Evaluating the Effectiveness of Orientation Indicators with an Awareness of Individual Differences

TINA ZIEMEK, SARAH CREEM-REGEHR, WILLIAM THOMPSON, and
ROSS WHITAKER, University of Utah

Understanding how users perceive 3D geometric objects can provide a basis for creating more effective tools for visualization in applications such as CAD or medical imaging. This article examines how orientation indicators affect users' accuracy in perceiving the shape of a 3D object shown as multiple views. Multiple views force users to infer the orientation of an object and recognize corresponding features between distinct vantage points. These are difficult tasks, and not all users are able to carry them out accurately. We use a cognitive experimental paradigm to evaluate the effectiveness of two types of orientation indicators on a person's ability to compare views of objects presented in different orientations. The orientation indicators implemented were colocated, which shared a center-point with the 3D object, or noncolocated with (displaced from) the 3D object. The study accounts for additional factors including object complexity, axis of rotation, and users' individual differences in spatial abilities. Our results show that an orientation indicator helps users in comparing multiple views, and that the effect is influenced by the type of aid, a person's spatial ability, and the difficulty of the task. In addition to establishing an effect of an orientation indicator, this article helps demonstrate the application of a particular experimental paradigm and analysis, as well as the importance of considering individual differences when designing interface aids.

Categories and Subject Descriptors: H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Graphical user interfaces (GUI); 3D graphics*; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Psychology*

General Terms: Experimentation, Measurement, Human Factors

Additional Key Words and Phrases: Mental rotation, cognitive support, methodology, visualization

ACM Reference Format:

Ziemek, T., Creem-Regehr, S., Thompson, W., and Whitaker, R. 2012. Evaluating the effectiveness of orientation indicators with an awareness of individual differences. ACM Trans. Appl. Percept. 9, 2, Article 7 (June 2012), 23 pages.
DOI = 10.1145/2207216.2207218 http://doi.acm.org/10.1145/2207216.2207218

1. INTRODUCTION

Three-dimensional (3D) visualization systems allow users to view the projections of three-dimensional objects on two-dimensional (2D) displays. Typically, information about object shape is conveyed using standard pictorial cues such as shading, perspective, and interposition, but may also involve animations depicting motion. Extracting relevant information from 3D visualizations in applications such as computer-aided design (CAD) and medical imaging is a difficult task for some users, and there is

This material is based on work supported by the National Science Foundation under grants 0745131 and 0914488.
Contact author: T. Ziemek; e-mail address: tziemek@cs.utah.edu.

DOI 10.1145/2207216.2207218 http://doi.acm.org/10.1145/2207216.2207218

evidence suggesting that not everyone may benefit from the advantages of 3D visualization as they are currently designed [Keehner et al. 2008; Khan et al. 2008]. This article addresses one aspect of the perception of 3D visualizations, namely the problem of accurately perceiving visualizations that are shown in multiple, simultaneous views. Multiple static views allow users to simultaneously view an object from different viewpoints and allow features to be seen that would otherwise be occluded [Tory 2003]. While technology advances have enabled widespread use of dynamic and interactive visualizations, there are numerous circumstances in which users rely on multiple static views. For example, designers, architects, and engineers rely on multiple views in several state-of-the-art 3D software applications, including Unity 3D. In Unity 3D, in order to accurately place or position a camera or in-scene object, users must correctly interpret the spatial relationships of a scene that are shown from multiple 2D projections. Interfaces of this sort, where projections at different angles must be reconciled against other views or a mental model, are found in a range of applications in design, engineering, and medicine. Given that the images are not inherently 3D, a series of cognitive tasks are needed for the viewer to form an accurate mental representation of the object or space. These tasks include establishing a correspondence between perspectives and keeping track of an object's features between views. It can also be difficult to see when object features have changed in morphology from one data set to another, and this can be extremely challenging if the viewpoints also have changed. Understanding the relative orientation in a pair of views may help with this complexity. Findings from basic research in perceiving object constancy show that individuals do understand the identity of multiple noncanonical views of objects, but that object recognition of different views can be difficult, particularly given rotations in depth that lead to occlusion of parts and foreshortening of the object [Lawson and Humphreys 1998].

The goal of the present work is to examine how two types of orientation indicators could be used to improve a user's ability to compare and comprehend multiple views of a 3D visualization. Orientation indicators are in-scene graphical aids that illustrate rotational changes of an object and the user's current viewing direction [Khan et al. 2008]. Prior research has raised awareness of the difficulties that individuals may have when working in a virtual space [Fitzmaurice et al. 2008; Tory 2003; Tory and Möller 2004], and orientation indicators are one solution to help viewers maintain orientation in a desktop environment [Khan et al. 2008; Stull et al. 2009]. Orientation indicators may provide users with cognitive support, which can be defined as assistance from an artifact to help a user to think and solve problems [Walenstein 2002], and free cognitive resources that modeling and visualization applications may unnecessarily impose. This view follows a theory of distributed cognition [Zhang and Norman 1994] that external representations in the environment are a critical counterpart to the internal processes of the mind.

We contribute to the existing work on graphical aids for visualization by focusing on the importance of evaluation, as there is no guarantee that users will benefit from even the most well-intentioned and technically developed tools [Tory and Möller 2005]. This article presents an evaluation of orientation indicators with an established cognitive experimental paradigm, *mental rotation* [Shepard and Metzler 1971; Vandenberg and Kuse 1978]. Decades of work in spatial cognition has demonstrated that visuospatial thinking and mental representation can be systematically evaluated [Shepard and Metzler 1971; Shepard and Cooper 1982; Zacks and Michelon 2005; Hegarty and Waller 2005]. Mental rotation tasks are most commonly used to examine the mechanisms underlying spatial reasoning and the internal construct of mental imagery, however many researchers have had success in using the mental rotation paradigm to evaluate the perception of computer graphics (e.g. [Barfield et al. 1988; Bülthoff and Edelman 1992; Rizzo and Buckwalter 1997; Tarr 1995; Hubona et al. 1997; Ziemek et al. 2008]). This methodology allows us to use objective, controlled experimentation to evaluate the influence of an orientation indicator on the perceived orientation of a 3D object. This paradigm also allows

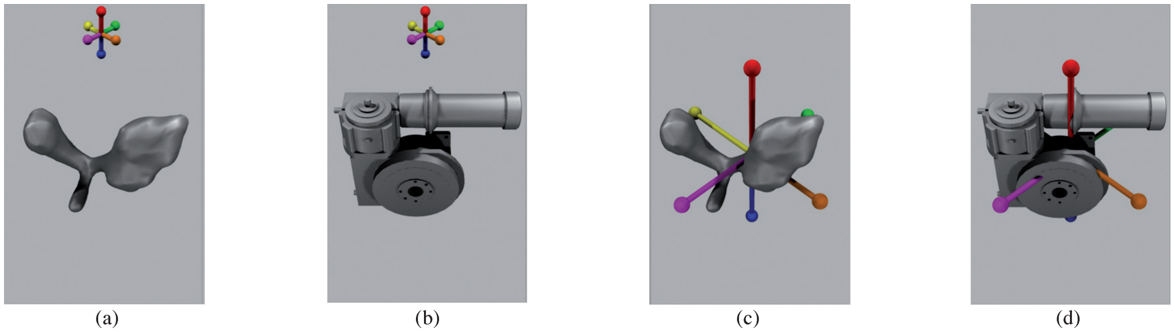

Fig. 1. (a) Noncolocated orientation indicator, applied to an anatomical-like object; (b) noncolocated orientation indicator, applied to a mechanical part; (c) and (d) colocated orientation indicators applied to the same objects.

us to examine several task factors that may influence the effectiveness of an orientation indicator. We specifically examine the effectiveness of orientation indicators that can be noncolocated (displaced) from the object or colocated with the target object (see Figure 1).

2. BACKGROUND

Research has shown that techniques can be implemented to address the challenges that arise when using computer graphics software, such as the development of haptic displays [Brooks, Jr. et al. 1990]; the addition of multiple viewpoints [Tory 2003]; and the augmentation of existing navigation tools [Fitzmaurice et al. 2008]. The use of orientation indicators as aids to 3D visualization has received some, but limited, attention. Khan et al. [2008] implemented an orientation indicator in CAD software, but did not compare task-performance between conditions when the indicator was present and when it was not present. More recently, Stull et al. [2009] found that orientation references helped individuals learn anatomy from a 3D visualization. Orientation indicators have been implemented in medical imaging software in the forms of bounding boxes, virtual human figures, and aids that depict the left, right, anterior, and posterior sides of an object [Cox 1996; Gering et al. 2001]. However, these orientation indicators have not been quantitatively assessed.

2.1 Variables that May Affect 3D Visualizations

Despite the enthusiasm regarding the use of computer generated images, research on how users perceive digital representations and the information users gain from a 3D visualization is to date limited and inconsistent [Huk 2006; Keehner et al. 2008]. Some research has shown that the presence of a 3D visualization is beneficial [Garg et al. 2001; Keehner et al. 2008], whereas other research shows that sometimes 3D visualizations do not provide extra information that is used effectively, such as in anatomy learning tasks [Garg et al. 1999; Huk 2006]. Previous research has indicated that when individuals from a broad population were assigned a shape-related task that entailed interaction with a 3D visualization, failure correlated with the inability to find an appropriate orientation from which to view the data [Keehner et al. 2008]. It is possible that some individuals cannot access information from a visualization because they get disoriented when working in an abstract virtual environment. The concepts and tools needed to maintain orientation in a computer-generated environment may be difficult to learn, and some users have even rejected using 3D computer graphics applications [Fitzmaurice et al. 2008; Khan et al. 2008]. Keehner et al. [2008] found that not all users were able to find the most "informative view"; that is, the view that gives key information within a visualization.

Similarly, Velez et al. [2005] reported that some individuals thought the most "informative view" was always the back projection of the object even if it was a side or bottom view.

There are a variety of possible factors that influence task-performance with a 3D visualization. Task-goals may influence the cognitive processes used. For example, recalling a view of a familiar object may involve different mechanisms of memory and attention compared to a task requiring the detection or comparison of a specific feature of a complex object. The nature and complexity of the models used is also likely an important factor. While most research that has evaluated visualization and human performance has tested only one category of model, such as carpal bones [Garg et al. 1999, 2001] or vertebrae [Stull et al. 2009], Velez et al. [2005] found that complexity of a 3D object (quantified by counts of object surfaces, object edges, and object vertices) affected 3D object comprehension. Finally, user interaction with an application is a factor that has received some attention. While it may seem intuitive that interactivity would help in 3D visualization tasks, several studies have found that individuals with interactive control perform the same or worse than individuals with an animated visualization which they cannot control [Hegarty 2004; Keehner et al. 2008; Sando et al. 2009].

Individual spatial abilities may also affect the utility of a 3D visualization. The term *spatial abilities* refers to a broad range of skills involving the mental representation and manipulation of information about geometric entities. Research has shown there is a natural variation between people in their spatial abilities; individual differences have been found in a variety of tests of visuospatial abilities [Hegarty and Waller 2005]. There are several possible alternatives for how spatial ability might influence the use of a 3D visualization [Hegarty 2004; Huk 2006; Keehner et al. 2008]. First, it may be that high spatial ability is a necessary prerequisite to using a 3D visualization. There is some evidence that high spatial ability is correlated with accuracy with a 3D visualization [Velez et al. 2005]. A three-dimensional tool improved learning for high spatial ability individuals [Garg et al. 2001], but put low spatial ability individuals at a significant disadvantage [Garg et al. 1999]. Huk [2006] found that only students with high spatial abilities benefited from 3D objects. Those with low spatial abilities also had more difficulty with complex geometric objects than those with high spatial abilities. Velez et al. [2005] reported that low spatial ability participants could only solve simple geometric objects such as cubes and cones. A second alternative is that those with low spatial ability would benefit most from 3D visualizations; if these individuals find it difficult to construct internal models, they might benefit from the structure of an external model. The work of Stull et al. [2009] supports this claim, showing that spatial understanding of 3D objects by low spatial individuals can be improved to near that of high spatial individuals with the use of cognitive aids [Stull et al. 2009]. A third possibility is that 3D visualizations will benefit everyone equally.

Several researchers have argued for the importance of considering individual differences in the design of human computer interaction systems [Dillon and Watson 1996; Chen et al. 2000; Modjeska and Chignell 2003; Dünser et al. 2006]. For instance, an investigation of gender differences in 3D computer graphic applications concluded that the purported poor performance of women compared to men in navigating virtual environments disappeared if users were provided with a wide field of view display [Czerwinski et al. 2002]. Moreover, it has been found that features added to make 3D computer graphic applications more accessible are not only popular with novice users but experienced users as well [Fitzmaurice et al. 2008]. Work to date on designing other forms of software with an awareness of the effects of individual differences is also limited [Huff 2002; Beckwith et al. 2005; Beckwith et al. 2006].

2.2 Tasks Where Multiple Views Could Challenge Users

Multiple views are both common and useful, but they also present the user with additional challenges by introducing additional complexity over a single view [Baldonado et al. 2000]. Previous research

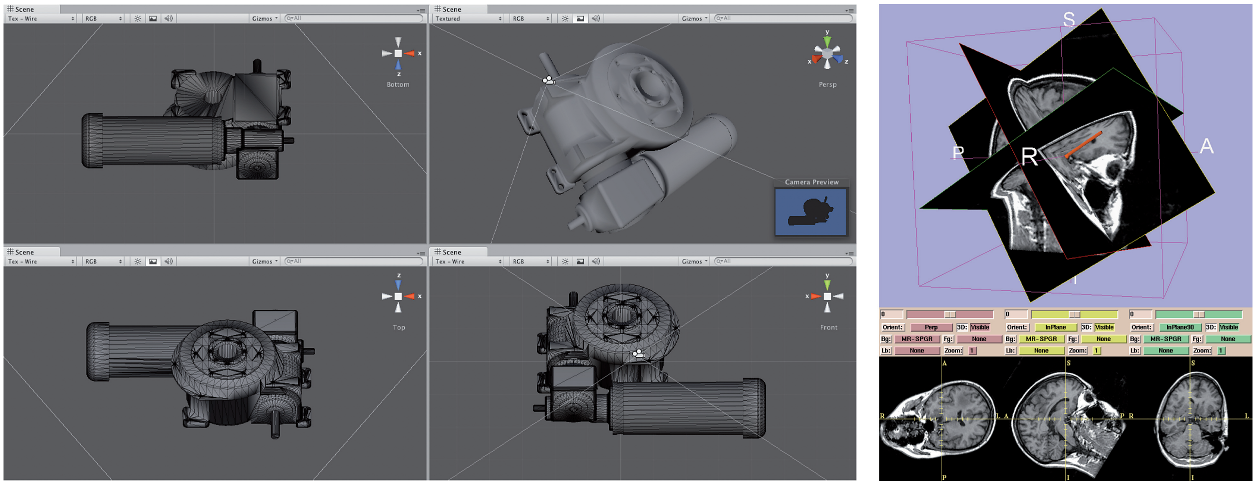

Fig. 2. Computer-aided design and visualization applications are often done using multiple views of a 3D model. Noncolocated orientation indicators used in the application Unity 3D, shown in left image, and colocated orientation indicators in application 3D Slicer, shown in right image, to indicate object's orientation. Unity 3D image courtesy and copyright Unity 3D. 3D Slicer image courtesy and copyright of David Gering.

has evaluated multiple views for information visualization [North and Shneiderman 2000; Baldonado et al. 2000], whereas the present research specifically addresses multiple views of the sorts of 3D objects common in CAD systems and in scientific and medical visualizations. Examples of some specific applications are described below.

2.2.1 *Mechanical CAD.* An important trend in mechanical CAD is the move towards 3D design software. Prior to the advent of CAD software, mechanical designs of individual parts and objects were typically specified by drafting on paper a set of *orthographic views* (sometimes called *multiview drawings*), representing the parallel projection of the object from various viewing directions. Viewing directions were typically separated by 90° and aligned in some natural way with the object. Modern mechanical software automates this drafting process, but the base representation can now be a 3D object (as shown in Figure 2).

These applications allow users to create, edit, and view 3D objects and scenes on traditional two-dimensional displays [Khan et al. 2008]. Viewpoints that are difficult to achieve in the real world, such as a bird's eye view, are easily attainable in 3D computer graphic applications. Users can view objects from any angle and orient the part in any position. However, controlling the virtual viewpoint and understanding the position of the virtual camera in relation to an object is a challenging task for those new to virtual environments [Khan et al. 2008]. Another consequence is that users may have to maintain an association and track object features between views. This problem may become even more complicated when designers are working with complex objects. Multiple views are used in CAD and modeling software such as Unity 3D and Autodesk's AutoCAD, Maya, and 3D StudioMax.

2.2.2 *Medical Visualizations.* Certain subsets of the medical community have adopted 3D visualizations into clinical practice [Gering 1999; Golland et al. 1998]. Several medical imaging and visualization software packages allow for multiple views of data. Multiple views are used in visualization software such as OsiriX, 3D Slicer, Anatomy Browser, Seg3D, and ImageVis3D. Figure 2 shows one way a volume can be oriented in 3D Slicer [Gering 1999]. These tools can be used for education, image-guided therapy, and also presurgical planning and reference [Anderson et al. 1998; Golland et al. 1998;

Hata et al. 2001; MacLeod et al. 2008; Tokuda et al. 2010]. For instance, the application 3D Slicer is used in image-guided surgery and allows the surgeon to view 3D surface models of key anatomical and functional structures [Gering et al. 2001] from preoperative data in the interventional context.

Medical education has already made a dramatic shift towards using 3D visualizations and digital representations of anatomy in academic curricula [Keehner and Khooshabeh 2002]. Educators are recommending digital representations for the study of anatomical structure, function, and spatial relationships [Silen et al. 2008]. Medical professionals rely on a detailed understanding of spatial structures in the human body [Hegarty et al. 2007], but medical students have difficulties in achieving this level of understanding [Silen et al. 2008]. It is believed that realistic 3D objects will enhance a student's learning experience [Garg et al. 1999]. There have been major initiatives such as the Visible Human Project to acquire spatial data from human organs and create 3D objects which are used for teaching and learning gross anatomy [Jastrow and Vollrath 2006].

3. EXPERIMENTAL EVALUATION

The purpose of providing users with an orientation indicator is to help them understand, identify, or recognize a viewed object when presented in different orientations. Our main goal was to test the effectiveness of the orientation indicator using two versions of the mental rotation paradigm [Shepard and Metzler 1971; Vandenberg and Kuse 1978]. To assess whether orientation indicators affect participants' performance, we created a series of four computer-based experiments. We tested whether mental rotation performance would change as a function of the orientation indicator and whether effects would differ depending on spatial ability and several task factors including the class of object (mechanical or biological) and the axes of rotation.

3.1 General Methods

Across all experiments, we measured accuracy (and response time in Experiment 3 and 4) on congruency judgments of 3D objects presented in different orientations, as a function of (a) two types of orientation indicators; (b) two classes of 3D objects; (c) two axes of rotation; (d) the influence of individual differences in visuospatial abilities; and (e) two different tasks. The first type of orientation indicator did not share a center point with the 3D object, and is referred to as *noncolocated*; the second type of orientation indicator did share a center point with the 3D object, and is referred to as *colocated*. Both the orientation indicators and 3D objects were always displayed as static pictures. In each experiment, the class of objects, axes of rotation, and the presence/absence of the orientation indicator were manipulated within-subjects, whereas visuospatial ability was necessarily a between-subjects factor. The type of orientation indicator and the type of task performed was constant within a single experiment and between-experiment comparisons of the effect of the orientation indicator type were performed later.

We predicted that both types of graphical aids could improve an individual's ability to make same/different judgments on 3D objects shown in different orientations, however, we did not have *a priori* predictions about which might be more effective. It could be that a noncolocated aid would provide a frame of reference for the object that could be used to facilitate the object judgments without distracting from the object features themselves. It could also be that the colocated aid would be effective because, when superimposed on the object, users could more directly perceive the correspondence between objects. We predicted that spatial ability may also influence an individual's performance on the tasks. Spatial ability has been a factor in prior research regarding the effectiveness of 3D visualizations, showing mixed results depending on the task. Given the varied findings, we considered whether the orientation indicator might be used in a different way by low versus high spatial groups. For example, those lower in spatial ability may be able to take advantage of a colocated aid because it does not require an additional transformation of the aid to the object, but would not be able to effectively use the noncolocated aid. Or, the additional complexity associated with the aid could prohibit low spatial users from benefiting from the aid at all. For those higher in spatial ability, the aid could facilitate their mental transformations for instances that are difficult for them, for example a complex object or large degree of rotation, or it might go unused if these individuals adequately perform the task without it.

Two different experimental designs were used, both of which have been employed extensively in past studies of mental rotation. The first of these, which we refer to below as the *choose-two-of-four task*, presented a target object and four possible matches. Participants had to pick the two correct matches from the four possibilities [Vandenberg and Kuse 1978]. Although the choose-two-of-four task is time-limited overall, there is no time limit on individual trials. This design was used in Experiments 1 and 2. The second design was a *same/different task* [Shepard and Metzler 1971] in which participants decided on each trial whether a pair of objects was the same or different. The same/different task was time-limited per trial, which allows response time to be measured on individual trials. This design was used in Experiments 3 and 4. The choose-two-of-four task emphasized accuracy, while the same/different task emphasized both accuracy and response time. Performance on the same/different paradigm can be used to make inferences about applications where there is time pressure, and can also provide evidence as to whether the orientation indicator increases response time. Lastly, it can be used to evaluate whether response time is linear with respect to degrees of rotation; this data cannot be determined using the choose-two-of-four mental rotation task. Across all experiments we assessed two axes of rotation, the vertical axis: parallel to the image plane and the horizontal axis parallel to the image plane.

In the experiments involving colocated aids, small object translations were applied to the object models when needed to insure that the center of the aid was within the interior of the object. Half of the rendered objects in Experiments 2 and 4 were affected. This had the unintended consequence of causing slight shading differences compared to the same object in the same orientation when presented in Experiments 1 and 3, with the main effect being that certain shadows were darker in Experiments 1 and 3. In a few cases, these translations also resulted in differences in where the objects were placed in the display window between equivalent colocated and noncolocated conditions. However, the critical within-subjects variable was the presence/absence of the aid, and this comparison was unaffected by the way in which objects were centered on the colocated aids.

Participants in each experiment were University of Utah students who were given either psychology course credit or compensation of $10 for their participation. Participants were not allowed to participate in multiple experiments. Because the orientation indicators rely on color perception, participants self-reported color blindness. Only one participant in all five experiments reported an inability to discern colors. The participant reported that his color blindness did not effect his ability to do the task or use the aid because there were enough colors he could perceive accurately.

Participants' spatial visualization ability was measured using two paper-and-pencil tests: the *Paper Folding Test* [Ekstrom et al. 1976] and the *Mental Rotation Test* [Vandenberg and Kuse 1978]. Each test had 20 questions, and consisted of two parts that were timed for three minutes each. The paper-folding test was scored by awarding one point for every correct answer minus three-quarters of a point for every incorrect answer. The mental rotation test was scored by awarding two points for every correct answer minus two points for every incorrect answer. Penalizing participants for incorrect answers was done to prevent them from guessing in both spatial ability tests, and is a standard scoring method. Standardized scores (z-scores) were calculated for the two paper-and-pencil tests, which were combined to create an aggregate measure for each participant across all four experiments (160 participants total: 88 females, 72 males). Participants were classified as high spatial ability or low spatial ability based on a natural break in the distribution of scores ($z = .151$) that was very close to the median

(z = .221). This resulted in 78 participants in the low spatial ability group (z-score range −6.022 to .085) and 82 participants in the high spatial ability group (z-score range .216 to 4.150). A Pearson correlation showed a strong positive relationship between the scores on the two paper-and-pencil tests, $r = .609$, $p < .01$ (two-tailed). Additional analyses for each experiment were conducted with spatial ability classified by median splits for each of the tests separately. These analyses did not qualitatively differ in the main effects and interactions reported below for each experiment, and so are not reported.

3.2 Experiment 1

Experiment 1 tested the presence of noncolocated orientation indicators to on-screen 3D geometric objects. A mental rotation task, the *choose-two-of-four task*, was used to assess the effectiveness of the orientation indicators. Participants were shown a target object and four possible matches. Participants had to pick the two correct matches from the four possibilities [Vandenberg and Kuse 1978].

3.2.1 *Participants.* Experiment 1 tested 40 participants (21 females, 19 males). The high spatial ability group (z-score range .216 to 2.823) had 23 participants (8 females, 15 males, mean age = 21.9 years, SD = 3.9). The low spatial ability group (z-score range −5.621 to .038) had 17 participants (13 females, 4 males, mean age = 27.5 years, SD = 9.5). Refer to Section 3.1 for how the participants' spatial visualization ability was measured.

3.2.2 *Stimuli.* Object stimuli were shown either with a noncolocated orientation indicator, or no orientation indicator. Noncolocated orientation indicators were placed above the object stimuli with the origin of the rotation displaced from the objects that they corresponded with. The indicator was shown rotated in the same axis and amount as the object, and both remained static images.

We based the look of the orientation indicator off of the coordinate system icons often used in CAD programs, but felt additional colored markers would help users who are not experienced with 3D CAD and visualization applications. The indicator had six markers, each with a unique color. Participants were not given instruction on how the aid could help to solve the tasks; they were only told the aid rotated the same amount and direction as the object.

Two classes of objects were used, since object complexity has been shown to affect task performance with 3D visualizations. One class of objects was mechanical parts that were constructed of distinct pieces; the other class of objects was anatomical structures that represent blood vessels, an aneurysm, or organism that is composed of abstract parts. These two classes of objects stem from the 3D object perception experiment conducted by Cole et al. [2008], which used models in which people could easily infer shape, did not have a lot of self-occlusion, were not too familiar, and were somewhat simple without much fine-scale detail [Cole et al. 2008]. We believe this criteria is well suited for both the mental rotation paradigm and also the application areas of 3D CAD software and medical visualizations; see Figure 4 for the ten 3D object stimuli. All object stimuli were positioned such that the object fit within the clipping window. The anatomical structures were assembled using digital embryos [Brady and Kersten 2003]. The mechanical structures were assembled using 3D models from 3D Content Central (`http://www.3dcontentcentral.com`). All models were modified and rendered with Autodesk's Maya software version 8.5. The type of lighting used were area lights with Blinn shading. The visual extent of the image (target object plus four object choices) presented on each trial had a width subtending 27.6° and a height subtending 8.3°.

3.2.3 *Experimental Design and Procedure.* Experiment 1 tested a noncolocated orientation indicator with the choose-two-of-four task using object rotations about the vertical axis parallel to the image plane and rotations about the horizontal axis parallel to the image plane. Although participants were told to perform as quickly and accurately as possible, the choose-two-of-four task emphasizes accuracy.

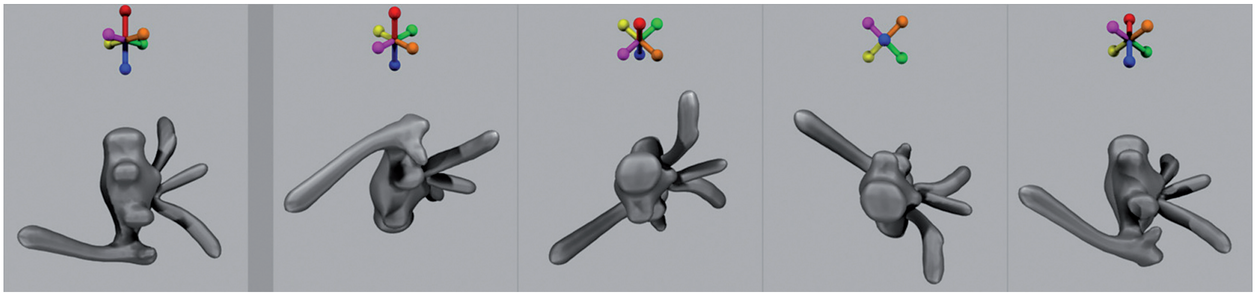

Fig. 3. Example trial for Experiment 1, involving a noncolocated orientation indicator: Choose which two of the four objects on the right match the target object on the left.

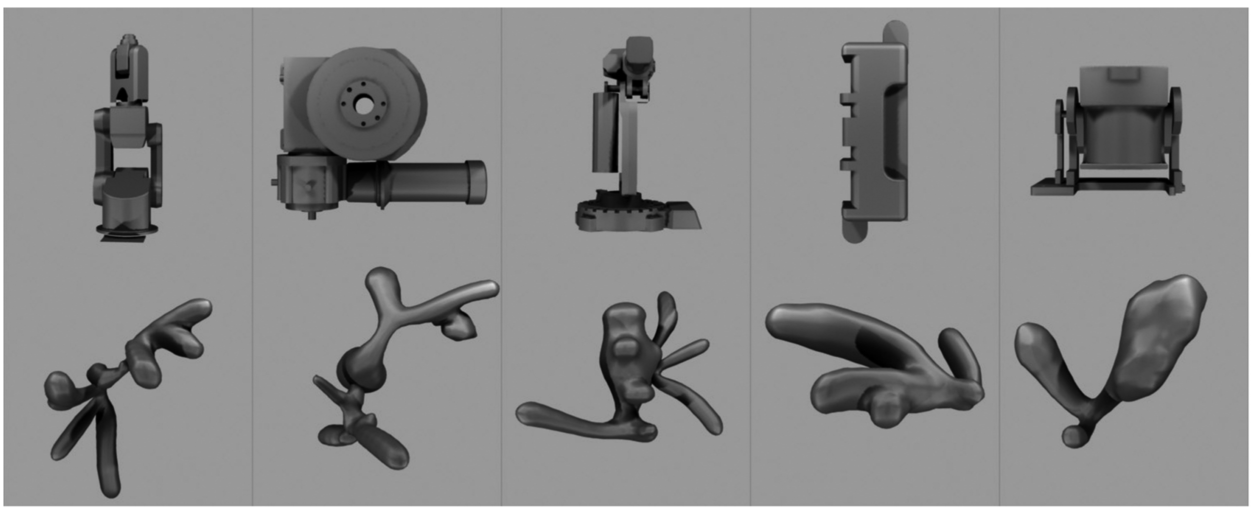

Fig. 4. Ten stimuli used in the experiment. Mechanical parts on top, anatomical structures below. Each stimulus shown in 0° orientation.

In this task, participants were shown four objects and they were to decide which two of the four objects matched a target object (see Figure 3). Two of the four objects were mirror images of the target object, and thus were not congruent in shape to the target object.

Experiment 1 included rotations about the vertical axis parallel to the image plane, hereafter vertical axis, and rotations about the horizontal axis parallel to the image, hereafter horizontal axis. In trials with rotation about the vertical axis, mirror objects were made by reflecting the object about the horizontal axis such that the left and right of the object were reversed. In trials with rotation about the horizontal axis, mirror objects were made by reflecting the object about the vertical axis such that the top and bottom of the object were reversed. Mirror objects were made in this manner to prevent participants from being able to use strategies other than mental rotation to solve the task. Figure 4 shows each object in its original 0° position from which reflections and rotations were based. Mirrored objects were also rotated from the initial position.

There were 40 trials total, of these trials, half presented an aid and half did not. Within each aid/no aid condition, half of the trials had objects rotated about the vertical axis and half had objects rotated about the horizontal axis. Within each axis condition, half of the trials had mechanical object stimuli and half had anatomical object stimuli. The target objects were positioned at 0°, 15°, 345°, 30°, and 330°. The four objects were positioned at 285°, 300°, 315°, 330°, 345°, 0°, 15°, 30°, 45°, 60°, and 75°. We tested five different degrees of disparity between the target object and the four objects: 15°, 30°, 45°, 60°, and 75°. Each degree of disparity between the target and four possible target object choices was used two times in each condition (i.e., presence of aid, class of object, and axis of rotation). The trials for anatomical objects were setup exactly the same as the trials for mechanical objects except the order of the four target object choices was rearranged. Furthermore, each of the ten objects was used in four trials (i.e., horizontal rotation no aid, vertical rotation no aid, horizontal rotation aid, vertical rotation aid). Trials with the orientation indicator were setup using the same rotational disparities between the target image and the object choices as the trials without the orientation indicator. The degrees of disparity between the four choices were also the same between orientation indicator and no indicator conditions.

Participants performed four blocks of trials; two blocks showed the orientation indicator, and two blocks did not have the indicator present. Two blocks were mechanical parts and two blocks were anatomical structures. We counterbalanced the order of the aid condition and object type condition across participants and gender to prevent performance differences attributed to practice effects. Participants were given four minutes to complete each block of trials, with three short breaks in between blocks. The breaks came after the first, second, and third blocks and lasted one minute, two minutes, and one minute, respectively. Each block had 10 trials. It was possible for a participant to time out and not finish a block of trials. They were also permitted to skip trials, and if time allowed they were given another chance to answer skipped trials. Instructions emphasized the importance of accuracy over response time and did not explicitly mention the manipulation of the class of objects or the axes of rotation.

Accuracy on the task was scored by giving two points for every correct answer and subtracting two points for every incorrect answer. This scoring method corrects for guessing and follows the conventional scoring method for Vandenberg and Kuse mental rotation tests [Vandenberg and Kuse 1978; Stull et al. 2009]. These scores were then normalized on a scale of zero to one. The orientation indicator variable, along with the class of objects and axis of rotation was varied within subjects to test for differences within the individual. Spatial ability was variable between subjects.

Participants performed the experiment individually in a controlled experiment room where lighting was held constant. The experiment was controlled using E-Prime and viewed on a 19-inch monitor. Viewing position was also held constant, with the observer's head located approximately 31 inches from the monitor. Although participants were instructed to remain seated in one location, head movement was not controlled for. Participants responded with a button box. For the four object choices, buttons were spatially mapped to the object choices on the monitor. During the breaks, participants read articles from the popular press to prevent them from devising cognitive strategies to solve the task. At the end of the computer portion of the experiment, participants were given a written survey regarding the experimental task similar to the one given in Peters et al. [1995]. Participants' spatial visualization ability was measured using the two paper-and-pencil tests mentioned before.

3.2.4 *Results and Discussion.* A 2(orientation indicator) × 2(class of objects) × 2(axis of rotation) × 2(spatial ability) ANOVA was performed on the mean scores for Experiment 1. Partial eta squared (η_p^2) is used as an indication of effect size, reflecting the amount of variance accounted for by the independent variable. Experiment 1 showed no effect of the noncolocated indicator, $F(1,38) = .5$, $p = .5$, $\eta_p^2 = .01$. Participants performed nearly the same with the aid (.77) versus without the aid (.76). One possibility for the lack of overall effect of the orientation indicator is that it was not attached to the

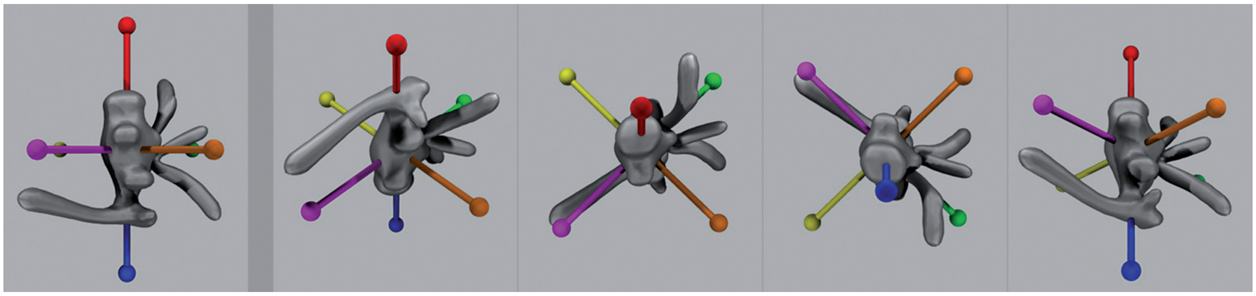

Fig. 5. Example trial for Experiment 2, involving a colocated orientation indicator: Choose which two of the four objects on the right match the target object on the left.

object. The noncolocated indicator could have led the user to solve the task first for the aid, and then transform the information about rotation to the object. Thus, users may have had difficulty recovering information from the aid and translating it to the object.

The indicator effect was modulated by the axis of rotation, as shown in the indicator × axis interaction, $F(1,38) = 8.4$, $p < .01$, $\eta_p^2 = .18$, with the presence of the noncolocated indicator leading to increased accuracy for objects rotated about the horizontal axis (.77 versus .73, $p < .05$), but no difference for the vertical axis (.76 versus .79, $p = .20$). Some specific instances of objects may have had characteristics making them more difficult to comprehend given rotations around the horizontal axis, and so more likely to benefit from an aid. While no obvious differences in self-occlusions or other appearance effects occurred between cases rotated around the horizontal and around the vertical axes, some subtle properties of object shape may have had an effect [Lawson and Humphreys 1998]. Intrinsic differences in that difficulty of comprehending rotations around the two different axes may also have played a role [Parsons 1987].

In addition, there was a significant effect of the class of objects, $F(1,38)= 26.6$, $p < .01$, $\eta_p^2 = .41$. Participants performed better on mechanical objects (.81) versus anatomical objects (.72). Participants may have improved performance with mechanical parts because they were comprised of distinct shapes, whereas anatomical parts were not. Participants may also have been able to identify a particular feature on the mechanical objects and use that feature to determine a correspondence between two objects. In addition, the mechanical objects had distinct vertical or horizontal axes that were aligned with the axis of rotation; in contrast, the anatomical objects did not have this property, which may have led to more difficulty in the decisions about these objects. Finally, there was an overall effect of spatial ability group, $F(1, 38) = 14.15$, $p < .01$, $\eta_p^2 = .27$, showing superior performance for the high versus low ability group.

3.3 Experiment 2

Because the noncolocated orientation indicator in Experiment 1 did not increase participants' accuracy in the mental rotation task, Experiment 2 tested the effect of a colocated orientation indicator. A colocated indicator could eliminate the need for a user to transfer the information about rotation from a noncolocated aid to the object. Thus the information about rotation from a colocated aid may be more accessible. The same mental rotation task, the choose-two-of-four paradigm, was used to assess the effectiveness of the colocated orientation indicators. The only difference between Experiments 1 and 2 was the orientation indicators used (see Figure 5).

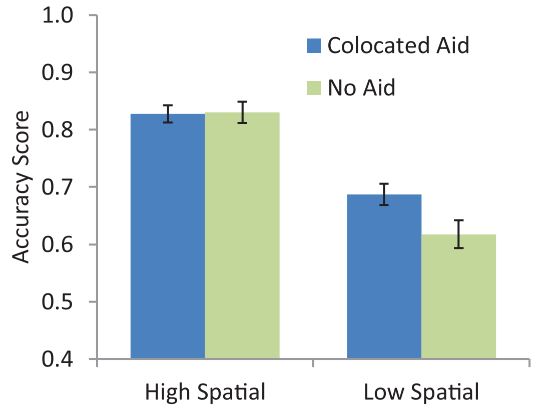

Fig. 6. Mean accuracy score, normalized from 0 to 1, as a function of presence/absence of a colocated aid and spatial ability in Experiment 2. The choose-two-of-four paradigm was used to evaluate performance. Error bars are +/- 1 SE.

3.3.1 *Participants.* Experiment 2 tested 40 participants (24 females, 16 males). The high spatial ability group (z-score range .220 to 3.312) had 21 participants (9 females, 12 males, mean age = 22.5 years, SD = 4.7). The low spatial ability group (z-score range -3.082 to $-.415$) had 19 participants (15 females, 4 males, mean age = 23.7 years, SD = 7.3). Refer to Section 3.1 for how the subjects' spatial visualization ability was measured.

3.3.2 *Stimuli.* Object stimuli were shown with either a colocated orientation indicator or no orientation indicator. The colocated indicator was a larger scaled version of the noncolocated indicator (see Section 3.2.2). Colocated indicators rotated in the same coordinate system as the objects that they corresponded with, with the center of that coordinate system at the center of the object. Each indicator was shown rotated in the same axis and amount as the object. The 3D object stimuli were identical to those used in Experiment 1.

3.3.3 *Experimental Design and Procedure.* The experimental design and procedure was identical to Experiment 1; refer to Section 3.2.3.

3.3.4 *Results and Discussion.* A 2(orientation indicator) × 2(class of objects) × 2(axis of rotation) × 2(spatial ability) ANOVA was performed on the mean scores for Experiment 2. Experiment 2 showed a facilitatory effect of the colocated indicator on performance, $F(1,38) = 8.5$, $p < .01$, $\eta_p^2 = .18$. Participants showed an increase in accuracy with the aid (.76) versus without the aid (.73). Participants appeared to be able to better utilize the information about rotation that the colocated aid provided.

Furthermore, the colocated indicator effect was modulated by spatial ability, as shown by the indicator × spatial ability interaction, $F(1,38) = 10.1$, $p < .01$, $\eta_p^2 = .21$. As Figure 6 shows, the effect of the indicator in Experiment 2 was driven by low spatial learners. The low spatial group showed an increase in accuracy with the aid (.69) versus without the aid (.62), $p < .01$, whereas the high spatial group showed no change (.83) for both conditions. Given the overall superior performance in the high spatial group, $F(1,38) = 51.06$, $p < .01$, $\eta_p^2 = .57$, this interaction provides initial support for the hypothesis that supplemental tools would be particularly effective for individuals who have difficulty with a task.

As with the results in Experiment 1, the indicator effect was modulated by the axis of rotation, as shown by the indicator × axis interaction, $F(1,38) = 18.3$, $p < .01$, $\eta_p^2 = .33$, with presence of the indicator leading to increased accuracy for objects rotated about the horizontal axis (.79 versus .71, $p < .01$), but no difference for the vertical axis (.73 versus .75, $p = .34$).

Similar to the results of Experiment 1, there was a significant effect of the class of objects, $F(1{,}38) = 47.1$, $p < .01$, $\eta_p^2 = .55$, showing that participants performed better with mechanical objects (.79) versus anatomical objects (.70). Finally, there was an overall effect of spatial ability, $F(1, 38) = 51.05$, $p < .01$, $\eta_p^2 = .57$, showing superior performance for the high versus low ability group.

3.3.5 *Combined Analysis of Experiments* 1 *and* 2. A 2(experiment) × 2(orientation indicator) × 2(class of objects) × 2(axis of rotation) × 2(spatial ability) ANOVA was performed on the mean scores for Experiment 1 and Experiment 2 combined. This analysis was conducted to determine if the colocated indicator used in Experiment 2 increased accuracy more than the noncolocated indicator used in Experiment 1. There was no main effect of Experiment, $F(1{,}76) = 1.6$, $p = .2$, $\eta_p^2 = .02$, nor any interactions with Experiment and the other variables. Across the two experiments, there remained an effect of the orientation indicators on accuracy, as demonstrated in Experiment 2, $F(1{,}76) = 5.6$, $p < .05$, $\eta_p^2 = .07$. However, the reduced size of the effect compared to Experiment 2 alone suggests that the effect was driven by the larger effect of the colocated indicator in Experiment 2.

3.4 Experiment 3

Experiments 3 and 4 were carried out to provide additional information about participants' performance when shown orientation indicators using a different mental rotation task, the same/different task, emphasizing both response time and accuracy. Response time in the same/different task is measured on each individual trial; thus we can assess whether the degree of rotation and/or the orientation indicator leads to increased response time. Although similar conceptually to the previous task, the binary forced-choice and emphasis on speeded responses is distinct and could possibly lead to differences in strategies or task performance.

Experiment 3 assessed performance using a noncolocated orientation indicator using the same/different paradigm.

3.4.1 *Participants.* Experiment 3 tested 40 participants (21 females, 19 males). The high spatial ability group (z-score range .325 to 3.571) had 19 participants (4 females, 15 males, mean age = 23 years, SD = 4.7). The low spatial ability group (z-score range −2.823 to .085) had 21 participants (17 females, 4 males, mean age = 22.4 years, SD = 7.6). Refer to Section 3.1 for how participants' spatial visualization ability was measured.

3.4.2 *Stimuli.* Object stimuli were shown either with a noncolocated orientation indicator, or no orientation indicator. The stimuli were identical to those used in Experiment 1; refer to Section 3.2.2. The visual extent of the image presented on each trial (the pair of objects) had a width subtending 11.3° and a height subtending 8.3°. The choose-two-of-four images were wider than same/different images because there were more objects; an individual object's size did not differ between the two paradigms.

3.4.3 *Experimental Design and Procedure.* In this task, participants were shown two objects to decide whether these two objects were the same object, but shown in different orientations, or whether they were different objects. If they were different objects, one object was a mirror image of the other. Mirror objects were made using the same procedure as in Experiment 1 (refer to Section 3.2.2), see Figure 7 for example trials.

There were 160 trials total. Of these trials, half presented an aid and half did not. Within each aid/no aid condition, half of the trials had objects rotated about the vertical axis and half had objects rotated about the horizontal axis. Within each axis condition, half of the trials had mechanical object stimuli and half had anatomical object stimuli. Within each object stimuli condition, half of the trials had

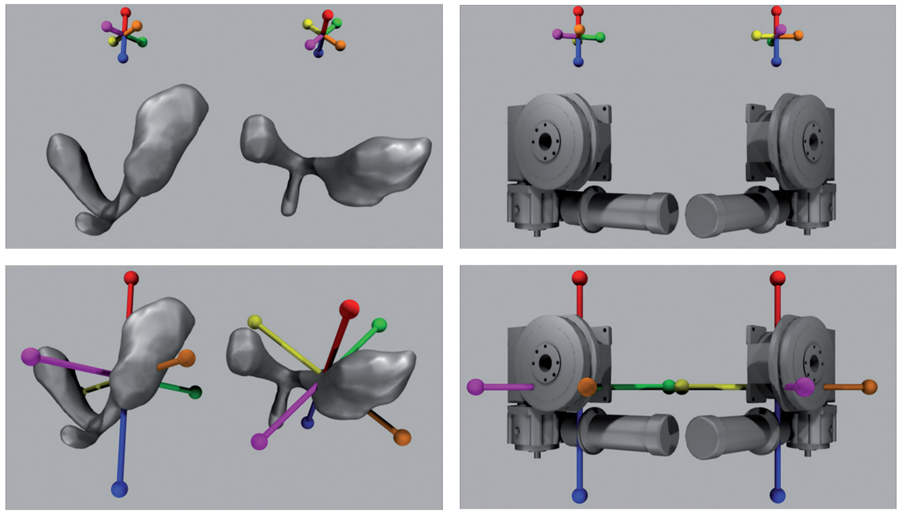

Fig. 7. Four different example trials: Are the objects the same object shown in different orientations, or are they different objects?

same objects and half had different objects. Objects were positioned at 285°, 300°, 315°, 330°, 345°, 0°, 15°, 30°, 45°, 60°, and 75°. Between the two objects there were five possible degrees of disparity, 15°, 30°, 45°, 60°, and 75°. Each degree of disparity was used two times in each condition (i.e., presence of aid, class of object, axis of rotation, and the same/different pair). The trials for anatomical objects were setup similarly to the trials for mechanical objects, except the order of the two objects was rearranged. Trials with the orientation indicator were setup using the same rotational disparities between the objects as the trials without the orientation indicator. Furthermore, each of the 10 objects was used 16 times.

Participants were given two blocks of trials, and within these two blocks, aid and no aid trials were presented randomly. We counterbalanced the two blocks across participants and gender. Participants were given 12 seconds per trial; if they exceeded this time limit they were not given a chance to respond and were presented with the next trial. Participants were not allowed to skip a trial. Instructions emphasized the importance of both accuracy and response time. Participants were given a two minute break between blocks of trials. During this break they read articles from the popular press to prevent them from devising cognitive strategies to solve the task. At the end of the computer portion of the task, participants completed the same written survey taken in Experiment 1 (see Section 3.2.3). The experiment took place in the same controlled experiment room used in Experiment 1.

Accuracy on the same/different task was scored by awarding one point for every correct answer. These scores were then normalized on a scale of 0 to 1. The orientation indicator, class of objects, and axis of rotation were varied within subjects. Spatial ability was a between subjects variable. Note that Experiment 3 scores should not be directly compared to Experiment 1 scores because of the intrinsic difference in the how the two experiment tasks are scored.

3.4.4 *Results and Discussion.* A 2(orientation indicator) × 2(class of objects) × 2(axis of rotation) × 2(spatial ability) ANOVA was performed on the mean accuracy scores for Experiment 3. Experiment 3 showed a facilitatory effect of the noncolocated indicator, $F(1,38) = 18.1$, $p < .01$, $\eta_p^2 = .32$. Participants increased in accuracy with the aid (.74) versus without the aid (.69).

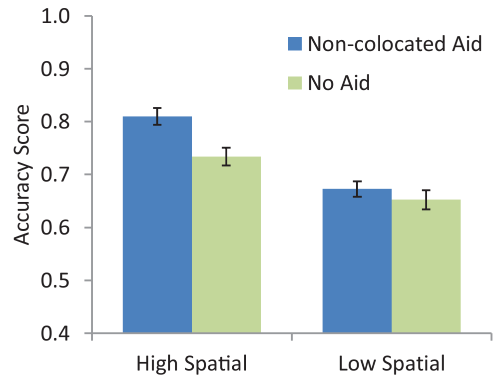

Fig. 8. Mean accuracy score on Experiment 3, normalized from 0 to 1, as a function of presence/absence of a noncolocated aid and spatial ability. The same/different paradigm was used to evaluate performance. Error bars are +/− 1 SE.

Furthermore, the effect of the noncolocated indicator was driven by the high spatial ability group, as shown in the spatial ability × indicator interaction, $F(1{,}38) = 6.1$, $p < .05$, $\eta_p^2 = .14$. As Figure 8 shows, the high spatial group showed an increase in accuracy with the aid (.81) versus without the aid (.73) $p < .01$, whereas the low spatial group showed a smaller and insignificant increase with the aid (.67) versus without the aid (.65) $p = .19$. It is interesting to note the effect of the noncolocated aid in this experiment compared to the lack of effect in Experiment 1. It may be that the increased difficulty associated with the time-pressure or the forced-choice paradigms in this task led to the effective use of the indicator, particularly for the high spatial ability participants.

The axis of rotation also had an effect on accuracy, $F(1{,}38) = 4.4$, $p < .05$, $\eta_p^2 = .10$. Participants performed better with object rotation about the vertical axis (.73) versus the horizontal axis (.70). In Experiments 3 and 4, trials involving rotation around the horizontal axis had somewhat more instances of self-occlusion than did trials involving rotations around the vertical axis.

Lastly, there was a significant effect of the class of objects, $F(1{,}38) = 38.7$, $p < .01$, $\eta_p^2 = .51$. Similar to the results of prior experiments, participants performed better with mechanical objects (.76) versus anatomical objects (.67). Finally, there was an overall effect of spatial ability, $F(1{,}38) = 29.0$, $p < .01$, $\eta_p^2 = .43$, showing superior performance for the high versus low spatial ability group.

A 2(orientation indicator) × 5(degree of rotation) × 2(spatial ability) ANOVA was performed on response time from the correct-only "same" trials. The class of objects and axis of rotation could not be analyzed because the majority of participants did not get at least one of these trials correctly for each degree of rotation. Participants showed increased response time with the aid (4.79 sec) versus without the aid (4.17 sec), $F(1{,}38) = 31.3$, $p < .01$, $\eta_p^2 = .45$. Participants also showed increased response times with a greater disparity, $F(4{,}152) = 8.6$, $p < .01$, $\eta_p^2 = .19$. A planned linear contrast confirmed a linear increase in response time with disparity, $F(1{,}38) = 19.13$, $p < .001$, $\eta_p^2 = .45$. This result could be interpreted as increasing time needed to mentally rotate one object to match the other, or it could be an indication of an overall increase in difficulty with greater disparity. A main effect of spatial ability, $F(1{,}38) = 5.6$, $p < .05$, $\eta_p^2 = .13$, showed that the high spatial group (4.93 sec) was overall slower to respond compared to the low spatial group (4.09 sec). It is noteworthy that the high spatial group performed more accurately overall (and with even greater accuracy given the aid), and were correspondingly *slower* than the low spatial group. These results suggest that the use of the noncolocated aid led to longer processing time, and could only be used effectively by those with high spatial abilities.

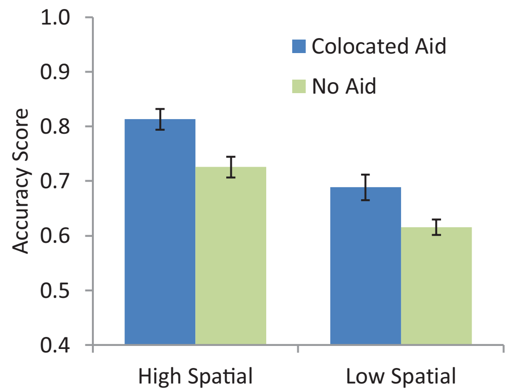

Fig. 9. Mean accuracy score on Experiment 4, normalized from 0 to 1, as a function of presence/absence of a colocated aid and spatial ability. The same/different paradigm was used to evaluate performance. Error bars are +/- 1 SE.

3.5 Experiment 4

Experiment 4 was created to test whether the improvement in accuracy with the colocated aid seen in Experiment 2 and with the noncolocated aid in Experiment 3; would replicate with the colocated aid used in the same/different task. Furthermore, we could assess the effects on response time concurrently.

3.5.1 *Participants.* Experiment 4 tested 40 participants (22 females, 18 males). The high spatial ability group (z-score range .221 to 4.150) had 19 participants (9 females, 10 males, mean age = 22.8 years, SD = 3.9). The low spatial ability group (z-score range -6.022 to $-.036$) had 21 participants (13 females, 8 males, mean age = 23.8, SD = 7.7). Refer to Section 3.1 for how participants' spatial visualization ability was measured.

3.5.2 *Stimuli.* Object stimuli were shown either with a colocated orientation indicator, or no orientation indicator. The stimuli were identical to those used in Experiment 3, refer to Section 3.4.2.

3.5.3 *Experimental Design and Procedure.* The experimental design and procedure was identical to Experiment 3, refer to Section 3.4.3.

3.5.4 *Results and Discussion.* A 2(orientation indicator) × 2(class of objects) × 2(axis of rotation) × 2(spatial ability) ANOVA was performed on the mean accuracy scores for Experiment 4. Experiment 4 showed a strong facilitatory effect of the colocated indicator, $F(1{,}38) = 41.5$, $p < .01$, $\eta_p^2 = .52$. As Figure 9 shows, the colocated indicator helped both spatial ability groups. Participants showed an increase in accuracy with the aid (.75) versus without the aid (.67). The colocated aid may have helped more generally in this experiment because information about rotation is more accessible when superimposed on top of the actual object. In addition, as part of the object, the colocated aid may be more automatically attended to and more difficult to ignore.

As with the results in Experiments 1, 2, and 3, the indicator effect was modulated by the axis of rotation, as shown by the indicator × axis interaction, $F(1{,}38) = 5.4$, $p < .05$, $\eta_p^2 = .12$, with presence of the indicator leading to increased accuracy for objects rotated about the horizontal axis (.75 versus .64, $p < .001$, $\eta_p^2 = .48$), but a smaller difference for the vertical axis (.75 versus .69, $p = .001$, $\eta_p^2 = .24$). As in the previous experiments, this difference in the effectiveness of the aid may be explained by overall differences in difficulty for different axes of rotation.

As in the other experiments, there was a significant effect of the class of objects on accuracy, $F(1{,}38) = 67.3$, $p < .01$, $\eta_p^2 = .64$. Participants performed better on mechanical objects (.77) versus anatomical objects (.65). Finally, there was an overall effect of spatial ability, $F(1{,}38) = 23.76$, $p < .01$, $\eta_p^2 = .38$, showing superior performance for the high versus low spatial ability group.

A 2(orientation indicator) × 5(degree of rotation) × 2(spatial ability) ANOVA was performed on response time from correct-only "same" trials. As in Experiment 3, this analysis collapsed over class of objects and axis of rotation. Participants showed increased response time with the aid (5.64 sec) versus without the aid (4.50 sec), $F(1{,}38) = 147.5$, $p < .01$, $\eta_p^2 = .80$. Participants also showed increased response times with greater degree of rotation, $F(4{,}152) = 9.2$, $p < .01$, $\eta_p^2 = .20$. A planned linear contrast confirmed a linear increase in response time with degree, $F(1{,}38) = 21.5$, $p < .001$, $\eta_p^2 = .36$. Similar to Experiment 3, this result could be due to increased time to mentally rotate larger degrees of disparity or an overall increase in difficulty. There was no overall effect of spatial ability on response times ($p = .89$). These results are consistent with those of Experiment 3 in terms of the relationship between the effect of the aid on accuracy and response time; accuracy improves with the aid, but response time also increases. Experiment 4 shows that with the colocated aid, an increase in accuracy and the corresponding increase in response time occurred for both high and low spatial ability groups.

3.5.5 *Combined Analysis of Experiments* 3 *and* 4. A 2(Experiment) × 2(orientation indicator) × 2(class of objects) × 2(axis of rotation) × 2(spatial ability) ANOVA was performed on the mean scores for Experiment 3 and Experiment 4 combined. Since both the noncolocated and colocated indicators showed effects with the same/different paradigm, as expected, an overall main effect of indicator was found, $F(1{,}76) = 58.3$, $p < .01$, $\eta_p^2 = .43$. There were no effects or interactions with the experiment.

Response time was also analyzed across Experiment 3 and Experiment 4. A 2(experiment) × 2(orientation indicator) × 5(degree of rotation) × 2(spatial ability) ANOVA was performed on response time from correct-only "same" trials. As seen in both individual experiments, participants showed increased response time with the orientation indicator (5.22 sec) versus without the indicator (4.34 sec), $F(1{,}76) = 146.0$, $p < .01$, $\eta_p^2 = .66$. There was a small effect for the experiment, $F(1{,}76) = 5.2$, $p < .05$, $\eta_p^2 = .06$, showing overall faster response times for Experiment 3 (4.49 sec) versus Experiment 4 (5.07 sec). There was an indicator × experiment interaction, $F(1{,}76) = 12.3$, $p < .01$, $\eta_p^2 = .14$, showing that response time with the orientation indicator was longer in Experiment 4 compared to Experiment 3 ($p < .01$), but there was no difference in response time for trials without the indicator across the two experiments ($p = .18$). There were no other interactions with the experiments. Since increased accuracy relates to slower response time in Experiments 3 and 4, the overall slower response time in Experiment 4 (compared to Experiment 3) for trials with the indicator is consistent with greater use of the aid overall. However, since this is a between-subject comparison across the two types of aid, it is possible that the response time difference could be attributed to differences between groups rather than differences between the use of a specific type of aid.

3.6 Spatial Ability and Gender

As seen in the individual experimental results above, there was an overall difference in accuracy in the high spatial ability versus low spatial ability groups in every experiment, $p < .01$. Given the prior claims of gender differences in spatial abilities (see Halpern and Collaer [2005]), we ran a 2(orientation indicator) × 2(class of objects) × 2(axis of rotation) × 2(gender) ANOVA on the mean scores for each experiment with gender instead of spatial ability as the between-subject variable. There was a significant overall difference in accuracy as a function of gender in Experiment 2, which used the choose-two-of-four task with a colocated aid, $p < .05$. The scores of this experiment showed that males (.81) outperformed females (.72). There was also a significant overall difference in accuracy between males and females in Experiment 3, which used the same/different task with a noncolocated aid, $p < .01$. The scores of this experiment showed that males (.75) outperformed females (.68). While related to the effects of spatial ability, these results indicate that gender is not as consistent in predicting performance across the four experiments as spatial ability.

4. GENERAL DISCUSSION

As anticipated, the type of orientation indicator implemented and individual spatial ability influenced users' abilities to make object orientation judgments of 3D geometric objects shown as multiple views. Overall, there was not much effect found with noncolocated orientation indicators (Experiment 1), although high spatial ability individuals did show some benefit (Experiment 3). Colocated aids helped both low and high spatial ability participants. Low spatial ability users improved in accuracy with a colocated aid in the choose-two-of-four task. Furthermore, both spatial groups showed improved performance with a colocated aid in the same/different task. These results provide evidence for several of the different predictions about how spatial ability could affect performance with 3D computer graphic applications. We found situations where an orientation indicator helped low spatial learners, high spatial learners, and all learners. We also found an instance where a noncolocated aid was not that beneficial to anyone (Experiment 1).

Task-difficulty may affect how much benefit there is from an aid, and the difficulty of the task can be a function of individuals' spatial ability, time pressure, and axis of rotation. Overall, the same/different tasks with time pressure revealed facilitory effects of the aid for all users. This may have been a result of increased difficulty related to speed of processing, as well as other task differences. Although the colocated aid increased accuracy, it also increased response time, showing evidence for a speed-accuracy trade-off. Speed-accuracy trade-offs are to be considered when designing an interface and are application dependent. In some applications, the increase in response time with an aid may be justified by the increase in accuracy. Orientation indicators particularly helped with rotation about the horizontal axis. This result could be because rotation about the vertical axis was easier for participants, either due to idiosyncratic aspects of the objects or intrinsic differences in comprehending rotations in depth around the two axes. Our experiments do not provide the data needed to determine the actual cause of this effect, and more research will be required to understand it.

Previous research indicates that imagining an object's rotation is difficult when only the object's initial position is given and no other information is provided [Parsons 1995; Wraga et al. 1999]. The experiments in this article expand on these findings and illustrate that people have difficulties determining a 3D object's orientation in space when the object is presented as multiple static views. It is possible that users have difficulty orienting objects in desktop environments because many of the cues commonly used to maintain a frame of reference in the real world are absent in these virtual spaces. In the real world we can orient ourselves via cues from our bodies and the environment, including the horizon, lighting, and other objects. In 3D CAD and visualization applications, objects are often presented in a vacuum of space, and users may become easily disoriented with camera perspectives that are from unfamiliar points of view. Users may have to make a conjecture about the spatial relationships within a virtual workspace [Glueck et al. 2009]. Ware and Arsenault [2004] found that frames of reference can impact the task-performance of making two virtual objects parallel (i.e., rotating one object until it matches the orientation of a target object).

Our orientation indicator may have provided viewers with a frame of reference by providing an in-scene coordinate frame. People appear to rapidly and automatically assign a major axis, and hence a top to objects [Zacks and Michelon 2005], and such axes play an important role in how we perceive their orientations in space [Shiffrar and Shepard 1991; Parsons 1987; Hinton and Parsons 1981]. The frame of reference provided by the orientation indicators could have been used to orient an object in one view to an object in another view. Zacks and Michelon [2005] suggest that in a same/different mental rotation task, requiring judgments about whether two images are identical or mirror images, viewers use a object-based transformation to rotate the reference frame. Our findings support this notion, showing that, in some circumstances, an orientation indicator facilitates users in maintaining an object-based frame of reference.

It is important to acknowledge the possibility that the orientation indicator was used to help participants define the relationship between various parts of an object without facilitating a mental transformation of the object. Using the colored axes of the indicator as relative features, participants may have compared the features of one object to another object to determine whether the objects were the same or different. This strategy would enable a comparison of multiple views without using mental rotation to determine correspondence. While this alternative cannot be ruled out given our current experiments, the response time functions found in Experiments 3 and 4, showing increasing response time with increasing rotational disparity, suggest that mental transformations of objects may have occurred.

We note some limitations of the present work which open up several areas of further study. First, the results found in this series of experiments are likely to be influenced by the specific parameters of objects that we chose. While our intent was to test two very different categories of objects relevant to 3D visualization applications, these objects did have different properties such as the major axis of alignment and its correspondence (or lack of) with the axes of rotation, and the particular resulting foreshortening or self-occlusion in the objects given each rotational axis. Future work is needed to test different object models and additional axes of rotation in order to generalize these findings further. Second, the design of this series of studies allowed for only between-subject comparisons of the type of orientation indicator. While each experiment allowed for comparisons of the exact same objects with and without the aid within each individual, we do not have as direct a comparison within the same person across types of aid (colocated and noncolocated). Despite these limitations, our results provide insight into the effectiveness of a given type of aid (compared to its absence) and demonstrate that performance effects may also be dependent on individual differences.

The distinction between the noncolocated aid and the colocated aid is important both for understanding how the aid facilitates performance and for advancement of visualization aids in 3D computer graphic applications. Since the noncolocated aid was displaced from the object, it added an additional transformation to correspond with the object. This may have prohibited lower-spatial individuals from attempting to use it, and overall, increased the cognitive effort involved in the spatial task. In contrast, the nature of the colocated aid makes it more intuitive in its spatial mapping to the object. Furthermore, the fact that it is superimposed on the object makes it more difficult for users to ignore. Importantly, some of the aids that are implemented in current applications rely on a displaced indicator. We have shown here that this may not be the most effective approach to facilitate spatial reasoning for all users.

The quantitative effects of the aid are consistent with the qualitative data collected. Low spatial participants stated that they tried to use the noncolocated aid, but that they were not sure whether it helped them. They made comments such as, "The aid added another level of complexity to the task, I first had to solve for the aid and then the object, and that was difficult for me," and "The aid was just another thing to compare." High spatial users said that they were "able to use the aid, but felt the task was the same level of difficulty throughout the experiment". In general, users who were presented with the colocated aid felt it helped them. Participants from both spatial ability groups commented that the colocated aid helped "because the colors showed me how to rotate the object". Some stated that they "always wanted the aid to be there". Lastly, participants from both spatial ability groups said that the "abstract" shapes were more difficult than the objects with "distinct pieces," supporting the quantitative findings of higher accuracy with the mechanical versus anatomical objects. While the type of object affected overall performance on the tasks, it did not generally determine the effectiveness of the orientation indicator.

5. CONCLUSION

Performance on visualization tasks involving multiple static views will not necessarily improve simply because cognitive aids such as orientation indicators are implemented. The utility of three-dimensional environments and orientation indicators can be increased by the controlled investigation of users' perceptions of the 3D objects. We have demonstrated the use of a cognitive experimental paradigm for a comparison of static images, but a similar methodological approach could be extended to systematically evaluate noninteractive and interactive dynamic displays. Future research could also evaluate oblique rotations, since in some circumstances they have been shown to be particularly difficult for both high and low spatial ability groups [Pani 1993; Parsons 1995; Peters et al. 1995]. This research provides a basis for an approach to future studies that draws on the theory and methods in spatial cognition to both inform the creation of graphical visualization aids and to evaluate them. By identifying users' difficulties with an abstract virtual space and the benefits of additional information we can make desktop environments more effective.

An awareness of individual spatial abilities is likely to help in creating and evaluating the presence of in-scene cognitive aids in 3D computer graphic applications. It has been shown that people differ widely in their spatial ability, not only in the general population, but also within specialized populations such as practicing surgeons [Keehner et al. 2004; Hegarty and Waller 2005]. An understanding of these differences may lead to more accessible visualization applications for all users. Our results not only provide evidence for the influence of individual differences on the effectiveness of a graphical aid, but also more generally highlight the need to quantitatively assess users' performance in desktop environments.

REFERENCES

ANDERSON, J., UMANS, C., HALLE, M., GOLLAND, P., JAKAB, M., MCCARLEY, R., JOLESZ, F., SHENTON, M., AND KIKINIS, R. 1998. Anatomy browser: Java-based interactive teaching tool for learning human neuroanatomy. *Radiological Soc. North Amer. – Electron. J.*

BALDONADO, M. Q. W., WOODRUFF, A., AND KUCHINSKY, A. 2000. Guidelines for using multiple views in information visualization. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI'00)*. ACM, New York, 110–119.

BARFIELD, W., SANDFORD, J., AND FOLEY, J. 1988. The mental rotation and perceived realism of computer-generated three-dimensional images. *Int. J. Man-Mach. Stud. 29,* 6, 669–684.

BECKWITH, L., BURNETT, M., WIEDENBECK, S., COOK, C., SORTE, S., AND HASTINGS, M. 2005. Effectiveness of end-user debugging software features: Are there gender issues? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 869–878.

BECKWITH, L., KISSINGER, C., BURNETT, M., WIEDENBECK, S., LAWRANCE, J., BLACKWELL, A., AND COOK, C. 2006. Tinkering and gender in end-user programmers' debugging. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 231–240.

BRADY, M. J. AND KERSTEN, D. 2003. Bootstrapped learning of novel objects. *J. Vision 3,* 6, 413–422.

BROOKS , F. P., JR., OUH-YOUNG, M., BATTER, J. J., AND KILPATRICK, P. J. 1990. Project GROPE Haptic displays for scientific visualization. In *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'90)*. ACM, New York, 177–185.

BULTHOFF, H. H. AND EDELMAN, S. 1992. Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proc. Nat. Acad. Sci. 89*. 60–64.

CHEN, C., CZERWINSKI, M., AND MACREDIE, R. 2000. Individual differences in virtual environments – introduction and overview. *J. Amer. Soc. Inf. Sci. 51,* 6, 499–507.

COLE, F., GOLOVINSKIY, A., LIMPAECHER, A., BARROS, H. S., FINKELSTEIN, A., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2008. Where do people draw lines? *ACM Trans. Graph. 27,* 3.

COX, R. W. 1996. AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res. 29*, 162–173.

CZERWINSKI, M., TAN, D. S., AND ROBERTSON, G. G. 2002. Women take a wider view. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 195–202.

DILLON, A. AND WATSON, C. 1996. User analysis in HCI – the historical lessons from individual differences research. *Int. J. Hum.-Comput. Stud. 45,* 6, 619–638.

DUNSER, A., STEINBUGL, K., KAUFMANN, H., AND GLUCK, J. 2006. Virtual and augmented reality as spatial ability training tools. In *Proceedings of the 6th ACM SIGCHI International Conference on Computer-Human Interaction*. ACM, New York, 125–132.

EKSTROM, R. B., FRENCH, J. W., HARMAN, H. H., AND DERMEN, D. 1976. *Manual for Kit of Factor-Referenced Cognitive Tests*. Educational Testing Service.

FITZMAURICE, G., MATEJKA, J., MORDATCH, I., KHAN, A., AND KURTENBACH, G. 2008. Safe 3D navigation. In *Proceedings of the Symposium on Interactive 3D Graphics and Games*. 7–15.

GARG, A., NORMAN, G. R., SPERO, L., AND MAHESHWARI, P. 1999. Do virtual computer models hinder anatomy learning? *Acad. Med. 74,* 10, 87–89.

GARG, A. X., NORMAN, G., AND SPERO, L. 2001. How medical students learn spatial anatomy. *The Lancet 357*, 363–364.

GERING, D. T. 1999. A system for surgical planning and guidance using image FUsion and interventional MR. M.S. thesis, MIT, Cambridge, MA.

GERING, D. T., NABAVI, A., KIKINIS, R., HATA, N., O'DONNELL, L. J., GRIMSON, W. E. L., JOLESZ, F. A., AND BLACK, P. M. 2001. An integrated visualization system for surgical planning and guidance using image fusion and an open MR. *J. Magnetic Resonance Imag. 13*, 967–975.

GLUECK, M., CRANE, K., ANDERSON, S., RUTNIK, A., AND KHAN, A. 2009. Multiscale 3D reference visualization. In *Proceedings of the Symposium on Interactive 3D Graphics and Games (I3D'09)*. ACM, New York, 225–232.

GOLLAND, P., KIKINIS, R., UMANS, C., HALLE, M., SHENTON, M., AND RICHOLT, J. 1998. Anatomy browser: A framework for integration of medical information. In *Proceedings of the 1st International Conference on Medical Image Computing and Computer Assisted Interventing (MICCAI 1998)*.

HALPERN, D. F. AND COLLAER, M. L. 2005. Sex differences in visuospatial abilities. In *Handbook of Higher-level Visuospatial Thinking*, P. Shah and A. Miyake Eds., Cambridge University Press, Cambridge, UK, 170–212.

HATA, N., JINZAKI, M., KACHER, D., CORMAK, R., GERING, D., NABAVI, A., SILVERMAN, S. G., DAMICO, A. V., KIKINIS, R., JOLESZ, F. A., AND TEMPANY, C. M. C. 2001. MR imaging-guided prostate biopsy with surgical navigation software: Device validation and feasibility. *Radiology 220*, 263–268.

HEGARTY, M. 2004. Dynamic visualizations and learning: Getting to the difficult questions. *Learn. Instruct. 14*, 343–351.

HEGARTY, M., KEEHNER, M., COHEN, C., MONTELLO, D. R., AND LIPPA, Y. 2007. The role of spatial cognition in medicine: Applications for selecting and training professionals. In *Applied Spatial Cognition: From Research to Cognitive Technology*, G. L. Allen Ed., Lawrence Erlbaum, Hillsdale, NJ, 285–315.

HEGARTY, M. AND WALLER, D. 2005. Individual differences in spatial abilities. In *Handbook of Higher-level Visuospatial Thinking*, P. Shah and A. Miyake Eds., Cambridge University Press, Cambridge, UK, 121–169.

HINTON, G. E. AND PARSONS, L. M. 1981. Frames of reference and mental imagery. In *Attention and Performance*, P. Shah and A. Miyake Eds., Vol. 9, Erlbaum, Hillsdale, NJ, 261–277.

HUBONA, G. S., SHIRAH, G. W., AND FOUT, D. G. 1997. 3D object recognition with motion. In *Extended Abstracts on Human Factors in Computing Systems (CHI'97)*. 345–346.

HUFF, C. 2002. Gender, software design, and occupational equity. *SIGCSE Bull. 34,* 2, 112–115.

HUK, T. 2006. Who benefits from learning with 3D models? The case of spatial ability. *J. Comput. Assist. Learn. 22*, 392–404.

JASTROW, H. AND VOLLRATH, L. 2006. Teaching and learning gross anatomy using modern electronic media based on the visible human project. *J. Comput. Assist. Learn. 22*, 392–404.

KEEHNER, M., HEGARTY, M., COHEN, C., KHOOSHABEH, P., AND MONTELLO, D. R. 2008. Spatial reasoning with external visualizations: What matters is what you see, not whether you interact. *Cognitive Sci. 32,* 7, 1099–1132.

KEEHNER, M. AND KHOOSHABEH, P. 2002. Computerized representations of 3D structure: How spatial comprehension and patterns of interactivity differ among learners. In *Proceedings of the AAAI Spring Symposium Series, Reasoning with Mental and External Diagrams*. AAAI Press, Menlo Park, CA, 12–17.

KEEHNER, M., KHOOSHABEH, P., AND HEGARTY, M. 2008. Individual differences among users: Implications for the design of 3D medical visualizations. In *User Centered Design for Medical Visualization*, F. Dong et al. Eds., Idea Group Inc., Hershey, PA, 1–24.

KEEHNER, M. M., TENDICK, F., MENG, M. V., ANWAR, H. P., HEGARTY, M., STOLLER, M. L., AND DUH, Q.-Y. 2004. Spatial ability, experience, and skill in laparoscopic surgery. *Amer. J. Surgery 188*, 71–75.

KHAN, A., MORDATCH, I., FITZMAURICE, G., MATEJKA, J., AND KURTENBACH, G. 2008. Viewcube: A 3D orientation indicator and controller. In *Proceedings of the Symposium on Interactive 3D Graphics and Games*. 17–25.

LAWSON, R. AND HUMPHREYS, G. W. 1998. View-specific effects of depth rotation and foreshortening on the initial recognition and priming of familiar objects. *Percept. Psychophys. 60*, 1052–1066.

MACLEOD, R., KHOLMOVSKI, E., DIBELLA, E., OAKES, R., BLAUER, J., FISH, E., VIJAYAKUMAR, S., DACCARETT, M., SEGERSON, N., AND MARROUCHE, N. 2008. Integration of MRI in evaluation and ablation of atrial fibrillation. *Comput. Cardiology 35*, 77–80.

MODJESKA, D. AND CHIGNELL, M. 2003. Individual differences in exploration using desktop VR. *J. Amer. Soc. Inf. Sci. Technol. 54,* 3, 216–228.

NORTH, C. AND SHNEIDERMAN, B. 2000. Snap-together visualization: A user interface for coordinating visualizations via relational schemata. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI'00)*. ACM, New York, 128–135.

PANI, J. R. 1993. Limits on the comprehension of rotational motion: Mental imagery of rotations with oblique components. *Perception 22,* 7, 785–808.

PARSONS, L. M. 1987. Visual discrimination of abstract, three-dimensional objects at many orientations. *Perception Psychophys. 42*, 49–59.

PARSONS, L. M. 1995. Inability to reason about an object's orientation using an axis and angle of rotation. *J. Exper. Psychol. Hum. Perception Perform. 21*, 1259–1277.

PETERS, M., LAENG, B., LATHAM, K., JACKSON, M., ZAIYOUNA, R., AND RICHARDSON, C. 1995. A redrawn van den Berg and Kuse mental rotations test: Different versions and factors that affect performance. *Brain Cogn. 28,* 1, 39–58.

RIZZO, A. A. AND BUCKWALTER, J. G. 1997. Virtual reality and cognitive assessment and rehabilitation: The state of the art. In *Virtual Reality in Neuro-Psych-Physiology: Cognitive, Clinical and Methodological Issues in Assessment and Rehabilitation*, G. Riva Ed., IOS Press, Amsterdam, 123–146.

SANDO, T., TORY, M., AND IRANI, P. 2009. Effects of animation, user-controlled interactions, and multiple static views in understanding 3D structures. In *Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization (APGV'09)*. ACM, New York, 69–76.

SHEPARD, R. N. AND COOPER, L. A. 1982. *Mental Images and Their Transformations*. MIT Press, Cambridge, MA.

SHEPARD, R. N. AND METZLER, J. 1971. Mental rotation of three-dimensional objects. *Science 171,* 3972, 701–703.

SHIFFRAR, M. M. AND SHEPARD, R. N. 1991. Comparison of cube rotations around axes inclined relative to the environment or to the cube. *J. Exper. Psychol.: Hum. Perception Perform. 17,* 1, 44–54.

SILEN, C., WIRELL, S., KVIST, J., NYLANDER, E., AND SMEDBY, O. 2008. Advanced 3D visualization in student-centered medical education. *Medical Teacher 30,* 5, 115–124.

STULL, A. T., HEGARTY, M., AND MAYER, R. E. 2009. Getting a handle on learning anatomy with interactive 3D graphics. *J. Edu. Psychol. 101,* 4, 803–816.

TARR, M. J. 1995. Rotating objects to recognize them: A case study on the role of viewpoint dependency in the recognition of three-dimensional objects. *Psychonomic Bull. Rev. 2,* 1, 55–82.

TOKUDA, J., FISCHER, G., DIMAIO, S., GOBBI, D., CSOMA, C., MEWES, P., FICHTINGER, G., TEMPANY, C., AND HATA, N. 2010. Integrated navigation and control software system for MRI-guided robotic prostate interventions. *Comput. Med. Imaging Graph. 34,* 1 (Jan.), 3–8.

TORY, M. 2003. Mental registration of 2D and 3D visualizations (an empirical study). In *Proceedings of the IEEE Visualization Conference*. 371–378.

TORY, M. AND MOLLER, T. 2004. Human factors in visualization research. *IEEE Trans. Visualization Comput. Graph. 10,* 1, 72–84.

TORY, M. AND MOLLER, T. 2005. Evaluating visualizations: Do expert reviews work? *IEEE Comput. Graph. Appl. 25,* 5, 8–11.

VANDENBERG, S. G. AND KUSE, A. R. 1978. Mental rotations: A group test of three-dimensional spatial visualization. *Perceptual Motor Skills 47,* 2, 599–604.

VELEZ, M., SILVER, D., AND TREMAINE, M. 2005. Understanding visualization through spatial ability differences. In *Proceedings of the IEEE Visualization Conference*. 511–518.

WALENSTEIN, A. 2002. Cognitive support in software engineering tools: A distributed cognition framework. Ph.D. dissertation, School of Computing, Science, Simon Fraser University.

WARE, C. AND ARSENAULT, R. 2004. Frames of reference in virtual object rotation. In *Proceedings of the 1st Symposium on Applied Perception in Graphics and Visualization (APGV'04)*. ACM, New York, 135–141.

WRAGA, M., CREEM, S. H., AND PROFFITT, D. R. 1999. The influence of spatial reference frames on imagined object-and viewer rotations. *Acta Psychol. 102*, 247–264.

ZACKS, J. M. AND MICHELON, P. 2005. Transformations of visuospatial images. *Behav. Cognitive Neurosci. Rev. 4,* 2, 96–118.

ZHANG, J. AND NORMAN, D. A. 1994. Representations in distributed cognitive tasks. *Cognitive Sci. 18*, 87–122.

ZIEMEK, T. R., CREEM-REGEHR, S. H., AND THOMPSON, W. B. 2008. Using mental rotation as a methodology to evaluate shape perception in computer graphics. In *Proceedings of the Symposium on Applied Perception in Graphics and Visualization*.

Received July 2010; revised October 2011; accepted November 2011