

Interacting with Molecular Structures: User Performance versus System Complexity

paper126

Abstract

Effective interaction in a virtual environment requires that the user can adequately judge the spatial relationships between the objects in a 3D scene. In order to accomplish adequate depth perception, existing virtual environments create useful perceptual cues through stereoscopy, motion parallax and (active or passive) haptic feedback. Specific hardware, such as high-end monitors with stereoscopic glasses, head-mounted tracking and mirrors are however required to accomplish this. Many potential VR users however refuse to wear cumbersome devices and to adjust to an imposed work environment, especially for longer periods of time. It is therefore important to quantify the repercussions of dropping one or more of the above technologies. These repercussions are likely to depend on the application area, so that comparisons should be performed on tasks that are important and/or occur frequently in the application field of interest.

In this paper, we report on a formal experiment in which the effects of different hardware components on the speed and accuracy of three-dimensional (3D) interaction tasks are established. The tasks that have been selected for the experiment are inspired by interactions and complexities, as they typically occur when exploring molecular structures. From the experimental data, we develop linear regression models to predict the speed and accuracy of the interaction tasks. Our findings show that hardware supported depth cues have a significant positive effect on task speed and accuracy, while software supported depth cues, such as shadows and perspective cues, have a negative effect on trial time. The task trial times are smaller in a simple fish-tank like desktop environment than in a more complex co-location enabled environment, sometimes at the cost of reduced accuracy.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics.]: Methodology and Techniques. Interaction Techniques I.3.7 [Computer Graphics.]: Three-Dimensional Graphics and Realism. Virtual Reality;

1. Introduction

A key challenge facing designers of virtual environments is how to effectively interact with objects in 3D virtual space. In order to enable users to effectively locate and manipulate the objects in such a space, it is especially important to depict the spatial relationships between these objects, in particular with respect to depth. The visual cues that allow humans to perceive depth have been extensively documented in the perception and computer graphics literature, see e.g. [Kel93, WFG92]. The use of these cues, and their influence on the performance of interactive tasks also remains an ongoing topic of research in the VR community, [ABW93, BB97]. One problem with providing users with a range of depth cues is that the implementation of many of these cues requires cumbersome and special-purpose hardware. For ex-

ample, providing active stereoscopic viewing requires the user to wear shutter glasses, providing motion parallax requires a sensor for tracking head positions, and providing solutions to problems caused by convergence/accommodation would even require an eye tracking device. Before deciding on such technologies, it is important for designers of virtual environments to understand the trade-offs involved. Indeed, the benefits of various depth cues come at a cost, both in terms of required hardware components and in terms of required user adjustment. Convincing (i.e., quantitative) arguments are therefore needed to weigh the pros and cons in any specific situation.

The motivation of our study is to understand these trade-offs in more detail for specific application scenarios. We have chosen structural biology as the application area, since

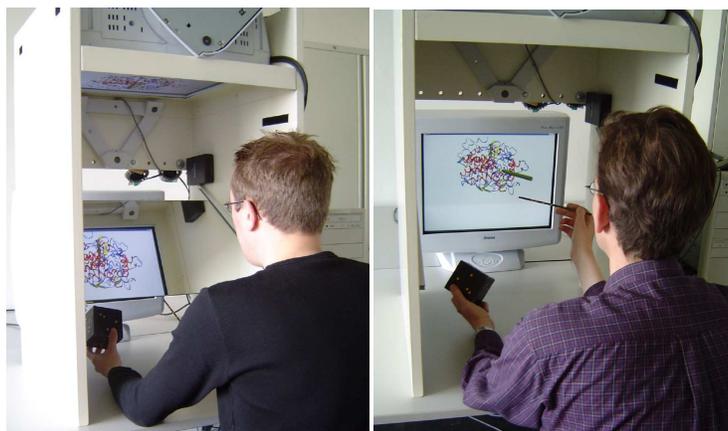


Figure 1: *The Personal Space Station (left) and the Personal Action Space (right).*

the feedback that we have received from biologists in the past is that virtual reality techniques offer great potential for perceiving structure [LLM*02]. However, almost all biologists refuse to wear the cumbersome VR equipment that is usually required today and most would prefer a VR environment that can be placed on their desktop, i.e., that would integrate well within their current work environment. This situation where potentially interesting technology is nevertheless not well accepted has motivated us to set up a formal experiment that can provide more quantitative arguments for the discussion. More specifically, we selected three application-specific 3D interaction tasks that are typical for exploring the structure of molecules: one atom selection task, and two steering tasks. We studied the performance with which these tasks could be executed in two different desktop environments: a co-located mirror-based virtual reality environment and a more traditional fish-tank virtual environment. The tasks were also evaluated under conditions of varying depth cue information. More specifically, we compared depth cues that require hardware support to implement (shutter glasses for active stereoscopic rendering, head tracking for motion parallax) and depth cues that can be mimicked by software rendering alone (shadows and perspective cues).

The paper is organized as follows. In the next Section we provide a short description of the Personal Space Station and Personal Action Space. These are the two desktop environments that we have used in our study. In Sections 3 and 4, we describe the formal experimental study and analyze the data, respectively. Finally, the experimental findings are summarized and directions for future research are explored.

2. Personal Space Station and Personal Action Space

From an apparatus point of view, there are three aspects that need to be distinguished when classifying and comparing

virtual environments; i.e. the tracking of the 3D interaction elements used as input, the display of the 3D virtual information, and the application software that links the observed inputs to the displayed outputs. In this study we compare two VEs that share the same tracking and application layer, and that differ only in the output layer. In this way, it becomes possible to assess the effect of the output layer on the user interactions.

2.1. Tracking

Both VEs aim at two-handed interaction in a 3D personal space that is within hands reach. Input devices, such as a pen for point selections and a cube to orient 3D models, allow users to manipulate 3D virtual objects. Input devices are tagged with infrared-reflecting markers, and are tracked in 3D by means of calibrated infrared stereo cameras, [LM03]. The tracking software can compute the pose of each device at approximately 55 times a second. The end-to-end latencies caused by the optical tracker are approximately 50-60 milliseconds, and the tracking accuracy is within 1 millimeter at the center of the interaction space. This accuracy gradually decreases near the corners of the workspace.

The advantages of optical tracking are that it allows for simultaneous tracking of multiple wireless input devices, and is not susceptible to electromagnetic interference. The obvious disadvantage is that of occlusion; i.e. hands manipulating the input devices may block the line-of-sight of a camera, resulting in a (temporary) loss of the position and orientation of a device. Our experience with many different subjects is that they do require some explanation of this system characteristic, but that they quickly adjust their behavior to minimize such problems.

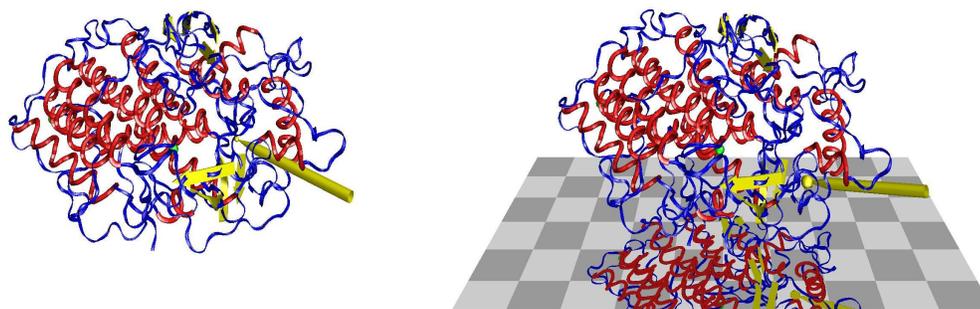


Figure 2: Visualization of the molecule without (left) and with colored shadows and stage (right).

2.2. The Personal Space Station (PSS)

The Personal Space Station (PSS), shown in the left panel of Figure 1, is a near-field VR system, [ML02]. Manipulation of virtual data is performed behind a mirror on which stereoscopic images of the 3D scene are reflected. The user's head is tracked to ensure perspectively correct images, which in turn is required to create the sensation of co-location between the physical interaction devices and the virtual objects that they control.

The PSS is similar to other mirror-based co-located systems as proposed by Schmandt, [Sch83]. The main difference is the usage of multiple passive haptic feedback devices, instead of a single force feedback device, [WSS99, AW04].

2.3. The Personal Action Space (PAS)

The Personal Action Space (PAS), shown in the right panel of Figure 1, is a fish-tank desktop (or Augmented Virtuality [MK94]) system. Manipulation of virtual data is performed in a small interaction volume, typically above a desktop surface. A vertical display surface is placed in front of the user onto which stereoscopic images are drawn. The size of the display can vary from large (e.g. a projector) to small (e.g. a desktop monitor). The user's head is tracked to ensure perspectively correct images.

The PAS is similar to other *fish-tank* like systems as proposed by Ware, [WAB93]. The main difference is the usage of multiple passive haptic feedback devices, instead of mechanical or electromagnetic tracked interaction elements.

2.4. Visual Molecular Dynamics Interface

Visual Molecular Dynamics (VMD) is a widely used molecular visualization program for displaying, animating, and analyzing large bio-molecular systems using 3D graphics and scripting [HDS96]. We have extended VMD to include the optically tracked interaction elements. Furthermore, we have included the option of showing the shadows of the molecules

and the interaction elements on a planar stage positioned somewhere in the scene. The modified version of VMD hence allows 3D interaction with molecules.

All experiments used the same visualization of a complex molecule, see Figure 2. The molecule, *Chloroperoxidase* from the fungus *curvularia inaequalis* [MW96], was provided to us by our colleagues from the University of Amsterdam, who also suggested many interactive tasks on the structure. The molecule is visualized as a combination of ribbon and cartoon representations: the backbone as a blue ribbon, the beta sheets as yellow arrows and the alpha-helices as red spirals.

The molecule is visualized at a fixed (center) position, while its orientation can be controlled by means of an interaction cube. A pen is used as selection/steering device. The tip size of the virtual pen can be varied to control the level of difficulty of the tasks.

3. Experiments

3.1. Conditions

The following experimental conditions were tested:

1. PSS versus PAS

The PSS provides users with access to virtual information by positioning the user's hands within the virtual environment. The 3D interaction elements are co-located with the virtual information they manipulate. For the PSS used in the experiments, the mirror is opaque. Hence, the user has only kinesthetic awareness of the interaction elements.

The PAS promotes the idea of using interaction devices for manipulating objects in virtual environment that is observed from the outside. The 3D interaction elements are remote from the virtual information they manipulate. The user can observe the interaction elements and the hands.

2. Presence (H) or absence (nH) of hardware-supported depth cues

Hardware supported depth cues include stereoscopic

viewing and motion parallax. In order to accomplish stereo rendering, subjects need to wear shutter glasses. Motion parallax is provided by direct coupling of a head tracker with the camera position within VMD. Moving the head changes the view on the molecule.

3. Presence (S) or absence (nS) of software-supported depth cues

Software-supported depth cues include a stage and projected shadows. The stage consists of a checkerboard at the bottom of the virtual environment. It serves as reference frame for the interaction space and should enhance depth perception. For shadows a virtual light source is positioned above the molecule. It projects colored shadows onto the stage, see Figure 2. Shadows of the input devices are also drawn.

Combining the H/nH condition with the S/nS condition leads to four conditions to be tested for each of the two systems (PSS/PAS).

In all conditions, three levels of difficulty are created by varying the radius of the tip of the selection pen. The following three radii are used: 2, 3.5 and 5 millimeter.

3.2. Tasks

We designed three different tasks with distinguishing characteristics.

Selection task The goal in the selection task is to select atoms in the molecule.

At the start of a trial one alpha helix (red spiral in the visualization) is highlighted by displaying two green spheres with a radius of 3 millimeter at both ends of the helix. Timing starts when the subject touches one of these spheres with the pen, upon which the selected sphere disappears. The subject now has to select the other sphere as fast as possible. When selected the second sphere disappears, time is recorded, and a new trial starts.

Only four of the helices in our molecule are used, resulting in distances between points to be selected from 8.3 to 10.1 centimeters. These four helices combined with the three levels of difficulty result in twelve trials (in semi-random order).

Sheet steering task The goal of the sheet steering task is to trace a curved path within the molecule.

At the start of a trial in this task one beta sheet (yellow arrow in the visualization) is highlighted by displaying two green spheres with a radius of 3 millimeter at both ends of the sheet. Timing starts when the subject touches one of these spheres, upon which the selected sphere disappears. The subject now has to trace a path over the beta sheet to the other end point, as accurately as possible. This means that at all times, the tip of the pen should stay in touch with the yellow arrow. As beta sheets are not planar, the path to

be followed is not a straight line, but contains one or more curves. When the sphere at the other end of the beta sheet is selected, it disappears, time is recorded, and a new trial starts.

Only four of the beta sheets in our molecule are used, resulting in (curved) path lengths varying from 4.3 to 5.6 centimeters. These four sheets combined with three levels of difficulty result in twelve trials (in semi-random order).

Helix steering task The goal of the helix steering task is to trace a straight path within the molecule.

At the start of a trial, one alpha helix is highlighted using the same mechanism as in the case of the selection task. Timing starts when the subject touches the sphere at one of the ends of the helix, upon which the selected sphere disappears. The subject now has to trace a path through the helix to the other end point, as accurately as possible. This means that at all times, the tip of the pen should stay completely within the red spiral. When the second sphere is selected, it disappears, time is recorded, and a new trial starts.

The same four helices as in the case of the selection task are used. These four helices combined with three levels of difficulty result in twelve trials (in semi-random order).

3.3. Subjects

Twenty subjects participated in the experiment. Ten subjects performed the experiments with the PSS and ten with the PAS. All subjects were experienced computer users (students).

Each subject performed all three tasks in all four conditions. The order of task was fixed: selection - helix - sheet. The order in which the subjects went through the four conditions (nH/nS, H/nS, nH/S, H/S) was randomized across subjects.

3.4. Procedure

Before the actual experiment started, subjects were instructed about what they were expected to do during the experiment. Subjects were shown how they could rotate the molecule and how they could use the pen to make selections. On indication that they knew how to manipulate the pen and rotate the molecule, a training session was started. In this session all three tasks were performed, and, whenever needed, instructions were provided on how to best perform the tasks. During the training session, the subject could explore how to perform the tasks most conveniently.

After the training session, the three tasks were performed in all four conditions. At the start of a task, the user was instructed whether to do the task as fast as possible (selection task) or as accurate as possible (helix and sheet steering tasks).

Times for each trial (time between selection of first and

second sphere) within a task were recorded. In addition, for steering tasks, we have defined an accuracy measure as the average error between the target and the pen during the trial. For the sheet steering task, the average error is defined as the average Euclidian distance between the surface of the sphere on the pen tip and the closest point on the beta sheet. In case of the helix steering task, average error is defined as the Euclidian distance between the point on the surface of the sphere on the pen tip that is furthest removed from the axis of the helix to the (cylindrical) outer surface of the helix (this distance is hence zero if the sphere is within the helix).

4. Results

In each of the three experiments (i.e., point selection, and sheet and helix steering), we collected 960 trials. Ten subjects performed the experiments on the PSS, while ten other subjects performed identical experiments on the PAS. The 48 trials of a single subjects consisted of 4 sets of 12 trials. The 4 sets corresponded to the four possible combinations of hardware-supported depth cues (stereo and head tracking) and software-supported depth cues (shadows and perspective). The 12 trials within a set contained 4 interaction tasks for each of the three different tip sizes (of 2, 3.5 and 5 mm). This tip size was the major factor influencing the difficulty of the task. The experimental data for different tip sizes will therefore be analyzed separately.

Statistical analysis of the time duration T revealed that $\log T$ could be approximated well by a Gaussian distribution for all tasks, in all conditions. From the estimated standard deviation of $\log T$ a 95% confidence interval could be constructed for the average $\log T$, using the t-distribution [DS01]. The average values and the confidence intervals on $\log T$ were subsequently mapped to duration T to allow for an easy interpretation of the results. The top panel of figure 3 summarizes the results for the "selection" task. You can observe 8 clusters of three data points, with 95% uncertainty intervals for the estimated means. The 4 leftmost clusters correspond to the 4 different combinations of depth cues in the PSS, while the 4 rightmost clusters correspond to similar combinations in the PAS. The three data points within a cluster are for the different tip sizes. The tip sizes of 2, 3.5 and 5 mm are represented by the leftmost, middle and rightmost points in a cluster, respectively. Corresponding plots for the "sheet" steering task and the "helix" steering task can be found in the other panels of figure 3.

The effect of the different factors on the task trial times can most easily be described by a linear regression model. More specifically, it was established that the following model

$$\log T = \log T_o + \log f_H \cdot H + \log f_S \cdot S + \log f_M \cdot M + \log f_{HS} \cdot HS,$$

submitted to IPT & EGVE Workshop (2005)

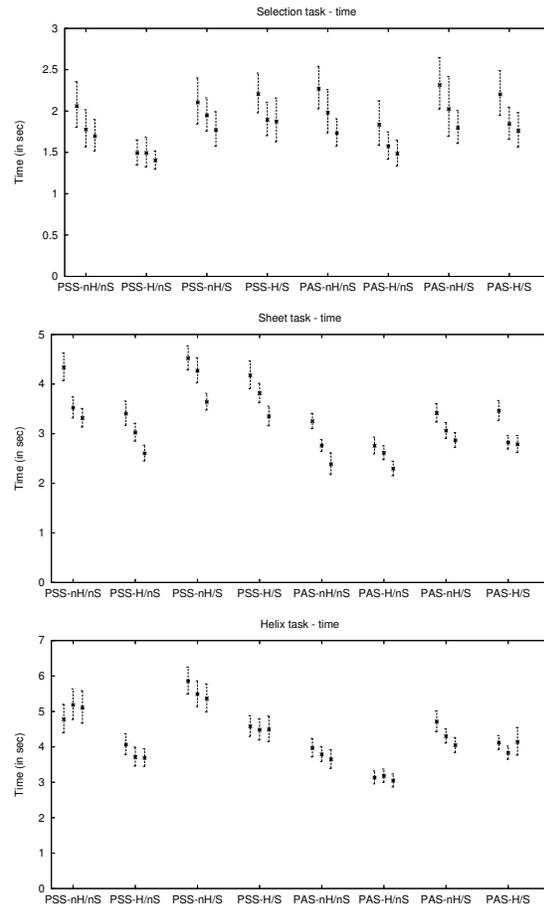


Figure 3: Task trial times with 95% confidence intervals. Top: selection task. Middle: sheet task. Bottom: helix task.

or equivalently,

$$T = T_o \cdot (f_H)^H \cdot (f_S)^S \cdot (f_M)^M \cdot (f_{HS})^{HS},$$

provided a good description for the trial times for all three different tip sizes. The parameters in this model are defined as follows:

- T_o is the average duration in the PSS condition without any depth cues,
- f_H is the trial time reduction factor resulting from hardware-supported depth cues,
- f_S is the trial time reduction factor resulting from software-supported depth cues,
- f_M is the reduction factor caused by moving from the PSS to the PAS,
- f_{HS} describes the interaction between both kinds of depth cues,
- H is 1 in case of hardware-supported depth cues (H), 0 otherwise (nH)

- S is 1 in case of software-supported depth cues (S), 0 otherwise (nS)
- M is 1 in case of the PAS, and 0 in case of the PSS.

We did not only estimate the regression variables themselves but also their 95 % confidence intervals (indicated by [low,high] in the tables). The results are summarized in table 1 for the "selection", "sheet" and "helix" experiments, respectively.

In case of the "selection" and the "sheet" experiment, the interaction task becomes easier as the tip size increases, and the corresponding trial time (T_o) hence decreases with tip size. In case of the "sheet" task, increasing tip sizes result in a more difficult task. Somewhat unexpectedly, varying tip size has however little or no effect on the duration of the interaction in the "helix" experiment. As expected, we obtain a significant effect of hardware-supported depth cues in all cases. The reduction factor f_H varies between 0.77 and 0.90. Only in a few cases do we observe a (slightly) significant effect of software-supported depth cues. In case of significant effects, the reduction factor f_S is however larger than 1, indicating that such depth cues increase rather than decrease the duration of the interaction. In cases where f_{HS} is significantly different from 1, it is also greater than 1. This demonstrates that software-supported depth cues are even more harmful when used in combination with hardware-supported depth cues than in cases where such cues are absent. Finally, there is also a significant difference between the PSS and the PAS systems for the steering tasks, where the reduction factor f_M varies between 0.77 and 0.83.

It is well known that people can trade time versus accuracy in an interaction task (as is for instance expressed in the familiar Fitts' law). Therefore, in order to better understand the performance of subjects within the different interaction tasks, we need to also analyze this accuracy. Since there is as yet no widely accepted metric for measuring the accuracy for 3D interaction tasks, we have adopted as accuracy measure the average error (D) between the virtual pen tip and the target.

Statistical analysis of the average error D revealed that $\log D$ could be more closely approximated by a Gaussian distribution than D itself. The average values and the 95% confidence intervals on $\log D$ were mapped to error D to allow for an easy interpretation of the results. Figure 4 summarizes the results for the "sheet" and the "helix" steering task, respectively.

It was established that the following model

$\log D = \log D_o + \log f_H \cdot H + \log f_M \cdot M + \log f_{HSM} \cdot HSM$,
or equivalently,

$$D = D_o \cdot (d_H)^H \cdot (d_M)^M \cdot (d_{HSM})^{HSM},$$

provides a good description for the accuracy measure for all three different tip sizes. The parameters in this model are defined as follows:

select	T_o	f_H	f_S	f_{HS}	f_M
2 mm					
estimate	2.06	0.765	1.020	1.305	1.305
low	1.87	0.678	0.904	1.099	1.099
high	2.27	0.864	1.152	1.550	1.550
3.5 mm					
estimate	1.84	0.817	1.058	1.151	1.044
low	1.67	0.724	0.938	0.971	0.959
high	2.02	0.922	1.194	1.365	1.137
5.0 mm					
estimate	1.71	0.841	1.040	1.211	1.008
low	1.57	0.755	0.934	1.040	0.934
high	1.86	0.937	1.158	1.410	1.087
sheet	T_o	f_H	f_S	f_{HS}	f_M
2.0 mm					
estimate	4.24	0.815	1.046	1.186	0.785
low	3.86	0.725	0.930	1.004	0.722
high	4.65	0.917	1.177	1.402	0.853
3.5 mm					
estimate	3.55	0.901	1.159	1.007	0.774
low	3.26	0.811	1.043	0.868	0.718
high	3.85	1.002	1.288	1.170	0.834
5.0 mm					
estimate	3.14	0.869	1.149	1.088	0.803
low	2.85	0.768	1.015	0.914	0.736
high	3.46	0.983	1.299	1.295	0.876
helix	T_o	f_H	f_S	f_{HS}	f_M
2.0 mm					
estimate	4.80	0.819	1.206	1.009	0.825
low	4.33	0.720	1.060	0.841	0.753
high	5.31	0.934	1.372	1.211	0.904
3.5 mm					
estimate	4.94	0.775	1.096	1.100	0.805
low	4.48	0.685	0.968	0.923	0.737
high	5.45	0.878	1.241	1.310	0.879
5.0 mm					
estimate	4.82	0.777	1.080	1.191	0.800
low	4.29	0.670	0.931	0.967	0.721
high	5.42	0.900	1.251	1.468	0.888

Table 1: Regression parameters for trial times. Top: selection task. Middle: sheet task. Bottom: helix task.

- D_o is the average error in the PSS condition without any depth cues,
- d_H is the error reduction factor resulting from hardware-supported depth cues,
- d_M is the error reduction factor caused by moving from the PSS to the PAS,
- d_{HSM} describes the interaction between both kinds of depth cues on the PAS,

We did not only estimate the regression variables themselves but also their 95 % confidence intervals. The results

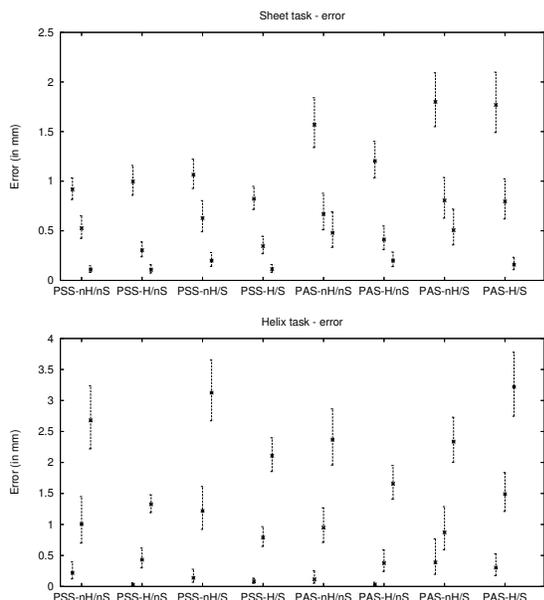


Figure 4: Average error with 95% confidence intervals. Top: sheet task. Bottom: helix task.

are summarized in table 2 for the "sheet" and "helix" experiments, respectively.

In case of the "sheet" experiment, the interaction task becomes easier as the tip size increases, and the corresponding error (D_o) hence decreases with tip size. In case of the "helix" task, increasing tip sizes result in a more difficult task. As expected, we obtain a significant effect of hardware-supported depth cues in all cases. The reduction factor d_H varies between 0.21 and 0.83. There is also a significant difference between the PSS and the PAS systems for most of the interaction tasks. The reduction factor f_M varies between 1.28 and 7.76, which illustrates that the accuracy is significantly less in the PAS than in the PSS. Note that there was no primary effect of software-supported depth cues (S) on the accuracy.

In order to visualize the trade-off between trial time and accuracy, we have plotted one against the other in Figure 5. Points in the plot represent the relation between time and accuracy with their confidence intervals (open boxed points are data taken from the PSS, diagonal crossed points are data taken from the PAS). The three different point symbols correspond to the three different pen tip sizes in each environment. Solid lines connecting corresponding points show the relation between speed and accuracy within one environment. Dashed lines connecting corresponding points show the relation between speed and accuracy across environments. Especially in the sheet task (left plot), the PAS seems to invite a faster interaction with a lower accuracy. The helix task seems to be more compelling, in the sense

sheet	D_o	d_H	d_M	f_{HSM}
2.0 mm				
estimate	1.02	0.828	1.547	1.336
low	0.85	0.659	1.231	0.920
high	1.23	1.040	1.945	1.941
3.5 mm				
estimate	0.55	0.532	1.278	2.034
low	0.39	0.354	0.850	1.046
high	0.76	0.799	1.920	3.956
5.0 mm				
estimate	0.04	0.254	4.542	0.492
low	0.02	0.081	1.458	0.077
high	0.09	0.789	14.148	3.143
helix	D_o	d_H	d_M	f_{HSM}
20 mm				
estimate	0.20	0.211	7.756	0.945
low	0.08	0.075	1.433	0.336
high	0.46	0.593	41.982	2.658
35 mm				
estimate	1.16	0.478	3.595	0.747
low	0.77	0.290	1.592	0.453
high	1.75	0.787	8.122	1.230
50 mm				
estimate	2.78	0.626	2.103	0.880
low	2.28	0.491	1.417	0.691
high	3.39	0.797	3.122	1.121

Table 2: Regression parameters for accuracy. Top: sheet task. Bottom: helix task.

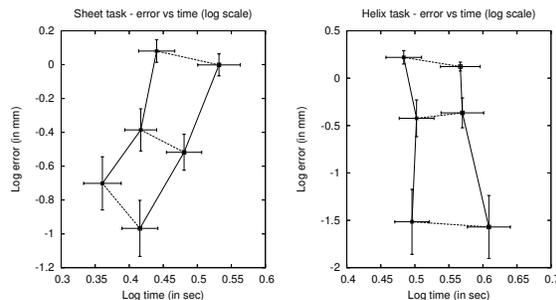


Figure 5: Trial time versus average error with 95% confidence intervals, under the condition H/nS. Left: sheet task. Right: helix task.

that the obtained accuracies are comparable in both systems. Despite this, the PAS system is still faster than the PSS system.

5. Discussion and Conclusions

Effective interaction in a virtual environment requires that the user can adequately judge the spatial relationships between the objects in a 3D scene. However, the benefits of

various depth cues come at a cost, both in terms of required hardware components and in terms of required user adjustment. It is therefore important to quantify the repercussions of dropping one or more of the above hardware technologies. In this paper, we have reported on a formal experiment in which the effects of different hardware components on the speed and accuracy of three-dimensional (3D) interaction tasks with molecular structures are established. We selected three application-specific 3D interaction tasks that are typical for exploring the structure of molecules: a point selection task (*atom picking*), a curved path tracing task (*sheet steering*), and a straight path tracing task (*helix steering*). From the experimental data, we have developed linear regression models to predict the speed and accuracy for the three interaction tasks.

We now discuss the most important conclusions:

- As expected, the hardware-supported depth cues (stereoscopic viewing, motion parallax) contribute significantly to task performance. The related reduction factor, f_H in tables 1 and d_H in table 2, indicate that hardware stereo and head tracking can reduce trial times by a factor of up to 25% and improve accuracy by a factor of 20% for steering tasks.

In this study, we have not examined the relative importance of stereo and head tracking.

- Surprisingly, the software-supported depth cues (shadows and perspective viewing) do not contribute to speed and accuracy ($f_S > 1.0$ and $f_{HS} > 1.0$ in table 1 whenever significant).

This agrees with the results of [HWSB99], who state that many shadows have negative impact compared with only small number of shadows. In the future we plan to simplify the shadow rendering, so that only shadows of the relevant objects, such as device, target objects, and molecular backbone representations, are projected.

- Users in the fish-tank like environment (PAS) performed no worse than users in the co-located environment (PSS) for steering tasks. Figure 5 shows that, for the sheet task, trial time is shorter in the PAS at the cost of reduced accuracy. For the helix task, PAS is faster without loss of accuracy.

References

- [ABW93] ARTHUR K., BOOTH K., WARE C.: Evaluating 3D task performance for fish tank virtual worlds. *ACM Transactions on Information Systems* 11, 3 (July 1993), 239–265. 1
- [AW04] ARSENAULT R., WARE C.: The importance of stereo and eye-coupled perspective for eye-hand coordination in fish tank vr. *Presence* 13, 5 (2004), 549–559. 3
- [BB97] BORITZ J., BOOTH K. S.: A study of interactive 3D point location in a computer simulated virtual environment. In *ACM Symposium on Virtual Reality Software and Technology* (1997), pp. 181–187. 1
- [DS01] DRAPER N., SMITH H.: *Applied Regression Analysis*. John Wiley and Sons, 2001. 5
- [HDS96] HUMPHREY W., DALKE A., SCHULTEN K.: Vmd - visual molecular dynamics. *Journal of Molecular Graphics* 14, 1 (1996), 3–38. 3
- [HWSB99] HUBONA G., WHEELER P., SHIRAH G., BRANDT M.: The role of object shadows in promoting 3d visualization. *ACM Transactions on Computer-Human Interaction* 6, 3 (July 1999), 214–242. 8
- [Kel93] KELSEY C.: *The Perception of Visual Information*. Springer-Verlag, 1993, ch. Detection of visual information, pp. 30–51. 1
- [LLM*02] LIERE R., LEEUW W., MULDER J., VERSCHURE P., VISSER A., MANDERS E., DRIEL R.: Virtual reality in biological microscopic imaging. In *Proceedings of the IEEE International Symposium on Biomedical Imaging 2002* (2002). 2
- [LM03] LIERE R., MULDER J.: Optical tracking using projective invariant marker pattern properties. In *Proceedings of the IEEE Virtual Reality Conference 2003* (2003), pp. 191–198. 2
- [MK94] MILGRAM P., KISHINO F.: A taxonomy of mixed reality visual displays. In *IEICE Transactions on Information Systems* (1994), pp. 1321–1329. 3
- [ML02] MULDER J., LIERE R.: The personal space station: Bringing interaction within reach. In *VRIC 2002 Conference Proceedings* (2002), pp. 73–81. 3
- [MW96] MESSERSCHMIDT A., WEVER R.: X-ray structure of a vanadium-containing enzyme: Chloroperoxidase from the fungus *curvularia inaequalis*. *Biochemistry* 93, 1 (January 1996), 392–396. 3
- [Sch83] SCHMANDT C.: Spatial input/display correspondence in a stereoscopic computer graphic work station. *Computer Graphics* 17, 3 (1983), 253–261. 3
- [WAB93] WARE C., ARTHUR K., BOOTH K.: Fish tank virtual reality. In *INTERCHI '93 Conference Proceedings* (1993), Ashlund S., Mullet K., Henderson A., Hollnagel E., White T., (Eds.), pp. 37–42. 3
- [WFG92] WANGER L., FERWERDA J., GREENBERG D.: Perceiving spatial relationships in computer-generated images. *IEEE Computer Graphics and Applications* 12, 3 (1992), 44–58. 1
- [WSS99] WIEGAND T., SCHLOERB D., SACHTLER W.: Virtual workbench: Near field virtual environment system with applications. *Presence* 8, 5 (1999), 492–519. 3