

Evaluating Stereo and Motion Cues for Visualizing Information Nets in Three Dimensions

COLIN WARE and GLENN FRANCK
University of New Brunswick

This article concerns the benefits of presenting abstract data in 3D. Two experiments show that motion cues combined with stereo viewing can substantially increase the size of the graph that can be perceived. The first experiment was designed to provide quantitative measurements of how much more (or less) can be understood in 3D than in 2D. The 3D display used was configured so that the image on the monitor was coupled to the user's actual eye positions (and it was updated in real-time as the user moved) as well as being in stereo. Thus the effect was like a local "virtual reality" display located in the vicinity of the computer monitor. The results from this study show that head-coupled stereo viewing can increase the size of an abstract graph that can be understood by a factor of three; using stereo alone provided an increase by a factor of 1.6 and head coupling alone produced an increase by a factor of 2.2. The second experiment examined a variety of motion cues provided by head-coupled perspective (as in virtual reality displays), hand-guided motion and automatic rotation, respectively, both with and without stereo in each case. The results show that structured 3D motion and stereo viewing both help in understanding, but that the kind of motion is not particularly important; all improve performance, and all are more significant than stereo cues. These results provide strong reasons for using advanced 3D graphics for interacting with a large variety of information structures.

Categories and Subject Descriptions: H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—*evaluation/methodology*; I.36 [Computer Graphics]: Methodology and Techniques—*ergonomics*

General Terms: Graphics, Human Factors.

Additional Key Words and Phrases: Head-coupled display, information visualization, network visualization, stereopsis, virtual reality

1. INTRODUCTION

A useful and interesting method for examining three-dimensional structures is to couple a perspective stereoscopic view of a 3D scene to the user's eye positions and update the view in real-time as the user moves. The key

Authors' address: Faculty of Computer Science, University of New Brunswick, P.O. Box 4400, Fredericton, NB E3B 5A3, Canada; email: (cware@UNB.ca).

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1996 ACM 0730-0301/96/0400-0121 \$03.50




Figure 1. Head-coupled stereo views resulting in a virtual 3D view of the network placed in the vicinity of the monitor.

elements of this are a high resolution monitor capable of running at a high frame rate, stereo glasses, and some method for tracking the user's head position (see Figure 1) [Deering 1992]. The positions of the user's two eyes are computed from the head position, and separate images are generated showing the correct perspective view of a set of virtual objects somewhere in the vicinity of the monitor screen. The result is a localized *Virtual Reality* (VR) environment that has a number of advantages over immersive virtual reality [Arthur et al. 1993; Deering 1992], not the least of these being that the everyday workspace of desk, filing cabinet, coworkers, and coffee mug are not excluded. In our previous work we have called this *Fish Tank VR* to characterize its localized nature and distinguish it from full immersive VR [Ware et al. 1993].

The advantages of Fish Tank VR representations of conventional objects are fairly obvious. For example, an industrial designer can see her design of a telephone handset as a truly three-dimensional object, or a trainee mechanic can see the engine part and how it is inserted as if he were looking at the real thing. But what about abstract information, such as database schemas, networks of human relationships, the structure of object-oriented code, or hypertext links? It is by no means clear that 3D diagrams offer any advantages over 2D ones in representing this kind of information. The present paper provides some strong evidence that 3D diagrams can offer significant advantages under the right viewing conditions. But before presenting this evidence we review some of the evidence as to whether advanced 3D viewing techniques should be used for viewing information networks. We discuss three types of evidence: perceptual, anecdotal, and experimental.

1.1 How Many Dimensions Are There in Visual Space?

In the following discussion, we use the term visual space in the everyday sense of perceived area or volume, ignoring the dimensions of, for example, color and texture. Using a naive view, moving from a 2D to a 3D display should vastly increase the amount of information that can be represented. Consider a 1000×1000 computer display. On a line we can perceive 1000 distinct pixels; on the plane we can perceive 1000^2 distinct pixels. Extending this logic we should be able to display 1000^3 distinct voxels in a 3D volume. These relations can be succinctly expressed by the following equations

$$I_{2D} = I_{1D}^2, \quad I_{3D} = I_{2D}^{3/2},$$

where I_{nD} represents the information that can be perceived given an n -dimensional display.

Clearly, there is a flaw in this logic; in general we do not perceive volumes of data; we do not perceive details of the insides of solids, only the layout of surfaces in space. This places an upper bound on the amount of information that can be represented in a 3D space.

Using a different naive view, we can argue that because stereo 3D viewing is based on input from two eyes this can at most double the amount of information perceivable, and we know that such effects as binocular rivalry mean that truly independent images from the two eyes cannot be perceived [Lack 1974]. In fact, it is only possible to extract depth information from highly correlated information presented to the two eyes. This alternative pessimistic view can be expressed by the equation

$$I_{3D} = C * I_{2D},$$

where C is some constant < 2.0 ; probably much closer to one than to two, due to this correlation.

This second view suggests that moving from 2D to 3D will yield only a small benefit in visualization, at least for the understanding of abstract data. If this view is correct, it will be hardly worthwhile using 3D representations of abstract multidimensional networks of information; a 2D view is likely to be equally effective and much simpler to produce.

There is a third view based on the ecological argument that because we have evolved in a 3D world, information presented in 3D will be processed more easily by the visual system. Networks of information do not have an inherent dimensionality in the geometric sense, but if our brains prefer 3D layouts, then a 3D layout may be more effective in conveying the information [Purcell and Stewart 1991]. We should also allow that the brain is capable of integrating information from the sequence of views of the world that are obtained as we move about and that this considerably increases the amount of information obtainable, although as with stereopsis these views must be highly correlated.

A truly 3D display is not simply a perspective picture, or a stereo perspective picture. It requires that the image be accurate from the actual viewpoint of the observer, especially if the observer is moving. A large part of space perception comes from the relative movement of objects; in fact, this seems to be considerably more important than stereopsis in helping us interpret spatial layout, whether it be from head motion, resulting in motion parallax [Gibson et al. 1959], or from object rotation giving the kinetic depth effect [Braunstein 1976; Bruno and Cutting 1988]. This movement allows the brain to integrate spatial information over time, and whereas a stereo display only gives two views to help understand a scene, a scene in which there is relative movement of the head and objects provides a whole continuum.

1.2 Anecdotal Evidence

There is growing evidence that representing diagrams in 3D can allow more complex information to be comprehended. An influential work has been the SemNet project [Fairchild et al. 1988]. This used a 3D representation to allow users to visualize large knowledge bases as nodes and arcs in a three-dimensional space. No attempt was made to quantify the advantages of this mode of viewing over 2D layouts, but much of the current interest in 3D abstract data visualization can be traced to this study.

The Cone Tree technique developed by Robertson et al. [1991] has the capacity to show considerably more tree-structured data than do more conventional 2D layouts. In a Cone Tree, all the children of a node are displayed around the base of a cone in 3D space, and this principle is applied recursively. The authors claim that as many as 1000 nodes may be displayable using Cone Trees without visual clutter—this is clearly more than could be contained in a 2D layout, although the Cone Trees require certain user manipulations to access some of the information. However, on the other side of the argument, there are 2D layout schemes that also claim to increase the amount of information displayable [Lamping and Rao 1994].

1.3 Empirical Evidence

Of particular relevance to the display of information networks is work that has shown that the number of errors in detecting paths through tree structures is substantially reduced if a 3D display method is used [Sollenberger and Milgram 1993; Ware et al. 1993]. Sollenberger and Milgram [1993; 1991] showed that both scene rotation and stereopsis helped reduce errors in a path tracing task. In their motion conditions the stimulus pattern rocked back and forth about a vertical axis (the perspective imagery was not coupled to head position). They found that motion was more valuable than stereopsis in reducing errors when subjects were required to discover to which of two tree roots a highlighted leaf node was attached. In a similar experiment, Ware et al. [1993] used a similar task with a head-coupled stereo display in which the perspective view was

coupled to the measured eye position of the observer. Although in Ware et al.'s display the motion was caused by head movement, the results were similar to those obtained previously by Sollenberger and Milgram. Because the geometric transformation of the image that may occur under head-coupling or scene rotation may be very similar, there is every reason to suppose that the same visual mechanisms are involved in both cases in constructing a 3D interpretation of the world.

A simple model for the integration of different depth cues is a weighted additive model [Bruno and Cutting 1988], according to which the brain computes a weighted sum of different information in determining the depth that will be perceived. Sollenberger and Milgram found their data suggested a superadditive model; however, the relationship between depth perception and percent correct in a path-tracing task was not made clear. Conversely, the data of Ware et al. is subadditive, although the same reservations about the link between the model and the data apply. Also, their experiment resulted in much lower error rates, and ceiling effects were undoubtedly present.

1.4 Motivation for New Studies

A number of unanswered questions are posed by the preceding studies. The first and most fundamental question is how much is gained by moving from a 2D to a 3D representation. Although a completely general answer to this question can never be expected because the answer must be task specific to some extent, any answer derived from an abstract task such as path tracing in networks of nodes and arcs will generalize to the large set of problems that can be represented in this way. There is also the more profound question related to the issue of how depth information from different cues is integrated in the brain. By measuring task performance with a series of graphs of different sizes under different viewing conditions, it should be possible to determine the function relating error rate to graph size for each condition. If these functions appear to belong to the same family, we can model their relationship.

A second question is whether the results of Sollenberger and Milgram and Ware et al. generalize to arbitrary graphs other than trees. This is important because tree layout is a relatively simple, well understood process and it is trivial to lay out a pair of trees in a plane so that they do not overlap. Hence the visualization problem posed by the previous studies could be solved easily without resorting to a 3D display. This is not the case for an arbitrary graph where the layout problem is more difficult and the advantages of 3D visualization may be more pronounced. Although much of the work on 2D layout of directed graphs has been directed at minimizing arc crossings, the algorithms are often complex [Eades and Xuemin 1989]. We hypothesize that 3D visualization will, to some extent, reduce the graph crossing problem because arcs will no longer appear in the plane of the screen.

A third question is what kind of motion is better for perceiving structure in information networks: (a) motion induced by perspective coupled to eye

position, (b) automatic rotation of the object, or (c) motion caused by linking the user's hand movements to the object. Sollenberger and Milgram, and Ware et al. showed advantages of motion in information perception, but the kind of motion was different in the two cases: scene rotation in the former case (causing the kinetic depth effect), and head motion in the latter (causing motion parallax). It should be noted that lateral motion of the viewpoint produces a very similar transformation to rotation of the scene about a vertical axis through the center of the scene. This leads to the hypothesis that both kinds of structured transformations may contribute to 3D space perception in similar ways to enhance network comprehension. However, it might be that the active search allowed by viewing modes (a) and (c) are critical to understanding because under these conditions the user has control over the view of the scene.

Experiment (1a) was designed to address the first and second of the preceding questions, and Experiment (1b) was designed to supplement (1a) which yielded only an incomplete answer. Experiment (2) was designed to address the third question.

2. EXPERIMENT 1a: HOW MUCH BETTER IS A 3D/STEREO DISPLAY FOR PATH TRACING IN A GRAPH?

The purpose of this experiment was to provide a quantitative estimate of the improvements in graph understanding offered by adding head coupling and stereo disparity cues to the display. We chose path tracing in a graph as the task because a graph is a very general data structure used to represent an information network. We chose error rates as the dependent variable because the previous studies described [Arthur et al. 1993; Sollenberger and Milgram 1991] have shown this to be highly sensitive to the mode of viewing. In addition, the ability to trace information paths accurately is fundamental to understanding information networks. The answer we were looking for was of the form "adding head-coupled stereo increases the size of a graph that can be understood by a factor of X." In order to make such an evaluation we designed an experiment with a range of graph sizes in the hope that the data would yield a family of curves (error rate as a function of graph size) to allow valid comparisons between viewing conditions. Two viewing conditions were used initially.

- (1) **2D:** No stereo, no rotation; the 3D graph was projected onto a 2D plane using an orthographic (parallel) projection by removing *Z* axis information, hence no depth information was presented.
- (2) **Stereo, head-coupled perspective:** The correct perspective view was generated for each eye position (continuously updated) using an apparatus as shown in Figures 1 and 2. The perspective was continuously updated based on the measured head position (and derived eye position) of the subject.




Fig. 2. Experimental setup

2.1 Hardware

The equipment used in this experiment consisted of StereoGraphics Corporation's CrystalEyes 3D LCD shutter glasses to provide the stereo, with Logitech Corporation's ultrasonic head-tracking built into the frame of the glasses. The computer used to produce the display was a Silicon Graphics Crimson VGX. In stereo mode, this system uses the top and bottom half of the frame buffer for the left and right images, respectively. These images are vertically expanded by repeating lines at display time, which results in half the vertical resolution and also causes lines to be drawn double thickness in a vertical direction. Another feature of stereo mode is that the monitor runs at 120Hz with each eye receiving a 60Hz update rate because of the synchronized shutter glasses. Figure 2 shows Glenn Franck wearing the CrystalEyes VR 3D LCD glasses. The triangle on top of the monitor is the reference sound source for the microphones embedded in the frames of the glasses.

2.2 Experimental Procedure

The subject's task was to decide whether there was a path of length two connecting two nodes that were highlighted in a randomly laid out graph. On each experimental trial there was either a path of length two (with an intervening node between the two highlighted nodes) or no path, with a 50% probability of each occurring.

The computer generated a random 3D graph consisting of different numbers of nodes, arranged in a simulated 17 cm³ volume. An example is shown in Figure 3. The nodes were divided into three equal-size groups.

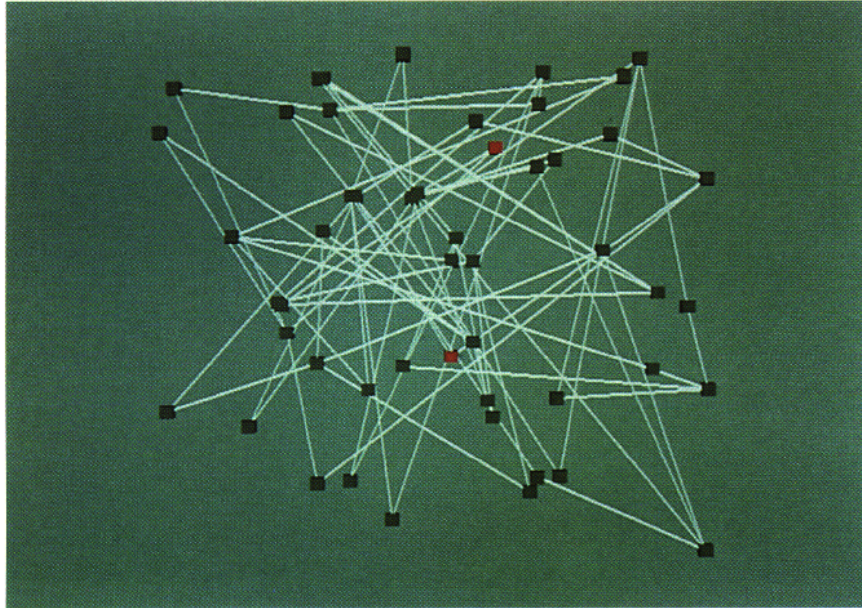


Fig. 3. Graph containing 78 nodes and 104 arcs as it was displayed on the monitor in the 2D condition. Two of the nodes are colored red. The subject's task was to determine if there was a path of length 2 between the two red nodes.

Two of these groups were leaf nodes, and the third was a group of intermediate nodes. Each node in one of the leaf groups was connected via arcs to exactly two different nodes in the intermediate group. For n nodes, this produced a total of $(\frac{2}{3} * n)$ connecting arcs. All nodes were placed randomly within the working volume.

Unhighlighted nodes were drawn in a dark gray color, and the highlighted ones were bright red; lighting was applied to all nodes to emphasize the three-dimensional effect. The size of each node was set to be 0.4 cm on each side. The arcs were white and drawn with double-width (two pixel) lines. Because the arcs were drawn as lines and not polygons, perspective sizing did not affect the apparent size of the lines. Due to the artifact described in Section 2.1, the vertical thickness (1.1 mm) was double the horizontal thickness (0.53 mm). The background was a flat light gray so as to minimize contrast and the ghosting effects associated with stereo. The screen image is shown in Figure 3. A stereo pair is illustrated in Figure 4 where the stimulus patterns consisted of a set of nodes and arcs laid out in a volume and rendered according to a variety of different methods, either with or without stereo and with or without head-coupled perspective. The images are set up for crossed eye viewing. If they are inverted they can be viewed either with a stereo viewing device or with the eyes parallel (something that requires considerable practice).




Fig. 4. This stereo pair is a line drawing simulation of the screen display with 51 nodes and 68 arcs.

Five graph sizes were used for each of the two viewing conditions.

Graph	# nodes	# arcs
1	24	32
2	51	68
3	78	104
4	105	140
5	132	176

This resulted in 10 size/condition combinations. Each of these was examined using two separate sets of 20 trials for each subject giving a total of 40 trials per condition. The graph remained the same throughout a trial set, but for each trial a different pair of nodes was highlighted. The highlighted nodes were actually connected on 50% of the trials, randomly determined.

For each subject, the experiment was divided into two experimental sessions administered on separate days with a complete set of conditions given on a pseudorandom order on each day. Before beginning the experiment each day, the subject was given a short warm-up that presented two trials under each of the conditions, but only using three of the five complexities (low, middle, and high) for each condition.

Prior to each block of trials, subjects were told which experimental condition to expect. On each trial, subjects were given as much time as required to respond, which they did by pressing a specified mouse button. The response time and the response validity were recorded.

Due to the different graph sizes and the relatively high numbers of polygons to draw in the more complex scenes, the update rate was reduced to 30 updates per second for all conditions. That is, the scene was redrawn into the frame buffer 30 times per second. However, the video refresh rate was 120 Hz (60 Hz for each eye). The subject wore the stereo glasses, regardless of whether the condition required them, in order to avoid effects related solely to the reduced luminance caused by the glasses.




Fig. 5. Error rate data from Experiment (1a). Vertical bars represent one standard error about the mean.

Eleven subjects took part in this experiment, eight of whom had used similar apparatus and/or graphics systems before.

2.3 Results and Discussion of Experiment 1a

The results from Experiment 1a are summarized in Figure 5. As can be seen, the errors in the head-coupled stereo condition were dramatically reduced compared to those obtained with the 2D condition. An analysis of variance revealed highly significant main effects, for 2D versus 3D, $F_{(1,100)} = 276.8$, $p < 0.01$; the number of nodes $F_{(4,100)} = 23.23$, $p < 0.01$, and an interaction between 2D versus 3D and the number of nodes $F_{(4,100)} = 7.76$, $p < 0.01$. These results reflect the very large difference in error rates between 2D and 3D conditions, the increase in errors due to graph size, and the fact that the error curve for 3D is relatively flat compared to that for 2D (leading to an interaction).

The time data is summarized in Figure 6. An analysis of variance for this data only reveals a significant main effect for the number of nodes, $F_{(4,100)} = 10.14$, $p < 0.01$. There was no effect of conditions and no interaction. Because times are clearly increasing with graph size, these data suggest that the important determinant of response time is the number of nodes.

A problem with this first experiment was that the results showed a considerably greater advantage to 3D viewing than had been anticipated;




Fig. 6. Time data from Experiment (1a). Vertical bars represent one standard error about the mean.

we had assumed that the condition with 132 nodes would yield large error rates even in 3D. However, the error rate in 3D was low throughout the range of graph sizes tested. Clearly, in order to make a meaningful estimate of the relative advantages of 3D over 2D, a considerably larger graph was needed in the 3D condition. Experiment 1b was designed to remedy the deficiencies of 1a, and we defer further interpretation of the results until after the presentation of Experiment 1b.

3. EXPERIMENT 1b

Based on the results of Experiment 1a, we redesigned the experiment with different ranges of graph complexities for the different conditions. We also added two conditions to help answer the question of the relative benefits of head coupling and stereo.

3.1 Conditions

Conditions 1 and 4 were the same as conditions 1 and 2, respectively, of Experiment 1a.

- (1) **2D**: No stereo, no rotation; the 3D graph was projected onto a 2D plane using an orthographic (parallel) projection by removing Z axis information, hence no overlap information was available.

- (2) **Stereo perspective:** No rotation; this condition made use of a pair of StereoGraphics CrystalEyes LCD shutter glasses to provide the disparity depth cues.
- (3) **Head-coupled perspective:** The scene's perspective projection changed continuously according to the subject's measured head position; the perspective projection was defined by a single viewpoint centered between the eyes.
- (4) **Stereo, head-coupled perspective:** Same as the preceding, except with stereo. The correct view was generated for each eye position (continuously updated).

The numbers of nodes used in the different conditions were as follows.

- (1) 21, 42, 63, 84, 105
- (2) 51, 81, 111, 141, 171
- (3) 81, 117, 153, 189, 225
- (4) 111, 156, 201, 249, 291.

To obtain the number of arcs multiply by $\frac{1}{3}$.

The procedure for Experiment 1b was in most respects identical to Experiment 1a, except that the number of trials per condition/size combination was reduced to 12 to keep the session time to a manageable size. This experiment involved 11 participants, 6 of whom had been exposed to a similar environment before. There was one person from Experiment 1a who took part in Experiment 1b as well.

3.2 Results and Discussion of Experiment 1b

Figure 7 summarizes the error data from this experiment, with data from Experiment 1a included for comparison. It shows a sequence of curves with varying gradients that appear to be roughly multipliers of each other with respect to the graph size. That is, error rate appears to be directly proportional to the number of nodes, with a different gradient for the different conditions.

To determine an appropriate model for this family of curves, we ran a series of polynomial regressions independently for each of the four curves. Remarkably, we found that a one-parameter fit, of the form

$$\text{Errors} = C * N \quad \text{where } N \text{ is the number of nodes}$$

accounted for 95% of the variance. Two and three parameter fits accounted for 98% of the variance. However, these do not allow for comparisons among the different conditions. In view of the success of the one parameter model, we decided to adopt it as a reasonable first approximation to the data. Because the lines for the different conditions are multiples of one another, it has the advantage of allowing for a direct comparison among the different conditions. The results of this fit are shown as the broad straight lines in Figure 7. The coefficients for the one parameter fit were as follows:




Fig. 7. Error data from Experiments (1a) and (1b). Vertical bars represent one standard error of the mean.

Condition	Coefficient value
2D	0.369
Stereo perspective	0.232
Head-coupled perspective	0.167
Stereo, head-coupled perspective	0.123

On this basis, we conclude that the graph that can be understood with head-coupled stereo is about 3.0 times as large as the 2D graph for any given error rate (taking the ratios of the gradients). Adding stereo alone appears to increase the comprehensible graph size by approximately a factor of 1.6, and adding head coupling alone appears to increase the comprehensible graph size by a factor of 2.2.

To test the significance of these differences, we divided the subject means by the number of nodes in each condition to remove covariance due to the different numbers of nodes. We then conducted pairwise analyses of variance. The results revealed that each of the four conditions was significantly different from the others with $p < 0.01$.

$$\text{Condition 2 vs. condition 1} \quad F_{(1,100)} = 34.54 \quad p < 0.01$$

$$\text{Condition 3 vs. condition 2} \quad F_{(1,100)} = 15.68 \quad p < 0.01$$

$$\text{Condition 4 vs. condition 3} \quad F_{(1,100)} = 8.69 \quad p < 0.01$$

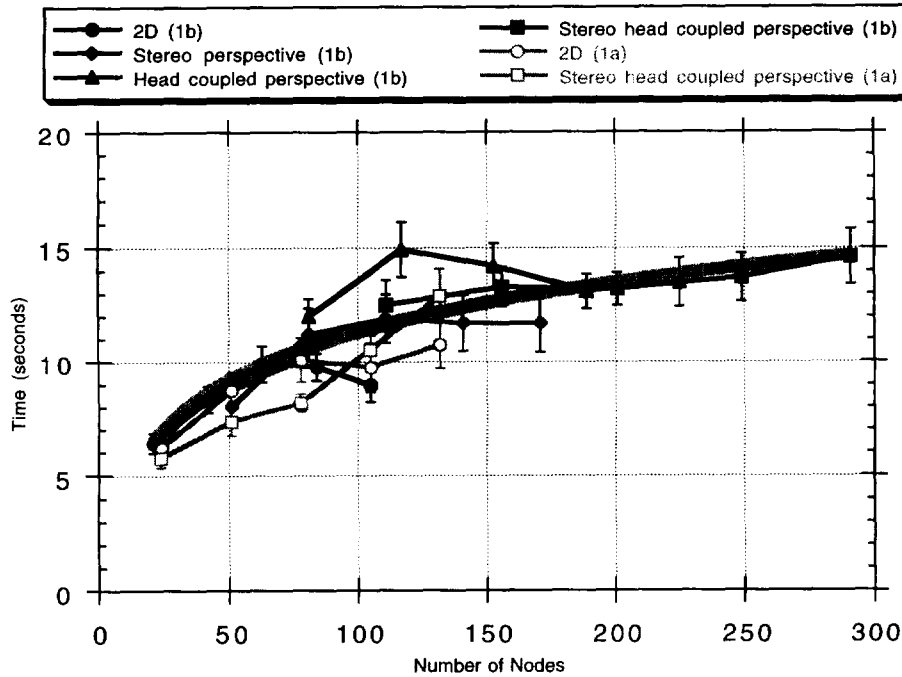


Fig. 8. Time to completion data from Experiments 1a and 1b. Vertical bars represent one standard error of the mean.

With differences as large as these it seems safe to assume transitivity. Thus, we conclude that all conditions were different with respect to error rates.

The average time for completion data is summarized in Figure 8. Completion time appears to depend more on the number of nodes and arcs in the graph than on the viewing mode as was shown in Experiment 1a, although there is evidence for different asymptotes for the different viewing modes. The times increase to about 13 seconds with approximately 100 nodes after which the curve levels off. The increasing part of the curve suggests that the time to process the data depends on the number of nodes, whereas the asymptote may represent the amount of time it takes for a subject to feel that they have extracted all they usefully can from the data. This asymptote is lower for 2D viewing, presumably, because the graph appears as a hopeless jumble sooner than it does for the 3D viewing modes. We found that the best fit to these data to be a function of the logarithm of the number of nodes:

$$\text{Time} = 3.21 \log_e(N/3.57) \quad r^2 = 0.77.$$

This function is represented by the broad line shown in Figure 8.

4. EXPERIMENT 2: WHAT IS THE BEST VIEWING MODE?

Experiments 1a and 1b provide compelling evidence that the application of head coupling and stereo allows greater information comprehension. Yet

we know from the previous work described in the introduction that different kinds of rotation can be used to enhance comprehension of a tree structure. This raises the issue of the relative merits of different viewing modes for data visualization. The nine conditions for this experiment included the four from Experiment 1b, two conditions with rotation added by hand motion (with and without stereo), and two other conditions with the rotation happening automatically (with and without stereo). To complete the set, we included a static perspective view condition that contained overlap and size perspective information. In order to keep this experiment manageable in length, we used only a single graph size for all conditions.

4.1 Method

Eighteen trials were conducted under each of nine conditions (described in the following). There were always 75 nodes and 100 arcs in every trial. Otherwise the method was the same as for Experiments 1a and 1b. Once again, eleven subjects were used. Five of these subjects had had exposure to a similar set-up before, and there were two participants who were involved in both Experiments 1b and 2.

4.2 Conditions

- (1) **2D:** No stereo, no rotation; the 3D graph was projected onto a 2D plane using an orthographic (parallel) projection by removing Z axis information, hence no overlap information was available.
- (2) **Static perspective:** No stereo, no rotation; essentially the same task as in (1), except that the graph is displayed using a perspective projection with the depth cues of relative size and overlap/occlusion. Conditions (2) through (9) all used a perspective projection.
- (3) **Stereo:** No rotation; this condition made use of a pair of StereoGraphics CrystalEyes LCD shutter glasses to provide disparity depth cues.
- (4) **Passive rotation:** No stereo; the scene rotated at a constant angular velocity of 20 degrees/sec about a vertical axis.
- (5) **Stereo, passive rotation:** Same as preceding except with stereo.
- (6) **Hand coupled:** No stereo, hand coupled; lateral movement of the mouse caused rotation of the scene about a vertical axis; mouse movement towards and away from the subject caused rotation of the scene about a horizontal axis. Movement was restricted to $\pm 128^\circ$ about the vertical axis and $\pm 49.1^\circ$ tilt (about the horizontal axis).
- (7) **Stereo, hand coupled:** Same as preceding, except with stereo.
- (8) **Head-coupled perspective:** The scene's projection changed continuously according to the subject's head position; the perspective projection was defined by a single viewpoint centered between the eyes.
- (9) **Stereo, head-coupled perspective:** Same as preceding, except with stereo; the correct view was generated for each eye position. Fish Tank VR.




Fig. 9. Results for Experiment 2. Vertical bars represent one standard error of the mean.

4.3 Results from Experiment 2

The results are summarized in Figure 9. The main difference among conditions is found in error rates. These ranged from a high of 26% in the 2D condition down to a low of 6.1% in the stereo hand-coupled condition.

We applied pairwise t-tests for all combinations of conditions and the results are given in Tables I and II. The results from large numbers of t-tests must be treated with some caution because 1 in 20 statistical differences (at the 0.05 level) will be a false positive.

The two 2D conditions gave significantly higher error rates than all other conditions although they did not differ significantly from each other. Other than this, we find that all three stereo conditions with motion were better than the stereo condition without motion. The stereo hand rotation condition was significantly better than all conditions except for the other rotation conditions. The head rotation condition was significantly worse than all three of the stereo rotation conditions.

These data confirm that the combination of stereo and motion is the most effective (average error 7.5%), and interestingly, they suggest that the method for producing the motion is not particularly important, inasmuch as no significant differences were found among the different combined stereo and motion conditions.

The average for the three motion alone conditions was 11.4%, whereas the average for the stereo alone condition was 15.4%. This agrees with the finding of Experiment 1b: that motion is more helpful than stereo. However, none of the differences among the three rotation and the stereo only conditions reached statistical significance. It is worth pointing out that

Table I. Error Rate Data: All Comparisons (* $p < 0.05$; ** $p < 0.01$)

	1	2	3	4	5	6	7	8	9
1 2D		**	**	**	**	**	**	**	**
2 Perspective	**		**	**	**	**	**	**	**
3 Stereo	**	**			*		**		*
4 Pass	**	**					*		
5 Stereo Pass	**	**	*					*	
6 Hand	**	**					*		
7 Stereo Hand	**	**	**	*		*		*	
8 Head	**	**			*		*		**
9 Stereo Head	**	**	*					**	

Table II. Time Data: All Comparisons (* $p < 0.05$; ** $p < 0.01$)

	1	2	3	4	5	6	7	8	9
1 2D						**		*	
2 Perspective						**		**	
3 Stereo						**		**	
4 Pass						**		*	
5 Stereo Pass						**		**	**
6 Hand	**	**	**	**	**		**		**
7 Stereo Hand						**		**	
8 Head	*	**	**	*	**		**		*
9 Stereo Head					**	**		*	

there are far more data involved in the comparison made for Experiment 1b, which may account for this discrepancy.

The response times were relatively uniform across conditions, and the fact that the times are approximately 10–12 seconds is entirely consistent with the data from Experiment 1b, given that the number of nodes was fixed at 75.

However, multiple *t*-tests revealed that there were significant differences. The hand and head motion conditions were slower than all other conditions, but they were not different from each other. There was only one other significant difference. The stereo head-coupled condition was slower than the stereo passive rotation condition.

5. CONCLUSION

We believe that our major contribution in this article has been to provide a first quantitative estimate of the benefits of stereo viewing with and without motion parallax for perceiving information nets. Our results surprised us, as we were prepared to believe that the second of the three theoretical views outlined in the introduction was the correct one and that adding 3D viewing would have only minimal benefits for an abstract data visualization task. The factor of three improvement that we actually measured (for the head-coupled stereo condition) should provide some strong encouragement to those who are developing experimental applica-

tions that use 3D viewing in information management systems [Koike 1993; Lieberman 1989; Mariani and Lougher 1992; Ware et al. 1993]. Our results suggest that this effort makes sense only if real-time rotation and/or stereo viewing is included in the data presentation. A static perspective image may add little in comparison with a 2D diagram and adding real-time rotation is more important than adding stereo. Unfortunately, structured motion is far more computationally expensive than stereo as it requires a minimum of 10 screen updates per second for a duration of at least 30 seconds, whereas a stereo pair only requires the computation of two images.

The practical conclusion from Experiment 2 is that although a variety of methods for introducing structured motion can help comprehension, the type of motion used should depend on the application. For example, if the selection of objects is important, then automatic rotation is not desirable because selecting moving objects is difficult. On the other hand, if head coupling is available, then this would probably not interfere with 3D selection. This is, presumably, because the motor control systems used for visually guided hand placement have evolved to work in conjunction with simultaneous head motion.

On the theoretical issue of the additivity of depth cues, we find a slight superadditivity in the cues of stereo and motion (60% improvement for stereo and 120% improvement for head coupling does not quite add up to the 200% improvement obtained with both). This agrees with the previous results of Sollenberger and Milgram [1993]. We do not wish to make much of this observation as we find it hard to believe that such a simple model can account for something as complex as space perception. In addition, there is an alternative interpretation of our result that does not rely on space perception at all. In our displays, more 2D patterns were inherently ambiguous than were the 3D patterns insofar as they contained paths that could not be resolved due to the densely overlapping network of nodes and arcs. It might be argued that the main advantage of the 3D view was simply the fact that motion allows the resolution of ambiguous arc destinations. Perhaps a mechanism that made all the nodes in a 2D display move at random would work just as well. Certainly, simple motion can be used to resolve patterns in abstract data [Limoges et al. 1989]. This is an interesting possibility which needs further investigation; however, it does not detract from the practical utility of our findings.

There does appear to be some penalty associated with using head-coupled viewing. A number of subjects stated that they found viewing the graph in the head-coupled stereo mode to be somewhat stressful. Part of this may be due to the difficulty of the task. Trying to perceive a path in a tangle of almost 300 nodes and 400 arcs is not easy. We also suspect that improvements in technology will help in this. Noise in the head-tracking system and the occasional briefly frozen image because of UNIX system functions are especially irritating in the head-coupled viewing modes. It is as if having adopted the real-world virtual reality metaphor, the brain expects the visual world to be stable, and if it is not we get a queasy feeling. This is

a problem that improvements in head tracking and display technology will help solve.

ACKNOWLEDGMENTS

The primary source of support for this research is an NSERC Canada Strategic Grant. We are grateful to Arthur Ryman at IBM Toronto Labs for his encouragement and support and for his assistance in providing code analysis tools. We also thank Tim Dudley at Bell Northern Research for providing the initial impetus to start this project and for ongoing enthusiastic support.

REFERENCES

- ARTHUR, K., BOOTH, K. S., AND WARE, C. 1993. Evaluating human performance for fishtank virtual reality. *ACM Trans. Inf. Syst.* 11, 3, 216–266.
- BRAUNSTEIN, M. L. 1976. *Depth Perception Through Motion*. Academic Press, New York.
- BRUNO, N. AND CUTTING, J. E. 1988. Minimodularity and the perception of layout. *J. Exper. Psychol. General* 117, 161–170.
- CARD, S. K., ROBERTSON, G. G., AND MACKINLAY, J. D. 1991. The information visualizer, an information workspace. In *CHI'91 Proceedings* (New Orleans, LA, April), 181–188.
- DEERING, M. 1992. High resolution virtual reality. *Comput. Graph.* 26, 2, 195–202.
- EADES, P. AND XUEMIN, L. 1989. How to draw a directed graph. In *IEEE Workshop on Visual Languages*. 13–17.
- FAIRCHILD, K. M., POLTROCK, S. E., AND FURNAS, G. W. 1988. SemNet: Three-dimensional graphic representations of large knowledge bases. In *Cognitive Science and Its Applications for Human-Computer Interaction*. Raymond Guindon, Ed. Lawrence Erlbaum, Hillsdale, NJ, 201–233.
- GIBSON, E. J., GIBSON, J. J., SMITH, O. W., AND FLOCK, H. 1959. Motional parallax as a determinant of perceived depth. *J. Exper. Psychol.* 58, 40–51.
- KOIKE, H. 1993. The role of another spatial dimensions in software visualization. *ACM Trans. Inf. Syst.* 11, 3, 266–286.
- LACK, L. C. 1974. Selective attention and the control of binocular rivalry. *Perception and Psychophysics* 15, 193–200.
- LAMPING, J. AND RAO, R. 1994. Laying out and visualizing large trees using a hyperbolic space. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, (Marina del Rey, CA, Nov.), 128–130.
- LIEBERMAN, H. 1989. A three-dimensional representation for program execution. In *IEEE Workshop on Visual Languages Proceedings* (Rome, Oct.), 111–116.
- LIMOGES, S., WARE, C., AND KNIGHT, W. 1989. Displaying correlations using position, motion, point size or point color. In *Graphics Interface Proceedings* (London, Ont., June), 262–265.
- MARIANI, J. A. AND LOUGHER, R. 1992. TripleSpace: An experiment in a 3D graphical interface to a binary relational database. *Interacting with Computers* 4, 2, 147–162.
- PURCELL, D. G. AND STEWART, A. L. 1991. The object detection effect: Configuration enhances perception. *Perception and Psychophysics* 50, 3, 215–224.
- ROBERTSON, G. G., MACKINLAY, J. D., AND CARD, S. K. 1991. Cone trees: Animated 3D visualizations of hierarchical information. In *CHI'91 Proceedings* (New Orleans, La., April), 189–194.
- SOLLENBERGER, R. L. AND MILGRAM, P. 1993. The effects of stereoscopic and rotational displays in a three-dimensional path-tracing task. *Hum. Factors* 35, 3, 483–500.
- SOLLENBERGER, R. L. AND MILGRAM, P. 1991. A comparative study of rotational and stereoscopic computer graphic depth cues. In *Proceedings of the Human Factors Society Annual Meeting* (San Francisco, CA, Sept.), 1452–1456.

WARE, C., ARTHUR, K., AND BOOTH, K. S. 1993. Fishtank Virtual Reality. In *INTERCHI'93 Technical Paper. Proceedings*. 37-42.

WARE, C., HUI, D., AND FRANCK, G. 1993. Visualizing object oriented software in three dimensions. In *Proceedings of CASCON'93*. (Toronto, Oct.) IBM Canada Ltd. and NRC. 612-620.

Received March 1994; revised April 1995; accepted November 1995