Information Visualization (2006) 5, 47–61

www.palgrave-journals.com/ivs

Human perception of structure in shaded space-filling visualizations

Pourang Irani[1]
Dean Slonowsky[2] and
Peer Shajahan[1]

[1]Department of Computer Science, University of Manitoba, Canada and [2]Department of Statistics, University of Manitoba, Canada

Correspondence:
Pourang Irani, Department of Computer Science, University of Manitoba, 556 Machray, R3T 2N2, Canada.
Tel: +1-204-474-8995;
Fax: +1-204-474-7609;
E-mail: irani@cs.umanitoba.ca

Abstract

Very early in the object recognition process the human visual system extracts shading information. 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 two experiments designed to investigate the effects of shading information on extracting the structure in space-filling visualizations. In the first experiment subjects performed better with the visualization tool with shading on structure-based tasks. Our results do not show that shading impairs users' judgment on size-based tasks. A subjective evaluation shows that users preferred interacting with the system when shading was available. The second experiment was designed to investigate further users' capacity to identify structural elements within the space-filling visualization. A sub-structure recognition task was employed in this experiment and results show that users are capable of identifying sub-structures quicker and with fewer errors when the visualization tool was equipped with shading information than without. The results of both experiments provide evidence that shading information can be used to effectively obtain structural information from space-filling visualizations.

Information Visualization (2006) **5,** 47–61. doi:10.1057/palgrave.ivs.9500113

Keywords: 2½-D Representation; structure-from-shading; shading and perception; space-filling visualization; TreeMap; CushionMap

Introduction

Hierarchies are common and abundantly available in our daily information managing activities. They describe the relationships among entities in file systems, organizations, and social structures. Hierarchies are organized into levels with a root at the top most level, and all other components at sub-levels. To adequately navigate or locate components within a hierarchy the structure needs to be evident to the user. A popular and common representation of hierarchical data is the 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.

Research in Information Visualization has developed a range of displays known as space-filling systems to make more efficient use of the display space available to the user. 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

Received: 15 August 2005
Revised: 21 December 2005
Accepted: 6 January 2006
Online publication date: 10 April 2006

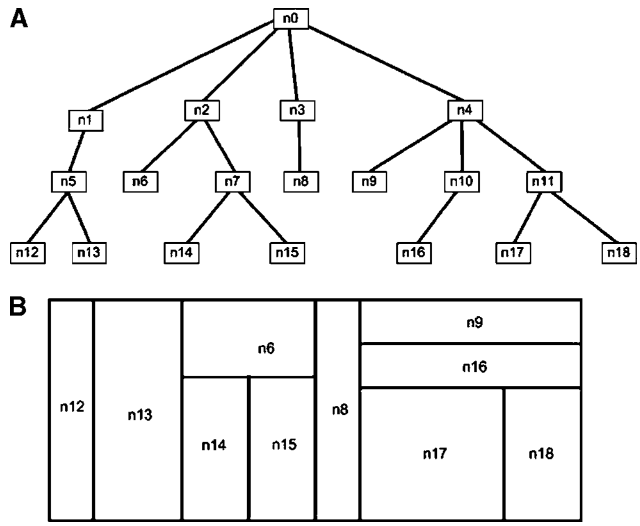

Figure 1 The node-link tree to the left clearly reveals the structure of the hierarchy. The TreeMap (an example of a space-filling visualization) to the right depicts the relative size of nodes but does not readily display the structure of the hierarchy.

into blocks of nested entities, with each block representing a node in the hierarchy. The hierarchical structure is usually created by nesting children nodes within their parents. As a result, boundaries between nodes are not always easily distinguishable and the structure of the hierarchy is not explicit.

An example of a space-filling technique is the TreeMap.[1] 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[2]). 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.[3] In Figure 1 parent nodes are not visible due to the nesting of components.

Several variations of the TreeMap have been developed.[4–6] One particular variation of the TreeMap relevant to the work here is the CushionMap.[7] The CushionMap was developed to make the hierarchy represented by a TreeMap visually more explicit. CushionMap uses shading algorithms to give a 2½-D impression, thereby facilitating the perception of node boundaries and making the structure more explicit. However, the claim that shading assists in seeing the hierarchical structure more discernable 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. Two experiments were conducted to test our hypotheses and the results of our investigation are reported here.

Related work

From the literature in information visualization, three particular evaluations of the TreeMap are relevant to our discussion. Turo and Johnson[8] 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, that is, 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.*[9] 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 radically 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' performed equally well with the TreeMap as with the Sunburst. 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[9] 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.

Kobsa[10] performed a comparative experiment with five well-known hierarchy visualization systems two of which were the TreeMap and CushionMap. Subjects performed structure-based and file attribute tasks. Kobsa's study did not place emphasis on whether shading facilitated interaction and task performance. The experiment was not designed with the hypotheses concerned with in this paper. The primary question in Kobsa's study relates to whether file management tasks can be performed better with hierarchy visualization tools other than Windows Explorer™. While subjects did not perform as well on the CushionMap as on the TreeMap, the results showed that Windows Explorer™ facilitated better most file management tasks over all the other tools.[10] Several reasons could be attributed to subjects poor performance with the CushionMap in comparison to the TreeMap. One explanation could be attributed to the lack of control with the interaction methods used in both visualizations. TreeMap was displayed using the TreeMap tool developed at the Human–Computer Interaction Lab at the University of Maryland whereas the CushionMap was tested using SequioView.[11] The difference in modes of interaction could have attributed to the differences. Another explanation could be due to the lack of control on the types of color mappings used for both hierarchies. The degree of control was not essential in Kobsa's study as mentioned earlier the experiment was not evaluating the effect of shading on the hierarchies.

The CushionMap system was designed to exploit the use of shading to make the hierarchical structure more explicit in the TreeMap.[7] The focus of our experiments has been to empirically test the effect of shading (such as that available for 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.

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.[12] In particular, there is evidence that such information is processed pre-attentively. A study by Enns and Rensink[13] 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 (3D) objects. The task consisted of locating single target items among one, six 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 *3D* objects.

Sun and Perona[8] extended the work of Enns and Rensink[13] by investigating the pre-attentive perception of elementary 3D shapes. To determine whether shading was more important than internal line crossings (which can contribute to determining the shape of a 3D 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[14] 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.

Although shading information can be processed pre-attentively, the human visual system makes certain assumptions about the direction of the light source. These effects, first published by Ramachandran,[15] are illustrated in Figure 4 below. In this figure, the shaded discs in the upper right and lower left appear as convex items. Their shadowed tops are the result of a perceptual bias that light is shining from above. The disc in the lower right appears concave, as the shadow resulting in its lower portion results from the imaginary-protruding surface. The effect is amplified when a top-shadowed disc is superimposed on a bottom-shadowed disc, giving the effect of a crater on a bump (upper left object in Figure 4).

The evidence of this bias is further supported by results from a recent physiological study by Hanazawa and Komatsu.[16] They suggest that neurons in a visual area of the brain (visual area V4) are adapted to the inclination where the light source is positioned somewhere above the viewer and the object or scene being viewed.

Beyond pre-attentive processing, shading information is also critical in structural object recognition. From a high-level view, object recognition is accomplished in a

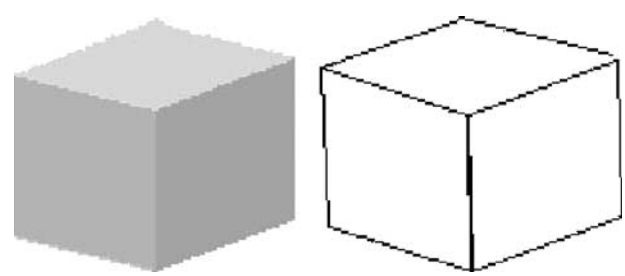

Figure 3 Sample targets used in the experiment by Sun and Perona.[8] A feature in common between the shaded item and the line drawing is the embedded Y-junction that assists in determining the shape of the object.

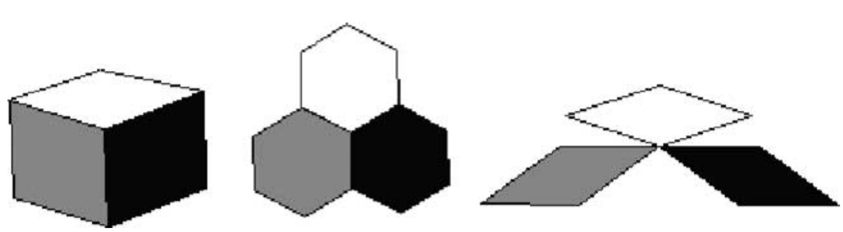

Figure 2 Targets used in Enns and Rensink's experiments.[13] The target to the left corresponded to projections of simple blocks under various lighting conditions. The pattern on the left was perceived faster than the 2D patterns on the right.

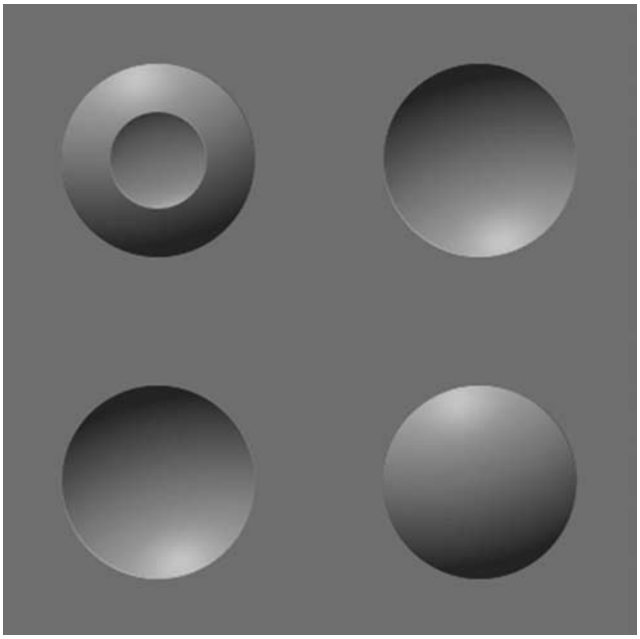

Figure 4 Our visual system makes assumptions about the directions of light source. Adapted from.[15]

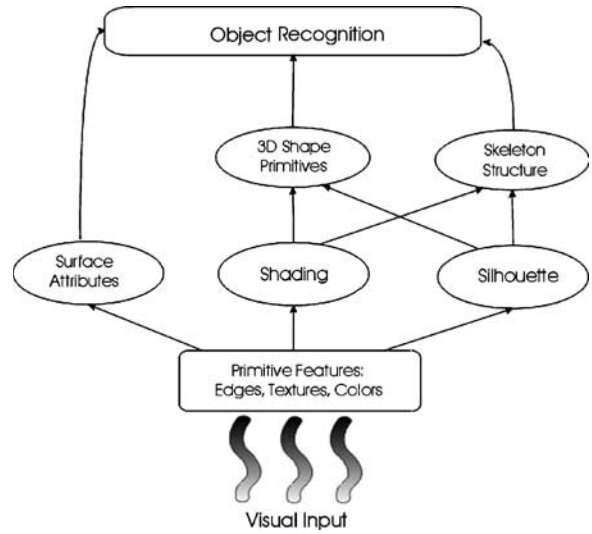

Figure 5 Theories of structural object perception propose a series of processing stages, culminating in object recognition. Shading plays a critical role in shape and structure extraction.

series of stages (Figure 5). 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[15] or 'blobs'.[24] At the same time it also facilitates the creation of a structural skeleton, containing information about how the components are interconnected[17,18] Ultimately, all of the information is combined which leads to object recognition.

Irani and Ware[19] 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 6). Their results show that using 3D shaded nodes and links resulted in more accurate sub-structure identification (11.4 *vs* 21% errors) and shorter times (4.1 *vs* 5.2 s). 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.

Furnas and Zhang[20] use shading information to depict scale levels in zoomable user interfaces. In their design, shading facilitates perceiving structures in hierarchies presented with nested boxes. Several shading tones are used to effectively show multiple zoom levels.

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

Figure 6 Diagrams used for comparing shading *vs* no-shading in explicit structures. Adapted from.[19]

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.

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*.[21] 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*.[22] 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 experiments. The first experiment compared CushionMap to the TreeMap on a set of generic tasks with respect to structural and size information. In the second experiment, a sub-structure identification task was designed to compare the two visualizations with respect to their ability of expressing structures.

Visualization tool

The experiment was conducted using the SequoiaView (version 1. 3) visualization tool, an application developed by van Wijk *et al.*[7] at the Department of Mathematics and Computing Science of the Technische Universiteit Eindhoven.[11] In SequoiaView users can view file hierarchies using both the TreeMap and CushionMap views (Figure 7). The tool is equipped with multiple configuration options including the mapping of file types onto color. We controlled all the options available in the tool and only created the two conditions: the presence and the absence of shading. All methods of interaction, that is, drilling down and up from hierarchies were similar. We used the default mappings of file extensions onto color and only used the following file types: '.bmp', '.jpeg', '.mp3', '.exe' and '.dll'. Similar mappings were used for both the CushionMap as for the TreeMap. For the experiments, 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.

Non-shaded treemaps have limited capability of showing structures. This is achieved by providing the entire path of files when the user hovers over a cell or by showing a border around a child node in yellow and its parent in red. For the experiments in our study we excluded these options to avoid any confounding effects with shading. SequoiaView is designed such that by moving the mouse over a node, displays a hint giving the absolute path of the file. We asked our participants to avoid using this feature. 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 sub-directory opens it so that the entire display is covered by the sub-directory. Browsing and navigating can be achieved by using the 'up-one-level' button but this was not offered as an option to the users. Use of the 'up-one-level' button was not permitted to better control the user's navigation throughout the experiment. Users were also asked to not modify any of the settings but instead to focus on completing the assigned tasks.

Experiment-1

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 sub-directories 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 sub-directories.

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).

Method

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.

Materials Participants performed the experiment on a 17 inch 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. The hierarchies used in this experiment contained 120 sub-directories and 2300 files. A sample hierarchy used in this experiment is shown in Figure 7. The resolution of the screen was set to 1024×768.

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 SequioaView. 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

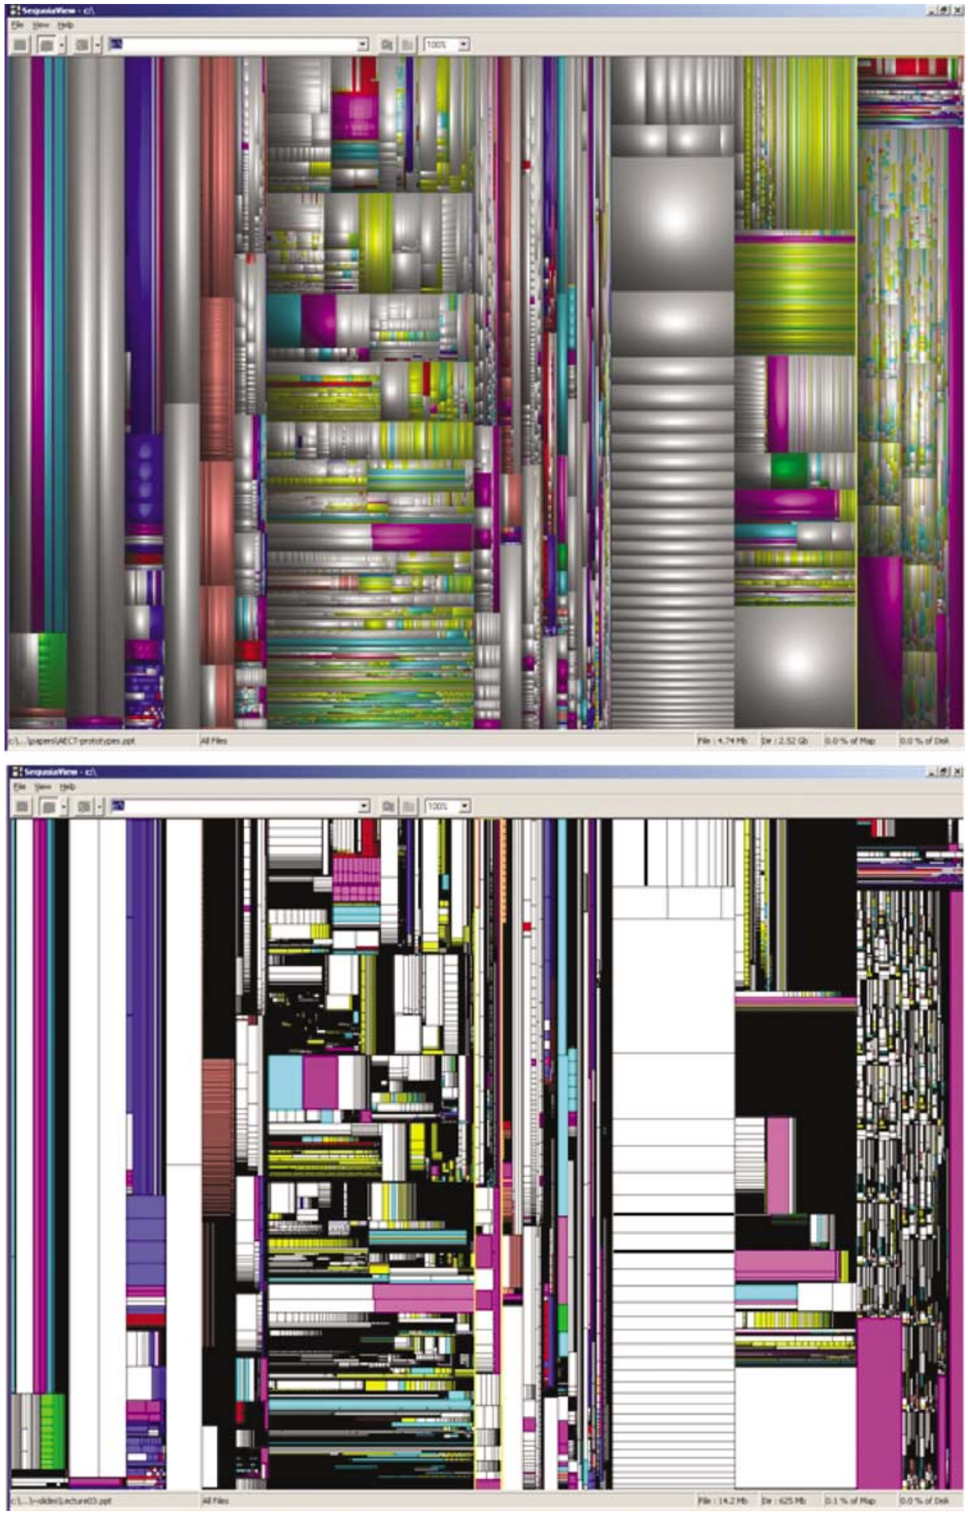

Figure 7 TreeMap and CushionMap representation of the local C: drive from SequoiaView.

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 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 s) 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 sub-directories in a given directory.
- Count the number of files in a given sub-directory.
- 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 sub-directory, 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.

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 (0–45 s) 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 $\times$ 2response 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 data sets, whereas average success rates were far from normal.

Following the methodology employed in an earlier and related study,[9] any task that was unsuccessfully completed or a timeout was excluded when calculating the average completion time. As a result, in the data set of completion times for structure-based tasks, two individuals had an average response of zeros 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 data set had only 18 pairs of responses.

The results are group by structure- or size-based tasks. Within each group subjects performed similarly therefore the analysis of the results is provided based on these two groups of tasks.

Effect of treatment and hierarchy order Of the subjects, 10 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, 10 subjects were randomly allocated to 'H1, then H2', the remaining 10 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- 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).

Analysis of structure-based tasks Figure 8 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 note 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 successfully 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.

Analysis of size-based tasks To test Hypothesis 2, we used a paired *t*-test to compare the mean average completion times of the four size-based tasks. Whereas the observed mean average completion times for the 20 subjects was 2.3 s 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.

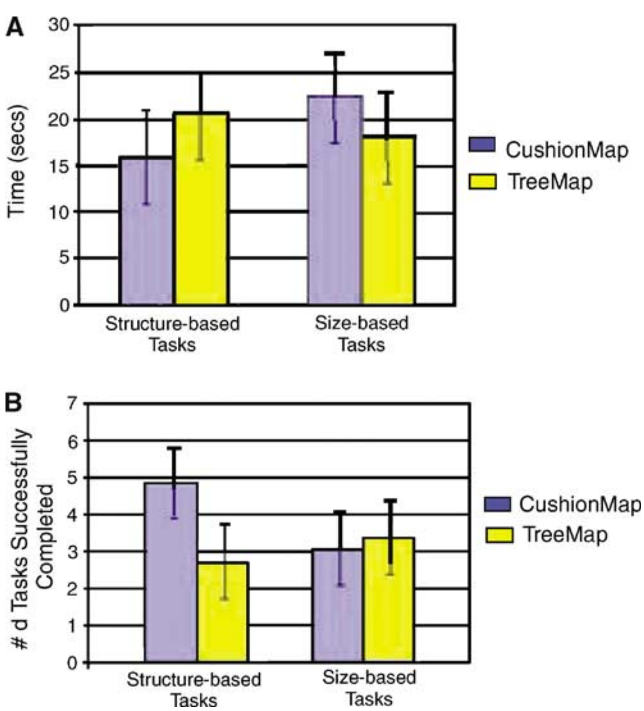

Figure 8 (a) Average completion times in seconds (for correct responses only) and (b) average number of tasks successfully completed.

Table 1 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.

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 2. 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 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.

Experiment-2

The results from the previous experiment do not conclusively determine whether shaded space-filling visualizations can reveal better the underlying hierarchical structure than their non-shaded counterparts. The results suggest that tasks requiring structural information is facilitated by the CushionMap but do not say anything about the user's ability to visually parse hierarchical structures. Therefore, a second experiment was designed to compare the effect of shading in conveying the sub-structural relationship of elements in space-filling hierarchies. In this experiment, the users performed tasks that required them to identify a sub-structure in a given directory, using both CushionMap (CM) and TreeMap (TM). Subjects were required to perform a sub-structure identification task, similar to the one described in Irani and Ware.[19] In order to perform the task successfully, theory suggests that subjects will store the target sub-structure with a set of structural description rules in

Table 1 Statistical significance of TM *vs* CM on structure- and size-based tasks

	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

visual working memory. In this task, subjects are presented with a hierarchical sub-structure for 15 s. After this exposure period, they are required to locate the sub-structure within several hierarchies. Two types of sub-structures were constructed (simple and complex).

Figure 9 shows an example of a 'simple hierarchy' (highlighted in red color). During the experiment, the subjects were given an unlabelled hierarchy (not the entire directory) and they were asked to find the location of the sub-structure in the directory. The simple hierarchy has a maximum depth of three levels, whereas the complex hierarchy has a maximum depth of four levels (Figure 12). Figure 9 shows the hierarchical representation of a directory. Figures 10 and 11 show the representation of the same directory using CM and TM. The blue colored block in Figures 10 and 11 represent the highlighted sub-structure presented in Figure 9.

Figure 12 shows an example of a 'complex hierarchy' (highlighted in red color) task. Figure 13 and 14 show the representation of the same directory using CM and TM. The blue colored sub-structure in Figure 13 and 14 represent the highlighted sub-structure of Figure 12.

In this experiment (Figure 12–14), 32 directories were created (16 were designed to do the test on the simple sub-structure and the other 16 were constructed for the complex sub-structure). In order to avoid learning effects, these directories contained the same general structure but with a different ordering and arrangement of sub-structure. The total number of files and folders were the same in all hierarchies.

In this experiment, the participants were equally divided into two groups. Half of the participants started the experiment with the TreeMap and the other half started on the CushionMap. The set of trials consisted of {CM-Simple, TM-Simple}, {CM-Complex, TM-Complex}, {TM-Simple, CM-Simple}, and {TM-Complex, CM-Complex}. After completing the tasks on one tool the participants took a brief break and switched onto the other tool.

Based on results from earlier studies on perception (as discussed in the section on visualization tool) and the experimental results described in previous section, the following effects were anticipated:

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

Hypothesis 2: The *shading* condition (CM) will result in higher performance on *tasks related to identifying complex hierarchies* than the *no-shading* condition (TM).

Table 2 **Subjective preferences averaged across subjects for both types of representations. 5 represents 'strongly agree' and 1 'strongly disagree'**

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

Method

Subjects A total of 20 undergraduate students from a local University participated in the experiment. These participants 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. However, none of the participants had

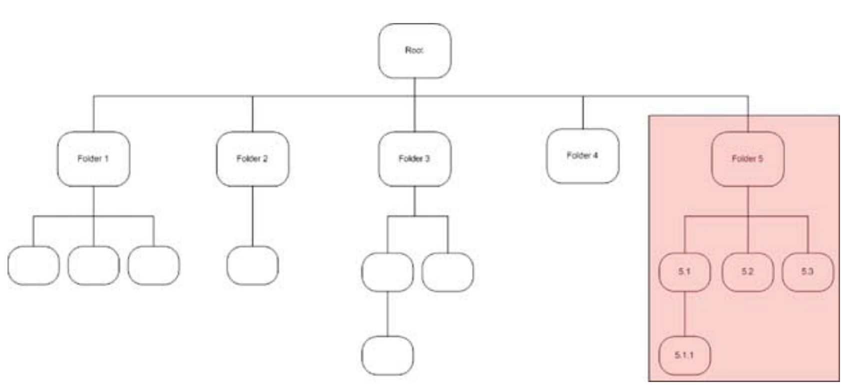

Figure 9 Hierarchical representation of a '*simple hierarchy*'.

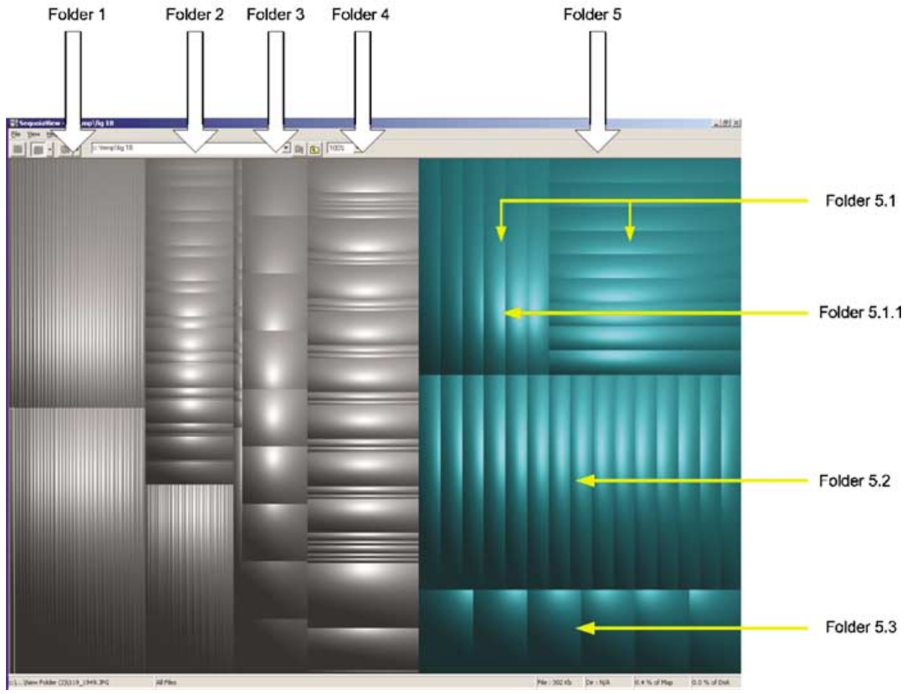

Figure 10 CushionMap containing the sub-structure in Figure 9.

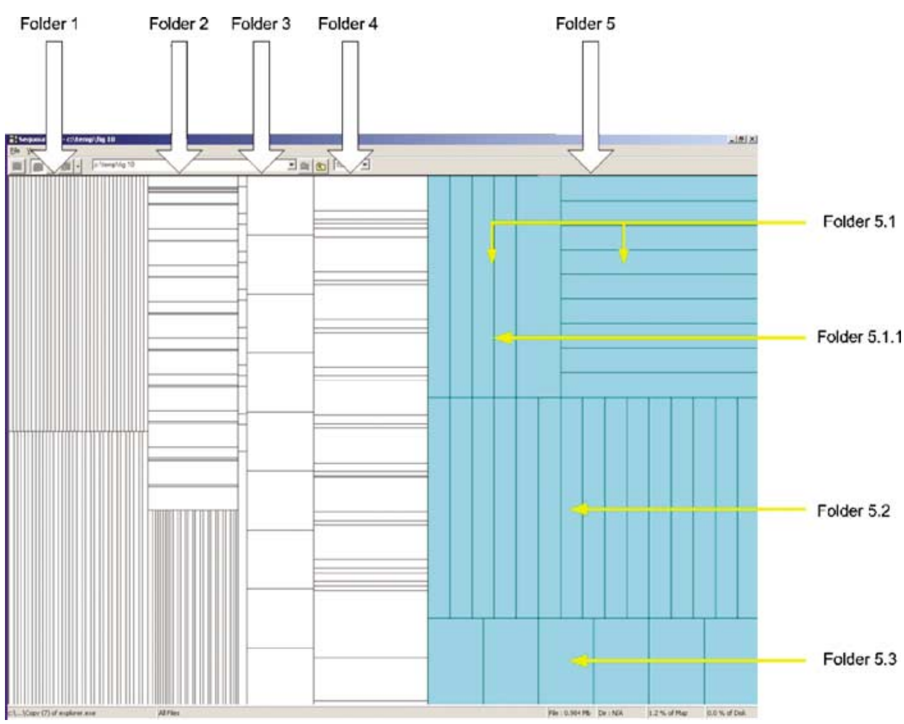

Figure 11 TreeMap containing the sub-structure in Figure 9.

previous experience using SequoiaView and the TreeMap or CushionMap visualization tools.

Materials Participants performed the tasks using SequoiaView on a 15.4 inch TFT with TruBrite LCD monitor, running over WinXP. The resolution of the screen was set to 1024×768.

Procedure Before the start of the experiment, a sample file directory was selected and shown to the subjects.

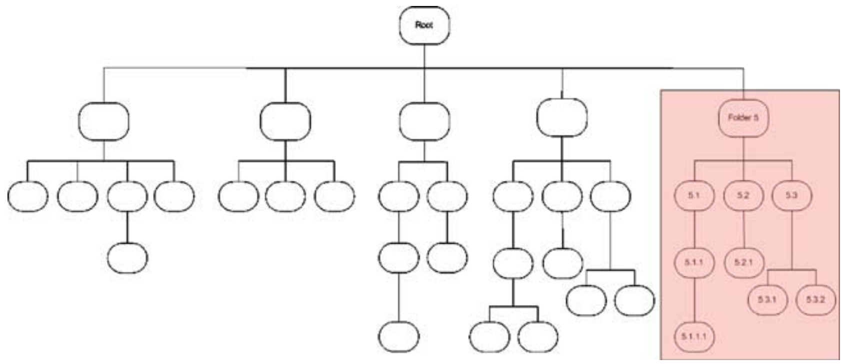

Figure 12 Hierarchical representation of a '*complex hierarchy*'.

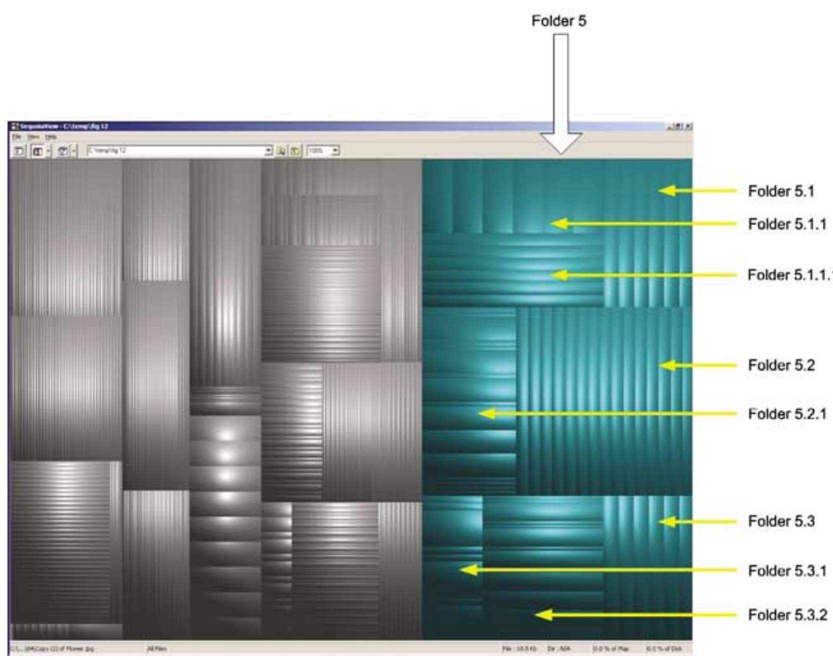

Figure 13 CushionMap containing the sub-structure in Figure 12.

This sample directory was first displayed using Windows Explorer™. The subjects were also explained the parent–child (hierarchical) relationship between the files and folders in the directory. Then the same directory was displayed using one of the space-filling visualizations. The subjects were also explained how the files and folders were arranged in the space-filling visualization. This was done by showing both the Windows Explorer™ and the space-filling visualizations, in parallel. When the subjects indicated that they were comfortable with the space-filling system and they understood how related items were presented, the other space-filling visualization was shown. The order of presentation of these space-filling visualizations varied between the subjects (e.g. group-1 received training with CM first, where as group-2 received training with TM first).

The experiment began when participants felt that they were comfortable with both the tools. After the description session, each participant performed the sub-structure identification task using both the tools, on both types of hierarchies. The participants were first shown the sub-structure and then were presented with a set of hierarchies half of which contained the sub-structure. The participant's task involved identifying whether the hierarchy is present in the shown (displayed) directory. If the hierarchy is present in the displayed directory, then they were asked to locate the position of the hierarchy in that directory. We measured participants' performance as a success (if they correctly identified the presence or absence of the sub-structure in the hierarchy) or a failure (if they incorrectly identified the presence or absence of the sub-structure in the hierarchy). The time to execute the task was also recorded. An upper limit of 150 s was set

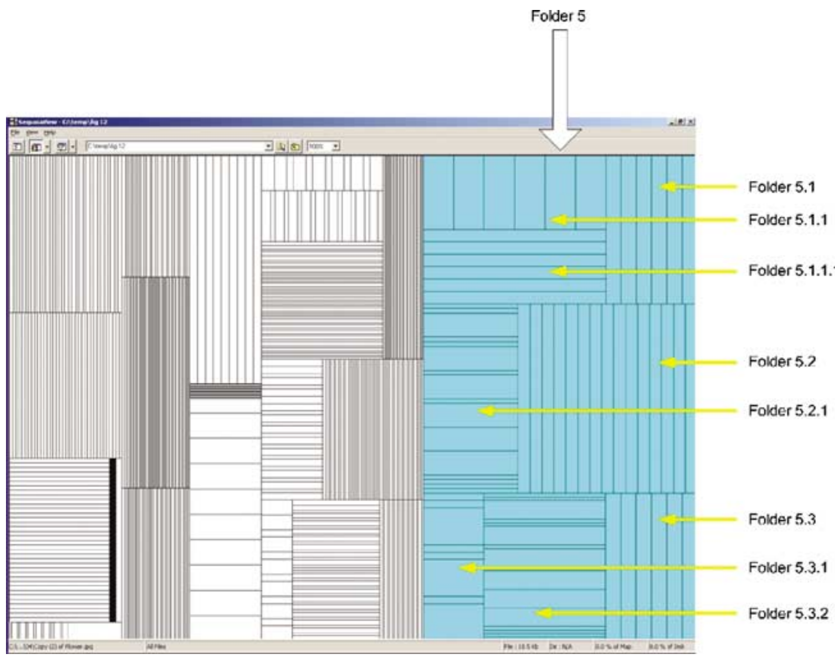

Figure 14 TreeMap containing the sub-structure in Figure 12.

for identifying the presence or absence of the sub-structure, after which a timeout was recorded.

Results and discussion

To test the two hypotheses stated in the beginning of this section, we measured subjects' performance with respect to two variables: time until completion (0–150 s) and successful/unsuccessful completion (1/0). For both hypotheses, we recorded the average response for both of these variables. These measurements were taken for each of the 20 subjects each containing 20 pairs (CM, TM). For the dichotomous response variable, timeouts were classified as failures (0). Following the methodology employed in an earlier and related study,[9] any task that was unsuccessfully completed or a timeout was excluded when calculating the average completion time. Average completion times were consistent with the normality assumption in all data sets, whereas average success rates were far from normal.

Of the subjects, 10 were randomly allocated to 'TM, then CM' and the remaining 10 were allocated to 'CM, then TM' (these two groups remained the same in all stages of the experiment). To justify our pooling of the subjects, we first verified that treatment order did not have 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 the sub-structure identification task, we did not detect a significant difference between these two groups with respect to the variable CM−TM ($F_{1,38} = 0.05$, $P = 0.825$). A non-parametric Kruskal–Wallis test on the average success rate data detected no

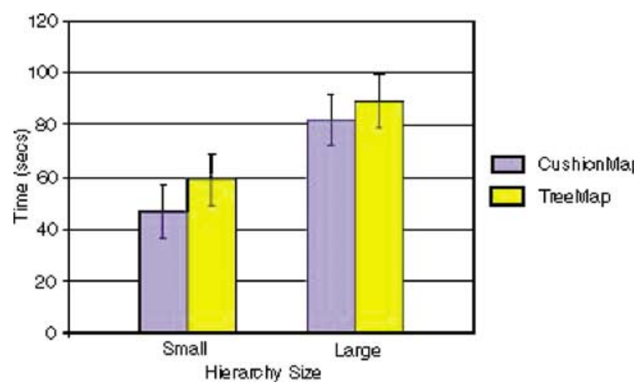

Figure 15 Average completion times for small and large hierarchies with both methods.

significant effects of treatment order (CM or TM first) on the sub-structure identification task ($\chi^2(1) = 0.003$, $P = 0.955$).

The average performance results are summarized in Figure 15. The time to locate target data were analyzed by means of a 2×2 (Type of Method $\times$ Sub-structure Size) one-way analysis of variance (ANOVA), with both Type of Method (CM *vs* TM) and Sub-structure Size (simple *vs* complex) serving as repeated measures. An alpha level of 0.05 was used for all statistical tests. Type of Method was found to be significant ($F(1, 159) = 7.88$, $P = 0.006$) with the CushionMap method group's mean task time (64.33 s) being faster than the TreeMap method group's mean (73.95 s). The main effect for sub-structure was statistically significant ($F(1, 159) = 84.05$, $P < 0.001$).

Table 3 Average error rates for small and large sub-structures with both methods

	Small (%)	*Large (%)*
CushionMap	12.50	19.38
TreeMap	23.75	28.13

However, no significant interaction effect was found between Type of Method and Sub-structure Size, $F(1, 159) = 0.550$, $P = 0.459$.

The average error rate is summarized in Table 3 below. A non-parametric Wilcoxon rank-sum test does not show a significant difference in average error rates for the CushionMap over the TreeMap. Overall the average error rate with the CushioMap is 18.13 *vs* 25.65% for the TreeMap. The overall average error rate is the average for all subjects across both conditions.

The results suggest that participants were faster roughly 1.5 times faster with the CuhionMap than with the TreeMap for visually parsing hierarchical structures. The results do not show significant difference in error rates. This suggests that while participant were careful in providing the correct response, it took them longer to complete the task with the tool in the absence of shading. Overall, these results provide some evidence that shading can facilitate the viewing of sub-structures in space-filling representations.

Conclusions

In this paper, we have presented the results of two experiments 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. In addition, we postulated that shading will facilitate visual segmentation of elements sharing common borders and thereby facilitate visual hierarchical parsing. By enhancing the user's ability to visually parse elements in the space-filling representation, we anticipated that sub-structure identification would be facilitated by the shaded representation.

Our results confirmed the first hypothesis. Users were faster and more accurate in completing directory management tasks with the shaded hierarchies. On the other hand, we did not obtain any conclusive results on the unfavorable effect of shading for size-based tasks, which is not in support of our second hypothesis. This warrants a more discerning follow-up experiment where the effect of shading on size judgments is better controlled. The results of the second experiment reveal that participants were faster in performing a sub-structure identification task when shading is available and are equally accurate with or without shading information. To successfully perform the second experiment participants are required to perform a mental operation of internally translating a node-link hierarchical sub-structure to a space-filling structure. It is remarkable to notice that such transformations were possible within relatively short time limits and with relatively few errors. Subjective responses suggest that the participants preferred interacting with the system when shading was available. Altogether, the results of both experiments support and add to the previous body of literature on the nature and benefits of shading. This affirms the intuition of the designers of the CushionMap.

Although not conclusive, our data suggest 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.

Based on the results of both experiments, the degree to which each set of tasks (structure-based or size-based) is supported by either type of visualization is shown in Figure 16. 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.

One reason why shading may facilitate structural parsing in space-filling representation could be due to the contrast and delineation between adjacent elements created by shading effects. In space-filling visualizations, the nodes share common borders. To highlight the segmentation in these regions shading seems to be an effective method. Beyond providing good segmentation cues, shading also creates sharp contrast between elements in the scene. Such contrasts are particularly useful for space-filling representations or any other types of visualizations that require adjoining elements for conserving space or simply for depicting underlying data or structures.

The following design suggestions can be adopted by practitioners:

- Space-filling representations should use shading information to depict structural components.
- Visualizations in general that rely on common borders between elements and that require attention to specific elements could take advantage of shaded scenes for assisting the visual system in segmenting elements in the scene.

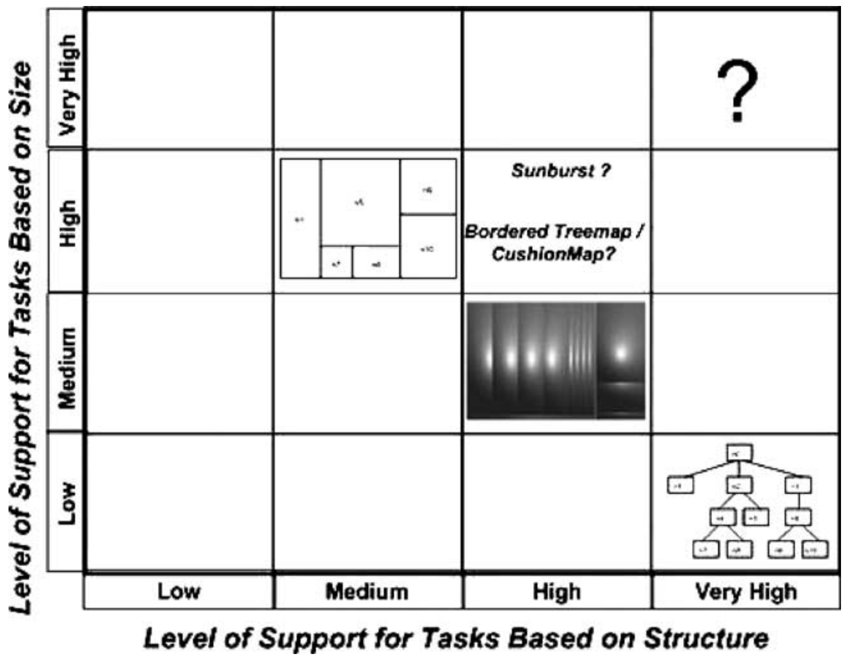

Figure 16 Space-filling techniques can be approximately positioned according to the degree to which they support structure-based and size-based tasks. Note that the top-right corner is empty – none of these systems fully support both types of tasks.

- In most interfaces, shading properties create visual effects that lead users to naturally click on items (as indicated by subjects in our experiments). Therefore, designers need to carefully balance the amount of shading visualizations. For instance, single item and disjoint elements need not be shaded as these could allude to a button or other clickable item.

The investigation contributes to the growing body of literature in information visualization related to evaluation methods and techniques. Such empirical studies could be beneficial to designers in building systems which will be effectively utilized by real-users[9,22,23]

Acknowledgments

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

References

1 Shneiderman B. Tree visualization with TreeMaps: a 2d space-filling approach. *ACM Transactions on Graphics* 1990; **11**: 92–99.
2 SmartMoney. www.SmartMoney.com.
3 Eick SG. Visualization and interaction techniques. *CHI'97 Tutorial Notes on Information Visualization*. ACM SIGCHI: Atlanta, Georgia, 1997.
4 Bederson BB, Shneiderman B, Wattenberg M. Ordered and quantum treemaps: making effective use of 2D space to display hierarchies. *ACM Transactions on Graphics* 2002; **21**: 833–854.
5 Bruls DM, Huizing C, van Wijk JJ. Squarified Treemaps. *Proceedings of the Joint Eurographics and IEEE TCVG Symposium on Visualization*. Springer: Vienna, 2000; 33–42.
6 Vernier F, Nigay L. Modifiable Treemaps containing variable-shaped units. *IEEE Visualization* 2000 (Late Breaking Hot Topics); 4.
7 Wijk JJ.van, Wetering H.van de. Cushion TreeMaps. *Proceedings of the IEEE Symposium on Information Visualization*. IEEE Computer Society Press: Silver Spring, MD, 1999; 73–78.
8 Sun JY, Perona P. Preattentive perception of elementary three-dimensional shapes. *Vision Research* 1996; **36**: 2515–2529.
9 Stasko JT, Catrambone R, Guzdial M, McDonald K. An evaluation of space-filling information visualizations for depicting hierarchical structures. *International Journal of Human-Computer Studies* 2000; **53**: 663–694.
10 Kobsa A. User experiments with tree visualization systems. *IEEE Symposium on Information Visualization* 2004; 9–16.
11 SequoiaView. Technische Universiteit Eindhoven, available at www.win.tue.nl/sequoiaview.
12 Freeman WT, Adelson EH, Pentland AP. Shape-from-shading analysis with shadelets and bumplets. *Association for Research in Vision and Ophthalmology* 1990; **31**: 410.
13 Enns JT, Rensink RA. Influence of scene-based properties on visual search. *Science* 1990; **247**: 721–723.
14 Braun J. Shape-from-shading is independent of visual attention and may be a 'texton'. *Spatial Vision* 1993; **7**: 311–322.
15 Ramachandran VS. Perception of shape from shading. *Nature* 1988; **331**: 163–166.
16 Hanazawa A, Komatsu H. Influence of the direction of elemental luminance gradiennts on the responses of V4 cells to textured surfaces. *The Journal of Neuroscience* 2001; **21**: 4490–4497.
17 Biederman I, Gerhardstein PC. Recognizing depth-rotated objects: evidence and conditions for three dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance* 1993; **19**: 1162–1182.

18 Marr D. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W.H. Freeman, San Francisco, 1982.
19 Irani P, Ware C. Diagramming information structures using 3D perceptual primitives. *ACM Transactions on Computer Human-Interaction* 2003; **10**: 1–19.
20 Furnas GW, Zhang X. *Illusions of Infinity: Feedback for infinite worlds, User Interface Software and Technology (UIST)*. ACM: New York, 237–238.
21 Carswell MC, Frankenberger S, Bernhard D. Graphing in depth: perspectives on the use of three-dimensional graphs to represent lower-dimensional data. *Behaviour and Information Technology* 1991; **10**: 459–474.
22 Zacks J, Levy E, Tversky B, Schiano DJ. Reading bar graphs: effects of depth cues and graphical context. *Journal of Experimental Psychology: Applied* 1998; **4**: 119–138.
23 Ware C. *Information Visualization: Perception for Design*, Morgan Kaufmann Interactive Technologies Series Morgan Kaufmann Publishers: San Francisco, 2004.
24 Biederman I. Recognition-by-components: a theory of human image understanding. *Psychological Review* 1987; **94**: 115–147.