

Combining 2D and 3D Views for Orientation and Relative Position Tasks

Melanie Tory

Torsten Möller

M. Stella Atkins

Arthur E. Kirkpatrick

School of Computing Science
Simon Fraser University
{mktory, torsten, stella, ted}@cs.sfu.ca

ABSTRACT

We compare 2D/3D combination displays to displays with 2D and 3D views alone. Combination displays we consider are: orientation icon (i.e., side-by-side), in-place methods (e.g., clip planes), and a new method called ExoVis. We specifically analyze performance differences (i.e., time and accuracy) for 3D orientation and relative position tasks. Empirical results show that 3D displays are effective for approximate navigation and relative positioning whereas 2D/3D combination displays (orientation icon and ExoVis) are useful for precise orientation and position tasks. Combination 2D/3D displays had as good or better performance as 2D displays. Clip planes were not effective for a 3D orientation task, but may be useful when only one slice is needed.

Author Keywords

2D and 3D visualization, display design, empirical study, experiment, orientation and relative position tasks.

ACM Classification Keywords

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

Both 2D and 3D visualizations are useful for analyzing 3D spatial data. Springmeyer et al. [6] observed that 2D views (i.e., slices or orthographic front/back, right/left, or top/bottom projections) are often used to establish precise relationships, whereas 3D views (other types of orthographic or parallel projections) are used to gain a qualitative understanding and present ideas to others.

Although displaying both 2D and 3D views is becoming more common, little research to compare and evaluate

different methods of combining 2D and 3D views has been done. Our research addresses the issue of which display technique(s) are best for specific situations and tasks.

In general, 2D views are good for seeing details of a particular part and navigating or measuring distances precisely (since only one dimension is ambiguous) [7, 8]. Three-dimensional displays are good for gaining an overview of a 3D space, understanding 3D shape, and navigating approximately in 3D [8, 14]. Since 3D and 2D views serve different purposes, having both visible may benefit certain tasks such as orienting and positioning objects relative to one another.

For example, radiologists typically view 3D medical scans as 2D slices to make details more apparent. They may also use a 3D view to gain an overall qualitative picture of the data, to explain ideas to other physicians, or to place slicing planes in non-standard orientations. Similarly, parts of a CAD model near the front can occlude parts at the back. For this reason, CAD models are usually displayed from several viewpoints at once, often from three standard orthogonal directions (2D orthographic views) plus one or more oblique viewpoints (to give an impression of 3D structure).

RELATED WORK

Methods to Combine 2D and 3D Views

In-Place Techniques

Clip planes show 2D slices in their correct position relative to the 3D view (the slice is “in-place”), so understanding relationships between views is easy. However, a clip plane removes all data between itself and the viewer, so information can be hidden. The “planar brush” [15] does not remove sections of the 3D view, but limits the 3D view to an outline or semi-transparent surface. Another alternative is to open up a volume along a cutting plane (e.g., [1, 2, 3]), so 3D view information is pushed aside but not removed.

Orientation Icons

With orientation icons (OI), 2D and 3D views are side-by-side. The 3D view “orients” users so they understand positions of 2D views. Because 2D views are physically separated from the 3D view and can be translated and rotated from their original location, understanding relationships between views can be challenging [10].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2004, April 24–29, 2004, Vienna, Austria.

Copyright 2004 ACM 1-58113-702-8/04/0004...\$5.00.

ExoVis

With ExoVis, the 3D view is placed in the centre of the display and 2D views are shown in the surroundings. 2D views may be translated but not rotated from their original orientation, so understanding the relationship between the 2D and 3D views requires more effort than in-place techniques, but less than orientation icons [10].

Comparisons of 2D/3D Displays

Many studies have compared 2D and 3D displays (e.g., [4, 7, 12]). However, they consider widely different applications and tasks, and vary greatly in terms of display parameters. A few research studies have elucidated more general principles about 3D and 2D displays. The “Proximity Compatibility Principle” (PCP) [14] suggests that tasks requiring integration of spatial dimensions (i.e., 3D knowledge) will benefit from 3D displays, whereas tasks requiring focused attention on one or two dimensions will benefit from 2D displays. Wickens et. al. determined experimentally that 3D (perspective) representations were better than 2D representations only for more integrative questions requiring knowledge of several dimensions [14]. Nonetheless, not all experimental evidence agrees with the PCP. Based on their experiments, St. John et al. [8] concluded that 3D displays were best for shape understanding tasks, whereas 2D displays were superior for relative positioning tasks, even when the positioning tasks required 3D knowledge. However, their relative positioning results may have limited value in practice because:

- Resolving ambiguity in their “3D” display required relating two 3D views via mental rotation, a challenging cognitive activity [5].
- Although shadows are known to be important for 3D positioning [13], no shadows were drawn.
- In one version of the experimental task, the “2D” display included a partial 3D view for reporting purposes, so the display was really a 2D/3D combination. In the other version of the task, participants simply had to click on the circle in the top 2D view, so the task did not really require 3D understanding of the scene.

Although many studies have compared 2D and 3D displays, few have considered combination displays or how to best create them. St. John et al. [9] compared (1) 3D views, (2) 2D views, and (3) a side-by-side combination of 2D and 3D views for a 3D route planning task. Time to complete the task was fastest with the side-by-side (i.e., orientation icon) display, indicating that combinations of 2D and 3D views are valuable. However, St. John et al. only looked at orientation icon displays and therefore did not consider the question of how to best combine 2D and 3D views.

In previous work [11], we heuristically compared OI, ExoVis, and in-place techniques, forming hypotheses about when each would be useful. We then showed that mentally relating 2D and 3D views was easiest with in-place

techniques, moderate with ExoVis, and most difficult with OI [10]. These studies were limited since they did not consider realistic tasks or complex displays, and did not consider 2D or 3D views alone. Our paper addresses these issues. We also present a variation of St. John et al.’s [8] relative position experiment that addresses its limitations and extends the results to include 2D/3D combinations.

STUDY SCOPE AND OBJECTIVES

We compared 2D, 3D, and several 2D/3D combinations for two tasks. We selected tasks predicted to benefit from combination displays: relative positioning and orientation.

Although the eventual goal is to study complex situations (e.g., many different views with many orientations), we limited the current study to one 3D view and three 2D views (oriented along the major axes of the space). In addition, 2D views could not be reoriented and the displays could not be rotated or otherwise changed. By strictly controlling interaction, we prevent it from becoming a confounding factor and avoid complicating the analysis with many variables. Also, axis-aligned 2D views are very common in many applications. Future studies with more complex visualization tasks will consider more complex displays.

METHOD

In the relative position task, participants estimated the height of a ball relative to a block shape (a purely perceptual task). In the orientation task, they adjusted the orientation of a plane relative to a torus (a perceptual and motor task). These tasks were selected for several reasons:

- They are common tasks in many application domains.
- They were predicted to benefit from a combination of 2D and 3D views.
- Objectives could be precisely defined and potentially confounding factors could be controlled.

Our design was between-subjects for display type and within-subjects for the position task ball height and the orientation task trial type. We measured time and accuracy.

Relative Position Task

This task was based on St. John et al.’s “over different” task [8]. Participants determined the position of a ball relative to a block shape (see Figure 1 on the colour plate). We made several changes to the “over different” task. We used a different block style so the number of possible ball positions was constant for all trials. In the St. John task, participants determined which block was directly beneath the ball, and reported their answer by clicking on a separate 3D view. Because a 3D view was available, the “2D” condition was really a 2D/3D combination. Using this reporting method would not allow us to compare a strict 2D condition to a 2D/3D combination. Instead, participants determined which block was directly beneath the ball, estimated the amount of empty space (vertically) between the ball and block shape, and reported this height as their answer. Correctly

identifying the height required understanding the ball's position relative to the block shape. This task also involves a truly 3D spatial relationship between the block and ball (rather than simply the 2D layout from a top-down view).

Stimuli

We compared five displays, as illustrated in Figure 1:

- A) **2D**: orthographic projections from the top, front, and right. These were arranged in a box shape so that their orientation matched the direction of projection.
- B) **3D rotated**: two 3D views, rotated 90° relative to one another (similar to St. John et al.'s [8] "3D" condition).
- C) **Orientation icon (OI)**: side-by-side 3D and 2D views. 2D views were arranged as in the 2D condition. This is the closest condition to St. John et al.'s "2D" condition (but their 2D views were not arranged in a box shape).
- D) **ExoVis**: 3D view with 2D projections surrounding it.
- E) **3D shadow**: one 3D view with a directional light centered above the block shape, so that a shadow of the ball projected directly downwards onto the block beneath it.

Scenes were modeled in Trispectives Technical v. 2.0. Following our previous work [10], blocks were generated by removing 2, 5, or 8 cubes from a base shape containing 27 cubes (3 x 3 x 3). Example block shapes are shown in Figure 1. Removal of cubes was purposely constrained so that all shapes would have similar overall form and complexity. Specifically:

- The resulting object remained as a single connected component and was not allowed to lose its 3 x 3 x 3 structure (i.e., no 3 x 3 slab was completely removed).
- Blocks were removed contiguously from either one or two locations, but not from more than two locations.
- Few blocks were removed from the "back" of the object where the shape's geometry would be hidden.
- Blocks were removed from the top down, and no blocks were removed from the bottom slab. Thus there were always 9 sub-blocks that could be beneath the ball.

A red sphere (the "ball") was positioned directly above one of the sub-blocks. The sphere's diameter was the width of a sub-block. Empty vertical space between the block shape and ball was 0, 0.5, 1, 1.5, or 2 sub-block sized units.

Stimuli were presented as static images. Like St. John et al. [8], 3D views were rendered with orthographic projection to make the figures appear as small objects (e.g., toys) viewed up close. 2D views were rendered from the top, right, and front sides of the object (third-angle projection). 2D views were kept in their original orientations, forming a box shape – to help participants understand and remember views.

Procedure

Participants completed 20 trials (4 with each of the 5 ball heights) with one of the 5 displays. Heights were in pseudo-

random order. The same sequence of block/ball scenes was shown for each display so only the type of view differed.

Participants reviewed instructions that explained the task and views and gave examples with answers. "2D" participants viewed instructions that included a 3D view (to help them understand how the 2D views were constructed), but were warned that the 3D view would not be present during the trials. Participants completed five practice trials (with answers provided) followed by 20 experimental trials. The experimenter helped participants through the practice trials to ensure they understood the task. Participants were left alone to complete the experimental trials.

Participants were instructed to be as accurate as possible. Breaks were permitted only between trials. Trials began when participants pressed a "Ready" button and ended when they answered with a button labeled "0.0", "0.5", "1.0", "1.5", or "2.0", as shown in Figure 1 (2D image). No time limit was imposed. Answers could not be changed. After each trial, participants reported their confidence in their previous answer. In a post-trial questionnaire, participants rated the difficulty of task components and gave their opinions of the display. Due to space limitations, only time and error data will be presented in detail here; other results will appear in a later publication.

Hypotheses

We predicted:

H1: 2D would be faster and more accurate than 3D rotated as found by St. John et al.[8].

H2: 3D shadow would be fastest, but estimating height would be difficult so there would be many errors.

H3: OI and ExoVis would have fewer errors than all other displays and be faster than 3D rotated and 2D displays.

Orientation Task

In the orientation task, participants used a 3 degree-of-freedom (DOF) input device to orient a plane relative to a torus, such that the torus was cut into two identical parts (as if slicing a bagel in half). This task was modeled after slicing plane orientation tasks in medical imaging and other volume data applications. For example, medical images of the chest area are usually aligned with the main axes of the body. However, because the heart is at an oblique angle, physicians often need to orient an oblique slice through the region in order to get a useful view. Orientation tasks are also common in other 3D graphics applications such as CAD (e.g., to set the angle of the roof of a house).

We chose an abstract task (a torus and plane) rather than a specific domain (e.g., medical imaging or CAD) to provide generic results that are hopefully applicable to many fields. In addition, this task requires only a simple display (minimizing conflicting factors), does not require domain knowledge, and has a clearly defined correct answer.

Stimuli

We compared five displays, as illustrated in Figure 2:

- A) **2D**: slices through the middle of the torus, perpendicular to the three major axes. Slices were arranged in a box shape so that their relative orientations were correct.
- B) **3D**: Projection of a 3D scene containing the torus and plane. Orthographic projection was used so the objects appeared as if they were viewed up close.
- C) **Orientation icon (OI)**: side-by-side 3D and 2D views. 2D views were arranged as in the 2D condition.
- D) **ExoVis**: 3D view with 2D slices surrounding it.
- E) **Clip Plane**: 3D view that could be cut through the middle with clipping planes perpendicular to the three major axes. Clip planes removed everything in front of them.

Torus orientations changed for each trial. Orientations were constrained such that they could not be around only one of the standard axes (so that the orientation was incorrect in more than one 2D view) and could not be more than 55° from horizontal (to prevent awkward physical positions) or less than 12° (so a minimum amount of work was required). These constraints were determined by trial and error.

Input Device

Figure 3 illustrates our custom input device. We used a 6 DOF Polhemus Fastrak to input plane orientation. Position data from the Fastrak was discarded. To improve stimulus/response compatibility between the display and input device, we attached the Polhemus sensor to a square piece of plywood. The orientation of the plywood directly mapped to the orientation of the red plane.

With the clip plane, having all the slices on at once hides most of the display, so we believed participants would want to turn the slices on and off. In addition, participants needed a way to easily start and end the trials. To accomplish this, we positioned a 3-button mouse on the plywood; the three slices could be turned on and off using the three buttons. Mouse buttons were labeled with colours to match the colours of slices on the display. The mouse ball was removed to make room for the Polhemus sensor and make only the buttons functional. Mouse buttons were wired to the same buttons of a 2nd 3-button mouse; the regular mouse could be used for ordinary mouse actions while the custom input device could record mouse clicks.

Procedure

Each participant used one of the 5 displays and completed 16 trials. Orientations were in pseudo-random order. Trials were identical for all displays. Participants reviewed instructional materials that explained the task and the views. Instructions for the 2D condition included a 3D view to help participants understand how the 2D views were constructed; these participants were warned that the 3D view would not be present during the trials. A demonstration of the software and input device were performed. Participants then

completed five practice trials (with help if necessary) followed by 16 experimental trials (where the experimenter observed but did not offer help).

Participants were asked to be as accurate as possible. Breaks were permitted only between trials. Trials started with a single click and ended with a double-click. No time limit was imposed. After each trial, participants estimated their error (in degrees) and typed this number at a prompt. In a post-trial questionnaire, participants rated the difficulty of task components and gave opinions of the display.

Hypotheses

We predicted:

H4: 3D would be fastest but least accurate.

H5: 2D and 2D/3D combination displays would be equally accurate, but combination displays would be faster since 3D supports approximate navigation.

Experimental Set-Up

Custom experimental software was run on a dual Pentium II computer with 128 MB of memory and 1024 x 768 display resolution. Participants interacted with the computer using a mouse (relative position task), and our specialized input device (orientation task). The keyboard was used only to type estimated error values for the orientation task.

Participants

Forty volunteers were recruited from various levels of the computing science and engineering student populations at our university (from first year to graduate level). Participants had varied levels of experience with computer graphics. Mean age was 27.

Overall Procedure

Participants were randomly assigned to one of five displays (A through E). Each participant completed both tasks with one of the five display types. The position task “3D rotated” display was paired with the orientation task “3D” display, and the “3D shadow” display was paired with the “clip plane”. Between-subjects design was chosen to prevent carry-over effects (e.g., learning or fatigue) from one display to the next. Order of the two tasks was counterbalanced. An equal number of males and females participated in each display group and task order. Participants completed all parts of the first task (including follow-up questionnaires) before beginning the second task. Experimental sessions lasted approximately 45 minutes.

RESULTS

We analyzed our results by analysis of variance (ANOVA). When necessary, we used the Huynh-Feldt correction. We used Tukey HSD post-hoc tests for between-subjects variables and pairwise comparisons for within-subjects variables.

Relative Position Task

We used 5 x 5 (ball height x display) ANOVA.

Position Task Timing Data

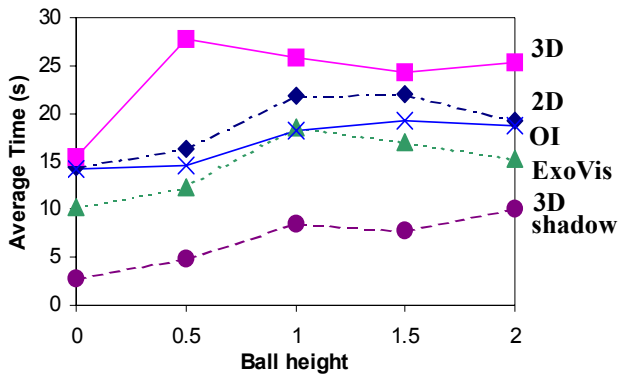


Figure 4: Timing data for the relative position task

Average trial time is shown in Figure 4. ANOVA found significant main effects for height ($F(4,140)=21.9$, $p<0.001$, $\eta_p^2=0.385$) and display ($F(4,35)=8.3$, $p<0.001$, $\eta_p^2=0.485$) and a significant interaction between height and display ($F(16,140)=2.5$, $p=0.002$, $\eta_p^2=0.225$). 3D shadow was significantly faster than all other displays ($p<0.015$). Height 0 was significantly faster than all other heights ($p<0.003$), and height 0.5 was faster than heights 1 and 2 ($p<0.01$).

Position Task Errors

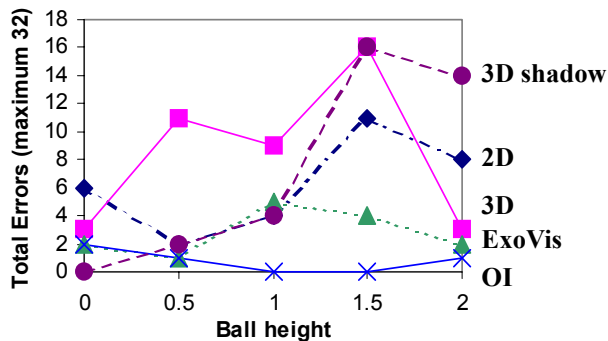


Figure 5: Errors for the relative position task

Total errors (sum over all participants) are given in Figure 5. ANOVA found significant main effects for height ($F(3.7,128.5)=6.27$, $p<0.001$, $\eta_p^2=0.152$) and display ($F(4,35)=5.1$, $p=0.002$, $\eta_p^2=0.367$) and a significant interaction between height and display ($F(14.7,128.5)=3.3$, $p<0.001$, $\eta_p^2=0.273$).

OI had significantly fewer errors than 3D and 3D shadow ($p<0.041$) and marginally significantly fewer than 2D ($p=0.056$). ExoVis had significantly fewer errors than 3D ($p=0.048$). Height 1.5 had significantly more errors than heights 0 and 0.5 ($p=0.004$), and marginally significantly more errors than 2.0 ($p=0.089$). There were no significant differences between displays for height 0.0. At height 0.5, 3D had significantly more errors than all other displays ($p<0.026$). At 1.0, 3D had more errors than OI ($p=0.03$). Displays were most different at height 1.5, where OI was better than 3D and 3D shadow ($p=0.002$) and marginally significantly better than 2D ($p=0.062$), and ExoVis was

marginally significantly better than 3D and 3D shadow ($p=0.09$). At height 2.0, OI was significantly better than 3D shadow ($p=0.035$) and ExoVis was marginally significantly better than 3D shadow ($p=0.054$).

Position Task Results Summary

Our results support H1, H2, and H3.

Orientation Task

From observations and participants' comments, we suspected that in the 3D condition, trials where the side of the torus was visible were easiest. Participants knew the plane was aligned when it became a simple line and/or aligned with the symmetry of the torus. Thus, we divided our trials into 3 types: (1) side trials (the torus hole was not visible), (2) top trials (full extent of the hole was visible), and (3) other trials (the hole was partially visible). (Note that the three trial types were distributed relatively evenly over the duration of the experiment.) We then analyzed our results by 3 x 5 (trial type x display) ANOVA.

Orientation Task Timing Data

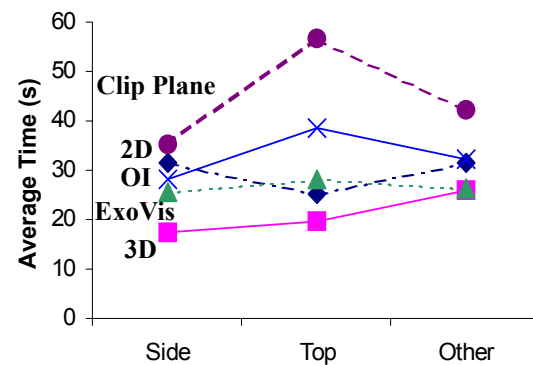


Figure 6: Timing data for the orientation task

Average trial time is summarized in Figure 6. ANOVA found significant main effects for trial type ($F(1.9,67.8)=4.9$, $p=0.011$, $\eta_p^2=0.124$) and display ($F(4,35)=3.9$, $p=0.01$, $\eta_p^2=0.308$) and a significant interaction between trial type and display ($F(7.7,67.8)=4.5$, $p<0.001$, $\eta_p^2=0.340$).

The clip plane took significantly longer than the 3D display ($p=0.008$), and marginally significantly longer than ExoVis ($p=0.059$). Side trials were faster than other trials ($p=0.012$). Differences between displays were only significant for top trials, where clip planes took significantly longer than 2D, 3D, and ExoVis ($p<0.034$). The difference between OI and 3D was marginally significant ($p=0.052$). There were marginally significant differences between 3D and clip plane ($p=0.051$) and 3D and 2D ($p=0.095$) for side trials. For 2D displays, top and other trials were significantly different ($p=0.046$). For 3D displays, other trials were different from both top and side trials ($p<0.008$). For clip planes, top trials were different from side and other trials ($p<0.009$). No significant trial type differences were found for OI or ExoVis displays.

Orientation Task Error

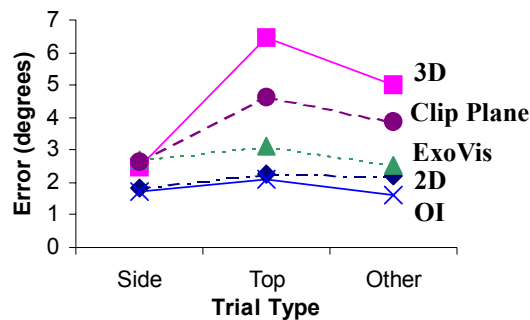


Figure 7: Error for the orientation task

Error data are given in Figure 7. ANOVA found significant main effects for trial type ($F(2,70)=15.4$, $p<0.001$, $\eta_p^2=0.305$) and display ($F(4,35)=7.3$, $p<0.001$, $\eta_p^2=0.454$) and a significant interaction between trial type and display ($F(8,70)=3.6$, $p=0.002$, $\eta_p^2=0.291$).

The 3D display had significantly more error than 2D and OI ($p<0.006$). Clip plane had more error than OI ($p=0.008$) and marginally significantly more error than 2D ($p=0.052$). Side trials had less error than top and other trials ($p<0.002$). No significant differences between displays were found for side trials, suggesting that 3D can be just as accurate as 2D and combination displays for these trials. However, for top and other trials, 3D displays were significantly worse than all displays except clip planes ($p<0.05$). Clip planes were worse than OI for other trials ($p=0.002$) and marginally significantly worse for top trials ($p=0.069$). Trial types were only significantly different for 3D and clip plane displays. For 3D displays, side trials had significantly less error than all other types ($p<0.001$), and for clip planes, side trials had less error than top trials ($p=0.021$) and marginally significantly less error than other trials ($p=0.075$).

Orientation Task Results Summary

Our results supported H4. The 3D display was consistently fastest but significantly less accurate than other displays. H5 was only partially supported. We found no significant accuracy differences between 2D, OI, and ExoVis, as expected, but times were also not significantly different.

DISCUSSION

Position Task

Combination 2D/3D and 3D shadow displays had higher likeability and lower difficulty ratings than 3D rotated or 2D displays. 3D rotated was difficult because estimating height was difficult and mental rotation required effort. Also, in each view, the ball position was ambiguous. One possible ball position could visually dominate over other possibilities, and was not always correct. Hence, we agree with St. John et al.[8] that 3D rotated displays are not effective for relative positioning. However, our results indicate that shadows can make 3D displays quite effective. Almost all errors for 3D shadow were only 0.5 units from the correct answer, indicating that people misjudged the ball

height, but not which block the ball was above. With the addition of a measuring tool, we believe people could be very accurate with 3D + shadow displays.

Nevertheless, shadows would not always be effective because the light must be placed in a very specific location relative to the objects of interest (e.g., if the light was slightly off to the side or one object was not directly above the other the light may be less effective). In addition, shadows can be difficult to interpret and costly to render in scenes with very complex geometry, such as the airplane model in Figure 8. Thus it is useful to consider 2D displays as an alternative way of resolving position ambiguity. Our results indicate that combination 2D/3D displays are a better choice than 2D alone for relative positioning tasks and should be chosen when 3D + shadow displays are not practical and/or 3D measurement tools are unavailable.

Orientation Task

Our results indicate that 3D displays are best for navigating approximately (within about 5 degrees) but not precisely. Interestingly, accuracy differences did not appear on side trials. It appears that 3D can be just as accurate as 2D and 2D/3D combinations when a “good” view can be found. However, finding a good view requires extra time and users may not know what view is best, so we are unsure whether 3D would outperform 2D and combination displays if this action were allowed. In addition, we expect a “good” view would only be possible when the model has symmetry or other defining features that can be seen from one side and used to align the slicing plane.

Although there were no significant time or accuracy differences between 2D and 2D/3D combinations, we observed that OI and ExoVis participants found an approximate solution more quickly. Participants using the 2D display did not seem to have a natural understanding of how to progress towards their goal. Thus (at least for novice users) we believe 2D/3D combinations are easier and better than 2D for approximate navigation, even though they do not show performance improvements for precise navigation.

Clip planes were worse than 2D/3D combination and 2D displays (slower overall and less accurate for top and side trials). Most clip plane participants could only work with one slice at a time (more than one slice was confusing or left too little information). Turning slices on and off was distracting and time consuming, and adjusting one dimension often accidentally changed others.

Overall Discussion

Because of our small group size (8 participants), a few spatially gifted or challenged participants could strongly influence our results. To check for this, we first compared the groups' experience with 3D computer graphics. All groups were similar, except the ExoVis group had slightly lower experience. We also checked overall trends after removing outliers (subjects with high or low performance).

For the position task, 2D became slower than 3D rotated, but there were no other changes. Therefore, we believe our conclusion that 2D/3D combination displays are better than 2D or 3D rotated displays is sound. If anything, with larger subject groups we expect participants to perform better with ExoVis and worse with 2D displays. For the orientation task, removing outliers gave a faster time for ExoVis (similar to the 3D display). This matches our prediction that ExoVis participants might perform more poorly than other groups because of less 3D graphics experience.

We had hoped our experiments would show a significant difference between ExoVis and orientation icon displays, to help designers decide when to use each type of combination. The fact that no such difference was found indicates that having both 3D and 2D together is more important than the method of organizing them on the screen. Presumably, any difficulty understanding the display organization is overshadowed by the difficulty of the task. Differences between these displays may be more pronounced when more freedom to interact with the display is allowed (e.g., the ability to rotate the 3D view). Running the experiments with larger sample sizes (for greater statistical power) may also elucidate clear differences.

We should note that time/accuracy trade-offs are a possible confound in our experiment. To reduce this risk, we encouraged all groups of participants to be as accurate as possible, and gave feedback on practice trials so they would know the correct answers. Our results also indicate this was not a major factor since accurate displays were often not slowest and inaccurate ones were often not fastest.

CONCLUSIONS AND FUTURE WORK

We conclude that 3D displays with appropriate cues (e.g., shadows) can be most effective for approximate navigation and relative positioning. However, precise navigation and positioning are not possible with 3D displays except in specific circumstances. Combination 2D/3D displays (orientation icon and ExoVis) were useful for precise orientation and position tasks. Compared with 2D displays, combination displays had as good or better performance. Clip planes were not effective for a 3D orientation task, but may be useful when only one slice is needed.

We are planning studies to compare the displays with more complex interactivity. In addition, our current experiment used generalized tasks so the results could be relevant to more than one application domain. Our future experiments will consider more domain-specific tasks to verify that the results really are generalizable.

ACKNOWLEDGEMENTS

Thanks to the Imager lab at the University of British Columbia for loaning us the Polhemus. Colin Swindells developed our custom input device. This project was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

1. Carpendale, M. S. T., Cowperthwaite, D. J., Tigges, M., Fall, A., Fracchia, F. D. The tardis: A visual exploration environment for landscape dynamics. *Proc. Visual Data Exploration and Analysis VI*, SPIE vol. 3643 (1999).
2. Kurzion, Y., Yagel, R. Interactive Space Deformation with Hardware-Assisted Rendering. *IEEE Computer Graphics and Applications*, 17, 5 (1997), 66-77.
3. McGuffin, M.J., Tancau, L., Balakrishnan, R. Using Deformations for Browsing Volumetric Data. *Proc. IEEE Visualization*, 2003, 401-408.
4. Park, S.H., Woldstad, J.C. Multiple Two-Dimensional Displays as an Alternative to Three-Dimensional Displays in Telerobotic Tasks. *Human Factors*, 42, 4 (2000), 592-603.
5. Shepard, R.N., Metzler, J. Mental rotation of three-dimensional objects. *Science*, 171, (1971), 701-703.
6. Springmeyer, R.R., Blattner, M.M., Max, N.L. A Characterization of the Scientific Data Analysis Process. *Proc. IEEE Visualization*, 1992, 235 -242.
7. Smallman, H., St. John, M., Oonk, H., Cowen, M. Information Availability in 2D and 3D Displays. *IEEE Computer Graphics and Applications*, 21, 5 (2001), 51-57.
8. St John, M., Cowen, M.B., Smallman, H.S., Oonk, H.M. The use of 2D and 3D displays for shape-understanding versus relative-position tasks. *Human Factors*, 43, 1 (2001), 79-98.
9. St. John, M., Smallman, H.S., Bank, T.E., Cowen, M.B. Tactical Routing Using Two-Dimensional and Three-Dimensional Views of Terrain. *Technical Report 1849*, SSC San Diego Technical Reports, 2001.
10. Tory, M. Mental Registration of 2D and 3D Visualizations (An Empirical Study). *Proc. IEEE Visualization*, 2003, 371-378.
11. Tory, M. and Swindells, C. Comparing ExoVis, Orientation Icon, and In-Place 3D Visualization Techniques. *Proc. Graphics Interface*, 2003, 57-64.
12. Van Orden, K.F., Broyles, J.W. Visuospatial Task Performance as a Function of Two- and Three-Dimensional Display Presentation Techniques. *Displays*, 21, 1 (2000), 17-24.
13. Wanger, L.R., Ferwerda, J.A. and Greenberg, D.P. Perceiving spatial relationships in computer-generated images. *IEEE Computer Graphics and Applications*, 12, 3 (1992), 44-58.
14. Wickens, C.D., Merwin, D.H., Lin, E.L. Implications of Graphics Enhancements for the Visualization of Scientific Data: Dimensional integrality, Stereopsis, Motion, and Mesh. *Human Factors*, 36, 1 (1994), 44-61.
15. Wong, P.C., Bergeron, D. Brushing Techniques for Exploring Volume Datasets. *Proc. IEEE Visualization*, 1997, 429-432.




Figure 1: Displays for the relative position task. Participants judged the vertical distance between the block and ball. The correct answer is 1.0. The 2D image shows a partial screen shot. Participants answered via a button at the bottom of the screen.

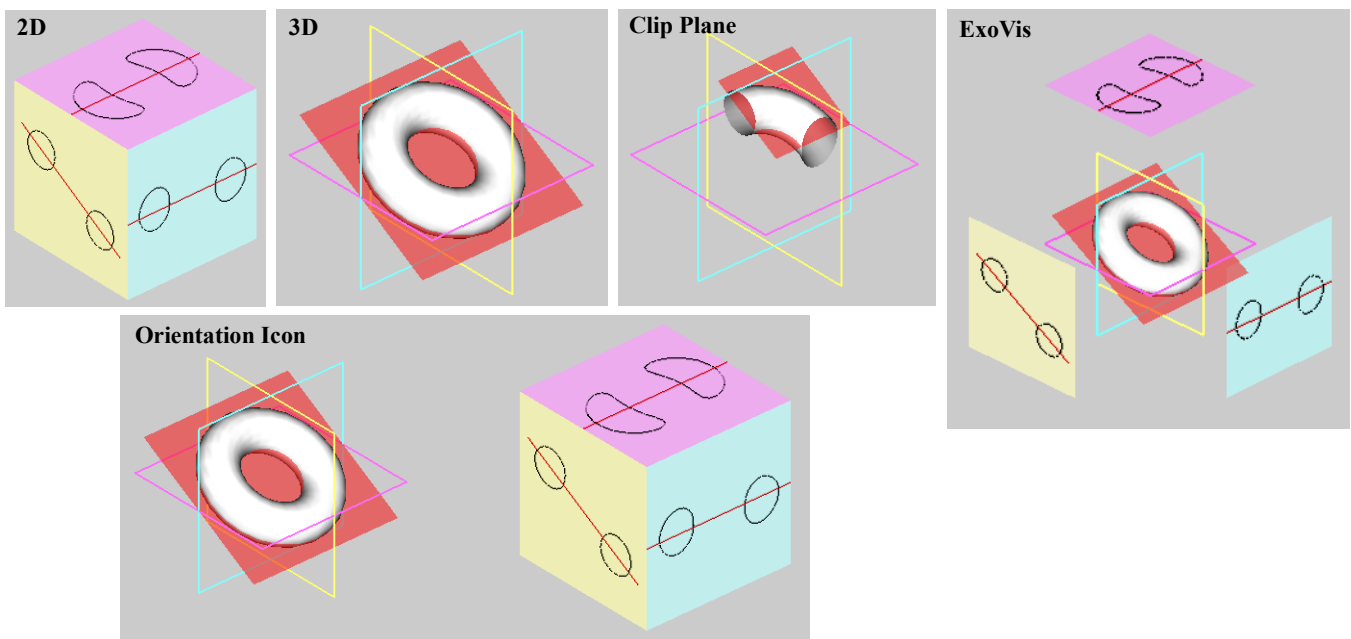


Figure 2: Displays for the orientation task. Participants aligned the red plane with the middle of the torus (correct solutions shown). 2D Views show cross-sections through the center of the torus and red plane. Coloured lines in 3D views indicate slice positions.




Figure 3: Custom input device for the orientation task. Tilting the board tilted the red plane. Buttons started / ended trials and turned slices on and off.



Figure 8: Virtual reality environment displaying airplane components. Complex geometry makes shadow generation costly and shadow interpretation difficult.