were faster and more accurate in completing directory management tasks with the shaded hierarchies. This supports and adds to the previous body of literature on the nature and benefits of shading. On the other hand, we did not obtain any conclusive results on the unfavorable effect of shading for size-based tasks. This warrants a more discerning follow-up experiment where the effect of shading on size judgments is better controlled.

Subjective responses suggest that the participants preferred interacting with the system when shading was available. This affirms the intuition of the designers of the CushionMap.

Although not conclusive, our data suggests the need to improve the CushionMap so that it will facilitate fast and accurate comparison of sizes of elements in a hierarchy. We may potentially be able to improve performance of size judgments and not affect visibility of the structure. This might be accomplished by modifying the type of shading, by using various forms of texture or by dynamically adjusting the display based on a pre-specified task. We are in the process of investigating and implementing these alternatives.

Our investigation also contributes to the growing body of literature in information visualization related to evaluation methods and techniques. As emphasized by Stasko et al. [11], such empirical studies could be beneficial to designers in building systems which will be effectively utilized by real-users.

7. Acknowledgements

We would like to thank Colin Ware (University of New Hampshire, USA) for his insight and support. This research was supported by grants from NSERC.

8. References

[1] Biederman, I. Recognition-by-Components: A Theory of Human Image Understanding, Psychological Review, 94:2, 115-147, 1987.

[2] Biederman, I. and Gerhardstein, P.C. Recognizing depth-rotated objects: Evidence and conditions for three dimensional viewpoint invariance, Journal of Experimental Psychology: Human Perception and Performance, 19:1162-1182, 1993.

[3] Braun, J. Shape-from-shading is independent of visual attention and may be a "texton", Spatial Vision, 7:311-322, 1993.

[4] Enns, J. T. and Rensink, R. A. Influence of Scene-Based Properties on Visual Search, Science 247:721-723, 1990.

[5] Carswell, C., S. Frankenberger, S. and Bernhard, D. Graphing in depth: perspectives on the use of three-dimensional graphs to represent lower-dimensional data, Behaviour and Information Technology 10, pp. 459–474, 1991.

[6] Irani, P. & Ware, C. Diagramming Information Structures using 3D Perceptual Primitives, ACM Transactions on Computer Human-Interaction, 10:1, 2003.

[7] Marr, D. Vision: A computational investigation into the human representation and processing of visual information, San Fransisco, CA: Freeman, 1982.

[8] Ramachandran, V. S. Perception of shape from shading, Nature, 331:163-166, 1988.

[9] SequoiaView, Technische Universiteit Eindhoven, available at www.win.tue.nl/sequoiaview.

[10] Shneiderman, B. Tree visualization with TreeMaps: a 2d space-filling approach. ACM Transactions on Graphics, 11(1):92–99, September 1990.

[11] Stasko, J.T., Catrambone, R., Guzdial, M. and McDonald, K. An Evaluation of Space-Filling Information Visualizations for Depicting Hierarchical Structures International Journal of Human-Computer Studies, vol. 53, pp. 663-694, 2000.

[12] Sun, J.Y. and Perona, P. Preattentive perception of elementary three-dimensional shapes, Vision Research, 36(16):2515-2529, August 1996.

[13] Turo, D. and Johnson, B. Improving the visualization of hierarchies with TreeMaps: design issues and experimentation, Proceedings of the IEEE Visualization, pp. 124-131. Boston, MA, 1992.

[14] Wijk, J.J. van, Wetering, H. van de. Cushion TreeMaps, Proceedings of the IEEE Symposium on Information Visualization, pp. 73-78, 1999.

[15] Zacks, J., Levy, E., Tversky, B. and Schiano, D. J., Reading bar graphs: effects of depth cues and graphical context. *Journal of Experimental Psychology: Applied* **4**, pp. 119–138, 1998.

Proceedings of the Eighth International Conference on Information Visualisation (IV'04)
1093-9547/04 $ 20.00 IEEE