

Visualization Task Performance with 2D, 3D, and Combination Displays

Melanie Tory, Arthur E. Kirkpatrick, M. Stella Atkins, and Torsten Möller

Abstract—We describe a series of experiments that compare 2D displays, 3D displays, and combined 2D/3D displays (orientation icon, ExoVis, and clip planes) for relative position estimation, orientation, and volume of interest tasks. Our results indicate that 3D displays can be very effective for approximate navigation and relative positioning when appropriate cues, such as shadows, are present. However, 3D displays are not effective for precise navigation and positioning except possibly in specific circumstances, for instance when good viewing angles or measurement tools are available. For precise tasks in other situations, orientation icon and ExoVis displays were better than strict 2D or 3D displays (displays consisting exclusively of 2D or 3D views). The combined displays had as good or better performance, inspired higher confidence, and allowed natural, integrated navigation. Clip plane displays were not effective for 3D orientation because users could not easily view more than one 2D slice at a time and had to frequently change the visibility of individual slices. Major factors contributing to display preference and usability were task characteristics, orientation cues, occlusion, and spatial proximity of views that were used together.

Index Terms—H.5.2 User Interfaces - Graphical User Interfaces (GUI), Screen Design, Evaluation/Methodology, I.3.3 Picture/Image Generation - Display Algorithms, J. Computer Applications (e.g., CAD, Medical Imaging)

INTRODUCTION

Spatial data can be presented in 2D views that provide information about only 2 dimensions or 3D views that provide information about 3D structure. The two view styles are effective for different tasks. In qualitative field observations, Springmeyer et al.[8] noted that 2D views are often used to establish precise relationships, whereas 3D views are used to gain a qualitative understanding and to present ideas to others. Other studies with various users and tasks have found that 2D views can enable analysis of details and precise navigation and distance measurements (since only one dimension is ambiguous)[7],[9] whereas 3D views facilitate surveying a 3D space, understanding 3D shape, and approximate navigation[9],[15].

Although many studies have compared displays consisting only of 2D or 3D views, few have compared displays combining 2D and 3D views. In this paper, we present a series of experiments that compare 2D, 3D, and combined 2D/3D displays for different tasks. Our objective is to identify the tasks for which each view is best suited.

We focus on tasks that require both overview and detail. During such tasks as orienting and positioning objects relative to one another, a 3D view may help gain an overall understanding of the space, whereas a 2D view may assist precise judgments and actions. An example is the analysis of volumetric medical scans. To see details without occlusion, radiologists typically view such data in 2D slices. They may also use a 3D view to gain an overall qualitative understanding, to explain ideas to other physicians, or to place slicing planes in non-standard orientations. Occlusion issues also arise in the display of computer aided de-

sign (CAD) models: parts near the front can occlude parts at the back. For this reason, CAD models are usually displayed from several viewpoints simultaneously, with three standard 2D orthographic views (used for precise editing and measurement) supplemented by one or more oblique viewpoints (to assist understanding of 3D structure).

Throughout this article we distinguish between “view” and “display”. We define a **view** as a single projection of a three-dimensional object. A **2D view** shows a thin slice through a scene or a front, back, right, left, top, or bottom orthographic projection. A **3D view** is any other type of orthographic or perspective projection that shows 3D spatial structure. We define a **display** as an arrangement of one or more views on the screen. A **strict 2D display** consists entirely of 2D views, and a **strict 3D display** consists entirely of 3D views. A **combined 2D/3D display** has at least one 2D view and at least one 3D view.

We describe three experiments comparing combinations of 2D and 3D views:

- Experiments 1 and 2 compared displays consisting of strict 2D and 3D views and displays combining 2D and 3D views for relative positioning (Experiment 1) and orientation (Experiment 2) tasks. In these experiments, the orientations of the views relative to the dataset were fixed, and the tasks and datasets were abstract. The controlled, abstract conditions permitted accurate estimation of the time and errors when performing the tasks with each display.
- Experiment 3 qualitatively explored design issues for combined 2D/3D views. In this experiment, the user could modify the orientations of the views relative to the dataset, and the task and dataset approximated conditions of actual use.

Together, these experiments provided a robust, multifaceted evaluation of combined 2D/3D views.

Melanie Tory is with the Computer Science Department, University of British Columbia, 2366 Main Mall, Vancouver, BC, Canada, V6T 1Z4. Other authors are with the Computing Science Department, Simon Fraser University, Burnaby, BC, Canada, V5A 1S6. Email: melanie@cs.ubc.ca, {ted, stella, torsten}@cs.sfu.ca.

MOTIVATION

Understanding spatial structure from a physically two-dimensional display

The display types considered in this paper are designed to render 3D structures on a screen consisting of two physical dimensions, requiring simulation of the third dimension. The mental processes of recovering the spatial structure of a displayed object are beyond the scope of this paper. For our purposes of evaluating some variations of 2D and 3D displays, we focus on two mental operations the user must perform to effectively use the display: determine distances and angles, and register multiple views of the same object.

Each display type represents a different compromise among support for the activities. A display consisting of 2D slices on the face of a cube (see Fig. 1(a)) makes it convenient to determine distance along any line parallel to a face. However, to estimate distance along a line intersecting more than one face, the user must first mentally register the views, correlating them to locate their common points. In contrast, the 3D Shadow display (see Fig. 1(c)) has only a single view, eliminating the need for registration. The projection of the display permits convenient estimation along the object faces, but the distance between the floating ball and the box is ambiguous if no further cues are given. Adding a directional light source that produces a shadow under the ball resolves the ambiguity.

View registration is particularly difficult when the views show different rotations of the same object. Registration requires mentally rotating the scene, which is a challenging cognitive activity for two 3D views [6], such as those shown Fig. 1 (b), as well as a 2D view plus a 3D view [11].

Analyses such as these naturally lead to questions of comparison: Are there tasks for which 2D or 3D views are uniformly superior?

Comparisons of Strict 2D and 3D Displays

Many studies have compared displays consisting exclusively of 2D or 3D views. The bulk of these studies have been domain- and task- specific. For example, Smallman et al. [7] showed that visual search was faster with 2D air traffic control displays. Van Orden and Broyles [14] found that 2D displays were as good as or better than 3D displays for aircraft speed and altitude judgments, but 3D volumetric displays were best for collision avoidance tasks. Park and Woldstad [5] reported that 2D and 3D displays were equally good for telerobotic positioning, provided the 3D displays were enhanced by extending reference lines from the face of the robot gripper.

Identifying overall principles about 2D/3D display design is difficult with domain-specific studies. The studies each consider different users and tasks and vary greatly in terms of display parameters. To provide more general guidance, Wickens et al. [15] proposed the “Proximity Compatibility Principle” (PCP). This principle states that tasks requiring integration of spatial dimensions benefit from 3D views, whereas tasks requiring focused attention on one or two dimensions benefit from 2D views. In support of this claim, Wickens et al. [15] presented data showing that 3D representations were better than 2D representations for

more integrative questions requiring knowledge of several dimensions. However, St. John et al. [9] published data partially contradicting the PCP. Although their finding that 3D displays were better for shape understanding tasks was predicted by the PCP, they also found that 2D displays were superior for relative positioning tasks, even when understanding the position required 3D knowledge. This last conclusion contradicts the PCP; however, it may be specific to their form of strict 3D display. In particular, registering the two views in their “3D” display required mental rotation. Our Experiment 1 replicated their task with a 3D display that did not require mental rotation.

Displays Combining 2D and 3D Views

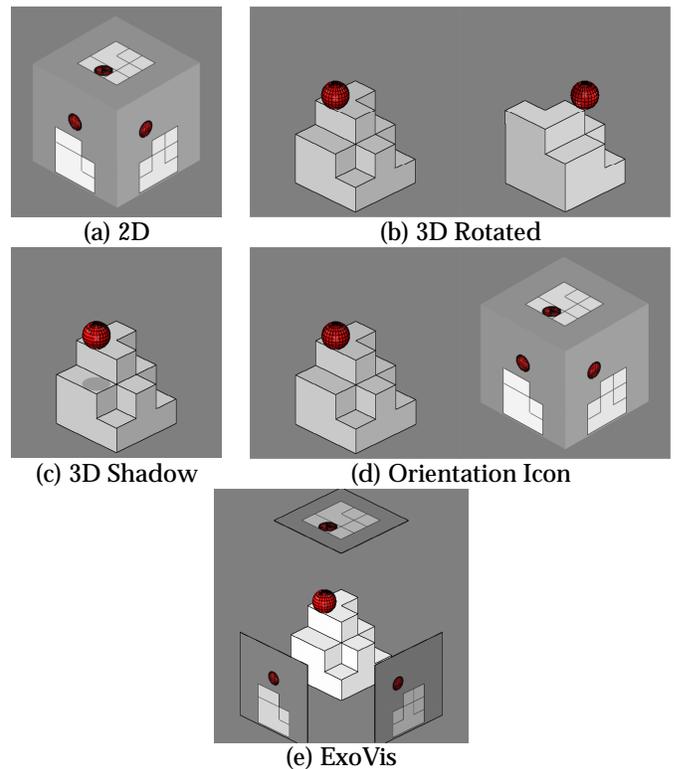


Fig. 1 (a) 2D, (b) 3D Rotated, (c) 3D Shadow, (d) Orientation Icon, and (e) ExoVis displays used in Experiment 1 (position estimation). Participants estimated the height of the ball relative to the block shape. In this example, the ball is at height 1.5 diameters above the block shape.

The evidence of differing strengths for 2D and 3D views suggests that displays combining both view types may provide superior performance. While this logic has appeal, the benefits of combining them may be offset by the greater complexity of the display. Increasing the number of disjoint views of the object increases the amount of view registration that the user must perform. Furthermore, some organizations of the different views, such as the ExoVis display described below, introduce the possibility of one view occluding other views, over and above any self occlusion in the displayed object. Consequently, adding more views may not always lead to higher task performance.

In an earlier paper, we defined three categories of combined 2D/3D displays, and compared them based upon factors such as flexibility, occlusion, and expected difficulty

of view registration [13]. The orientation icon (OI) display places 2D and 3D views side-by-side (see Fig. 1(d)). The 3D view “orients” users by indicating the positions of the 2D views relative to the object and each other. However, because 2D views are physically separated from the 3D view and can be translated and rotated from their original location, view registration can be challenging. In this paper, we use two types of OI displays, one where the 2D views are displayed in a box shape (Experiments 1 and 2), and another where the 2D views are displayed flat on the screen (Experiment 3). View registration is expected to be easier in the box version because mental rotation is not required.

The ExoVis display (see Fig. 1(e)) is an extension of “interactive shadows” [2] to incorporate 2D slices. ExoVis comprises a 3D view in the centre of the display with 2D views surrounding it in close proximity. The 2D views are translated but not rotated from their original orientations. Registering the 2D and 3D views requires less transformation than orientation icon displays with flat 2D views, but more than the clip plane displays described below.

Clip plane displays contain invisible planes that hide everything between themselves and the camera. In this way, clip planes show 2D slices in their exact position within the 3D view. For example, Fig. 2 shows a torus that has been cut by two orthogonal clip planes, leaving only the back quadrant visible. Variants include the “planar brush” [16], which shows a slice of a 3D object surrounded by an outline or semi-transparent surface of the 3D object’s outer contour, and opening the volume like a book (e.g., [1],[3],[4]). View registration is easier with a clip plane display than with orientation icon or ExoVis. However, because a clip plane clips off everything between itself and the viewer, it is less generally useful than other displays.

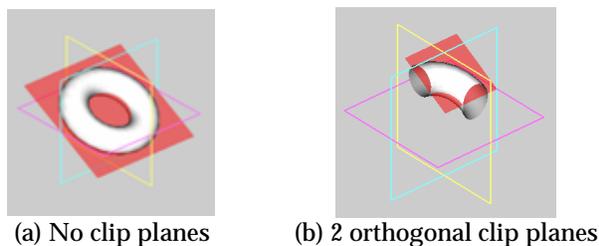


Fig. 2 Clip Plane display used in Experiment 2. The 3D scene is shown (a) without clip planes and (b) with 2 orthogonal clip planes. Coloured outlines represent possible clip planes. Blue and yellow clip planes have been turned on in (b). In the experiment, participants oriented the red (dark) plane so it cut the torus exactly in half.

Comparisons of Combined 2D/3D Displays

St. John et al. [10] compared strict 2D, strict 3D, and an orientation icon display for a 3D route-planning task. Task completion was fastest with the orientation icon display, indicating that combined 2D/3D displays are valuable. However, this experiment did not consider other combined 2D/3D displays such as ExoVis or clip plane displays.

In previous work, we evaluated the three combined 2D/3D displays according to heuristics of cognitive difficulty [13]. For mental registration of 2D and 3D views, a subcomponent of the larger process of interpreting the dis-

play, the heuristics predicted that clip planes would pose the smallest cognitive load, ExoVis a moderate amount, and orientation icon the greatest. A subsequent empirical study confirmed these predictions [11].

Contribution

The research reported in this paper extends the above studies in both task complexity and range of displays compared. Experiments 1 and 2 compared the combined 2D/3D displays as in our previous work [11], but included more than one 2D view in each 2D/3D display and applied the combined 2D/3D displays to higher-level cognitive tasks. These experiments also considered strict 2D and 3D displays in addition to the three combined 2D/3D displays. Experiment 3 increased the complexity further, allowing users to interactively rotate the scene and adjust positions of 2D slices within the display. Experiment 3 also used a more realistic dataset than Experiments 1 and 2 (a 3D medical image rather than a simple block shape).

Timing and error data from Experiments 1 and 2 were reported in [12]. In this paper, we summarize the timing and error results, and then include additional data and extend the discussion. Specifically, we report error sizes, participants’ ability to predict their error, and subjective ratings. Experiment 3 has not been previously published.

OVERVIEW OF EXPERIMENTS

In Experiments 1 and 2, we compared strict 2D and 3D displays to combined 2D/3D displays. Four displays (2D, 3D, orientation icon, and ExoVis) were included in both experiments. Each experiment also included a fifth display. In Experiment 1, we included St John et al.’s 3D rotated display (see Fig. 1(b)), to compare that display with another strict 3D display, 3D Shadow (see Fig. 1(c)). In Experiment 2, the fifth display was a clip plane. The clip plane was not included in Experiment 1 because that experiment used 2D orthographic views rather than slices, so clip planes were not possible. In Experiment 3, we qualitatively explored design issues in 2D/3D combination displays by comparing ExoVis and orientation icon display designs. Experiment 1 considered a relative position estimation task, Experiment 2 considered relative orientation, and Experiment 3 considered volume-of-interest (VOI) positioning.

In Experiment 1, participants used 3D views and/or 2D orthographic views to estimate the position of a ball relative to a block shape (see Fig. 1). This task is representative of many domains (e.g., positioning objects relative to one another in CAD or determining a tumour’s position relative to a blood vessel in medical imaging). These tasks could be performed using interactive measurement tools, but it is likely faster to approximate the relative position before (or instead of) using tools to determine exact positions.

In Experiment 2, participants used 3D views and/or 2D slice views to adjust the orientation of a plane relative to a torus (see Fig. 2). The task was modeled after slicing plane orientation tasks. For example, medical images of the chest usually align with the main axes of the body. Because the heart is at an oblique angle to this grid, physicians often need to orient an oblique slice. Orientation tasks are also

common in other domains (e.g., to set the angle of the roof of a house in CAD).

In Experiment 3, participants positioned a box-shaped volume-of-interest around an anomaly in a volume data set, similar to outlining a tumor or lesion for medical image analysis. Figures and details for Experiments 2 and 3 are given later.

These tasks were chosen for several reasons. First, they are common in applications. Second, our primary research objective was to investigate 2D/3D combination displays, so we chose tasks that were expected to benefit from both 2D and 3D views. (Strict 2D and 3D displays were then included in Experiments 1 and 2 to test this hypothesis.) For Experiments 1 and 2, we chose abstract tasks and simple shapes rather than domain-specific tasks and data sets to make the results generalizable to many fields. In addition, these tasks required only simple displays (minimizing conflicting factors), did not require domain knowledge, and had clearly defined answers. In Experiment 3, we used medical image volume data for contrast with the simple shapes in Experiments 1 and 2.

Experiments 1 and 2 were run sequentially with the same participants. Order was counterbalanced. Experiment 3 was run at a later time with different participants.

EXPERIMENT 1: POSITION ESTIMATION TASK

Method

We used a 5 x 5 (display x ball height) mixed design. In pilot work, we found an asymmetrical transfer effect across display types when several displays were used in quick succession; therefore display type was made a between-subjects factor. Ball height was a within-subjects factor. Height order was randomized, although slightly restricted to avoid having many similar heights in succession. We measured time, accuracy, and subjective ratings of difficulty, likeability, and confidence that answers were correct.

Participants

A total of 40 volunteers (20 male, 20 female) were recruited from university computing science and engineering student populations (ranging from first year to graduate level). Participants had varied levels of experience with computer graphics. Mean age was 27. Four students of each gender were randomly assigned to each display condition.

Task and Stimuli:

We used a variation of St. John et al.'s relative position estimation task [9]. Example trials are shown in Fig. 3. Participants determined the approximate empty vertical space (height) between a ball and a block shape. The empty space was always 0, 0.5, 1, 1.5, or 2 sub-block sized units. Although this may seem like a 1D task, determining the height actually requires a 3D understanding. Because the top surface of the block shape is not flat, estimating the height requires knowing the ball's horizontal position.

Scenes were presented as static images. Block shapes were generated by removing 2, 5, or 8 cubes from a base shape containing 27 cubes (3 x 3 x 3). Orthographic 2D views were rendered from the top, right, and front. As in

St. John et al.'s [9] study, 3D views were rendered with orthographic projection.

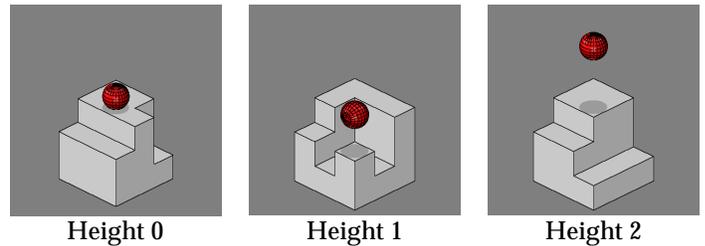


Fig. 3 Sample trials for Experiment 1. The 3D shadow display is shown with 3 different block shapes and ball heights. Heights 0.5 and 1.5 were also included in the experiment.

Display Conditions

We compared 5 displays: 2D, 3D rotated, 3D shadow, ExoVis, and orientation icon (OI). Example trials with each of these displays are shown in Fig. 1. 2D views were presented at oblique angles to improve the correspondence with 3D views. 3D rotated was included to replicate the corresponding condition in St. John et al. [9] and 3D shadow was included to assess the impact of a shadow cue. To create the shadow, a directional light source was placed directly above the ball, so a shadow projected directly downwards onto the block. Shadows were not shown in 3D views for the other displays so that participants would be forced to use the alternate views. Therefore variations in these views could be meaningfully assessed.

Procedure

Participants first reviewed instructions and examples with answers. They then completed 5 practice trials and 20 experimental trials (4 with each of the 5 ball heights) with one of the 5 displays. Answers were not provided for experimental trials. The same sequence of block/ball scenes was shown for each display (to ensure consistency).

A screen with one button labeled "Ready" appeared before each trial. A trial began when the participant pressed the "Ready" button, and ended when s/he pressed one of 5 buttons at the bottom of the screen to specify their answer.

Participants were instructed to be as accurate as possible and take as much time as they needed. No time limit was imposed since trials were not expected to take a long time. Answers could not be changed. Following each trial, participants rated their confidence in their answer for that trial (on a 7-point rating scale). In a post-trial questionnaire, participants gave opinions of the display and rated the ease of task components (overall ease, understanding block shape, understanding which cube the ball was above, and estimating the ball's height). Experimental sessions lasted approximately 20 minutes.

Analysis

We analyzed quantitative results by ANOVA (on means for each subject) followed by pairwise comparisons or Tukey HSD tests (for within- or between-subjects variables, respectively). When Mauchly's test of sphericity indicated it was necessary, we used the Huynh-Feldt correction. Non-

parametric tests were used for rating scale data. Error trials were excluded from the timing analysis.

Results

Fig. 4 summarizes mean time and error. For both measures, we found main effects for display (Time: $F(4,30)=6.1, p=0.001$, Error: $F(4,35)=5.1, p=0.002$) and height (Time: $F(3.6,106.9)=31.4, p<0.001$, Error: $F(3.7, 128.5)=6.3, p<0.001$), and a display/height interaction (Time: $F(14.2,106.9)=2.0, p= 0.024$, Error: $F(14.7,128.5)=3.3, p<0.001$).

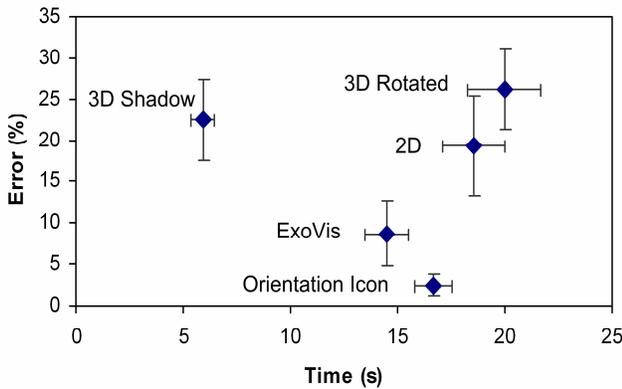


Fig. 4 Mean time and error (percentage of trials that were incorrect) for Experiment 1. Error bars show standard error.

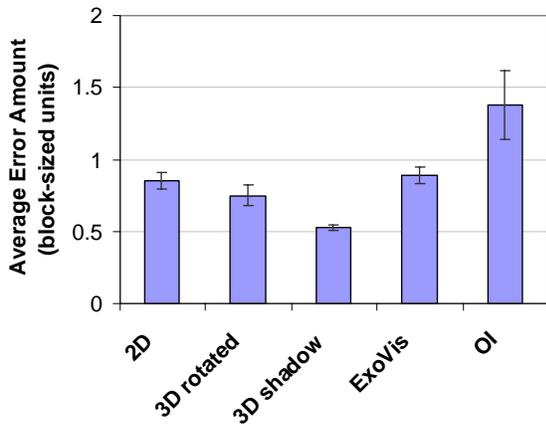


Fig. 5 Average error size in Experiment 1. Correct answers are not included in the averages. Error bars show standard error.

3D shadow was significantly faster than all other displays ($p \leq 0.035$), but had a large number of errors. As shown in Fig. 5, most of these errors (34 out of 36) were incorrect by only 0.5 height units, indicating that participants understood the ball’s position, but slightly misjudged the height. OI and ExoVis had moderate times (though not significantly different from 2D and 3D rotated). Although the size of their errors tended to be large (see Fig. 5), OI and ExoVis had fewer errors than the other displays (see Fig. 4). Specifically, OI had significantly fewer errors than 3D rotated and 3D shadow ($p \leq 0.04$) and marginally significantly fewer than 2D ($p = 0.056$). ExoVis had significantly fewer errors than 3D rotated ($p = 0.048$). 2D and 3D rotated dis-

plays (see top right corner of Fig. 4) were not very effective, having the longest times, many errors, and large error sizes.

Higher ball heights required more time for all displays, and produced more errors for 2D, 3D rotated, and 3D shadow displays. Ball height details and additional timing and accuracy results may be found in [12].

Participants rated the displays differently. Subjective ratings of overall ease of use, confidence that answers were correct, and display likeability are given in Fig. 6. The Kruskal Wallis test showed marginally significant differences between displays for confidence ($\chi^2 (df=4, N=40) = 9.2, p = 0.055$) and significant differences for likeability ($\chi^2 (df=4, N=40) = 17.2, p = 0.002$) and overall ease ($\chi^2 (df=4, N=40) = 13.4, p = 0.011$). OI and ExoVis were well liked, easy to use, and evoked the highest confidence that answers were correct. 3D shadow was next best, then 3D rotated, and lastly 2D. One interesting observation is that participants reported fairly high confidence in their answers for all displays, even those on which they performed poorly. This has potentially troubling implications for 2D, 3D rotated, and 3D shadow displays, where the error percentage was quite high. 3D rotated was considered easy for understanding the block shape, but difficult for understanding ball position and height. By contrast, understanding block shape was difficult with the 2D display.

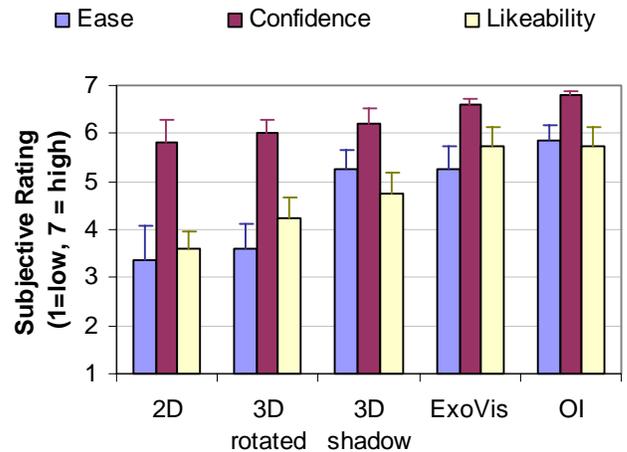


Fig. 6 Subjective ratings of Ease, Confidence, and Display Likeability for Experiment 1. Error bars show standard error.

We also examined learning effects. Fitting power laws of practice to successive trials indicated that trial time was fairly consistent during the actual experimental trials. A possible exception is the 2D display, which had a much steeper learning curve than the others. This suggests that 2D displays are quite difficult for novice users, but practiced users might be able to use them effectively.

Overall, our results indicate that 3D rotated displays are not effective for relative position estimation, replicating the results of St. John et al. [9]. Based on observations and participants’ comments, we determined that common problems with this display were difficulties estimating ball height (especially for high heights and 1/2 unit heights) and difficulty relating the views. Half the participants who used the 3D rotated display felt the views sometimes con-

flicted with one another, even though they knew otherwise. This occurred because the ball position was ambiguous in any one view, and one possible ball position sometimes visually dominated the other. It is possible that this type of 3D display may be substantially improved by allowing interactive rotation rather than only mental rotation, since relating the views should be easier. However, the time cost associated with rotating the view could be large.

Unlike 3D rotated, the 3D shadow display was very fast and received moderate ratings. Taken together with the St. John et al. [9] results, these show that the performance of strict 3D displays is sensitive to height cues and whether viewers must perform mental rotation to register views. Almost all participants who used the 3D shadow display commented that the major difficulty was estimating ball height, not understanding the ball's horizontal position. Error data also indicated this difficulty. Several participants either requested a ruler or were observed using their hand as a measuring tool, indicating that 3D shadow displays could be very effective for relative position estimation with the addition of measurement cues. Alternatively, designers could use a point light instead of a directional light so the shadow size would indicate height. Because estimating height was users' biggest complaint with this display, we expect likeability, confidence, and ease of use would increase if stronger height cues were added. We expect such cues would also benefit the other displays, but to a lesser extent, as their 2D views already provide good height cues.

Nevertheless, shadows would not always be effective because the light must be placed in a specific location relative to the objects of interest (e.g., if the light were slightly off to the side, it could be less effective). In addition, shadows can be hard to interpret and costly to render in scenes with complex or dense geometry.

As computing power increases, we can compute, collect, and store larger and larger amounts of data, such that complex and dense information spaces are becoming increasingly common. Thus it is useful to consider 2D views as an alternative way of resolving position ambiguity. Since ExoVis and OI had few errors, moderate time, and high ratings on all scales, we conclude that combination 2D/3D displays are a better choice than strict 2D displays for relative position estimation tasks and should be chosen when 3D + shadow displays are not practical and/or 3D measurement tools are unavailable. Note that this recommendation applies to novice users; with practice, users may become comfortable working with strict 2D displays.

Only 8 participants were in each group. This allowed us to examine several displays; however, significant differences could have been missed, and differences in spatial ability could be confounding. To check for this confound, we compared the groups' computer graphics experience. All groups were similar, except the ExoVis group had slightly lower experience. Furthermore, we checked timing data trends after removing outliers (exceptionally slow or fast participants). The only difference was that 2D became slower than 3D rotated. Therefore, we believe our conclusion that 2D/3D combination displays are better than 2D or 3D rotated displays is sound.

EXPERIMENT 2: ORIENTATION TASK

Method

We used a between-subjects design with a single factor, display type, having 5 levels. We measured time, accuracy, estimated error, and subjective difficulty ratings.

Participants:

Participants from Experiment 1 also took part in Experiment 2. Order of the experiments was counterbalanced.

Task and Stimuli

Participants used a 3 degree-of-freedom (DOF) input device (see Fig. 7) to orient a plane such that it cut a torus into two identical parts, as if slicing a bagel in half (see Fig. 8). Tori were rotated between 12° and 55° from horizontal. The axis of rotation varied for every trial.

Interaction Technique

Fig. 7 illustrates our input device. We used a 6 DOF Polhemus Fastrak to input plane orientation. Position data from the Fastrak was discarded. We attached the Polhemus sensor to a square piece of plywood. Orientation of the plywood directly mapped to orientation of the red plane. The camera position could not be changed.



Fig. 7 Custom input device for Experiment 2.

With the clip plane display, having all the slices on at once hides most of the display, so we believed participants would want to turn slices on and off. In addition, participants needed to easily start and end trials. To accomplish this, we positioned a 3-button mouse on the plywood; the 3 slices could be turned on and off using the 3 buttons. Mouse buttons were labeled with colours to match the colours of slices on the display. Mouse buttons were wired to the same buttons of a second 3-button mouse; the regular mouse was used for ordinary mouse actions while the custom input device recorded mouse clicks.

A "Ready" screen appeared between trials. Participants started each trial with a single click on the custom input device, and ended with a double-click. After each trial, participants estimated their error (in degrees) and typed this number at a prompt.

Display Conditions

We compared 5 displays (see Fig. 8). Three were comparable to those in Experiment 1: 2D, orientation icon, and ExoVis. The 3D display in this experiment did not have a

shadow (it would have been uninformative because the dataset had only one part), but was otherwise similar to the 3D shadow display in Experiment 1. The fifth display was a clip plane, a combined 2D/3D display that only applies to 2D slices. (Experiment 1 used 2D orthographic views, so it did not include the clip plane display.)

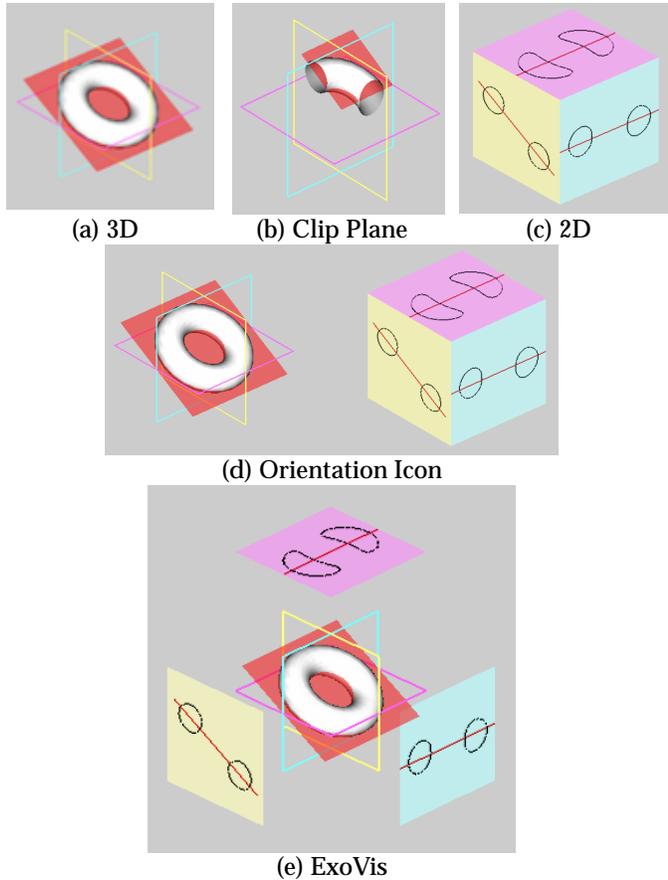


Fig. 8 Displays for Experiment 2 (orientation task). Participants oriented the red (dark) plane relative to the torus using the custom input device. Individual 2D views could be hidden, but the camera could not be moved. Examples show approximately correct solutions. The clip plane image (b) shows a 3D view cut by 2 orthogonal clip planes.

Procedure

Participants completed 5 practice trials and 16 experimental trials. The same sequence of stimuli was presented to every participant. The sequence was randomly generated, subject to the restriction that consecutive stimuli had to be somewhat different. Participants were asked to be as accurate as possible. No time limit was imposed since trials were not expected to take a long time. In a post-trial questionnaire, participants rated the difficulty of task components and gave opinions of the display. Experimental sessions lasted approximately 30 minutes.

Analysis

Based on participants’ comments, we suspected that trials with the 3D display were easier when the side of the torus was visible. Participants knew the plane was aligned when it became a simple line and/or aligned with the to-

rus. Therefore, for our analysis, we divided our trials into the 3 types shown in Fig. 9: side (torus hole not visible), top (full extent of the hole visible), and oblique (hole partially visible). Classification of each stimulus was done by visual inspection. If there was doubt about which category a stimulus belonged in, it was considered oblique. In total, there were 4 side, 4 top, and 8 oblique trials. Post-hoc analysis showed that all trial types were distributed over the duration of the experiment. Given this assurance that the mean time by trial type was not confounded with participant learning, we analyzed the data as a 5 x 3 mixed-effects ANOVA, with trial type as a within-subjects factor. Statistical procedures were the same as Experiment 1. Time was analyzed only for trials with error < 5 degrees.

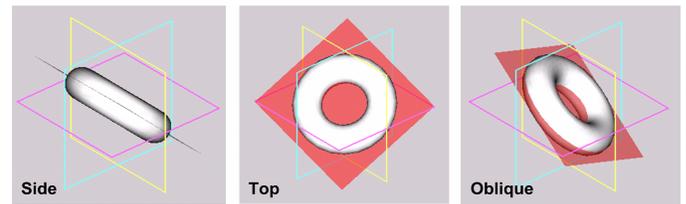


Fig. 9 Trial types for Experiment 2.

Results

Mean time and error data are summarized in Fig. 10. For both measures, we found a main effect for display (Time: $F(4,32)=4.1, p=0.009$, Error: $F(4,35)=7.3, p<0.001$) and a display/trial type interaction (Time: $F(7.5,60.3)=3.5, p=0.002$, Error: $F(8,70)=3.6, p=0.002$). There was also a main effect of trial type for the error data ($F(2,70)=15.4, p<0.001$).

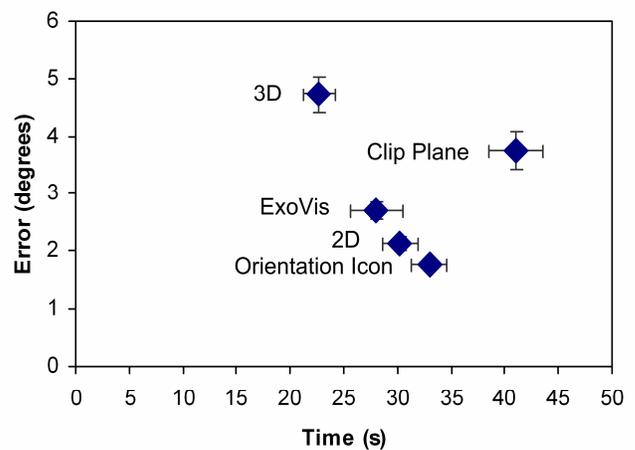


Fig. 10 Mean time and error data for Experiment 2. Error bars show standard error.

ExoVis, OI, and 2D performed well overall, having the lowest error and moderate times. Although the 3D display was fastest (significantly faster than clip plane ($p = 0.007$) and marginally significantly faster than OI ($p = 0.052$)), 3D had significantly more error than 2D and OI ($p < 0.006$). The clip plane performed poorly, having high error and requiring long times. Clip planes took significantly longer than 3D (p

= 0.007), and marginally significantly longer than ExoVis ($p = 0.099$). Also, clip planes had more error than OI ($p = 0.008$) and marginally significantly more error than 2D ($p = 0.052$).

However, significant differences between displays only occurred for top and oblique trials, not for side trials. All displays had low error levels for side trials. This suggests that 3D can be just as accurate as 2D and combination displays when a good view direction is used. For these and other timing and error statistical details, see [12].

By subtracting the actual error on each trial from the participant's estimated error (typed at a prompt following each trial), we obtained a measure of how well participants could predict their own accuracy, as summarized in Fig. 11. For this measure, we found significant differences between displays ($F(4, 35) = 2.7, p = 0.045$) and trial types ($F(2, 70) = 4.5, p = 0.014$) and a significant interaction between display and trial type ($F(8, 70) = 3.8, p = 0.001$).



Fig. 11 Ability to predict orientation accuracy.

As shown in Fig. 11, 2D and OI participants consistently overestimated their error, whereas 3D participants underestimated their error. This poses potential problems for 3D displays (because users may be overconfident in their accuracy) as well as 2D and OI displays (because users may take excessive amounts of time before they feel confident in their performance). We believe the problem with 3D displays was that participants could not always see changes in plane orientation relative to the torus, particularly for top trials. For 2D, participants may have overestimated their error because they did not feel confident in their 3D understanding of the scene; however, this does not explain why OI participants had similar results. ExoVis and clip plane participants were best able to predict their own accuracy. Overall, there was a significant difference between 3D and 2D ($p = 0.041$), and a marginally significant difference between 3D and OI ($p = 0.079$). However, these differences only existed for top and oblique trials. There were no significant differences between displays for side trials, suggesting that people can closely predict orientation accuracy with 3D displays when a good view is available.

We also examined learning effects. Fitting power laws of practice to successive trials indicated that time and accuracy became fairly consistent following the practice period.

Our rating scales did not show significant differences between displays, so results are not shown in this paper.

However, participants' comments and our observations provided interesting insight. Most people using the 2D display did not appear to naturally understand how to move the input device to progress towards their goal. Progress was generally made by trial and error, and by focusing on one dimension at a time. Availability of the 3D view produced more directed and coordinated movements.

OI and ExoVis participants tended to move quickly to an approximate solution using the 3D view, and then fine-tune individual dimensions using the slices. Some clip plane participants used a similar strategy, turning off all slices to get an approximate solution, and then using one slice at a time to adjust each dimension. Other clip plane participants started with the default view (2 slices); however, this view was only liked by one participant and we believe the others used it only because it did not require changing any settings. Most participants found it difficult to work with more than one clip plane at a time, and found switching between dimensions difficult and annoying, for at least two reasons. First, to switch dimensions users either had to randomly try input device buttons to find the correct one or move their eyes from the screen to the input device to match the clip plane colour to the input device button colour. Second, users would often correctly orient the plane in one dimension and then find that this action had altered the orientation in other dimensions.

As for Experiment 1, we analyzed time and error data without outliers in case differences in spatial ability strongly influenced our results. Removing one very slow participant resulted in a faster average time for ExoVis (similar speed to the 3D display), but produced no other changes in timing or error trends. As with Experiment 1, our results in this experiment seem robust to the influence of outlying levels of individual performance.

EXPERIMENT 3: QUALITATIVE EXPLORATION

Experiment 3 was a qualitative exploration of design issues for interactive combined 2D/3D displays. Our objectives were to:

- Contrast ExoVis and orientation icon for qualitative differences that might not have affected the overall times and errors in Experiment 1.
- Consider a more specific domain (volume visualization) to see whether our results would be similar to those for the abstract data sets.

We chose the flat OI display rather than the box OI display, permitting a contrast between 2D views at an angle (ExoVis) and 2D views straight on (flat OI).

Method

ExoVis and OI displays were compared (see Fig. 12). A within-subjects design was used so participants could compare the displays. Display order was counter-balanced.

Task and Stimuli

Participants positioned a box-shaped "volume of interest" (VOI) around an anomaly in a volume data set. Such VOI tasks are common in 3D imaging; they allow users to study interesting subregions separately from the volume as a whole. We chose a tomato data set so that detailed do-

main knowledge (e.g., medical knowledge) would not be necessary and university students could be participants.

Each display consisted of a 3D view and 3 orthogonal 2D slices, as shown in Fig. 12. The 3D view showed a semi-transparent white isosurface of the outside of the tomato plus a solid red isosurface of the anomaly. It also contained 3 “placeholder” planes that indicated slice positions. Participants could hide the placeholders or change their rendering style (solid, semi-transparent, or wireframe). The 3D view (for the orientation icon display) and the entire ExoVis scene (for the ExoVis display) could be rotated via mouse input. 2D views showed grayscale images of the current slices, where the anomaly appeared as a bright white spot. Slices could not be reoriented, but could be translated back and forth and hidden. A green wireframe box represented the VOI. Its position and size could be altered via sliders. Interaction methods were identical for both displays.

Our interaction techniques were largely separated from the display to make the interaction consistent and allow us to focus on display organization. We did not expect these interaction methods to be well-liked, but felt they would be sufficient to compare the two displays. Furthermore, we used participant’s discontent with the interaction method to initiate discussions about how they would like to interact with the displays they had experienced.

Procedure

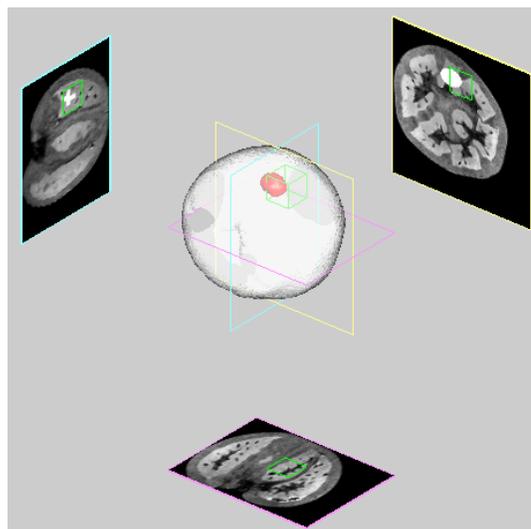
While participants performed the task, the experimenter asked questions to clarify what strategy participants were using, what parts of the display they were viewing, and any problems they were having. After participants tried the task with both displays, we conducted a semi-structured interview where we asked open-ended questions and how much participants agreed or disagreed with the following statements about each display:

- The display was...
 - Easy to learn
 - Clearly organized
 - Frustrating to use
- With the display it was easy to...
 - Complete the assigned task
 - Relate 2D and 3D views
 - Move the slices
 - Determine whether the anomaly was enclosed by the box
 - Adjust the box size

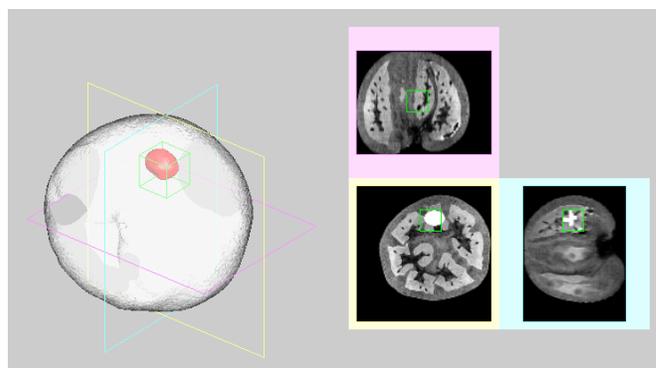
Participants

Six computer science or engineering graduate students (3 male and 3 female) and 2 computer science professors (1 male and 1 female) participated in the experiment. Half of the participants had taken part in Experiment 1 or a pilot study. All participants were previously known to the experimenter and were selected for their strong communication skills, to ensure the interview would be informative. Two of the participants were medical imaging specialists and at least one other participant had CAD experience. One of the participants is an author of this paper; at the time of the experiment, this participant was familiar with the over-

all goals of the research program, but not with the specific objectives of this study.



(a) ExoVis



(b) Orientation Icon

Fig. 12 Displays for Experiment 3. Placeholders are shown as wireframe planes. Participants positioned the green box around the VOI.

Results

Rating scale results are shown in Fig. 13. ExoVis was rated better on average than OI for all rating scale questions. In addition, 5 out of 8 participants said they preferred ExoVis overall. Of the other 3, 2 preferred OI and the other had no preference. Participants liked ExoVis because relating the 2D and 3D views was easier. Relying on colours alone (in the OI display) was possible but required more effort. One participant said ExoVis was especially helpful for relating views when the 3D view was rotated. A second major advantage of ExoVis was that 2D and 3D views were physically closer. Participants reported less eye movement between 2D and 3D views using ExoVis than using the orientation icon display. This was especially helpful when translating slice planes through the dataset. In addition, one participant commented that ExoVis was “more natural to use”, and a second participant said ExoVis gave a better feeling of control over the actions. The participant who had no display preference claimed that s/he did not use the 3D

view. Notice that viewing 2D slices obliquely (in ExoVis) was not considered a detriment by this participant.

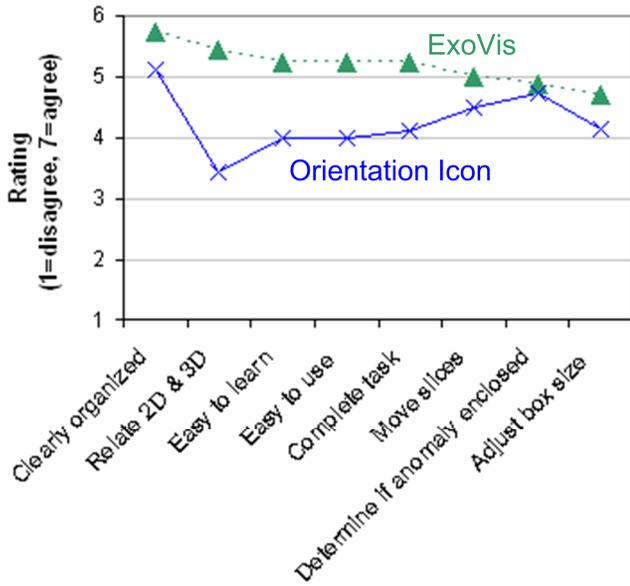


Fig. 13 Experiment 3 rating scale results, sorted by decreasing average rating for ExoVis.

We observed several different approaches to complete the VOI task, including:

- *Widest slice strategy*: first translate the 3 slices back and forth until each slice shows the largest possible white spot. Then adjust the box position and size so that it encloses the anomaly in these slices.
- *3D approximation strategy*: first use the 3D view to get the box in approximately the correct position. Then use the slices for fine-tuning and confirmation.
- *Slice approximation strategy*: first translate the slices so they show part of the anomaly. Translate the box to approximately the correct position. Rotate the 3D wireframe view and/or translate the slices to confirm box placement.
- *3D only strategy*: Rotate the 3D wireframe view so the camera points directly down one of the major axes. Adjust the box position so it is correct in the other 2 dimensions. Then repeat with the camera pointing down a different axis to complete the task.

Strategies that focused primarily on 2D views (e.g., widest slice strategy) worked quite well with both displays. Strategies that only used the 3D view a little also worked reasonably well (e.g., slice approximation and 3D approximation strategies). However, ExoVis caused serious problems for the 3D only strategy because of occlusion. Specifically, when the camera was positioned so it pointed along a major axis, either the 3D view occluded one of the slices or vice versa, making the task very difficult. The other two slices were seen from the side so they appeared as lines. Seeing slices from the side or from oblique angles was annoying to participants who used both 2D and 3D views and wished to view the slice contents. However, the participant

who chose the 3D only strategy actually found the lines helpful because the colours identified which slider would move the box in a particular direction (the placeholders also served this function but the participant found both together useful). This meant that s/he did not want to turn the slices off permanently. At the same time, s/he did not want to manually move slices or turn them on and off every time s/he changed the view orientation, and instead wanted them to move automatically as the camera was moved to reduce occlusion. Such an automated placement algorithm is an interesting topic for future work.

Displaying slices straight on was the main advantage of the OI display and the main disadvantage of ExoVis. Flat slices were considered useful for precise positioning (because of higher resolution and lower distortion) and for comparing more than one data set (because slices could be placed side by side). Three participants suggested that the best display would be ExoVis with an option to view slices straight on. However, participants disagreed on how to specify when straight-on viewing should be used; some participants suggested a mouse click, but others wanted a less intrusive mechanism. One participant suggested having both oblique and non-oblique slices visible simultaneously to reduce the need for mouse clicks; however, this would require extra screen space that may not be available. Hence, future studies should consider the best method of switching between ExoVis slices and non-oblique versions.

One important problem with the OI display was that slices were distant from the 3D view, so users had to make large eye movements. We had placed the slices in an L-shape to represent an open box (as in CAD multi-view projections); however, most participants did not realize this or find it useful. Instead, they suggested placing the slices in a vertical row or surrounding the 3D view (as in the ExoVis display except with slices flat on the screen).

Another important factor was interaction technique. Participants wanted direct manipulation for all scene components. Mode buttons (to specify which object was being manipulated) were annoying and distracted users from their task. Hence, we believe the best interaction technique would allow users to specify which object to manipulate by simply pointing at the object or a specific part (e.g., an edge). Implementing this type of direct interaction may be more challenging with ExoVis because there are more objects in the scene. Additional research is needed to determine how many objects could be placed in a scene before this interaction technique would no longer be manageable.

Our VOI was box-shaped to make user interaction simple. If we had compared strict 2D and 3D displays, this may have provided an advantage to the 2D display. However, since Experiment 3 considered only 2D/3D combinations, we do not believe this limitation would affect the choice of which display was better. Nonetheless, future experiments with oblique-shaped VOIs could be interesting.

DISCUSSION

To achieve experimental control and enable testing with novices, Experiments 1 and 2 used shapes and tasks much simpler and abstract than the CAD applications from which

they were derived. Furthermore, interactivity was restricted, which would be rare in any non-experimental situation. Experiment 3 more closely matched the task and data of actual applications and included greater interactivity. Together, all three experiments provide valuable insight into usability of combined 2D/3D displays. Note that several differences between the experiments (e.g., stimuli shape, volumetric vs. surface data, and orthographic views vs. slices) limit our ability to directly compare them.

As seen by comparing Fig. 4 and Fig. 10, most display types showed consistent performance in Experiments 1 and 2. The strict 3D display produced rapid but inaccurate performance. Performance of OI and ExoVis was approximately equivalent and was consistent for both tasks. They were either slower or the same speed as 3D, but were quite accurate. Performance of the 2D display had equivalent speed to OI and ExoVis displays but its accuracy was sensitive to the task. The 2D display was as inaccurate as the 3D display for the position estimation task but became as accurate as the combination displays for the orientation task.

The clip plane display performed uniformly poorly on the one task for which we evaluated it, orientation. The display forced users to physically and cognitively switch between individual slices and a complete 3D view. Because the task required integrating information from several slices, this switching was cumbersome and prohibited good performance. Clip planes may be more useful when only one slice is needed, when slices can be used sequentially, or when complete slices are unnecessary (e.g., when users can work with a small box cut out of a 3D scene). Given the clip plane's current widespread use, we recommend evaluation of this display for other tasks. If it is consistently outperformed by other displays, it may be productive to replace it by a display supporting better performance.

These results are subject to several caveats. Performance of 3D displays is apparently sensitive to the cues presented. In the position estimation task, the 3D display with shadows was fast but the version with two rotated 3D views produced slow performance. Likewise, for the 2D orthographic projections, additional cues such as hidden contours could improve performance. Future comparisons should specify the cues included in all display types.

Our second caveat concerns trial type in the orientation task. Although performance was stable for the ExoVis and orientation icon displays, performance of the strict 3D and 2D displays was sensitive to the orientation of the object displayed in the trial. We suggest that the stable performance of the two combination displays (orientation icon and ExoVis) demonstrates an advantage: By including elements of both 2D and 3D displays, the combination displays permit users to vary their strategy according to the needs and possibilities of a specific visualization.

Given that combination displays are useful, what is the best type of combination? Lack of large quantitative differences between ExoVis and orientation icon displays indicates that having both 2D and 3D may be more important than the method used to organize them on the screen. Nonetheless, our qualitative results provide insight into when each technique is valuable, as described below.

Orientation cues provided by ExoVis and OI box (i.e., OI display where the 2D views were arranged in a box shape, as in Experiments 1 and 2) methods were considered valuable, especially while participants were learning the tasks and displays. This was particularly important for users with little 3D graphics or CAD experience. Orientation cues were also important for understanding projections (e.g., for the block shapes in the position estimation task), relating the 3 DOF input device to the display in the orientation task, and for rapidly switching attention between 3D and 2D views.

At the same time, viewing slices at an angle sometimes made precise judgments challenging. This problem was more pronounced with interactive rotation because viewing angles could be very oblique. Interactive rotation also caused occlusion problems with ExoVis.

Personal preferences for the displays varied. A system that allows users to switch between ExoVis, OI box, and OI flat (i.e. OI display where the 2D views are flat on the screen, as in Experiment 3) displays may resolve these issues. ExoVis may also benefit from an automated placement algorithm that moves objects to reduce occlusion as the view angle changes.

Another important factor was proximity of views that were used simultaneously. For the orientation task, a few participants complained that with ExoVis they could not see more than one 2D slice without moving their eyes. This forced them to use a strategy that focused on one slice at a time. By contrast, for the VOI task, some participants felt ExoVis was better because slices were closer to the 3D view. These observations illustrate the importance of matching the display type to the task and strategy. For many participants, the orientation task was divided into two distinct phases: approximation with the 3D view followed by fine-tuning with the 2D views. Here, OI may be best because it separates the 3D and 2D views to match the strategy. Similarly, ExoVis may be better for tasks that require frequent switching between 3D and 2D views (e.g., to reposition slices and verify the box position in the VOI task).

Experiment 3 showed substantial variation in users' strategies and preferences. For example, some users could perform precise tasks with 3D views alone by using rotation, whereas others found this too confusing and preferred to use 2D views or a combination of 2D and 3D views. The wide variations in task strategy that we observed emphasize the importance of allowing personal customization.

CONCLUSIONS

Strict 3D displays with additional cues such as shadows can be effective for approximate relative position estimation and orientation. However, precise orientation and positioning are difficult with strict 3D displays except possibly in specific circumstances (e.g., with appropriate lighting, viewing angle, and measurement tools). For such precise tasks, combination 2D/3D displays (orientation icon and ExoVis) were better than strict 2D or 3D displays. Compared to strict 2D displays, combination displays performed as well or better, inspired higher confidence, and allowed more integrated navigation. Clip plane displays

were not effective for 3D orientation because it was difficult to use more than one slice at a time and challenging to integrate information from several slices. OI displays with flat 2D views were useful for some precise judgments, whereas OI box and ExoVis displays were useful for understanding projections, relating the display to a 3D input device, and for rapidly switching attention between 3D and 2D views. OI displays may be preferred when the task has distinct 2D and 3D phases, and ExoVis may be preferred when 2D and 3D are used closely together.

FUTURE WORK

Our results indicate that interactively switching between oblique ExoVis 2D views and non-oblique versions would be valuable. In future user studies, we would like to test several interaction techniques to achieve this. In addition, we would like to develop and test an automatic placement algorithm for ExoVis so that objects do not occlude each other upon interactive rotation.

Our experiments involved generalized tasks so the results could be relevant to many applications. Future experiments could examine how the results apply to a variety of tasks in different domains. Because of the high individual variability, further subjects should be studied to assess the impact of users' background and skills on their performance. For example, experiments with experts (e.g., physicians or CAD technicians) may yield different results than experiments with novice users. A within-subjects study could better control for individual variation in spatial ability and allow examination of learning effects, but with the cost that fewer displays could be included. Other future studies could match display types to interaction techniques and devices, since different interaction methods might work best with different displays. Furthermore, an eyegaze study would reveal the extent to which the 3D and 2D views were used in making the decisions.

ACKNOWLEDGMENT

We wish to thank members of the Imager laboratory at the University of British Columbia for loaning us the Polhemus, and Colin Swindells for customizing the input device. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

REFERENCES

- [1] M.S.T. Carpendale, D.J. Cowperthwaite, M. Tigges, A. Fall, and F.D. Fracchia, "The tardis: A visual exploration environment for landscape dynamics," *Proc. Visual Data Exploration and Analysis VI*, SPIE vol. 3643, 1999.
- [2] K.P. Herndon, R.C. Zelezniak, D.C. Robbins, D.B. Conner, S.S. Snibbe, and A. van Dam, "Interactive shadows," *Proc. User Interface Software and Technology*, pp. 1-6, 1992.
- [3] Y. Kurzion and R. Yagel, "Interactive Space Deformation with Hardware-Assisted Rendering," *IEEE Computer Graphics and Applications*, vol. 17, no. 5, pp. 66-77, 1997.
- [4] M.J. McGuffin, L. Tancau, R. Balakrishnan, "Using Deformations for Browsing Volumetric Data," *Proc. IEEE Visualization*, pp. 401-408, 2003.
- [5] S.H. Park and J.C. Woldstad, "Multiple Two-Dimensional Displays as an Alternative to Three-Dimensional Displays in Telerobotic Tasks," *Human Factors*, vol. 42, no. 4, pp. 592-603, 2000.

- [6] R.N. Shepard and J. Metzler, "Mental rotation of three-dimensional objects," *Science*, vol. 171, pp. 701-703, 1971.
- [7] H.S. Smallman, M. St. John, H.M. Oonk, and M.B. Cowen, "Information Availability in 2D and 3D Displays," *IEEE Computer Graphics and Applications*, vol. 21, no. 5, pp. 51-57, 2001.
- [8] R.R. Springmeyer, M.M. Blattner, and N.L. Max, "A Characterization of the Scientific Data Analysis Process," *Proc. IEEE Visualization*, pp. 235-242, 1992.
- [9] M. St. John, M.B. Cowen, H.S. Smallman, and H.M. Oonk, "The use of 2D and 3D displays for shape-understanding versus relative-position tasks," *Human Factors*, vol. 43, no. 1, pp. 79-98, 2001.
- [10] M. St. John, H.S. Smallman, T.E. Bank, and M.B. Cowen, "Tactical Routing Using Two-Dimensional and Three-Dimensional Views of Terrain," *Technical Report 1849, SSC San Diego Technical Reports*, 2001.
- [11] M. Tory, "Mental Registration of 2D and 3D Visualizations (An Empirical Study)," *Proc. IEEE Visualization*, pp. 371-378, 2003.
- [12] M. Tory, T. Möller, M.S. Atkins, and A.E. Kirkpatrick, "Combining 2D and 3D Views for Orientation and Relative Position Tasks," *Proc. CHI*, pp. 73-80, 2004.
- [13] M. Tory and C. Swindells, "Comparing ExoVis, Orientation Icon, and In-Place 3D Visualization Techniques," *Proc. Graphics Interface*, pp. 57-64, 2003.
- [14] K.F. Van Orden and J.W. Broyles, "Visuospatial Task Performance as a Function of Two- and Three-Dimensional Display Presentation Techniques," *Displays*, vol. 21, no. 1, pp. 17-24, 2000.
- [15] C.D. Wickens, D.H. Merwin, and E.L. Lin, "Implications of Graphics Enhancements for the Visualization of Scientific Data: Dimensional integrity, Stereopsis, Motion, and Mesh," *Human Factors*, vol. 36, no.1, pp. 44-61, 1994.
- [16] P.C. Wong and R.D. Bergeron, "Brushing Techniques for Exploring Volume Datasets," *Proc. IEEE Visualization*, pp. 429-432, 1997.

Melanie Tory is a post-doctoral fellow in the Imager lab at The University of British Columbia. Her research focuses on user-oriented design and evaluation of visualization tools, particularly for scientific applications. The work reported in this paper was part of her PhD thesis at Simon Fraser University. Melanie received a PhD in Computing Science from Simon Fraser University in 2004 and a BSc in Microbiology from the University of British Columbia in 1999.

Arthur E. Kirkpatrick is an assistant professor of computing science at Simon Fraser University. His research interests include haptic interfaces and interaction techniques. He is co-director of the Graphics, Usability, and Visualization Lab. He received his Ph.D. in Computer and Information Science from the University of Oregon in 2000.

M. Stella Atkins is a Professor in the School of Computing Science at Simon Fraser University, and Director of the Medical Computing Lab at SFU. Her research interests include medical image display and analysis, and computers in medicine. This includes medical image segmentation, brain atrophy studies, image registration, denoising and enhancement, and radiology workstation design. Additionally, she is interested in the use of eyetrackers, which provide a new imaging modality for seeing inside the brain. She received her PhD in Computer Science from the University of British Columbia in 1985, and her BSc degree in Chemistry from Nottingham University in 1966.

Torsten Möller is an assistant professor at the School Of Computing Science at Simon Fraser University. His research interests include Scientific Visualization and Computer Graphics, especially the mathematical foundations of visualization and graphics. He is co-director of the Graphics, Usability and Visualization Lab and serves on the Board of Advisors for the Centre for Scientific Computing at Simon Fraser University. He received his PhD in Computer and Information Science from Ohio State University in 1999 and a Vordiplom (BSc) in mathematical computer science from Humboldt University of Berlin, Germany.