

The Cubic Mouse

A New Device for Three-Dimensional Input

Bernd Fröhlich John Plate

GMD/IMK.VE

D-53754 Sankt Augustin, Germany

{Bernd.Froehlich, John.Plate}@gmd.de

ABSTRACT

We have developed a new input device that allows users to intuitively specify three-dimensional coordinates in graphics applications. The device consists of a cube-shaped box with three perpendicular rods passing through the center and buttons on the top for additional control. The rods represent the X, Y, and Z axes of a given coordinate system. Pushing and pulling the rods specifies constrained motion along the corresponding axes. Embedded within the device is a six degree of freedom tracking sensor, which allows the rods to be continually aligned with a coordinate system located in a virtual world. We have integrated the device into two visualization prototypes for crash engineers and geologists from oil and gas companies. In these systems the Cubic Mouse controls the position and orientation of a virtual model and the rods move three orthogonal cutting or slicing planes through the model. We have evaluated the device with experts from these domains, who were enthusiastic about its ease of use.

Keywords

User interface hardware, two-handed interaction, virtual reality

INTRODUCTION

Many graphics applications require the input of three-dimensional coordinates to position objects in a virtual world. Desktop applications typically use a mouse, trackball or a more exotic device like a dial box for such input. This works reasonably well as long as the coordinate system of the world is aligned with the screen. However, once the virtual world is rotated the mapping of mouse movements is often no longer intuitive, e.g. it might happen that the mouse is moved to the right, but the virtual object moves to the left. This is often quite confusing.

We present a new input device that allows intuitive input of three-dimensional coordinates in virtual environment applications. The device (Figure 1), which we named the Cubic Mouse, consists of a cube-shaped case, three rods,

and control buttons. Each rod passes approximately through the center of two parallel faces of the case. The rods are perpendicular to each other and movable. They represent the X, Y, and Z axes of a coordinate system. There is also a six degree of freedom (6DOF) tracker embedded in the cube-shaped case, which we use to orient and position the virtual world in three-space relative to the observer. In this way the rods stay aligned with the coordinate system axes. By pushing and pulling the rods we specify motions of virtual objects constrained along the X, Y, and Z axes. Typically users hold the device in their non-dominant hand to position and orient the world, while the dominant hand operates the rods and the control buttons.




Figure 1: The Cubic Mouse device

For the development of the Cubic Mouse we had two driving applications. Within a consortium of car manufacturers and crash software vendors we are developing an application prototype for steering and visualizing car crash simulations in virtual environments. Stereoscopic virtual environments like Caves [1], Responsive Workbenches [6,7] or large screen projections facilitate the understanding of complex three-dimensional deformations occurring during a car crash. The main tool for investigating simulated car crash results are cutting

planes and chair cuts. We use the Cubic Mouse for navigating around the car model, for positioning cutting planes inside the model (Figure 2), and for performing chair cuts.




Figure 2: A car crash visualization on the Responsive Workbench. The Cubic Mouse is aligned with the principal axes of the car model. The user moves one of the three orthogonal cutting planes from the top into the model using the appropriate rod.

Our second driving application stems from the geoscientific domain. We work with geologists and geophysicists from oil and gas companies to evaluate virtual environment technology for reservoir discovery and characterization. Their data is mostly based on seismic measurements and information acquired from the actual drilling of wells. This data is three-dimensional in nature and stereoscopic virtual environments allow users to explore and understand complex subsurface structures in three dimensions. One traditional method of exploring the seismic data is by moving three orthogonal slices through a 3D seismic volume. In our system these slices can be positioned intuitively with the three rods. We also experimented with the positioning of orthogonal slices in volumetric data sets from CT and MRI scans in a similar way. In all cases users were able to use the basic features of our device with only two or three sentences of introduction and they were immediately enthusiastic about its ease of use.

Our main contribution is the development of an intuitive device for the input of three-dimensional coordinates in interactive graphics applications. We describe the realization of the device, introduce a set of possible application domains, and discuss some experiences of experts from these domains and results from an initial user study. We also present some variations on the basic design of our device.

RELATED WORK

In most of our application scenarios, the Cubic Mouse serves as a coordinate system prop. From this perspective, the closest relative is the head prop in [5] for neurosurgical visualization. In their system, users hold a small rubber

sphere or a doll's head with an embedded tracker in one hand. This head prop is used to control the orientation of a head model on the screen. The other hand holds a second prop which, for example, is used to position a cutting plane relative to the head prop. This is in contrast to our system, where the dominant hand is used to manipulate controls located on the Cubic Mouse held in the non-dominant hand.

A variety of systems use two-handed interaction techniques based on hand-held widgets, e. g. in [8] users hold a virtual widget in one hand and operate it with the other. In [9] users hold a miniature model of the virtual world in one hand and manipulate objects in the miniature with the other. These systems do not employ real world props other than tracked wands or data gloves.

The TouchCube by ITU Research presented at the Siggraph'98 exhibition is a cube-shaped input device with touch sensitive faces. By applying certain gestures to the touch sensitive surfaces, objects are moved in a three-dimensional world. The version presented was not tracked. A tracking sensor could be easily added, but the device relies largely on being mounted on a stand. Otherwise it would be difficult to apply two finger gestures simultaneously to two opposite surfaces as they are suggested in the patent (US patent US-A-5 729 249). The key concept of the Cubic Mouse is that moving a rod into a certain direction results in the movement of a virtual object into exactly the same direction thus heavily relying on tracking.

At the University of North Carolina there was a system developed that used pairs of X, Y, and Z sliders mounted in the corresponding directions in a fixed world coordinate system to specify pairs of orthogonal clipping planes [Prof. Brooks, UNC, personal communication]. This setup is very similar to the Cubic Mouse, except that the Cubic Mouse can be used to specify clipping planes relative to an arbitrarily oriented object.

From another point of view, the Cubic Mouse is related to devices which allow the separate input of X, Y, and Z coordinates. For example, the dial box uses separate dials to specify constraint motions along the X, Y, and Z axis, but there is no intuitive connection between a given dial and the corresponding axis.

HARDWARE AND SOFTWARE IMPLEMENTATION

The Cubic Mouse was built using off the shelf parts. The cube-shaped case's edge length is 9 centimeters (3.5 inches) and was determined by the size of readily available potentiometers. The total weight of the device plus the 1.3 meters of cable is around 300 grams (2/3 of a pound). The case and the cables contribute most of the weight. They could be built with much lighter materials, which would reduce the weight considerably. The latest prototype has buttons mounted at both ends of each rod plus six application programmable buttons mounted on one of the Cubic Mouse's faces as shown in Figure 1. The two cables come off the opposite face. Most of our applications have a

natural up and down direction, so the face with the buttons becomes typically the “up” face and the cables come off the “down” face. Linear potentiometers are used to measure the positions of the rods. A Polhemus Fastrak sensor provides the spatial position and orientation information for the Cubic Mouse.

We built the device within a few days and integrated it into the Avocado graphics framework [11] without any difficulties. Avocado is based on SGI’s Performer toolkit and OpenGL. The system supported already an analog/digital (A/D) converter. After plugging the Cubic Mouse into the A/D box the corresponding button states and potentiometer values were immediately available in Avocado. Most of the programming was done in Scheme, Avocado’s scripting language. Avocado supports a variety of output devices, so we were able to experiment with the Cubic Mouse in a CAVE, on a two-sided Responsive Workbench system (a Responsive Workbench with an additional vertical screen at the back), and in a monitor based environment.

APPLICATION SCENARIOS

We experimented with the Cubic Mouse in four different application domains. Most of our experiences were collected from an application prototype developed for the oil and gas industry, where we use the Cubic Mouse for positioning three orthogonal slices in a seismic data set. We used the Cubic Mouse in a similar way for the visualization of CT and MRI data. Clipping planes and chair cuts are important tools in engineering visualization systems. We specify these cuts with the Cubic Mouse and presented our system to a crash engineer from a large automotive company.

Data Exploration for the Oil and Gas Industry

The oil and gas industry acquires enormous amounts of seismic data for the exploration of potential new reservoirs. This data has to be sighted by geologists and geo-physicists to discover the precious oil and gas containing subsurface structures. The raw seismic data is processed into regular three-dimensional arrays - the seismic cubes. These seismic cubes along with information obtained from drilled wells are the basis for building a subsurface model. Geo-engineers roam through the seismic cubes to find areas of interest and model subsurface structures within these areas. The whole process is much more complicated than this, but these two phases are very important and they are repeated many times.

Often geo-engineers visualize their data sets by rendering subsurface structures as polygonal models and by representing the seismic volume as three orthogonal slices. A typical data set is shown in Figure 3. The Cubic Mouse allows geo-engineers to look at the data set from different directions while moving the slices through the seismic volume. This is realized in the following way: The seismic cube’s orientation is always kept in sync with the Cubic Mouse’s orientation and the rods move the seismic slices.

This way the rods stay always perpendicular to the slicing planes, which makes it very easy and intuitive to find the desired rod without looking at the device.




Figure 3: A typical oil exploration data set containing subsurface structures, wells, and seismic slices. The subsurface model consists of two main structures: horizons and faults. Horizons separate two earth layers, and faults are breaks in the rocks, where one side is moved relative to the other. Horizons are typically horizontal structures and faults are typically vertical structures. Three orthogonal slicing planes are used to visualize the seismic volume. Typically, one slice, the so called inline-slice, is oriented perpendicular to the main fault direction. The time-slice is oriented horizontally and the crossline-slice is perpendicular to both.

Two buttons are used for scaling the model up and down. A third button serves as a clutch, which allows users to freeze the model in the current orientation and lay down the Cubic Mouse. In our system, which was originally developed for the two-sided Responsive Workbench, the Cubic Mouse’s translation was mapped 1:1 to the translation of the model. Releasing the clutch attaches the model to the Cubic Mouse’s current location. This allows users for example to pan the model on the workbench by releasing the clutch on the left side, moving the Cubic Mouse and therefore the model to the right side, pressing the clutch and moving the hand back to the left side, releasing the clutch again, and so on. This worked quite well on the workbench, since most of the visible parts of the model are within an arms length

reach. When we started using the application in a CAVE, the model was typically displayed at a much larger scale. To compensate we used a 1:3 or 1:4 hand to model movement ratio.

Another issue is how to define the center of the rotate and zoom operations applied to the model. At the start of our application, we display the model at a scale such that the entire model is visible. The origin for the zoom and rotate operations is defined as the origin of the model's bounding sphere. After zooming and panning the model, occasionally our users want to change the origin to a new location or feature within the model. Usually they are already investigating this area by moving the slicing planes back and forth through the feature. Thus, by simultaneously pushing the two zoom buttons, we let them define the new origin of rotation as the intersection point of the three slicing planes.

Visualizing Volumetric Medical Data

CT, MRI and PET scanners generate volumetric data sets similar to the seismic data sets used in our geological scenario. Traditionally, medical visualization uses three orthogonal slicing planes (Figure 4) to view human cross sections. We use the Cubic Mouse for this scenario in exactly the same way as for the geo-scientific visualization. This gives users the skull in their hand, similar to the system described in [4], and allows them to position the slicing planes using the rods.




Figure 4: Three cross sections through a human head: The transversal plane is oriented horizontally, the frontal plane is parallel to the front, and the sagittal plane is perpendicular to both and moves from ear to ear.

Cutting Planes and Chair Cuts

For engineering applications, cross sections are an important tool for viewing and understanding the structure of three-dimensional models. Often two or three orthogonal cutting planes are applied at once or a so called chair cut is

performed as shown in Figure 5. A chair cut removes one octant of the model, whereas three clipping planes cut away seven octants. These operations are complementary to each other and we allow users to switch between them. The definition of the octant being cut away or kept depends on the orientation of the clipping planes, which can be toggled using the buttons mounted on the two ends of each rod.

Since clipping planes in OpenGL can only be used to create convex cuts, we use a three pass rendering algorithm similar to [12] to create the chair cut. The first pass renders one half of the model using one clipping plane. The second pass renders an additional quarter of the model using a second clipping plane and the first clipping plane reversed. The third pass renders the last octant and uses a third clipping plane and the first and second clipping plane reversed. This approach sends the model three times down the graphics pipeline, which increases rendering times typically by a factor of two to three. Our system is based on SGI's Performer graphics tool kit, which uses a three process pipeline to render graphical objects. The application process positions objects in the virtual world, the culling stage removes polygons outside the viewing frustum, and the rendering stage feeds the remaining polygons into the graphics hardware. We mostly avoid the slow down for the chair cut by considering clipping planes already during the Performer culling stage similar to the already employed view frustum culling. This removes invisible polygons before they are handed over to the graphics hardware.




Figure 5: This chair cut clips away the left, upper, rear octant of a car model to permit a better view of the interior.

EXPERIENCES AND RESULTS

We observed people using the Cubic Mouse during a planned user observation session, while giving presentations to our consortium partners, and during a variety of demos.

User Observation

For a user observation session we recruited 12 subjects (8 male, 4 female) from staff and students of our research

organization. None of the subjects had ever used the Cubic Mouse before. The subjects were presented with the medical scenario at the two-sided Responsive Workbench shown in Figure 4. After handing over the Cubic Mouse to a subject we gave a short introduction of a few sentences explaining how to rotate the data set, describing the functionality of the rods, the two zoom buttons, and the clutch button. The subjects had to perform various tasks like finding a cross section through the nose and showing it from the side, finding a cross section through the eyes and showing it from the top, investigating the brain, finding the brain stem, and so on. After the subjects had performed these tasks, they received a questionnaire with twenty questions to assess the overall reaction to the input device, learnability, and the reaction to specific device characteristics. The questions were answered on a scale from 0 to 7.

On the whole, users found the Cubic Mouse natural and easy to manipulate (see video). They became proficient with the device in a few seconds and were able to fulfill our requests very quickly. Particular results from our questionnaire are:

- The test took 4 to 8 minutes for each person, one person being 4 minutes, one person being 8 minutes, the rest between 5 and 7 minutes. There was no correlation between the time used and the experience people had with virtual environments.
- Users felt very much in control of the application.
- Users with small hands found the device slightly too big and the rods slightly too long. The device can be comfortably held in a large hand, but should be smaller for small hands.
- Nobody complained strongly about the weight and surprisingly the cables were mostly not an issue, but our subjects used the device only for a short time.

Even though we did not tell our subjects how to hold the Cubic Mouse, it turned out that in general they held the device in their non-dominant hand and operated the rods with the dominant hand. These observations correspond closely with Guiard's findings [3] on how humans use their hands in asymmetric bimanual tasks. Both hands are used symmetrically to perform rotations of the model that can not be achieved by just twisting the wrist of the non-dominant hand.

The most surprising and convincing observation is that almost all users did not look at the device for switching rods from the first use. Instead they focus on the task and on the model shown on the screen. We observed that operating the buttons required sometimes a quick glance at the Cubic Mouse. These observations are also backed by observations done at the Siggraph 1999 exhibition, where the Cubic Mouse was used by a few hundred people.

An important reason for the intuitive operation of the Cubic Mouse can be found in the strong proprioceptive cues

provided. The cubic shape of the device in the non-dominant hand provides a strong cue for the dominant hand for finding the rods, which is additionally supported by visual cues from the application.

Experiences with Experts

When developing the prototype for the oil and gas visualization, we first implemented a virtual tools based approach similar to the one described in [10], [2]. Users had to pick up different tools for each task, e.g. a zoom tool, a rotation tool, and so on. To drag around a seismic slice, the user had to pick up a drag tool with a tracked wand, point to the slice, press the button on the wand, and move the slice by moving the wand. Sometimes the slices were hard to find since they were hidden behind faults or horizons. Among other demos we worked with one geologist for three days, did a one day evaluation session with three geologists, and presented the system for four hours to 20 experts from different oil companies. We showed both versions of our system and our users found that with the Cubic Mouse, the most common tasks were immediately available and easy to perform. They had to resort to the tool based interface only occasionally.

A few month ago a crash engineer from a large automotive company visited us to review our progress in a crash visualization project our group is involved in. We presented to him the simple cutting plane and chair cut scenario shown in Figure 5 on a Responsive Workbench system with his crash test data. Even though the demonstration ran at a low frame rate due to the complex finite element model, he stated that the Cubic Mouse is all he ever needed. He added that the intuitiveness of the device was "hard to beat". His single complaint was that reversing the direction of an individual clipping plane was non-intuitive, because there were no buttons mounted at the ends of the rods at that time, which has been fixed with our latest version of the Cubic Mouse.

The medical scenario has not yet been presented to medical doctors, but based on our experience with the oil and gas experts and the automotive engineers, we are highly motivated to explore this application domain in the near future.

CONCLUSIONS AND FUTURE WORK

Coordinate systems are an essential building block for most computer graphics applications. The Cubic Mouse literally puts an application or object coordinate system into the user's hand. The rods provide a strong affordance and make the device obvious to use. The Cubic Mouse's two-handed operation allows one hand to rest against the other which reduces fatigue and allows extended use.

When presenting the Cubic Mouse, people are often inspired to think about further enhancements of the device. In particular everybody wants to be able to turn the rods and we have just completed the development of a first prototype. These additional three degrees of freedom lead to an intuitive six degree of freedom input device. Turning

one of the rods specifies a rotation of a virtual object around the appropriate axis of a given reference frame represented by the Cubic Mouse. In contrast to a standard six degree of freedom tracker we have here the possibility to specify the six degrees of freedom independently from each other. This is often very desirable, since it allows a much more precise input much like a dial box, just in a more intuitive way. Another use for these additional degrees of freedom would be the fine adjustment of X, Y, and Z coordinates during a positioning task.

Currently the device is mostly used as an absolute positioning system. Moving a rod from one stop to the other moves the corresponding virtual object a given distance. This requires the frequent use of the clutch when moving objects over a larger distance. With a spring mechanism to reset the rods to a centered rest position the device turns into a well designed relative positioning device, without range limitations.

Our first prototype tracks the rods' movement with potentiometers, which is really a primitive approach. Obviously, optical tracking could be used to overcome the static friction problem of potentiometers and it would have higher resolution. Even more interesting is the use of step motors attached to the rods. In addition to tracking the rods' movements this would allow haptic feedback and turns the device into an *active* real world prop.

We presented an input device, which intuitively combines constrained translation and potentially rotation input with a 6DOF free-motion device. There is still more work to do to explore the full potential of this setup. However some geologists predicted already from their first experience that the Cubic Mouse will become a standard for their application domain within the next years.

ACKNOWLEDGMENTS

This work was partially supported by the VRGeo and SimVR consortia. We thank the members of the consortia for their valuable feedback during our meetings. Special thanks to Ernst Kruijff for help with the user testing, and to David Tonnesen for a thorough critique of an earlier version of the paper. Thanks also to Jakob Beetz, Hartmut Seichter, Jan Springer, Michael Tirtasana, Henrik Tramberend, and Jürgen Wind for help with the paper, and to the rest of the VE group at GMD for their support.

The first author is grateful to Pat Hanrahan and the Stanford Graphics Group as well as Frank Crow and Interval Research for providing a unique research environment at Stanford, where many of the foundations for this research were laid.

REFERENCES

1. Cruz-Neira, C., Sandin, D.J., and DeFanti, T.A. Surround-screen Projection-based Virtual Reality: The Design and Implementation of the CAVE. *Proceedings of SIGGRAPH '93*, 135-142, 1993.
2. Cutler, L. D., Fröhlich, B., and Hanrahan, P. Two-Handed Direct Manipulation on the Responsive Workbench. *1997 Symposium on Interactive 3D Graphics*, 107-114, 1997.
3. Guiard, Y. Symmetric division of labor in human skilled bimanual action: the kinematic chain as a model. *The Journal of Motor Behaviour*, 19(4):486-517, 1987.
4. Hinckley, K., Pausch, R., Goble, J.C., and Kassell, N.F. A Survey of Design Issues in Spatial Input. *In Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 213-222, 1994.
5. Hinckley, K., Pausch, R., Goble, J.C., and Kassell, N.F. Passive real-world interface props for neurosurgical visualization. *In Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems*, pages 452-458, 1994.
6. Krüger, W., Bohn, C.-A., Fröhlich, B., Schüth, H., Strauss, W., and Wesche, G. The Responsive Workbench. *IEEE Computer*, 42-48, July 1995.
7. Krüger, W., and Fröhlich B. The Responsive Workbench. *IEEE Computer Graphics and Applications*, 12-15, May 1994.
8. Mine, M.R., Brooks F. P. Jr., and H. Sequin, C. H. Moving Objects in Space: Exploiting Proprioception In Virtual-Environment Interaction. *In SIGGRAPH 97 Conference Proceedings*, pages 19-26. August 1997.
9. Pausch, R., Burnette, T., Brockway, D., and Weiblen, M.E.. Navigation and locomotion in virtual worlds via flight into Hand-Held miniatures. *In SIGGRAPH 95 Conference Proceedings*, pages 399-400, August 1995.
10. Poston, T., and Serra, L. The Virtual Workbench: Dextrous VR. *In Virtual Reality Software and Technology (Proceedings of VRST'94, August 23-26, 1994, Singapore)*, pages 111-122, August 1994.
11. Tramberend, H. Avocado: A Distributed Virtual Reality Framework. *Proceedings of VR'99 Conference, Houston, Texas*, 14-21, March 1999.
12. Viegas, J., Conway, M. J., Williams G., and Pausch, R. 3D Magic Lenses. *In Proceedings of the ACM Symposium on User Interface Software and Technology, Papers: Information Visualization*, pages 51-58, 1996.