The Effects of Changing Projection Geometry on the Interpretation of 3D Orientation on Tabletops

Mark Hancock[1], **Miguel Nacenta**[1,2], **Carl Gutwin**[2], **Sheelagh Carpendale**[1]

[1]University of Calgary
Department of Computer Science
{msh,sheelagh}@cpsc.ucalgary.ca

[2]University of Saskatchewan
Department of Computer Science
{nacenta,gutwin}@cs.usask.ca

ABSTRACT

Applications with 3D models are now becoming more common on tabletop displays. Displaying 3D objects on tables, however, presents problems in the way that the 3D virtual scene is presented on the 2D surface; different choices in the way the projection is designed can lead to distorted images and difficulty interpreting angles and orientations. To investigate these problems, we studied people's ability to judge object orientations under different projection conditions. We found that errors increased significantly as the center of projection diverged from the observer's viewpoint, showing that designers must take this divergence into consideration, particularly for multi-user tables. In addition, we found that a neutral center of projection combined with parallel projection geometry provided a reasonable compromise for multi-user situations.

INTRODUCTION

The range of applications that are being built for tabletop displays is rapidly increasing, and now includes 3D objects and environments [1, 13, 24]. Some application areas benefit greatly from the combination of 3D data presentation and the natural collaboration and face-to-face communication affordances of tabletop systems—examples include tasks such as the planning of surgical procedures with 3D body models, urban planning discussions using models of real-world structures, and collaboration over 3D visualizations. In addition, tasks in multi-display environments may require that individual displays be able to indicate other surfaces or data in the real world (e.g., an object may need to be oriented towards a partner object on another display).

Displaying 3D objects on tables, however, presents new problems for designers of tabletop applications. The representation of a 3D virtual scene on a 2D surface such as a tabletop requires the projection of the virtual 3D objects onto the display surface—the choices made in creating this 3D image, such as where center of projection (CoP) is for the image, or whether a perspective or parallel geometry is used, can have dramatic effects on the appearance of the resulting scene. If the projection is poorly designed, the resulting image on the tabletop appears distorted, and it becomes difficult for the viewer to determine the shape and orientation of objects in the 3D scene. To illustrate this problem, Figure 1 shows what the images from Figure 2 would look like when observed from different points of view at a tabletop display using two different types of projection geometries.

This problem is intensified when people work together with 3D data around a table. Usually, 3D projections on tabletops have only one virtual viewpoint, so that some of the collaborators around a table will see a distorted view of some of the 3D objects. In tasks where the group needs to discuss details of the model such as shape, orientation, and tilt, these distortions could cause misunderstandings and difficulty in communicating about the model. Multiple viewpoint exceptions [12] share the same issues when multiple people are discussing a single 3D object.

Little is known about the problem of interpreting 3D models on 2D tables—about what types of actions are most affected, about the severity of the errors that people make, or about how to choose a projection that minimizes the negative effects. To investigate these issues, we carried out a study in which people were asked to estimate the orientation of a 3D object, projected onto a tabletop display with different CoPs and different projection geometries. We found that as CoP moved further away from the observer, their error in estimating orientation angle significantly increased. However, we also found that when parallel projection geometry is used in combination with a neutral CoP (i.e., between the two viewers), accuracy was as good as with the egocentric projection. Therefore, optimizing the perspective for one person will cause major problems for the others in the group—but providing a neutral CoP and parallel projection geometry may help to mitigate the problem.

As 3D content becomes more common in tabletop systems, understanding how design decisions affect interpretation becomes critical. Our work identifies center of projection and projection geometry as important factors in the usability of 3D tabletop applications, and provides clear guidelines for the design of systems where people must be able to determine the orientation and tilt of 3D objects.

BACKGROUND AND RELATED WORK

In this section, we first describe the fundamentals of pictorial representation of 3D images on a 2D plane, then survey

ITS '09, November 23–25, 2009, Banff, Alberta, Canada.

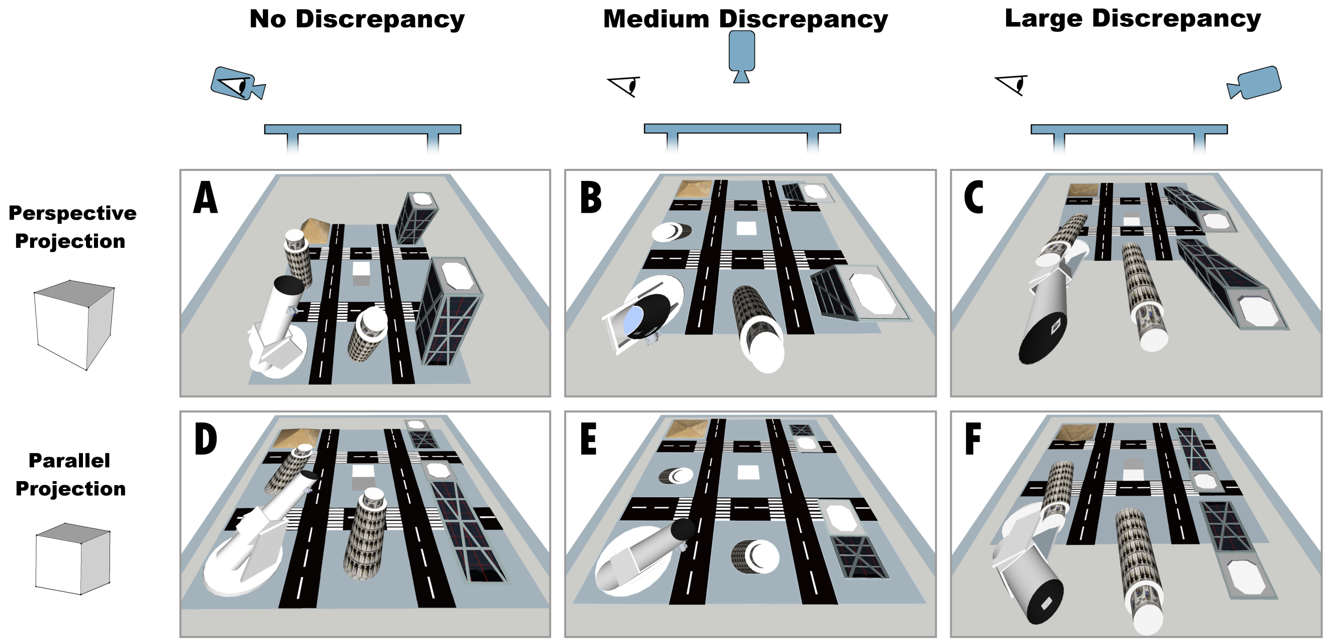

Figure 1: Appearance of 3D models rendered on a table with different levels of discrepancy between point of view (PoV) and center of projection (CoP) using parallel and perspective projection geometries.

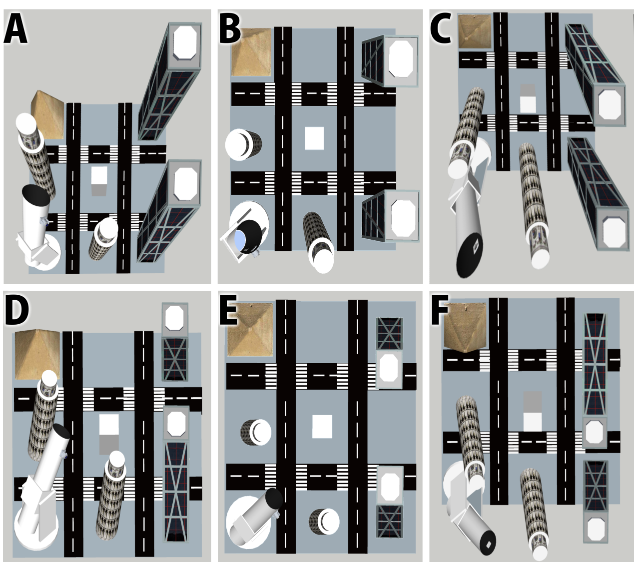

Figure 2: The images rendered to the tabletop display in Figure 1.

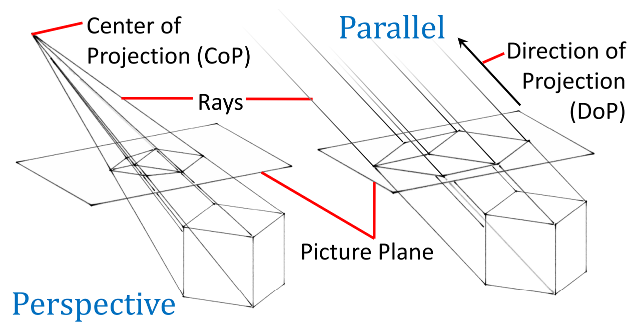

Figure 3: Projection geometries.

related phenomena and the theories that have been developed to explain our perception of those images. We then summarize the few examples in human-computer interaction (HCI) that deal with the perception of 3D on tabletop displays. Note that we also introduce several terms that are used throughout the remainder of the paper.

Basics of Pictorial Representation

To accurately represent a 3D scene in a picture we create straight lines (rays) that go from every point in the 3D scene to the center of projection (CoP) and intersect them with the plane of the picture (Figure 3, left). The picture is geometrically correct for the viewer if the CoP coincides with the point of view (PoV)—the location from where the picture is observed. This method, usually called perspective (or pinhole) projection, has been used by artists for centuries [14]. Currently, the same perspective projection fundamentals underlie most of 3D computer graphics and virtual reality. 3D elements represented in a picture through perspective projection appear smaller if they are further away from the CoP. Also, most lines that are parallel in the 3D scene converge to vanishing points in the picture.

An alternative to perspective geometry is to use rays that are parallel to each other and project in a fixed direction instead of converging to a CoP (Figure 3, right). Parallel projections result in pictures where parallel lines in the 3D scene are preserved as parallel in the picture. Parallel projections do not have a CoP because the projection lines do not converge, but instead the CoP determines a direction of projection. Parallel projections are often used in architecture and engineering because they preserve parallel lines and because they make direct measurement easier; however, parallel projections cannot generate impressions in the viewer's retina that are equivalent to what the real scene projects on the viewer's eyes.

The Perception of Pictorial Space

Scientists and artists have long investigated how pictures represented on a 2D plane are still able to create the impression of depth (what is called by psychologists and artists the *pictorial space*). Some researchers indicate that this is possible because pictures can accurately reproduce on the retina some of the necessary cues of 3D perception (e.g., the foreshortening of objects due to linear perspective and the visual change in receding textures) [2, 3, 7, 19, 20].

However, flat projections of 3D space only create retinal images equivalent to the 3D scene if the PoV of the observer is located at (or very close to) the CoP used to generate the image [20]. When PoV and CoP are at very different angles to the picture plane, or are at very different distances (e.g., when we look at photographs on a table, a painting from a lateral point of view, or a movie from the side aisle), the differences between correct and distorted retinal projections can be very large. If pictorial perception is dependent only on the geometry of the projected retinal image, this should result in the perception of a space that is deformed compared to the depicted space [5, 8, 20].

Regardless of the distortion, observers are remarkably good at still perceiving a relatively accurate pictorial space [21]. However, there is still controversy in the perception research community about the underlying processes that support correct space perception from geometrically incorrect retinal images (what we call space constancy). Some suggest that the visual system corrects distortions based on geometrical information from the represented scene (e.g., assuming certain angles are straight [18], objects are rigid [2], or certain converging lines on the picture are actually parallel in the real scene [19]), and others propose that information about the correct CoP can be recovered from perceptual information about the surface where the picture is projected (e.g., from accommodation and other 3D cues [4], or from the shape of the frame of the picture [16]).

Although the perception of pictorial space is relatively stable regardless of the discrepancy between the locations of the CoP and the PoV, the relationship between the pictorial space and the physical space is not equally stable. In particular, the perceived orientations towards the physical space of elements within the picture plane can vary depending on the position of the observer [3, 9, 20]. This effect is best exemplified by the famous U.S. recruiting poster of Uncle Sam, in which he points directly at the observer regardless of how far she is or how oblique she stands to the plane of the poster. For elements within the picture that point perpendicular to the picture plane (e.g., Uncle Sam's finger), the perceived orientation always follows the observer regardless of its position, and therefore the perceived orientation of the object with respect to the plane of the image can vary almost 180°. For objects that do not point perpendicularly to the picture plane the possible variation in the perceived angle is reduced; at the extreme (objects that are aligned with the picture plane) geometrical accounts of orientation perception [3, 5] predict that the pointing direction will not vary with changes in the PoV. This effect is referred to in the literature as the differential rotation effect (DRE) [8] or the *la Gournerie* effect [3]. The DRE is also subject to scientific controversy; experiments have shown that the geometrical predictions do not necessarily fit all data, especially for very oblique PoVs [7]. The possible causes might be found among the cues that cause the pictorial space constancy discussed above (e.g., frame and perceptions of the picture surface through binocular cues [21]).

In this paper, we investigate the effect of the discrepancy between the CoP and PoV locations on the perceived orientation of objects from the pictorial space into the physical environment (we will refer to this simply as "orientation"). In contrast with previous research in the area, we focus on the interactive tabletop scenario, and therefore we explore factors and conditions that are relevant for interactive tabletops. For example, although a large proportion of the studies mentioned above restrict participants to monocular perception, the use of a single eye to work on tabletops is not reasonable; all our tests are, instead, binocular. Previous research shows that binocular and monocular observation of pictures from non-coincident PoV and CoP is different [21].

We also compare perspective and parallel projections because, although parallel projections are incapable of generating a geometrically accurate retinal image, they have been shown to look more natural than their perspective counterparts [11], and are extensively used in architecture and engineering for their accuracy. Similarly, we include in our study conditions with motion parallax because it provides potentially strong cues [4,6] and has not generally been considered in pictorial research, which is usually more concerned with static pictures. Motion parallax is a depth cue that arises from the continuous change of the PoV with respect to the perceived objects, where many slightly different retinal images are composed to reconstruct the 3D structure of a scene. Motion parallax is extensively relied upon by fishtank virtual reality [22].

3D Perception in HCI and Tabletops

The perception and manipulation of shapes in oblique displays has received some attention in the HCI literature. For example, Wigdor and colleagues [23] studied how the slant of the surface affects the perception of several magnitudes (length, angle, area) for 2D data; Hancock and colleagues [12] looked at different alternatives for the representation and manipulation of 3D objects in Tabletops; and Nacenta and colleagues [17] studied how the correction of perspective distortion in oblique displays affects basic motor and cognitive processes. Finally, Grossman et al. surveyed different 3D technologies for horizontal surfaces, including the 3D cues that they provide [10].

STUDYING ORIENTATION IN TABLETOP DISPLAYS

While there have already been many studies exploring orientation in pictorial space, and these have generated some specific (if somewhat controversial) theories about the effect of CoP/PoV discrepancies, how these theories apply to tabletop display environments is still largely unknown. We do not know what the effects of discrepancy are on horizontal surfaces (experimental setups so far have been vertical), how

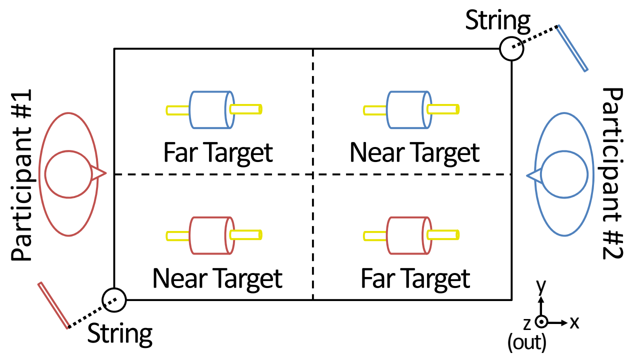

Figure 4: A diagram of the experimental setup.

discrepancy interacts with different types of projection geometries (perspective vs. parallel), or whether motion parallax can help compensate for the distortions created by the highly oblique PoVs typical for tables.

In our study, we focus on the analysis of CoP/PoV discrepancies to inform design choices of 3D tabletop applications. We explore three levels of discrepancy (Figure 1): when the CoP coincides with a person's PoV, when the CoP is directly above the table, and when the CoP coincides with someone else's PoV. We are also interested in the combination of discrepancy with the two main types of projection geometries (perspective and parallel), and with motion parallax, which was described in the previous section, and is reproduced through the real-time tracking of the participant's head. Our study is designed to answer three main questions:

- What are the effects of discrepancy on error?
- How do different projection geometries and motion parallax affect errors due to discrepancy?
- Are there any special cases that designers could use to alleviate errors?

Method

Participants

Twenty-four participants (11 female, 13 male) were recruited from the local community. Ages ranged from 19 to 36 (Mdn = 28, SD = 4.5). People were recruited in pairs (3 female, 4 male, and 5 mixed), but our design was symmetric, so each participant was analysed separately.

Apparatus

Figure 4 shows a diagram of the experimental setup. Participants stood at the ends of a 146 cm × 110 cm bottom-projected tabletop display with a resolution of 2800 × 2100 (19 pixels / cm). The viewpoint of each participant was tracked using a Vicon[1] motion tracking system and markers placed on hats that the participants wore throughout the experiment. Vicon markers were also placed at the end of a string which was also attached to the tabletop corner. These strings were manipulated by an attached wand and used to record participants' answers about the angle of the target (described below).

[1]Vicon Motion Tracking. http://www.vicon.com

Task

The experimental task asked participants to determine the orientation of target objects. For each trial, two 3D target objects were displayed on the screen (one per participant, always displayed on the half of the display at the participant's right). Targets were either in the 'near' half or the 'far' half of the table. Each object was a composite of a long thin cylinder inside a shorter thicker cylinder, each with the same axis. To provide their answer for each task, participants moved the wand until the string, stretched tight, was oriented at the same angle as the main axis of the target. Once both participants had indicated the angle and pressed a 'done' button on the table surface, the next trial would begin. To remind the participants about the projection geometry being used, two groups of four cube frames were shown in the empty quadrants. These cubes were rendered in the same 3D model as the target objects in all conditions.

Conditions and Design

The focus of our study was on how the degree of discrepancy between the center of projection (CoP) and the observer's point of view (PoV) affects perception of object orientation. We thus varied the discrepancy between these two points as the primary factor in our study. The three levels we chose correspond to likely choices when designing a 3D application for a tabletop display: no discrepancy (when the CoP and PoV are the same), medium discrepancy (when the CoP is directly above the table), and large discrepancy (when the CoP is set to someone else's PoV). We also tested the effects of both projection geometry (either perspective or parallel) and motion parallax—that is, whether the CoP dynamically followed the participant's PoV, resulting in small perspective changes as the participant moved their head.

Targets were shown in three different angular orientations, and in two locations. As shown in Figure 5, targets could be at either 0° (laying flat on the table and pointing towards the end of the table where the participant was located), 60° (pointing upwards towards the end of the table), or 90° (pointing straight up from the table). Targets never leaned to the left or right; that is, they always stayed coplanar with the longitudinal vertical plane. In addition, targets could appear either in the 'near' or 'far' halves of the table, as shown in Figure 4. The experimental factors and levels were thus:

- Discrepancy between CoP and PoV (none, medium, large)
- Motion parallax (absent, present)
- Projection geometry (parallel, perspective)
- Angle (0°, 60°, 90°)
- Location (near, far)

With a medium discrepancy (when the CoP is directly above the table), it is not sensible to introduce motion parallax, as there is no person to move with the CoP. Thus, the first two factors combine into five discrepancy-parallax conditions shown in Table 1. These five conditions were each tested with two projection geometries (perspective or parallel), and in each of these ten, participants performed all six combinations of target angle and target location, in random order. For each of the six combinations, participants carried out one practice trial, and two testing trials.

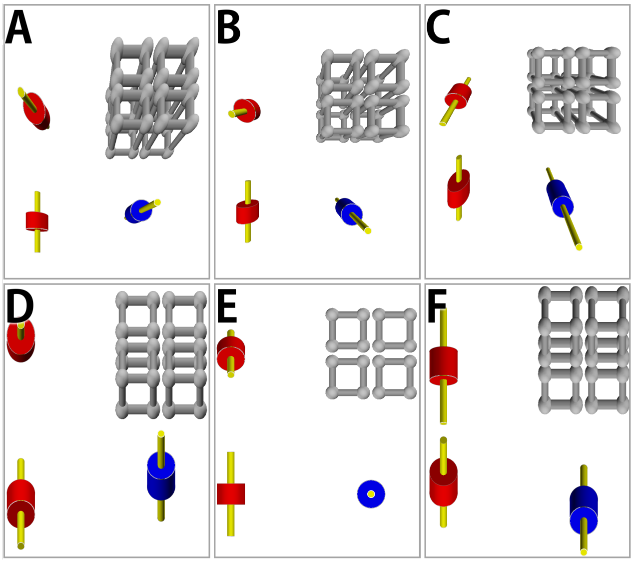

Figure 6: The three angles and visual hints rendered to the display in each discrepancy condition (matching the order of Figure 1).

		Discrepancy		
		None	Medium	Large
Motion	Absent	NA	MA	LA
Parallax	Present	NP		LP

Table 1: The discrepancy-parallax conditions.

The study therefore used a 5 discrepancy-parallax condition $\times$ 2 projection geometry $\times$ 3 angle $\times$ 2 location fully-crossed within-participant design. The 10 condition-projection pairs were counterbalanced between participants using a random Latin Square. Each participant performed 60 practice + 120 testing = 180 trials.

We measured error in angle and task completion time (TCT). The error in angle had two parts (Figure 5, right):

E_w: the error within the longitudinal vertical plane (e.g., if the target pointed towards the end of the table with an angle of 60°, and the participant held the wand such that the string had an angle of 65° in the longitudinal vertical plane, the error was 5°).

E_r: the error away from the longitudinal vertical plane (no targets leaned left or right in this way, but we measured this error since some projections can make targets appear to lean).

Hypotheses
Based on our study design we formulated the following primary hypotheses:

H1: As the discrepancy increases, E_w will increase.

H2: When there is no discrepancy, perspective geometry and motion parallax will reduce E_w; when there is discrepancy, perspective geometry and motion parallax will increase E_w.

H3: Medium discrepancy (CoP directly above the table) will be a special case that decreases E_w.

We also arrived at some secondary hypotheses:

H4: E_w will be least when the angle of the object is horizontal (0°) and most when the object is vertical (90°). This hypothesis will corroborate the differential rotation effect [8].

H5: The type of projection geometry will affect E_r (i.e., errors outside the plane in which we vary the angle).

H6: The use of motion parallax will require more time for the participants to determine the orientation.

RESULTS

We performed a full factorial repeated-measures analysis of variance (ANOVA) on our data and a series of planned comparisons for post-hoc analysis of an expected interaction between the discrepancy-parallax condition and projection geometry (H2). Our planned comparisons correspond to our primary hypotheses as follows:

- To test H1, we perform pairwise comparisons in the order of least to most discrepancy: NA to MA and MA to LA.
- To test H2, we additionally compare the two motion parallax conditions to the endpoints of discrepancy: NP to NA and LA to LP. An effect of perspective geometry would appear as a main effect of the ANOVA.
- To test H3, we additionally compare the MA condition with the remaining two conditions (NP and LP).

We performed these planned comparisons separately for each projection geometry and used a Bonferroni-corrected type I error threshold (i.e., $\alpha/12$).

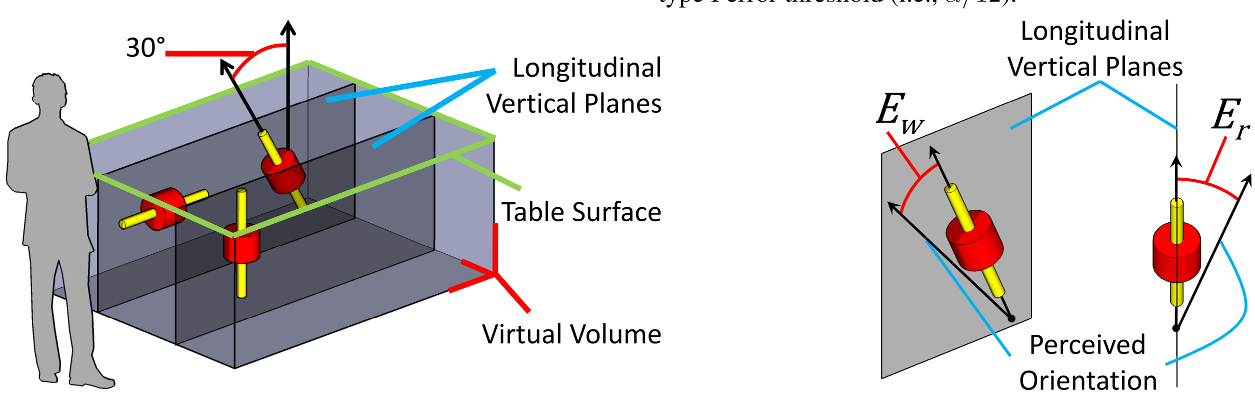

Figure 5: (left) The experimental setup and (right) the two types of error.

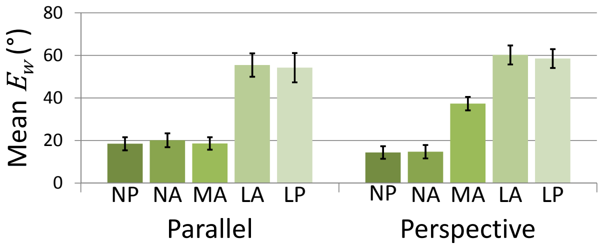

Figure 7: Mean error in angle (E_w) for the ten condition-projection pairs.

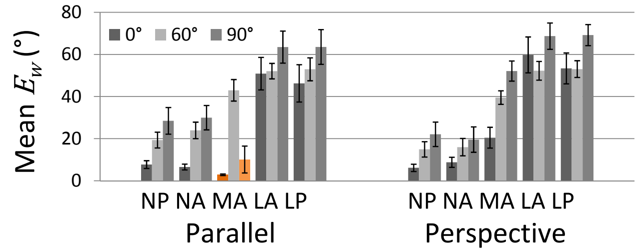

Figure 8: Mean E_w separated by angle of the object. For the MA-parallel condition, $0°$ and $90°$ were exceptions.

Error in Angle Within the Plane (E_w)

Table 2 shows the results of our planned comparisons:

		Parallel	Perspective
H1	NA to MA	$p = .59$	$p <$ **.001**
	MA to LA	$p <$ **.001**	$p <$ **.001**
H2	NP to NA	$p = .25$	$p = .80$
	LA to LP	$p = .80$	$p = .62$
H3	MA to NP	$p = .95$	$p <$ **.001**
	MA to LP	$p <$ **.001**	$p <$ **.001**

Table 2: Planned comparisons for E_w.

What was the effect of discrepancy on error? (H1)

There was a significant main effect of discrepancy-parallax condition on E_w ($F_{4,44} = 39.5$, $p < .001$). As can be seen in Figure 7, error increased overall as the discrepancy between CoP and PoV increased. Therefore, H1 is confirmed.

What were the effects of projection geometry and motion parallax on error? (H2)

There was no main effect of projection geometry on error in angle ($F_{1,11} = 2.6$, $p = .14$), with mean error for perspective conditions ($M = 37.1°$, $SD = 2.0°$) only slightly higher than for parallel conditions ($M = 33.4°$, $SD = 2.9°$). There was also no significant difference for the presence of motion parallax (Table 2 & Figure 7). Therefore, H2 is not confirmed.

Is it a special case to have CoP directly above the table? (H3)

There was a significant interaction between discrepancy-parallax condition and projection geometry ($F_{4,44} = 6.8$, $p < .001$). Our planned comparisons indicate that the interaction was due to the special case of medium discrepancy when using a parallel projection. As can be seen in Figure 7, error in MA with perspective geometry was substantially higher than error in MA with parallel geometry. Furthermore, in the perspective case, the medium discrepancy condition (MA) followed the expected trend of being in between the no discrepancy conditions (NP and NA) and the large discrepancy conditions (LA and LP), whereas in the parallel case, the MA condition was still less than the LA and LP conditions, but had as little error as the NP and NA conditions (Table 2 & Figure 7). Therefore, there is limited support for H3.

Did different target angles lead to different error? (H4)

There was a significant main effect of target angle on error ($F_{2,22} = 7.1$, $p < .01$). Post-hoc analysis showed that errors for all three angles were significantly different ($p < .05$) in the order we predicted with H4 ($M_{0°} = 26.3°$, $M_{60°} = 36.7°$, $M_{90°} = 42.7°$).

Target angle was also involved in several interactions, including significant two-way interactions between angle × condition ($F_{8,88} = 7.9$, $p < .001$) and angle × projection ($F_{2,22} = 9.6$, $p = .001$), a significant three-way interaction between angle × condition × projection ($F_{8,88} = 8.6$, $p < .001$) and a significant four-way interaction ($F_{8,88} = 2.4$, $p = .02$). Figure 8 shows that these interactions are largely explained by the special case of the medium discrepancy condition (H3). Specifically, for horizontal (0°) and vertical (90°) angles using a parallel projection, participants were able to determine the orientation with a high degree of accuracy, going against the trend predicted in H4, which held or was not significant for all other conditions.

Other analyses

All other main effects and interactions were not significant. Specifically, location did not have a significant effect on E_w, nor did it interact with other factors.

Left-Right Error in Angle (E_r)

There were significant main effects of condition ($F_{4,44} = 8.3$, $p < .001$) and angle ($F_{2,22} = 29.4$, $p < .001$), a significant condition × projection interaction ($F_{4,44} = 24.7$, $p < .001$), a significant condition × angle interaction ($F_{8,88} = 9.4$, $p < .001$), and a significant three-way condition × projection × angle interaction ($F_{8,88} = 13.3$, $p < .001$). No other main effects or interactions were significant.

Our planned comparisons showed only two significant differences. One in the parallel projection showed significantly more error in the NA condition than the MA condition ($p < .01$) and one in the perspective projection showed significantly more error in the MA condition than the NP condition ($p = .001$). These two cases are best explained through the three-way interaction (Figure 9). For the horizontal angle (0°), the left-right error was small for all conditions. For the other two angles, participants had more errors in the parallel projection when there was no discrepancy (NA and NP), as well as in the perspective projection with medium discrepancy (MA).

Task Completion Time

We performed the same factorial ANOVA on TCT. There was a significant main effect of condition ($F_{4,44} = 6.4$, $p < .001$). Using the same planned comparisons, aggregated

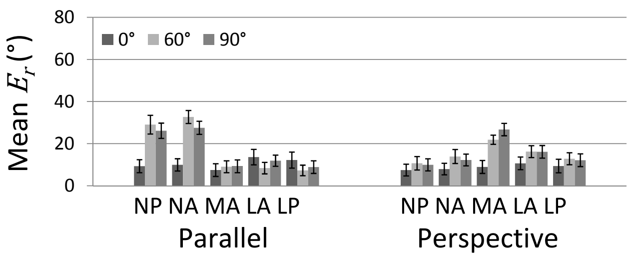

Figure 9: Remaining error (E_r) for each condition, projection, and angle.

across both projection geometries, revealed that the MA condition was significantly faster than the LP condition ($p < .001$). Though the other pairwise differences were not significant, there is a clear trend that both conditions involving motion parallax ($M_{NP} = 6.7$ s, $M_{LP} = 7.5$ s) were slower than the three without ($M_{NA} = 5.6$ s, $M_{MA} = 5.4$ s, $M_{LA} = 6.1$ s).

There was also a significant interaction between angle and location ($F_{2,22} = 16.1$, $p < .001$), but we did not investigate further, as it did not involve the discrepancy-parallax condition, nor projection geometry.

DISCUSSION & CONCLUSION

We summarize the main findings of our study as follows:

- As the discrepancy between CoP and PoV increases, so does the error in people's ability to judge the orientation of 3D objects.
- For parallel projections, the case where the CoP is directly above the table is special and may reduce the problem of this discrepancy.
- Motion parallax did not improve people's ability to judge orientation, nor did it make the error worse when the discrepancy was large.
- People will take more time to acquire and process information when motion parallax is available.
- We did not find effects involving location, suggesting that effects of discrepancy are minimally impacted by the location of objects on the screen.

First, our main result was that error increased with increasing discrepancy between CoP and PoV. This relationship follows our expectations based on the idea of the differential rotation effect (DRE). This effect (as described in the "Uncle Sam" example) causes objects oriented perpendicular to the picture plane to seem as thought they point toward the observer. This phenomenon can explain the increasing error that we saw for the three target angles: objects that were at 90° (fully perpendicular) were difficult for participants to judge, whereas objects at 0° were interpreted with high accuracy. These difficulties arising from DRE suggest that this effect must be considered when designing tabletop systems.

Second, our study showed that a neutral CoP and a parallel projection geometry is a special case. One possible explanation is that the frame of the table may suggest that objects are oriented relative to that frame, and not to the PoV. However, this explanation does not account for the difference between perspective and parallel projections. A perhaps more likely explanation is that objects rendered in this geometry lose their 3D appearance and become perceivable as 2D within the plane of the table. That is, a cylinder pointing out of the display becomes a circle, and when pointing horizontally becomes a rectangle. This explanation is also consistent with the differences we found for each object angle (Figure 8).

There are several practical recommendations that have emerged from our study. First and foremost, we have shown that decisions made about the projection geometry are important and that the designer cannot blindly use the 'defaults' from 3D graphics. In particular, attaching the CoP to one person's PoV can introduce errors of up to 60° in perception for another person at the table, and using a CoP above the table together with a perspective projection can introduce errors over 40°. Our study provides some evidence that a parallel projection with a CoP directly above the table may alleviate some of the problems introduced by this discrepancy (down to a 20° error). Note also that the use of perspective versus a parallel geometry is not necessarily a binary choice; a CoP that is very far above the table may reap some of the benefits of a parallel projection, while maintaining some of the perspective depth cues.

Other possible design solutions to the problem of discrepancy include dedicating parts of the screen to the different CoPs of the viewers [12], or to have different images projected to different people, either through polarized glasses [1] or as a result of their viewing angle [15]. Alternatively, the designer can introduce a method of switching between different CoPs corresponding to different people or discrepancies (e.g., with a button). However, it is still not known how these solutions will impact the applications in which they are used. Will switching between views make it difficult to refer to objects that you worked with previously? Does the need to wear glasses or a reduced screen size affect the collaboration in some way?

Despite our attempts at thoroughness, the problem of discrepancy has many unexplored aspects with respect to tabletop display environments. The study that we ran is specific to orientation, and further research is required to determine the full impact of CoP/PoV discrepancies. Space constancy or framing may eliminate the issue of discrepancy for other tasks, such as comparing 3D objects or determining an object's shape. Stereoscopic depth cues in combination with motion parallax may also prove more helpful. Further study is required to tease out these differences.

ACKNOWLEDGEMENTS

We would like to thank Natural Science and Engineering Research Council of Canada, Albertas Informatics Circle of Research Excellence, Alberta Ingenuity, and the Canadian Foundation of Innovation for research support. We also thank the reviewers and members of the iLab for their helpful comments on this work.

REFERENCES

1. Agrawala, M., Beers, A. C., McDowall, I., Fröhlich, B., Bolas, M., and Hanrahan, P. The two-user responsive

workbench: support for collaboration through individual views of a shared space. In *Proc. SIGGRAPH*, ACM Press/Addison-Wesley Publishing Co. (1997), 327–332.

2. Cutting, J. E. Rigidity in cinema seen from the front row, side aisle. *J. Exp Psychol Human 13, 3* (1987), 323–334.

3. Cutting, J. E. Affine distortions of pictorial space: Some predictions for goldstein (1987) that la gournerie (1859) might have made. *J. Exp Psychol 14, 2* (1988), 305–311.

4. Cutting, J. E. How the eye measures reality and virtual reality. *Behav Res Meth Instrum Comput 29* (1997), 27–36.

5. de La Gournerie, J. *Traité de Perspective Linéaire Contenant les Tracés pour les Tableaux, Plans & Courbes, les Bas Reliefs & les Décorations Théatrales, avec une Théorie des Effets de Perspective*. Dalmont et Dunod, 1859.

6. Dijkstra, T. M. H., Cornilleau-Peres, V., Gielen, C., and Droulez, J. Perception of three-dimensional shape from ego-and object-motion: Comparison between small-and large-field stimuli. *Vision Research 35, 4* (1995), 453–462.

7. Ellis, S. R., Smith, S., Grunwald, A., and McGreevy, M. W. Direction judgement error in computer generated displays and actual scenes. Pictorial communication in virtual and real environments (1991), 504–526.

8. Goldstein, E. B. Spatial layout, orientation relative to the observer, and perceived projection in pictures viewed at an angle. *J. Exp Psychol Human 13, 2* (1987), 256–266.

9. Goldstein, E. B. Perceived orientation, spatial layout and the geometry ofpictures. Pictorial communication in virtual and real environments (1991), 480.

10. Grossman, T. and Wigdor, D. Going deeper: a taxonomy of 3D on the tabletop. In *Proc. Tabletop* (2007), 137–144.

11. Hagen, M. A. and Elliott, H. B. An investigation of the relationship between viewing condition and preference for true and modified linear perspective and adults. *J. Exp Psychol Human 2, 4* (1976), 479.

12. Hancock, M. and Carpendale, S. Supporting multiple off-axis viewpoints at a tabletop display. In *Proc. Tabletop* (2007), 171–178.

13. Hancock, M., Carpendale, S., and Cockburn, A. Shallow-depth 3D interaction: design and evaluation of one-, two- and three-touch techniques. In *Proc. CHI*, ACM (2007), 1147–1156.

14. Jones, R. K. and Hagen, M. A. The perceptual constraints on choosing a pictorial station point. *Leonardo 11, 3* (1978), 191–196.

15. Kitamura, Y., Nakayama, T., Nakashima, T., and Yamamoto, S. The IllusionHole with polarization filters. In *Proc. VRST*, ACM (2006), 244–251.

16. Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., and Todd, J. T. Pointing out of the picture. *Perception 33* (2004), 513–530.

17. Nacenta, M. A., Sakurai, S., Yamaguchi, T., Miki, Y., Itoh, Y., Kitamura, Y., Subramanian, S., and Gutwin, C. E-conic: a perspective-aware interface for multi-display environments. In *Proc. UIST*, ACM (2007), 279–288.

18. Perkins, D. N. Compensating for distortion in viewing pictures obliquely. *Percept Psychophys 14* (1973), 13–18.

19. Saunders, J. A. and Backus, B. T. The accuracy and reliability of perceived depth from linear perspective as a function of image size. *J. Vision 6, 9* (2006), 933–954.

20. Sedgwick, H. A. *The effects of viewpoint on the virtual space of pictures*, 460–479. Taylor & Francis, Inc., 1993.

21. Vishwanath, D., Girshick, A. R., and Banks, M. S. Why pictures look right when viewed from the wrong place. *Nat Neurosci 8, 10* (2005), 1401–1410.

22. Ware, C., Arthur, K., and Booth, K. S. Fish tank virtual reality. In *Proc. CHI*, ACM (1993), 37–42.

23. Wigdor, D., Shen, C., Forlines, C., and Balakrishnan, R. Perception of elementary graphical elements in tabletop and multi-surface environments. In *Proc. CHI*, ACM (2007), 473–482.

24. Wilson, A. D., Izadi, S., Hilliges, O., Garcia-Mendoza, A., and Kirk, D. Bringing physics to the surface. In *Proc. UIST*, ACM (2008), 67–76.