Workshop on the Challenges of 3D Interaction

Kenneth P. Herndon, Andries van Dam and Michael Gleicher

1 Introduction and Motivation

3D computer graphics is becoming more and more popular due to the increased availability of 3D hardware and software on all classes of computers. However, despite this growing popularity and the existence of a number of successful 3D graphics applications, particularly in CAD, CAE, and medical and scientific visualization, the field is still very immature. There are no widely accepted standards for hardware or software platforms; learning to implement or use 3D graphics software is still extremely laborious; and the most effective ways for humans to interact with synthetic 3D environments are still not clear.

In the past few years, a number of research and commercial groups around the world have been studying these problems [1][3][9][11][16][20][23][29][33][34][37]. Most have realized that 3D graphics applications are significantly more difficult to design, implement and use than their 2D counterparts. The authors were interested in understanding more about these difficulties, and their many diverse sources – users' perceptual, cognitive, and motor skills and abilities, limitations of currently available input and output devices, the nature of 3D tasks, and the variety of implementation strategies and development environments. In order to get a reasonable understanding of these challenging issues, input from other researchers both in our own field and others, such as perceptual psychology, kinesiology and graphic design is necessary.

In order to draw together an sufficiently diverse collection of researchers to discuss these issues, we organized a workshop at CHI'94 called "The Challenges of 3D Interaction". The primary objective of this workshop, held on April 24th and 25th, 1994, was to identify the principles and techniques necessary to design, implement and evaluate specifically 3D user interfaces – that is, interfaces that exist in the same 3D environment as the 3D application objects they control. Despite the title of the workshop and the obvious bias of the organizers toward specifically 3D computer graphics, many of the issues discussed during the workshop turned out to be applicable to user interface research in general.

2 Workshop Structure

The seventeen workshop attendees were selected on the basis of position papers submitted by almost 40 applicants (Appendix B describes how to obtain copies of these documents). In selecting participants, we attempted to gather as diverse a set of experiences as possible.

Most of the first day was devoted to individual presentations by the three organizers and the seventeen attendees. Each person gave a ten-minute talk followed by a short question-and-answer period. The remainder of the two-day workshop was made up of the following panel discussions:

- Application space
- Fundamentals of 3D interaction – what is a 3D user interface?
- Psychology
  - Perception
  - User studies and evaluation of user interfaces
- Conceptual design
  - Borrowing from and designing for the "real world"
  - Borrowing from 2D user interface experience
- Current state of the art in 3D user interface research
- Non-traditional interfaces
- Agenda and directions for future research

These topics arose partly from our experience in the field and partly from the attendees' position papers. In preparation for the workshop, the organizers selected a team of two or three participants to lead each panel and gave each team a list of potential questions to help them organize a short introduction to the topic. Each team then moderated a group discussion. The following sections describe the main points of each of the panel topics.

3 Applications for 3D Graphics

Stuart Card, David Zeltzer

This panel began by considering the following list of application areas that have successfully used computer-generated 3D graphics and 3D user interfaces. There are, of course, non-3D applications within each of the following areas, but we were concerned with identifying those particular applications which have driven 3D user interface research to date.

- Envisionment of intrinsically 3D data: architecture, CAD/CAM [10], modeling and animation, medicine
- Fantasization of worlds: games/fantasy, simulation and training
- Information Visualization: financial data, scientific visualization, databases, libraries

- Communications: teleconferencing, meetings

We believe that these same application areas will continue to drive future 3D interface research. However, the hardware and software implementations of future applications must address the following issues:

Scale: Though many of the 3D applications developed in research labs are dealing with new interface and presentation ideas, they have generally been very small-scale, proof-of-concept implementations (commercial systems are generally more robust and complete and therefore of greater use to more people, but do not necessarily push the limits of user interface research). If these applications are to be useful they must be able to support higher density problems. Design applications must support much larger databases of information (e.g., the Boeing 777 database is measured in gigabytes). Likewise, information visualization applications must handle large amounts of data while still displaying it legibly. There will also be a great difference between consumer and higher-end applications for medicine and science: these latter applications will generally require higher degrees of accuracy and resolution than applications for entertainment.

Level of Detail: User interfaces must also provide ways to manage the relative size relationship between objects in a computer-generated 3D environment and the view of this environment as seen by a human. For example, some scientific datasets contain data at many different scales (e.g., large-scale flow of air past an aircraft fuselage versus the much more detailed flows around the jet engine turbines). An application for visualizing this data must provide ways for users to change the size of their view so that these data at varying scales can be observed.

Interaction model: The traditional interaction model used for 2D WIMP (Windows, Icons, Menus and Pointers) interfaces may not be appropriate for all 3D interfaces, especially those that use input and output devices with more than two degrees-of-freedom (DOF). Whereas systems employing WIMP-style interaction models generally enforce a distinct separation between application and user interface and limited communication between them, so-called “non-WIMP” [17] user interfaces must support:

- higher bandwidth input and output
- many degrees of freedom
- real-time response
- continuous response and feedback
- probabilistic input
- multiple simultaneous input and output streams from multiple users

Interaction techniques: The tasks for which a particular application is used will determine the interaction techniques its user interface must support. Applications for architectural design, for instance, must provide very precise controls for object placement and size, but a program to help you to position furniture in your house may need only gestural placement tools. Similarly, applications for which immersive displays are best must support very different navigation and manipulation tools from applications for desktop use.

While it is agreed that each application places different demands on the design of its user interface, a number of interaction tasks seem to be universal to 3D graphics applications:

- Object creation, model definition
- Object selection
- Object placement and editing: affine transformations (e.g., translate, rotate, scale, alignment, etc.); modifications to other parameters (e.g., color, shading, etc.)
- Viewpoint control
- Perception: extracting cognitive information from an environment
- Programming: defining behavior of objects and relationships between objects

Just to clarify the terminology, this is a list of *tasks* which are common in 3D applications. *Interaction techniques* are the interfaces one uses to complete a particular task. These *interaction techniques* are discussed in detail in Section 4, “Fundamentals of 3D Interaction”.

Utility of 3D UIs: The user interfaces of some 3D graphics applications will probably not migrate naturally into 3D. Architects, for instance, have been trained for centuries to use 2D drafting tools to design 3D buildings. This may indeed be the optimal method for this task and giving them 3D tools on a computer may not help them at all to design buildings. However, in order to communicate their work to clients, they must generally translate their 2D drawings into 3D – for this task, a fully 3D application may be appropriate.

4 Fundamentals of 3D Interaction

Jock Mackinlay, Lutz Kettner

The interaction techniques used in today’s 3D graphics applications to manipulate synthetic objects and navigate through synthetic worlds are generally *ad hoc* implementations of task-dependent designs. Little work has been done to unify the wide variety of techniques into a universally applicable set. There is, however, a feeling that a general set of 3D interaction techniques will eventually emerge. Some examples of current techniques are:

- ray intersection picking for object selection
- direct-manipulation object placement (screen-aligned translation, virtual sphere rotation, etc.)
- pan, zoom, and rotation of camera
- use of textures, stereo displays, shadows and other depth cues to aid perception
- widgets for constraining gestural user input to specific degrees of freedom (e.g., object handles, grids)

While incomplete, this list does point to some of the difficult problems of 3D interaction. For instance, performing object selection by ray intersection can be problematic when there are many objects in a scene, because 3D objects tend to occlude one another.

It is often difficult to implement direct-manipulation interaction techniques for 2D input devices (mouse, tablet, trackball, etc.) because correlating 2D hand movement in the real world to object movement in the synthetic world can be difficult, especially if the planes of motion are not similar. The chain of indirection typically used in mouse-driven 3D user interfaces (from a user's hand through the mouse, then through a widget, and finally to the object) only complicates the matter. Higher-DOF devices (e.g., 6D tracker, glove, etc.) may alleviate some of these problems, but many of these devices, at least in their current incarnations, are very fatiguing to both eyes and muscles and have relatively poor spatial and temporal resolution (though advances are being made in these respects). The use of physical constraints may help alleviate this strain and may also increase interaction precision, just as straight edges and tabletops, for example, increase the accuracy of drawing straight lines with pen and paper. Indeed, devices with three or more DOFs may ultimately provide a more direct interface to 3D manipulations than their 2D counterparts, but more work must be done to make them easier to use for both developers and users.

Navigation in computer-generated 3D worlds is also problematic because these environments often have relatively little detail (compared with the real world) and the techniques used to control the viewpoint are sometimes difficult to use. Adding markers and landmarks to virtual worlds may help prevent people from getting lost so easily. But even if users knows their current position and where they want to go, they must still be able to control the navigation technique.

The use of metaphor in 3D user interfaces is also important in teaching users how to translate their real-world actions into synthetic-world effects [27]. The popular desktop metaphor in 2D windowing systems is a particularly good example of this, even though it does not exactly mimic real desktop behavior. Rooms, offices and other spatial and physical metaphors have been proposed for 3D graphics applications to give users a sense of the space in which they can operate [21]; viewing metaphors include perspective and orthogonal views, lenses [4], windows and other ways of distorting space. One can distinguish between object or physical metaphors (e.g., folders or trashcan) and action metaphors (e.g., realistic actions such as pointing, grabbing and flying; non-realistic actions such as teleportation).

Metaphors are something of a mixed blessing, however [19]. For instance, if a particular metaphor is interpreted literally by users, they may expect that their real-world knowledge of that thing will transfer into the synthetic world and that the interface will thus require little learning. In all likelihood, however, it will not transfer in full, since metaphors are rarely implemented completely. The challenge to user interface designers is to develop interfaces that promote the maximum transfer of users' knowledge of physical objects and actions into virtual environments.

5 Psychology

Jim Ferwerda, Jack Loomis
Rich Gossweiler, Colin Ware, Christine MacKenzie

This topic was divided into two separate panels. The first focused on the different ways in which humans perceive the world around them, and the characteristics 3D user interfaces must have to exploit these perceptual and spatial reasoning skills. The second focused on how people perform 3D interaction tasks, problems posed by multiple-DOF input devices, and user studies and evaluation of 3D user interfaces.

5.1 Navigation

This panel began with a discussion of the various cues that humans use to navigate through the real world, asserting that we must first understand real navigation if we are to understand how to build interfaces for navigation in virtual worlds [31].

Navigation is the planning and execution of travel through space, real or virtual, carried out with reference to external or internal representations of the space being traveled. An appropriate conceptual framework for navigation by machines, humans and other species recognizes the three kinematic orders upon which updating of position and orientation is based: position, velocity, and acceleration [14].

- Position-based navigation: also called "pilotage" or "piloting". This form of navigation relies on external signals indicating position and orientation (e.g., in the form of signals from stationary landmarks).
- Velocity-based navigation: also called "dead-reckoning" or "path integration". This relies on external signals indicating speed and direction of travel. Linear and rotary displacements are obtained by integrating these signals. Proprioception[1], optic flow, and acoustic flow are all potential sources of such velocity information in both real and virtual environments.
- Acceleration-based navigation: also called "inertial navigation". With this form of navigation, linear and rotary accelerations of the traveler are doubly integrated to obtain linear and rotary displacements with respect to the initial position and orientation. Humans sense such accelerations on the basis of vestibular (inner ear) [24], cutaneous, and proprioceptive stimulation. Without a motion base providing acceleration signals, navigation with 3D UIs will have to rely on position and velocity information alone.

As yet, little is known about which sensory signals are most important in effective navigation by humans, but it seems likely that visual position and velocity information is most important. The vestibular system provides rotational information based on the frequency of the rotational movement: for low frequency movements it measures rotary acceleration, while for high frequency movements it measures rotary velocity [2].

Even in the field of perceptual psychology, human navigation is not very

1. Proprioceptive information derives from both the signals from muscles, tendons, and joints and from the commands issued by the brain to the muscles; it provides an awareness of body postures and changes in posture as well as of movement through space.[8]

well-understood. In most graphics systems today, vision is the only channel used to provide navigational cues to users, though some systems generate aural, haptic and vestibular cues (e.g., immersive systems like flight simulators). The extent to which pure optic flow can provide adequate information about one's position and orientation in space is still unknown. However, as mentioned earlier, landmarks can be very useful for determining location and are relatively easy to add to 3D graphics applications. Features such as horizons provide information about one's orientation relative to a ground plane, and textured surfaces and shadows can inform users about their distance from objects. All of these techniques can help people to understand their relation to an environment around them, though it is still unclear how exactly people build and maintain so-called "cognitive maps".

5.2 Perception[2]

The characteristics that 3D user interfaces must have to exploit the perceptual and spatial reasoning skills of users fall into five categories:

Functional fidelity: Taken together, the various sensory cues provided by an interface define the functional fidelity of a given synthetic environment. The functional fidelity of an application must be appropriate to the tasks being performed. The representation of the 3D objects in the scene should be useful, but not necessarily photorealistic. More research is needed to quantify the levels of functional fidelity needed for various tasks. The level will, of course, depend on the task itself, but may also vary for users with different expertise.

Responsiveness: 3D user interfaces must be very quick to respond to user input so that natural exploratory and manipulative behavior can occur. Studies have shown that lag in immersive environments greatly affects the accuracy of positioning tasks [6]. If a user interface is too slow, it becomes non-interactive. In most computer graphics systems, lag is generated at all points of the I/O pipeline, including input devices, software (application processing and user interface), and hardware (rendering).

Affordances: An affordance provides information about what can be done with an interface and what the consequences are of using it [30][36]. Affordances allow the creation of objects and environments that have meaningful properties, geometry, materials and behavior, and provide clues about how to interact with these objects and environments.

Appeal to mental representations: User interfaces must be organized in such a way that they are recognizable to users. For instance, metaphors like those discussed earlier can be used to transmit the structure and behavior of a user interface to reduce the cognitive load placed upon the novice.

Multiple/integrated input and output modalities: User interfaces should be able to use more than just the visual channel for communication. Voice and gesture can be used for input, sound and haptics for output.

It was noted during this panel discussion that most perceptual psychologists doing research today are not studying the specific problems of concern to user interface developers in the HCI context (they are investigating the sensory mechanisms that underlie our perceptual functioning).

5.3 3D Interaction Tasks

People are inherently active, generative, and intentional, and fundamentally biological and biomechanical. These qualities imply a perception-action feedback loop in every interaction a human makes, including user interfaces to computer applications. The physical constraints include constant forces such as gravity and time-varying forces like friction that become active when people touch or manipulate an object. Without thinking about them, people use these forces to simplify tasks. Current human-computer interfaces very rarely include haptic devices, and so are mainly limited to providing visual or audible feedback to user actions. Humans generally respond to visual stimuli by looking toward them and to sounds by turning toward the source. User interfaces may be made more effective by capitalizing on these basic orienting mechanisms.

5.4 Multiple-DOF Input Devices

Multiple-DOF devices such as 3D or 6D trackers and glove input devices are excellent choices for multiple-DOF tasks such as gestural viewpoint or object placement. Similarly, 2D devices are well suited for 2D tasks such as drawing or pointing on a CRT. Table 1 lists the degrees of freedom for a number of devices and representative tasks.

Device	DOF	Task
Slider/dial	1	volume control
Mouse	2	picking, drawing, 2D location
6D mouse, tracker	6	docking, view control
Glove, face	16+	hand/face animation
Body suit	100+	whole body animation

Table 1: Degrees of freedom of input devices.

The relationships between devices and tasks listed above are obvious and are generally straightforward to implement and use. Problems arise, however, when using devices for tasks that do not have the same number of degrees of freedom. For instance, specifying the 3D position and orientation of an object with a mouse is difficult because one can only modify two degrees of freedom at a time. When the task space has more degrees of freedom than the input device, the user interface to the task often must take the form of a complex dialogue of composed interactions rather than direct manipulation. Some of these problems were discussed in Section 4 above, "Fundamentals of 3D Interaction".

Conversely, using higher-DOF devices for lower-DOF tasks can also be confusing if the input device is not physically constrained to the same degrees of freedom as the task. Consider using a mouse to change the value of a one-dimensional horizontal Motif slider: vertical movements of the mouse have no effect on the slider,

2. Further reading can be found in [15][32][35].

thus impairing the illusion that the screen cursor and slider have a direct relationship. Even if the screen cursor is constrained to the space of the slider (i.e., in only one direction so that the cursor is always on the slider), the physical mouse can still move in the orthogonal direction. These same phenomena occur when modifying a one-dimensional slider in a virtual environment with a 6D tracker because the tracker is not constrained to the same degrees of freedom as the slider.

A possible solution to the problem of missing constraints is to use physical props to constrain or suggest to the user preferred degrees of freedom, even though the prop may have as many actual degrees of freedom as the input devices connected to them. In this way, users will tend to move according to the perceived constraints of the system [22]. An alternative approach is to develop ways to constrain the task space to fit the input device being used. More research must be done to find how best to constrain multiple-DOF 3D interaction tasks to more manageable sizes.

5.5 User Studies and Evaluation

User studies are as important for 3D user interfaces as they are for 2D, but since the range of tasks differs greatly in 3D, so also must the metrics used to evaluate 3D user interfaces. Some studies have been done to compare interaction techniques for placing objects in 3D [18][25], for instance, but there have been no comprehensive evaluations of entire 3D user interfaces.

6 User Interface Design

Paul Isaacs, Bruce Wilson
Dan Venolia, C. Overbeeke

This topic was also divided into two separate panels: the first panel considered borrowing interface ideas from the “real world”, and the second concerned borrowing ideas from 2D user interface design and the effects of input devices on UI design.

6.1 Borrowing from the “Real World”

It was generally felt that “real world” objects and ways of interacting with them can offer valuable insights for 3D user interface design. As in real-world design, there are a number of criteria by which to evaluate computer interfaces: layout/visibility, legibility/affordances, ergonomics, color, shape, material, feedback, etc. The relative importance of each of these criteria will vary depending on the input and output devices used, the user’s level of expertise, and the task at hand. Whether real world designs will work in a virtual world depends greatly on factors of the physical environment, such as screen resolution and lag in input devices, computation and graphics pipelines, which greatly affect usability. One must also consider whether a user interface design can be implemented with available hardware and software.

6.2 Borrowing from Other Disciplines

3D user interface designers can borrow ideas and techniques from a variety of other disciplines to help direct their work, including industrial design, architecture, 2D user interface design, toy design, cinematography, and psychology. In particular, the potential contributions of cinematography to 3D UI design were discussed. One important difference between these two fields is that cinema is mainly a passive experience for the viewer, whereas interacting with computers is active. Cinematographic techniques used to manipulate human viewers, such as cuts, establishing shots, and camera movements, might not be as useful in highly interactive applications as lighting and audio, for instance, which are used to set mood and provide feedback about what is happening or will happen in a scene.

Techniques used by industrial designers, however, such as affordances, color coding, and iconography, are very important to good user interface design. A simple rule of thumb is that one should never have to read a label to understand how to use an object – its affordances should convey this information. While this is an honorable goal, interfaces, even those to real-world objects, rarely achieve it in practice. Iconography can be used to communicate the effect of an interface object. For example, a button should invite pushing by its affordances, and should indicate its function via icons (though icons are often not self-disclosing). Color and size can be used to relate the button to other interface objects.

6.3 Borrowing from 2D UI Design

3D user interface designers can learn a great deal from the field of 2D UI design, but must be careful when attempting to use 2D techniques for 3D interfaces. Some techniques common to 2D interfaces may be much more difficult to both implement and use in 3D environments simply because they were not designed for 3D. General concepts like “direct manipulation” can certainly be applied to 3D, but WIMP-style widgets like menus may not transfer so easily.

6.4 User Community

User interfaces to 3D graphics applications must be tailored to suit their particular user communities (architects, surgeons, designers, etc.). The design of a user interface must also consider the different levels of skill of its user community. Any given application will probably be used by both novices and experts. Because novice users tend to spend more time exploring an environment than performing specific activities within it, novice user interfaces should be easy to understand and provide immediate feedback and guidance. An expert user interface can take more liberties since its users have a mental model of the operations that can be performed with it. An interface could constantly change to accommodate a user who is gaining (or losing) expertise by selectively adding and deleting components.

6.5 Input Devices

Input devices also have a great impact on user interface design [7]. Mouse-based interfaces have generally been limited to a point-and-click dialogue between user and computer, though a few interfaces have been developed that interpret both postures (static positions of the hand) and gestures (movements of the hand over time). Many-DOF input devices such as 6D trackers, gloves and others can provide input data with much more content than mice. Glove interfaces based on gesture or posture recognition have been found very hard to use because it is difficult to remember the gesture or posture corresponding to a given command. However, though they have not been widely used yet, gloves may yet become quite useful devices as better software techniques (e.g., speech and gesture) are developed to support them.

One possibility for future interfaces is to follow the real world's example closely and provide different input devices for different tasks. This solution obviously does not exploit the flexibility of the computer. However, it is quite difficult to build interfaces that make it easy to perform a wide variety of tasks with a single input device.

7 State of the Art in 3D UI Research

Bill Chapin, Mark Green

Current 3D user interface research involves mainly exploring the design space of 3D interfaces to existing 3D graphics applications and looking for new applications that 3D graphics will enhance. This exploration includes experimenting with different interaction models, interaction techniques, and software architectures for underlying support systems.

Despite the large amount of research in interactive 3D graphics over the past ten years, the field is still very immature with respect to providing integrated toolkits for constructing new 3D interaction techniques or user interfaces. There is no system that provides for 3D the functionality that state-of-the-art 2D interface builders, such as UIM/X, do for 2D. Some commercial products, like SGI's OpenInventor, provide programming environments at a higher level than basic graphics languages like OpenGL, XGL or PHIGS+, but still rely on programming in C++. Products like Autodesk's Cyberspace Developer's Kit and research systems like the MR Toolkit from the University of Alberta provide interfaces to low-level device drivers and simple object creation. The vast majority of research systems being used today were created from scratch, starting with low-level graphics libraries. Also, the range of hardware input and output devices used by 3D graphics systems is large and there are no standard interfaces for these devices. One consequence of this is that the particular input devices one chooses to use determine both the interaction model and style of the interface.

The current feeling is that there may in the future be generic toolkits for creating 3D interfaces for all application domains, but such toolkits will have to support a very wide variety of interaction tasks. Simulation applications will need 3D objects that behave according to sets of rules, such as the principles of Newtonian physics; design applications need techniques for manipulating and constructing 3D objects with precise control; visualization applications require navigation and orientation techniques as well as the ability to handle possibly large datasets.

Some of the factors which make 3D interface design difficult even for desktop systems (i.e., mouse and monoscopic CRT) have been discussed above. Some 3D interfaces, however, may also be fundamentally different from any existing 2D UI. Immersive environments, for instance, with many multiple-DOF input devices and multiple simultaneous users, will require software systems very different from those usually seen in 2D and 3D applications (these differences will likely be at the operating system, application and user interface library levels). Direct manipulation of objects becomes the primary mode of interaction in immersive 3D environments, as opposed to indirect control via 2D or 3D widgets.

Some preliminary solutions to device handling, software structure of underlying systems, and basic geometry and behavior editing have already been presented, but these are only first steps. Unsolved problems include:

- standardization of interaction primitives
- presentation technologies (LCD, haptic, aural, etc.)
- input techniques (speech and gesture recognition, gaze tracking, non-intrusive position-measuring devices, etc.)
- networking
- multimodal I/O
- relieving symptoms due to prolonged use of input and output devices
- scale, level of detail

8 Non-Traditional Interfaces

Yoshifumi Kitamura, George Fitzmaurice

This panel aimed at identifying some of the emerging input and display technologies and discussing their effects on the design of 3D user interfaces and software systems for interactive 3D graphics. The categories of relevant technologies include:

- tracking (hand, body, eye, mouth, etc.)
- recognition (speech, gesture, etc.)
- display (visual, haptic, aural, etc.)

Each of these can be applied to desktop, semi-immersive, immersive or augmented reality systems, as appropriate. Numerous applications of these various hardware technologies have been published in recent years. New technologies, such as holographic displays and esoteric input devices for special purpose applications, will be developed as needed. There are as yet no standard multi-DOF input devices or display devices for interacting with or visualizing 3D worlds. Until standards are adopted, however, software systems for interactive 3D graphics applications will have to be designed with specific hardware technologies in mind.

9 Agenda and Future Directions

Our final discussion focused on topics for future research in interactive 3D graphics applications. We believe the research community needs to:

- Develop efficient and useful interfaces to existing 3D graphics applications. Often, 3D user interfaces work by displaying elements of the real world in a virtual environment in a way that facilitates a better understanding of the virtual world by providing a means to interact with

it and possibly change it. 3D graphics can be used to represent real-world objects, but can also be effective for visualizing more abstract data [12].
- Identify new applications for which 3D user interfaces are needed.
- Assess the perceptual, cognitive and motor capabilities of humans engaged in 3D interaction in the physical world, as well as in the HCI context. 3D interfaces will benefit from a better understanding of human spatial cognition and of ways to portray 3D space, shapes and relationships among objects so that users can perceive and interact with them effectively.
- Assess the different demands placed on user interface design by desktop, semi-immersive, immersive, and augmented reality systems [5][26].
- Incorporate other senses, such as audition, touch and kinesthesis, that have not yet been widely used in 3D user interfaces.
- Find general techniques for certain interaction tasks that are independent of input and output devices (e.g., navigation and manipulation).
- Define a *metric of fidelity* by which we can determine how much realism a synthetic environment must have to be understandable. The current feeling is that the level of fidelity required is task-dependent and that there are few general guidelines to follow.
- Explore the possibilities of developing 3D user interfaces and 3D graphics applications that are appropriate for and available to a wider audience. For instance, American Disabilities Act (ADA) advocates are pushing for greater access to computing facilities by the blind [28], a community rarely considered by graphics-oriented researchers who are mainly concerned with exploiting the visual capabilities of the sighted to facilitate learning and productivity. The apparent contradiction here is worth some careful thought.
- Develop better tools for creating 3D applications. Present tools are not as well developed as their 2D counterparts.
- Define the human-computer dialogue for interactive 3D graphics applications so that concrete guidelines for 3D user interface design can be laid out.
- Address software engineering issues unique to interactive 3D applications. Time-critical rendering, modeling and simulation techniques must be developed to guarantee interactive frame rates though negotiated graceful degradation even in complex environments [13]. It is felt that current operating systems are not capable of supporting the real time demands of 3D interfaces. Other software engineering concerns include support for iterative design of interfaces and for distributed and networked environments.
- Address hardware limitations, such as resolution and accuracy, which are amplified in augmented reality systems for which accurate registration with the real world is crucial.
- Find ways to assess the usability of complex user interfaces.

Research in the above areas will be useful in determining design requirements for 3D user interfaces, applications and hardware.

10 Acknowledgments

We are very grateful to the people who helped out during the workshop by taking notes, operating audio and visual equipment, making copies and everything else: Dan Robbins, Tom Meyer, Scott Snibbe, Rick Zobel, Virginia Howlett, Susan Papallardo and Alan Wexelblat. We also thank those who made helpful comments on draft versions of this report.

Appendix A Participants

Workshop Organizers

Andries van Dam, Brown University
Kenneth Herndon, Brown University
Michael Gleicher, Carnegie Mellon University

Workshop Participants

Stuart Card, Xerox Palo Alto Research Center
Bill Chapin, Stanford University
Jim Ferwerda, Cornell University
George Fitzmaurice, University of Toronto
Rich Gossweiler, University of Virginia
Mark Green, University of Alberta
Paul Isaacs, Silicon Graphics
Lutz Kettner, University of Karlsruhe
Yoshifumi Kitamura, ATR Communication Systems Research Laboratories
Jack Loomis, University of California
Christine MacKenzie, Simon Fraser University
Jock Mackinlay, Xerox Palo Alto Research Center
Kees Overbeeke, Delft University of Technology
Dan Venolia, Apple Computer, Inc.
Colin Ware, University of New Brunswick
Bruce Wilson, Softimage
David Zeltzer, MIT

Appendix B Additional Information

Additional information about this workshop can be found via anonymous FTP at the site `wilma.cs.brown.edu:/pub/papers/graphics/chi94`, or via the WWW in the home pages for the Brown Computer Graphics Group (`http://www.cs.brown.edu/research/graphics/3d_ui/chi94/chi94-workshop.html`). In these locations can be found the text of this report, the position papers submitted by the above participants and other miscellany.

Appendix C Bibliography

This bibliography is by no means comprehensive, but includes references to a variety of work relevant to the field of 3D graphics systems, applications and user interfaces.

1 Akeley, K. RealityEngine graphics. In *Proceedings of SIGGRAPH'93*, pp. 109–116.
2 American Physiological Society (Bethesda, MD). *Handbook of Physiology*, Volume III, Part 2, Section 1: *The Nervous System*. 1984.
3 Bier, E. Snap-dragging in three dimensions. In *Proceedings of the 1990 Symposium on Interactive 3D Graphics*, pp. 193–204.
4 Bier, E., Stone, M., Pier, K.,

Buxton, W. and DeRose, T. Toolglasses and magic lenses: the see-through interface. In *Proceedings of SIGGRAPH'93*, pp. 73–80.

5 Bishop G., Fuchs, H., et al. Research directions in virtual environments: report of an NSF invitational workshop, March 23-24, 1992. In *Computer Graphics*, 26(3), pp. 153–177.

6 Bryson, S. Effects of lag and frame rate on various tracking tasks.

7 Card, S., Mackinlay, J. and Robertson, G. The design space of input devices. In *Proceedings of CHI'90*, pp. 117–124, April, 1990.

8 Clark, F.J. and Horch, K.W. Kinesthesia. In Boff, K.R., Kaufman, L., and Thomas, J.P. [Eds.], *Handbook of Perception and Human Performance*. v. 1: *Sensory processes and perception*, pp. 13/1–13/62. New York: Wiley, 1986.

9 Conner, D.B., Snibbe, S.S., Herndon, K.P., Robbins, D.C., Zeleznik, R.C. and van Dam, A. Three-dimensional widgets. In *Proceedings of the 1992 Symposium on Interactive Three-Dimensional Graphics*, pp. 183–188.

10 Coyne R., and Snodgrass. Rescuing CAD from rationalism. *Design Studies*, 14, 2, pp. 100–123, 1993.

11 Deering, M.F. and Nelson, S.R. Leo: A system for cost effective 3D shaded graphics. In *Proceedings of SIGGRAPH'93*, pp. 101–108.

12 Ellis S. [Ed.]. *Pictorial Communication in Virtual and Real Environments*, 2nd Ed. London, England: Taylor and Francis, 1993.

13 Funkhouser, T.A. and Séquin, C.H. Adaptive display algorithm for interactive frame rates during visualization of complex virtual environments. In *Proceedings of SIGGRAPH'93*, pp. 247–254.

14 Gallistel, C.R. *The Organization of Learning*. Cambridge, MA: MIT Press, 1990.

15 Gibson, J.J. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin, 1979.

16 Gleicher, M. and Witkin, A. Through-the-lens camera control. In *Proceedings of SIGGRAPH'92*, pp. 331–340.

17 Green, M. and Jacob, R. SIGGRAPH'90 Workshop Report: Software architectures and metaphors for non-WIMP user interfaces. *Computer Graphics*, v. 25, pp. 229–235, July, 1991.

18 Grissom, S. and Perlman, G. Step(3D): A portable discount usability evaluation plan for 3D interaction. Department of Computer Science and Information Science, Ohio State University. OSU-CISRC-2/93-TR7, 1993.

19 Halasz, F. and Moran, T.P. Analogy considered harmful. In *Proceedings of Human Factors in Computer Systems Conference*, pp. 383–386, March, 1982.

20 Harrell, C.B. and Fouladi, F. Graphics rendering architecture for a high performance desktop workstation. In *Proceedings of SIGGRAPH'93*, pp. 93–100.

21 Henderson, D. and Card, S. Rooms: the use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. In *ACM Transactions on Graphics*, v. 5(3), pp. 211–243, 1986.

22 Hinckley, K., Pausch, R., Goble, J.C., and Kassell, N.F. Passive real-world interface props for neurosurgical visualization. *Proceedings of CHI'94*, pp. 452–458, April, 1994.

23 Houde, S. Iterative design of an interface for easy 3D direct manipulation. In *Proceedings of CHI'92*, pp. 135–142.

24 Howard, I.P. The vestibular system. In Boff, K.R., Kaufman, L., and Thomas, J.P. [Eds.], *Handbook of Perception and Human Performance*. v. 1: *Sensory processes and perception*, pp. 11/1–11/30. New York: Wiley, 1986.

25 Katz, J. A user study for comparison of interaction techniques using 3D widgets across multidimensional input devices. Brown University Master's Thesis. May, 1994.

26 Kawalsky R.S. *The Science of Virtual Reality and Virtual Environments*. Wokingham: Addison-Wesley, 1993.

27 Lakoff, G. and Johnson, M. *Metaphors We Live By*. Chicago: The University of Chicago Press, 1980.

28 Loomis, J.M., Klatzky, R.L., Golledge, R.G., Cicinelli, J.C., Pellegrino, J.W., and Fry, P.A. Nonvisual navigation by blind and sighted: assessment of path integration ability. *Journal of Experimental Psychology, General*, 122, 73–91, 1993.

29 Molnar, S., Eyles, J. and Poulton, J. PixelFlow: High-speed rendering using image composition. In *Proceedings of SIGGRAPH'92*, pp. 231–240.

30 Norman, D. *The Psychology of Everyday Things*. Basic Books, 1988.

31 Schiesser, E.R. Principles of navigation. In C. Belove [Ed.], *Handbook of Modern Electronics and Electrical Engineering*. New York: Wiley, 1986.

32 Skarda C.A. Perception, connectionism and cognitive science. In Varela F. and J.P. Dupuy, *Understanding Origins: Contemporary Views on the Origin of Life, Mind and Society.* Deventer, The Netherlands: Kluwer, 1992.

33 Snibbe, S.S., Herndon, K.P., Robbins, D.C., Conner, D.B. and van Dam, A. Using deformations to explore 3D widget design. In *Proceedings of SIGGRAPH'92*, pp. 351–352.

34 Strauss, P.S., Carey, R. An object-oriented 3D graphics toolkit. In Proceedings of SIGGRAPH'92, pp. 341–349.
35 Varela F. Whence perceptual meaning? A cartography of current ideas. In Varela F. and J.P. Dupuy, *Understanding Origins: Contemporary Views on the Origin of Life, Mind and Society.* Deventer, The Netherlands: Kluwer, 1992.
36 Warren W.H. Environmental design as the design of affordances. In Flach, J.M., Hancock, P.A., Ciard, J. and Vincente K.J. [Eds.]. *The Ecology of Human-Machine Systems*. Hillsdale, NJ: Lawrence Erlbaum.
37 Zeleznik, R.C., et al. An object-oriented framework for the integration of interactive animation techniques. In *Proceedings of SIGGRAPH'91*, pp. 105–112.