Human-Computer Interaction -- INTERACT'03
M. Rauterberg et al. (Eds.)
Published by IOS Press, (c) IFIP, 2003, pp. 916-919

Empirical Evaluation of Performance in Hybrid 3D and 2D Interfaces

Sriram Subramanian, Dzmitry Aliakseyeu and Jean-Bernard Martens

User Centered Engineering group

Faculty of Industrial Design, Eindhoven University of Technology

s.subramanian, d.aliakseyeu, j.b.o.s.martens@tue.nl

Abstract: Experimental studies of spatial input devices have focused on demonstrating either the superiority of 3D input devices over 2D input devices, or the superiority of bimanual interaction over unimanual interaction. In this paper, we argue that hybrid interfaces that combine a 3D input device with a 2D input device have received little attention up to now and are potentially very useful. We demonstrate by means of an experimental evaluation that working with hybrid interfaces can indeed provide superior performance compared to strictly 3D and 2D interfaces.

Keywords: 3D input device, 2D input device, interaction techniques, allocation of control, evaluation methods, hybrid interfaces, Human-Computer Interaction.

1 Introduction

Technological improvements and cost reductions in computing power, display and sensor technology have resulted in an increasingly widespread use of 3D input devices. As a result there is a growing need to understand when and how these 3D input devices perform best. A common belief of 3D user interface design is that, because an application requires manipulating objects in 3D space, the best interface should necessarily also use 3D input devices.

In the last two decades several experiments on spatial input devices have been performed. One of the focal points of these investigations has been on understanding bimanual interaction. Experiments in bimanual interaction have mostly concentrated on showing that the Kinematics Chain (KC) model (Guiard, 87), which for example says that the non-dominant hand acts as a reference for, and precedes the dominant hand in a bimanual task, also holds for virtual object manipulations in interactive environments (Balakrishnan and Hinckley 99, Hinckley et al 98). There have also been experiments to show that bimanual interaction, when guided by the KC model, can be superior to unimanual interaction, both in the case of 2D input devices (Leganchuk et al, 98) and 3D input devices (Gribnau and Hennessey, 98). In this latter reference, the authors showed that two-handed 3D input achieved lower trial completion times than one-handed 3D input. However, in the one-handed case the subjects only had one 3D input device and were only allowed 6 Degrees of Freedom (DOF), whereas in the two handed case the subjects were allowed 12 DOF. This means that the benefit could either have been derived from bimanual interaction or from greater freedom of control (12 vs. 6 DOF).

An observation that arises from these experiments in bimanual interaction is that increasing the number of available DOF to more than what is strictly necessary can enhance performance. However, as noted by Hinckley (Hinckley, 94), it also makes good sense to exploit task-specific needs to reduce dimensionality. For example, if the user's task consists only of orienting an object, it makes little sense to allow simultaneous translation, since it only makes the user's task more difficult: the user must simultaneously orient the object and keep it from moving outside the field of view.

The emerging question that is addressed in this paper is whether it is better to increase or reduce the available DOF when performing a 3D interaction task? We except a priori that the outcome will point somewhere in the middle, i.e., that it can be advantageous to allow users more DOFs than what is strictly required for the task, while at the same time allowing them to constrain some of the DOFs as they progress with the task. This constraining of the DOF in a task implies that some sub-tasks will require only limited DOFs (say 1, 2 or 3 DOF). For such sub-tasks, 2D input devices are potentially more effective, in terms of precision and speed, than 3D input devices. The advantages of both approaches may even be combined by handling a 3D device in one hand and a 2D device in the other hand.

In Hinckley (Hinckley, 94) and Bowman (Bowman, 2001) the authors conform to this view that *hybrid interfaces*, that combine a 3D input device with a 2D input, should provide superior performance over strictly 3D or 2D interfaces. However, all experimental comparisons of bimanual interaction known to us either involve only 2D input devices or 3D input devices. We are not aware of any experiments so far with hybrid interfaces.

From the description above, the following relevant questions arise:

- Can we enhance the users performance by allowing them to introduce constraints in the control of an input device?
- Can a combination of a 3D input device with a 2D input device enhance user performance?

This paper summarizes the results of an experiment that we have conducted to start addressing these questions.

In order to investigate the above questions through empirical evaluation we needed a suitable experimental task. Two criteria played a role in selecting this task. First, we required the task to be practically relevant (for instance, for medical planning), so that the results of the experiments are potentially useful to the design of practical interaction techniques. Second, we wanted the task to be sufficiently powerful to bring out the differences (if any) between the alternative interaction techniques.

2 Experimental Platform and Task

The experiment was conducted on the Visual Interaction Platform (Aliakseyeu et al, 2000) (see Figure 1). The platform has a workspace (6 in Figure 1) where all the users actions are performed, and appropriate visual feedback is provided, and a communication space (5 in Figure 1) where additional visual feedback, such as perspective 3D rendering, is provided. Figures 2 and 3 provide more detailed layouts of these spaces in the specific case of our experimental task. The communication space displays the surface rendering of a solid 3D model (see 1 in Figure 2) and a bounded intersection plane (2 in Figure 2) that can be moved through the 3D model to create dynamically generated cross-sections (3 in Figure 2). In the workspace a Maximum Intensity Projection (MIP) (from the top) of the 3D model (3 in Figure 3) is provided at the location of the 3D model (1 in Figure 3). A dynamically updated outline of the top view of the intersection window (2 in Figure 3) is also overlaid on the projection. The intersection plane can be controlled using the Rigid Intersection Selection Prop (RISP) (Aliakseyeu et al, 2002). The intersection plane can be moved through the model by moving the RISP above the table at the location of the window (3 in Figure 1).

The workspace also supports an interaction through virtual *pins*. When a pin from the pin holder is placed at any point in the cross-section image (6 in Figure 3), the intersection plane is pinned to the location within the 3D model corresponding to that point. The pin can be positioned using a digital pen.

The experimental task requires the users to navigate the intersection plane through the solid 3D model in order to locate a dark-gray disk hidden inside the body of the model. The location of the disk is not visible in the surface rendering or MIP of the 3D model (see 1 in Figure 2). The goal of the user is to dock the intersection plane with the dark gray disk (visible in 3 in Figure 3). The task was completed when the conductor of the experiment can clearly see the dark-gray circular disk and the subject indicates that s/he has *finished*.

This task was chosen because of its similarity to the kind of tasks encountered in the medical domain, for example when trying to intersect with an aneurysm or blood vessel. The radius and thickness of the disk allow to control the difficulty of the task.

3 Interaction Techniques

Free Movement (F)

In this case the RISP is moved freely in 3D space and both its position and orientation have to match the disk position and orientation to successfully complete the task.

Combined 3D and 2D Techniques (FR, PR)

The obvious choices for constraining the users movement are to let the intersection plane respond to

- Only the rotational movements of the RISP, i.e., the positional changes in the RISP are ignored.
- Only the positional movements of the RISP, i.e. the changes in the orientation of the RISP are ignored.

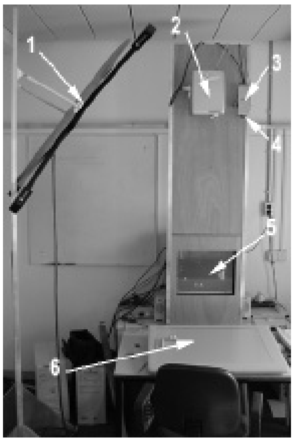

Figure 1: Visual Interaction Platform. 1–3D tracker, 2–LCD Projector, 3-Camera, 4-Infrared light source, 5-CRT screen (Communication Space) 6- Wacom Tablet with projection image (work space).

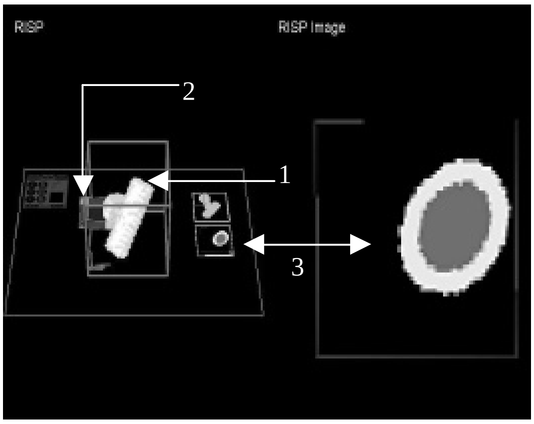

Figure 2: The Communication Space. 1-Solid 3D model, 2 – Intersection Plane, 3 – Cross-section generated.

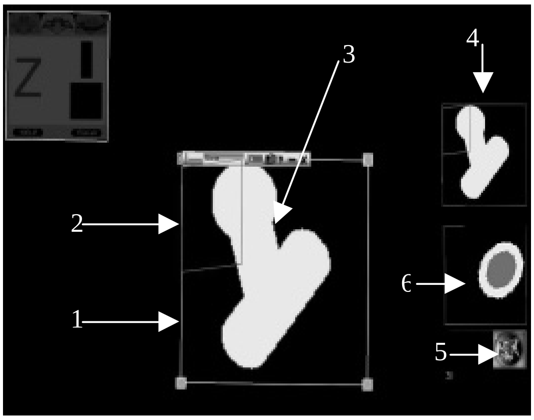

Figure 3: The Workspace. 1 – Location of the 3D model, 2 – Outline of the Intersection Plane, 3 – The Maximum Intensity Projection of the 3D model, 4 – A Copy of 2 and 3, 5 – Pins, 6 – Cross-section Image

Free Movement followed by Rotation (FR):

In this case the intersection window at first responds to all movements of the RISP. After the user has specified a point of rotation (by inserting a pin) the window responds only to the rotation of the RISP.

Positioning first followed by Rotation (PR):

In this case the intersection window first responds to the positional changes of the RISP. Once the user has specified a point of rotation (by inserting a pin), the window responds only to the rotation of the RISP.

In both above techniques, at the instant when the user wants to pin a certain point, s/he has to keep the RISP fixed in the desired position while dragging the pin to its destined location. This action actually tests the ability of the users to work simultaneously with a 3D prop and a 2D prop. Both techniques combine a 3D and a 2D prop along with the possibility of constraining the DOFs.

Menu based Navigation (MPR):

An advantage of *FR* and *PR* over *F* could be due to either the combination of 3D with 2D or to the ability to constrain the DOFs. Therefore, we introduced a fourth technique that exploits constraints, but only uses 2D interaction with the pen and tablet. In this method the user can first position the intersection plane using a menu function. The user can move the intersection plane up or down the axis perpendicular to the table using a 1D slider or move sideways, parallel to the table, using a 2D area pad. Once the user specifies a point of rotation (using the pin) the menu provides an Arcball controller (Shoemake, 1992) to modify the orientation of the intersection plane.

4 Experimental Design

Eight right-handed subjects (6 males and 2 females) participated in this experiment. Each participant performed all four techniques in a within-participant design. The technique order was counter-balanced using a Latin square design: 2 participants each performed in the order (*F, FR, PR, MPR),* (*FR, PR, MPR, F*), (*PR, MPR, F, FR)* and (*MPR, F, FR, PR).*

The subjects were given written instructions for the experiment. Following this, they were given verbal instructions outlining the properties of each technique, after which they attempted each technique in a practice block of 4 trials. This instruction phase took 30 minutes. After practicing, the participants performed 4 blocks of trials. The first two blocks consisted of 4 trials with each interaction technique and the second two blocks consisted of 5 trials with each interaction technique, resulting in a total of 18 trials per interaction technique. Participants took approximately one hour to complete the 72 trials. At the end of the experiment the participants had to fill an exit questionnaire to elicit subjective feedback on the ease of use and fatigue of each interaction technique.

5 Results and Conclusion

The total trial time was used as the primary measure of performance. Figure 4 shows the cumulative probability distribution of time for the different interaction techniques. Table 1 shows the means and standard deviations of the completion times for the different interaction techniques. Techniques FR and MPR did not differ significantly. Significant differences occurred between *PR* and *MPR* ($F_{1,\ 142} = 12.16$, $p<0.001$), and between *FR* and *F* ($F_{1,\ 142} = 10.042$, $p<0.003$).

Table 1: **Mean and Standard Deviation (SD) of Trial Completion times for different interaction techniques**

	F	**FR**	**PR**	**MPR**
Mean	20.09	16.37	12.65	16.17
SD	8.58	7.85	6.52	7.42

The unconstrained free movement was the slowest of the different interaction techniques, whereas starting with positioning followed by rotating was the fastest. At the time of conceiving the experiment we expected a priori that the FR technique would perform best, since the participant is free to get as precisely as possibly to the target and then drop off excess DOFs to close in on the target. However, this expectation is contradicted by the experiment.

The experiment did highlight that exploiting 2D and 3D input devices as well as introducing suitable constraints into the task can be an effective strategy for improving user performance.

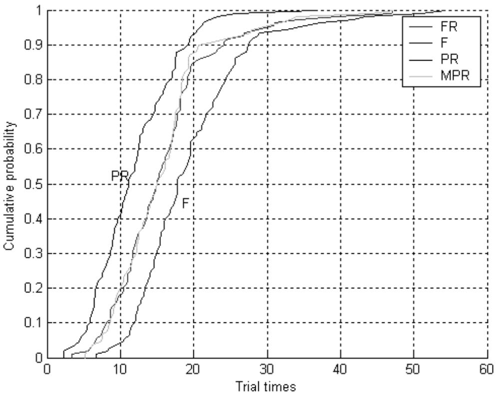

Figure 4: Cumulative distribution of trial times for different interaction techniques (PR - fastest, MPR, FR – intermediate, F – slowest)

References

Aliakseyeu D., Martens, J.-B., Subramanian, S., Vroubel, M., and Wesselink, W. (2001) Visual interaction platform. *Proceedings of Interact 2001,* Tokyo, Japan. July 2001, pp. 232-239

Aliakseyeu, D., Martens, J.-B., Subramanian, S., Rauterberg, R. (2002) Interaction Techniques for Navigation through and Manipulation of 2D and 3D Data, i*n Proc. of Eigth Eurographics Workshop on Virtual Environments*, Barcelona, pp. 179-199, 2002.

Balakrishnan, R., & Hinckley, K. (1999). The role of kinesthetic reference frames in twohanded input performance. *In Proceedings of the 12th annual acm symposium on user interface software and technology*. ACM Press, pp. 171–178

Bowman, D., Kruijff, E., LaViola, J., and Poupyrev, I. (2001) An Introduction to 3D User Interface Design. *Presence: Teleoperators and Virtual Environments*, vol. 10, no. 1, 2001, pp. 96-108.

Gribnau, M. W., & Hennessey, J. M. (1998). Comparing single- and two-handed 3d input for a 3d object assembly task. *In Chi 98 conference summary on human factors in computing systems*, ACM Press, pp. 233–234.

Guiard, Y. (1987). Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *In Journal of motor behavior* , pp. 486-517.

Hinckley, K., Pausch, R., Goble, J. C., & Kassell, N. F. (1994). A survey of design issues in spatial input. In *Proceedings of the Symposium on User Interface Software and Technology,* ACM Press, pp. 213–222.

Hinckley, K., Pausch, R., Proffitt, D., and Kassell, N. (1998). Two-handed virtual manipulation. *ACM Transactions on CHI 5*, pp 260-302.

Leganchuk, A., Zhai, S., & Buxton, W. (1998). Manual and cognitive benefits of twohanded input: an experimental study. *ACM Transactions on Computer-Human Interaction (TOCHI),* 5 (4), pp. 326–359.

Shoemake, K., (1992) “ARCBALL: A User Interface for Specifying Three-Dimensional Orientation Using a Mouse”, *Graphics Interface*, 1992, pp. 151-156