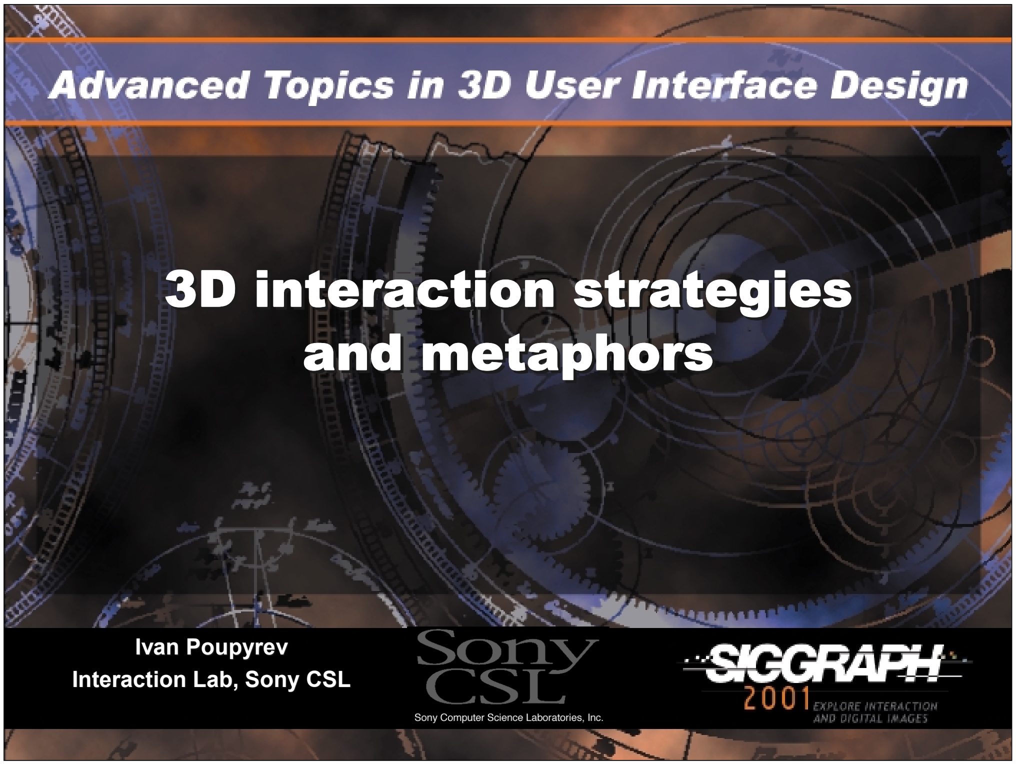

Ivan Poupyrev, Ph.D.
Interaction Lab, Sony CSL

E-mail: poup@csl.sony.co.jp

WWW: http://www.csl.sony.co.jp/~poup/

Address:
Interaction Lab, Sony CSL
Takanawa Muse Bldg.,
3 – 14 – 13 Higashigotanda
Shinagawa-ku, Tokyo 141-0022
Japan

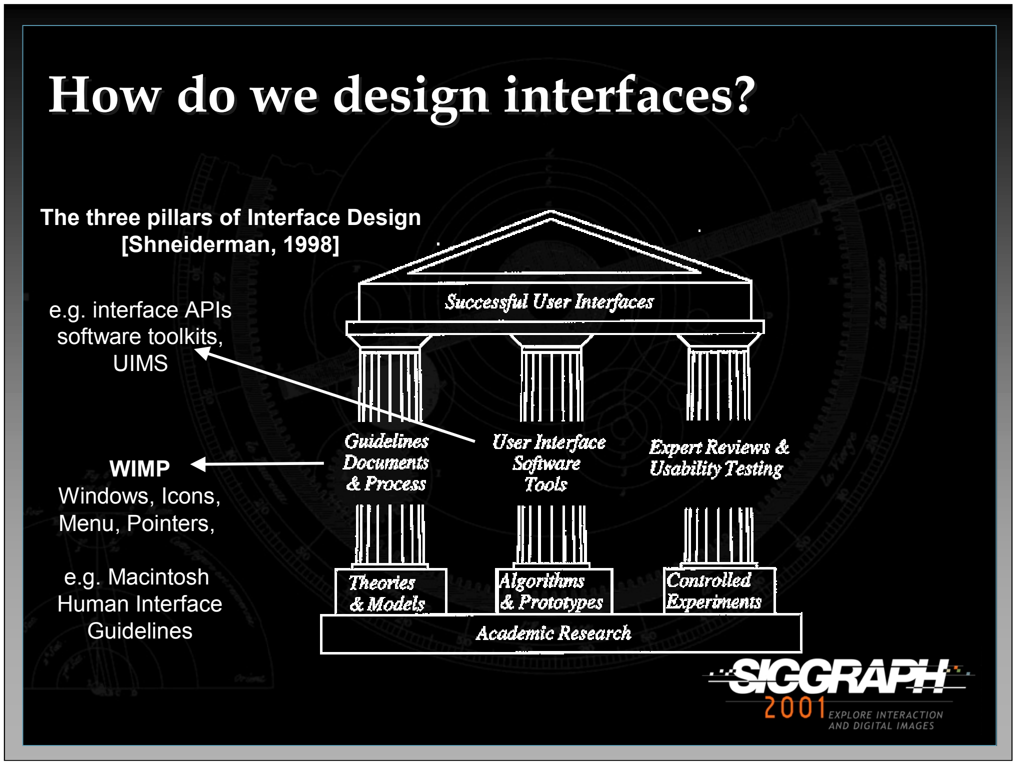

Because 3D user interfaces are special types of human-machine interfaces, the results of many years of research and development of traditional 2D interfaces can be also applied to the design of 3D interfaces.

These are the three pillars of successful user interface design, according to Shneiderman (1998): *guideline documents*, *user interface software tools* and *expert review and usability testing*. An example of the guidelines that are often used in designing 2D interfaces is the Macintosh Human Interface Guidelines, which outline the basic elements of the user Macintosh 2D interface, their functionality, purpose, layout, and visual appearance (Apple, 1992). These and other similar 2D interface design guidelines provide designers with basic building blocks of the user interface. Thus, interface designers do not have to invent user interfaces themselves, but can construct interfaces out of instances of icons, menus, dialog boxes, windows and others interface elements, as well as assign them various properties, names and functionality.

Furthermore, interface designers and developers do not even have to implement these basic interface elements: the user interface API (i.e. Application Programming Interface) provides access to libraries of already implemented behavior and functionality of the interface elements, which has become a standard part of the operating system. High-level interface design tools such as graphical editors allow designers to "draw" an interface for application, while even higher level tools such as UIMS (user interface management systems), provide ever more powerful tools in designing interfaces.

-continued on the next page

Certainly, even with the help of all these tools, designing high-quality interfaces still remains a complex and challenging task, requiring multiple iterations and usability studies to evaluate and refine designed interfaces (a third pillar of the interface design in Shneiderman's diagram).

While there has been a lot of criticism of the dominant desktop WIMP paradigm (e.g. Norman, 1999), it cannot be denied that, in spite of all their shortcomings, desktop graphical user interfaces have been a major step toward interfaces that can be effectively used by large numbers of users across different computing platforms.

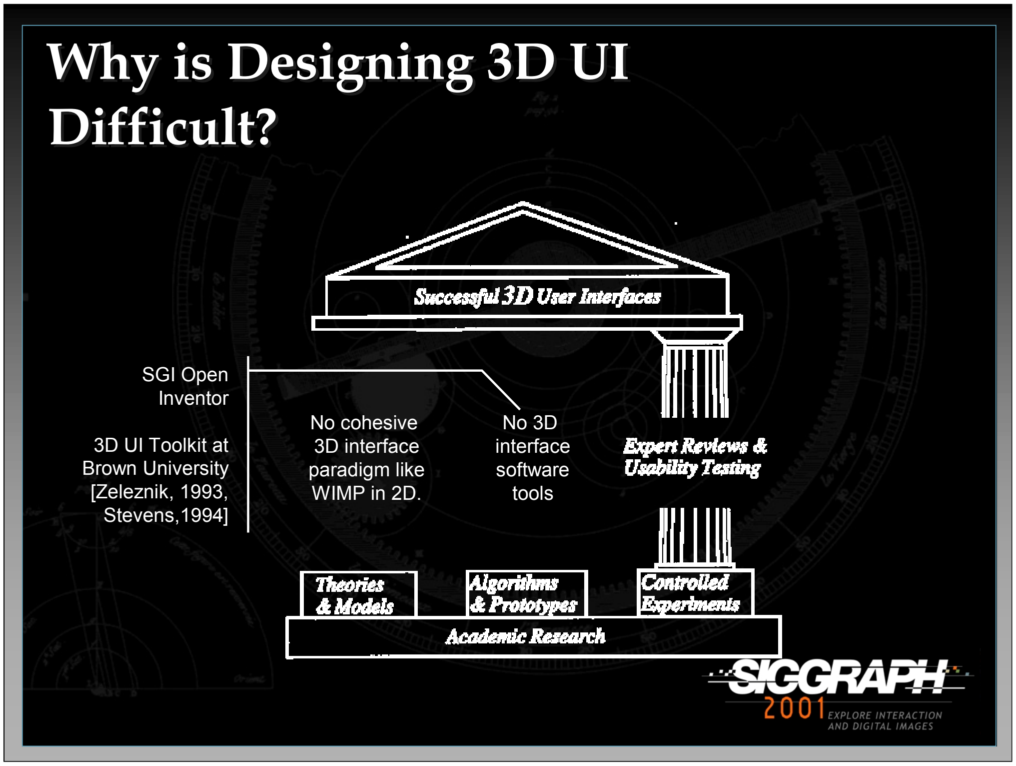

Designing 3D interfaces is still an art because first, there no cohesive 3D interface paradigm exists. What are the most basic classes of elements for 3D user interfaces? How do they relate to each other? There have been many 3D interaction techniques reported in the literature, some of them with guidelines for their use. However, it is not clear how they all relate and compare to each other, or how we should approach design of complex interaction sequences to do complex tasks. Consequently, there are few if any 3D interface design guideline documents that the designer of 3D user interfaces can rely on.

Second, there are currently no tools to support the design of 3D user interfaces, beyond the most simple, for example in Open Inventor. Currently, if designers need to use certain interaction techniques or tools they must either implement them themselves from scratch or invent new techniques. As a result designing interactive 3D user interfaces is a very time consuming task. In addition the produced interfaces are rarely formally evaluated and are usually designed mostly on the basis of designer intuition and common sense. Consequently, today's 3D interfaces are incompatible from application to application: each has its own look and feel.

One of the reasons for this is that the design space in 3D user interfaces is significantly larger then in 2D and large portions of it remains unexplored. 3D interface technology is still rapidly developing, and the new input and display devices, interaction techniques that appear require consequent revaluation. Furthermore, it is not uncommon that for certain application tasks or devices there have been no interaction techniques constructed, i.e. the design space of 3D user interfaces still has many empty spots.

-continued on the next page

In conclusion, the design of 3D interfaces is currently based on a mixture of intuition, common sense, informal rules of thumb, previously reported or ad-hoc designing techniques, and few general human factors guidelines.

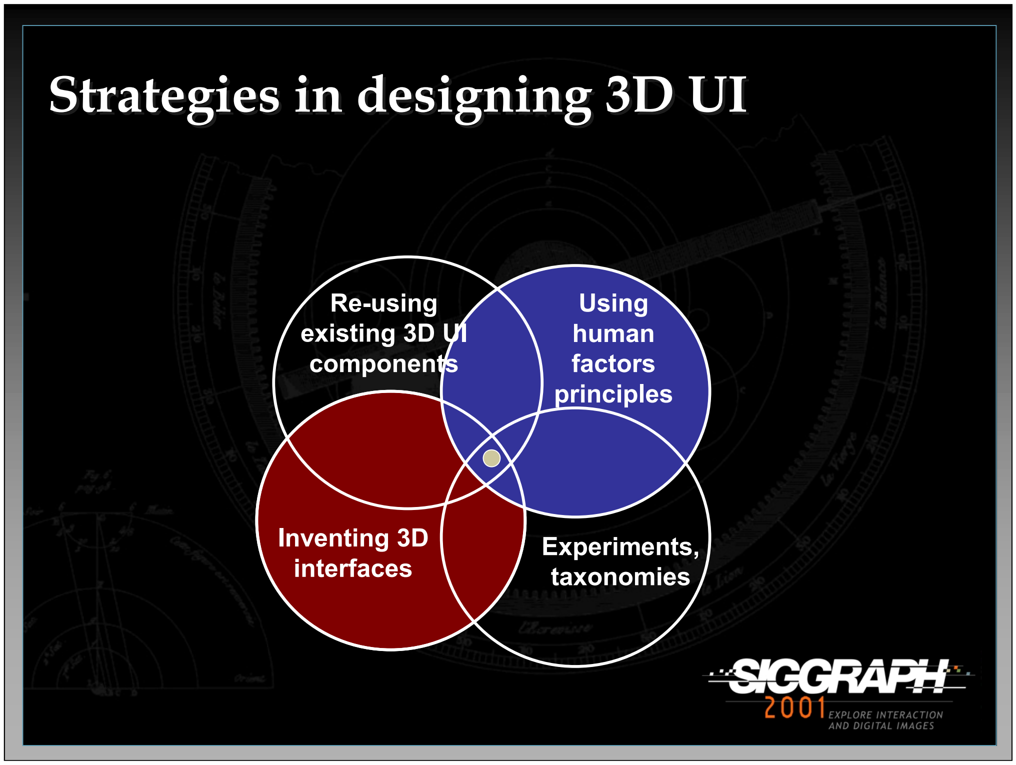

Today there are basically four major “pillars” in designing 3D user interfaces. First of all, we can reuse some of the interaction techniques developed before. The previous lectures in this tutorial have presented some of the interaction techniques reported in the literature.

Second, the human factors literature, related both to general principles of human-machine interaction as well as to specific principles of interaction with computers, is also highly relevant in the design of 3D user interfaces.

Third, the design of complex interaction sequences often require developers either to invent new interaction techniques or to adopt existing interaction techniques. This is because 3D interaction is a rapidly developing field and new input and output devices as well as new applications that employ 3D input are constantly being developed.

Finally, the design can be guided by some of the 3D user interface taxonomies and results of experimental studies reported in the literature. Some of the taxonomies and systematic design techniques will be covered in the other lectures of this course.

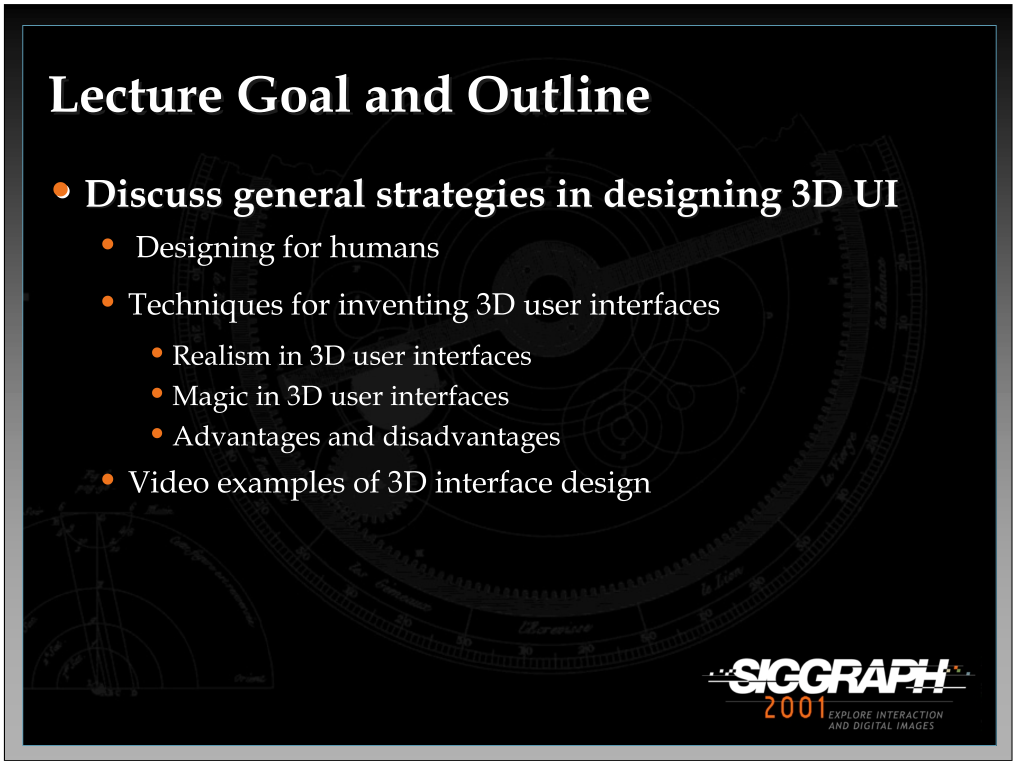

In this lecture I will discuss some of the informal and general approaches for designing 3D interfaces.

The lecture starts with discussing some of the basic human factor principles that can help to design better 3D user interfaces. The area of human factors engineering is vast and any in-depth discussion is far beyond the scope of this lecture. However, a few of the often used and reported principles will be briefly discussed.

The lecture then continues by presenting simple techniques for inventing 3D user interfaces. In order to present these principles in a more or less cohesive manner, I will separate them into two categories: first are methods based on the "realism" (or isomorphism) in 3D UI design, an approach which tries to borrow ideas from the real world. The second category is "magic" or non-isomorphism. In this approach we are trying to design interfaces that are significantly different from the real world and that allow the user to interact with 3D computer graphics environments in a very different manner than in a real world.

Methods and ideas that will be discussed in this lecture are general in the sense that they apply not to a single task or application, which was the case of the techniques that we talked about before, but to any interaction technique, interface or 3D system. Most of the design principles discussed here are informal, based on rules of thumb, esthetics, and stealing from other areas of human activity. Nevertheless, many of today's successful 3D interfaces have been designed based on these ideas.

Most of the interface design principles from human factors research can be directly applied to designing 3D user interfaces. Simplicity, consistency, feedback from operations, error prevention and aesthetic appeal are as important in 3D interaction as in any other human-machine interfaces.

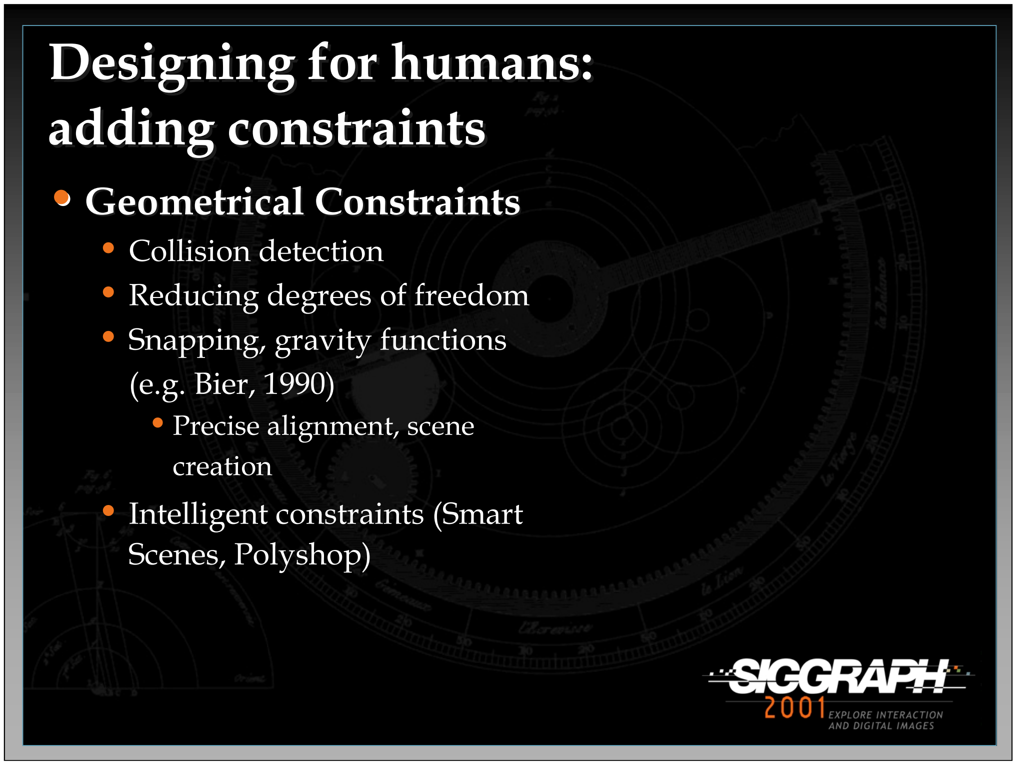

The simplest example of constraints is using collision detection: the users freedom is limited by not allowing them to go freely through virtual objects, which in many cases makes interaction much easier, especially in navigation techniques. Constraints have been used in 3D user interfaces in two main forms. First, they were used to reduce the number of degrees of freedom to make manipulation simpler. For example, an object can be constrained to only move on the surface of a plane, making it positioning simpler.

Second, constraints we used to snap objects to a 3D grid or special guiding objects, e.g. surfaces, lines and planes with which manipulated objects aligns (e.g. Bier, 1990). Snapping can make selection and object arrangement significantly easier.

Using both hands while interacting with computers allows us to transfer our everyday manipulation skills into interaction with virtual environments which makes it easier and more effective. That us why, two-handed or bi-manual input has been investigated extensively in 2D interfaces (e.g. Buxton et al. in 1986). In 3D interaction two-handed input has also been asuccessfully used to design compelling 3D user interfaces (e.g. Mapes and Moshell, 1995).

There are two basic ways to incorporate two hands in 3D interaction. The first is *symmetric* bi-manual manipulation, where each hand can be used to perform different, separate tasks. Interaction is symmetrical in the sense that input from both hands is equal, for example typing on a keyboard.

A second approach is to allow the user to use both hands to perform a single task, for example selecting from a hand held menu or rotating an object with one hand while fixing the center of rotation with the other. In this case of bi-manual input the use of the hand is asymmetrical in the sense that each hand assumes a certain role that depends on the action of the other hand. Hinckley (1997) demonstrated that in the cooperative two-handed manipulation the left (non-dominant) hand defines a general spatial frame of reference for precise actions of the right (dominant) hand. This property can be explicitly used to design interesting interaction techniques for object manipulation, for example the Voodoo Dolls technique (Pierce et al., 1999).

Feedback plays crucial importance in designing user interfaces. Our ability to self-regulate body movements, e.g., manipulating objects, depends on spatial and temporal correspondence between a large variety of sensory feedbacks: visual, tactile, kinesthetic, proprioceptive and others. If the 3D user interface response, e.g., visual feedback, conflicts with kinesthetic or proprioceptive feedback produced by the human motor system, then the user performance degrades (e.g. Smith and Smith, 1987). Hence, the quality of 3D user interfaces depends on whether they preserve *compliances* between the multiple-dimensions of sensory feedback, i.e., a stimulus-response (S-R) compatibility (Fitts, 1953).

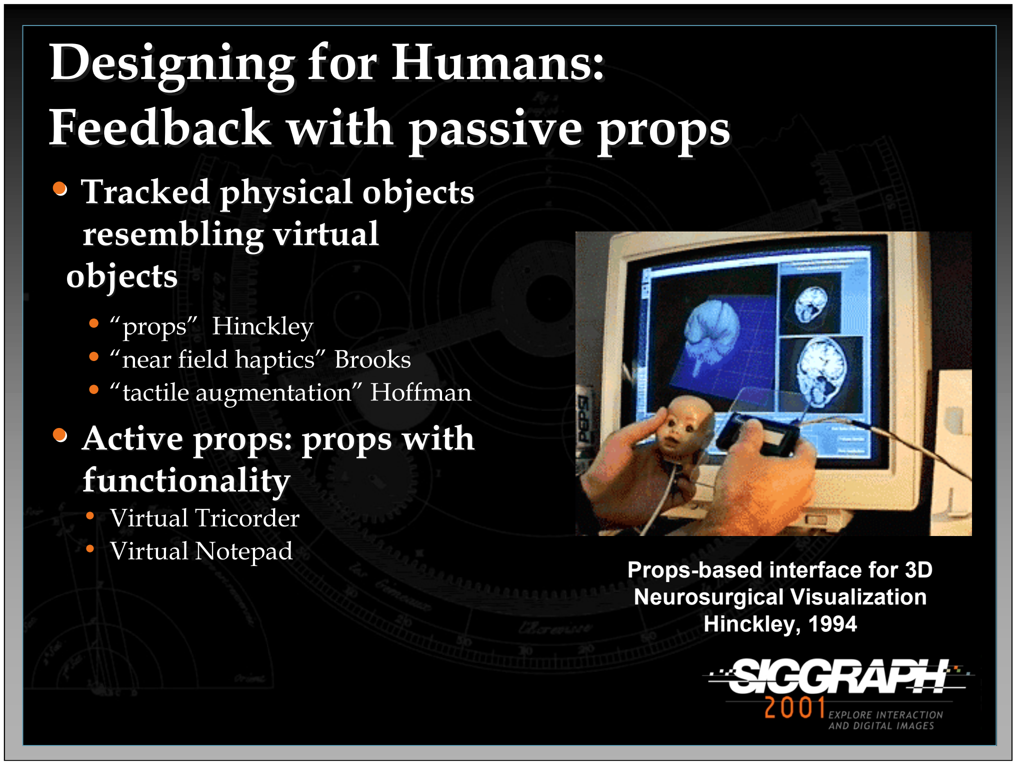

One popular technique in 3D user interface design is to consider and control the physical shape of input devices, providing the user with passive tactile feedback. The main idea here is to match the shape and/or appearance of a physical object with a virtual one. The term that is used to refer to this technique is "*passive physical props*" and they were first introduced by Hinckley when he described a 3D interface for visualizing and interacting with neursurgical data (Hinckley et al., 1994). The input device in his interface was a doll's head fitted with a 6DOF sensor. By manipulating the dolls head, the user was able to quickly and reliably relate the orientation of the input device to the orientation of the brain data on the screen, resulting in efficient and enjoyable interaction.

The interface designed by Hinckley was non-immersive, in immersive environments the passive props can also be spatially registered with virtual objects providing inexpensive physical feedback to the user. Hoffman refers to the technique as "tactile augmentation" (Hoffman et al., 1998).

-continued on the next page

Virtual tricorder by Wloka et al. 1996

Extending this technique, we can design *active physical props* by matching shape and functionality of input devices with a corres-ponding virtual object. Active physical props were first introduced by Wloka (1995), where a physical mouse was registered with a virtual mouse of the exact same shape, allowing the user to easily interact with it while immersed in a virtual environment. Another example of an active physical prop is the Virtual Notepad, where a virtual tablet is matched with a real pressure-sensitive tablet, allowing the user to write while immersed in VE (Poupyrev et al, 1998).

Designing for Humans: Feedback with passive props

- **Advantages**
  - Haptic and tactile feedback
  - Increases presence (Hoffman, 1998)
  - Common reference frame (Hinckley, 1994)
  - Ease of learning
  - Direct control of realism
    - e.g. treatment of phobias
- **Disadvantages**
  - Tracking multiple objects
  - Does not increase performance (Hinckley 97, Ware and Rose 99)

Using props in treatment of spider phobia Carlin, Hoffman et al., 1994

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Using passive and active physical props is an extremely useful design technique for 3D user interfaces. Props allow us to add inexpensive physical and tactile feedback, significantly increasing presence for immersive environments (Hoffman et al., 1998) and establishing a common frame of reference between the device and desktop 3D user interfaces. The introduction of tactile augmentation allows us to explicitly control the realism of virtual environments, which can be useful in such applications as the treatment of phobias (Carlin et al. 1997).

The disadvantages of props are that they require tracking of multiple physical objects, which might be expensive. Also, experimental studies done so far have not shown any improvement in user performance for motor tasks when using props (Hinckley 97, Ware and Rose 99).

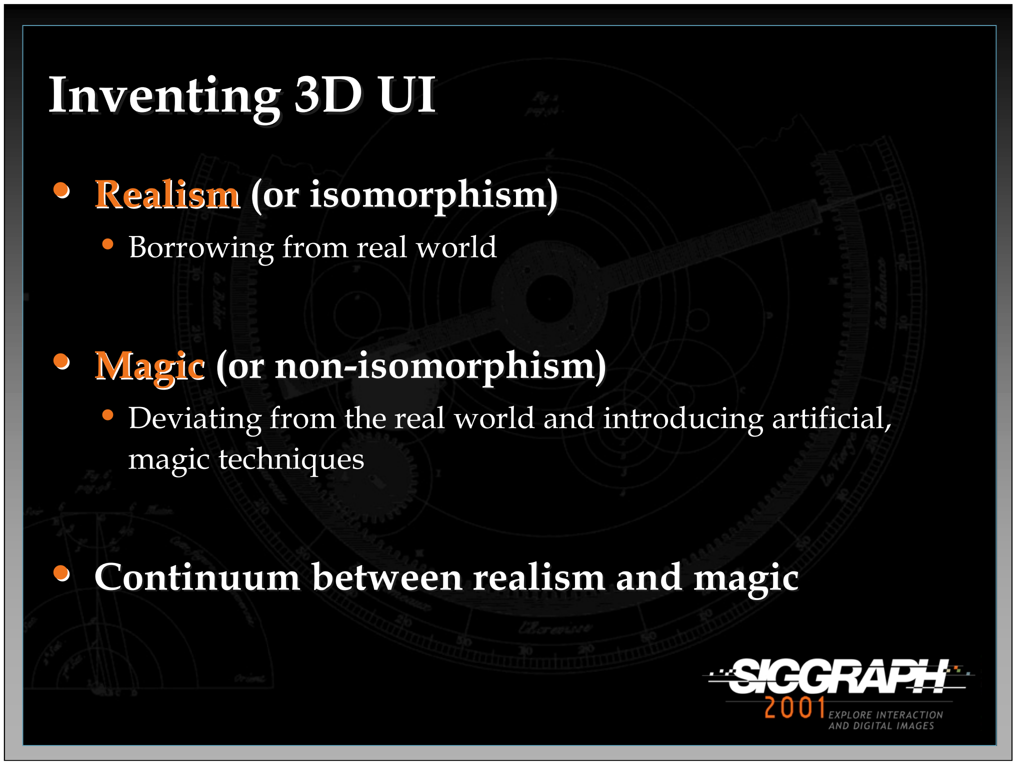

While human factors principles offer important design guidelines for 3D interface design, they do not really suggest how we can create compelling 3D interfaces or invent effective techniques. In this section, I survey some of the informal, rule-of-thumb approaches that have been taken in designing 3D user interfaces.

The approaches that I will discuss are categorized into two categories. The first includes approaches based on "realism" (or isomorphism), an approach that tries to borrow ideas for 3D interface design from the real world. The second is the "magic" approach, or non-isomorphism, in which we are trying to design interfaces that are significantly different from the real world and allow the user to interact with 3D computer graphics environments in a very different manner then in a physical world.

While this categorization is useful, it is not very strict – usually there is a continuum between realism and magic.

The basic approach to design 3D interfaces is to simply imitate the real, physical world as closely as possible. This approach is important for all simulation applications, such as training, battle field simulation, some entertainment applications, and evaluations of the usability of complex human-controlled mechanisms such as cars and tractors.

The advantages of this approach is that the user already knows how to use the interface from everyday experience, and an interface can be implemented either on the basis of the designer's intuition or clearly specified technical design requirements, such as in simulation applications.

The problem, however, is that the simulation is never exact due to the limitations of the technology. In non-simulation applications this approach introduces the same limitation as we have in the real world, which might be annoying and inefficient for the user. Furthermore, even in a simulation application there are often tasks that do not directly relate to the simulation itself, e.g. system control, and require use of 3D interaction techniques.

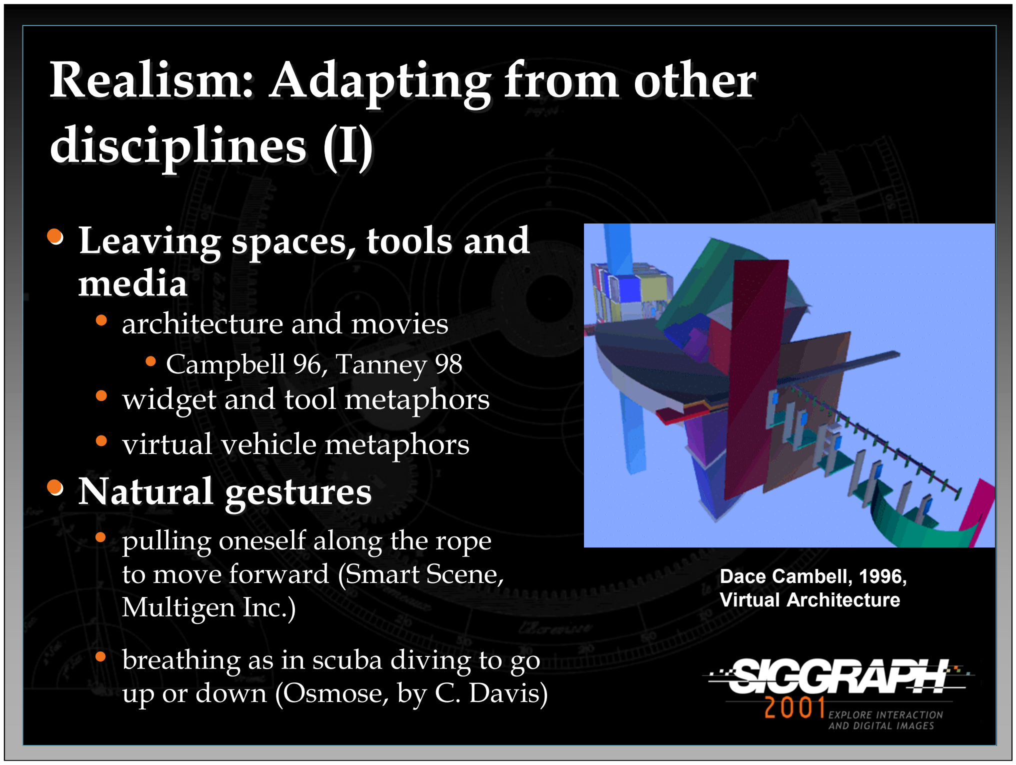

Instead of mimicking a real world we can steal and adopt ideas and/or existing artifacts from the real world. Indeed, movies and architecture has been a source of inspiration for much of VE design (e.g. Campbell 1996, Herndon et al., 1994). The virtual vehicle metaphor has been probably one of the most used techniques for navigation. Virtual widgets and tools have been often adopted from real-world physical tools and objects: for example in the dVise system from Division Inc, a lamp widget was used to set up lighting and an egg widget was used to create new objects.

Another way to adapt a real world for the virtual is to borrow natural physical gestures to perform interaction tasks. For example, in the Smart Scenes by Multigen Inc. (see also Mapes, 1995) the user moves in the environment by pulling himself or herself along an invisible rope. An even more radical method was used in the Osmose environment, where the user navigated by using breathing and balance control, a technique borrowed from the scuba diving technique of buoyancy control. The user is able to float upward by breathing in, to fall by breathing out, and to change direction and by altering the body's center of balance. The intent was to create an illusion of floating rather than flying or driving (http://www.softimage.com/Projects/Osmose/).

While borrowing from the real world, these techniques do not simply mimic but rather adopt real world tools and gestures, which can make the interface rather intuitive.

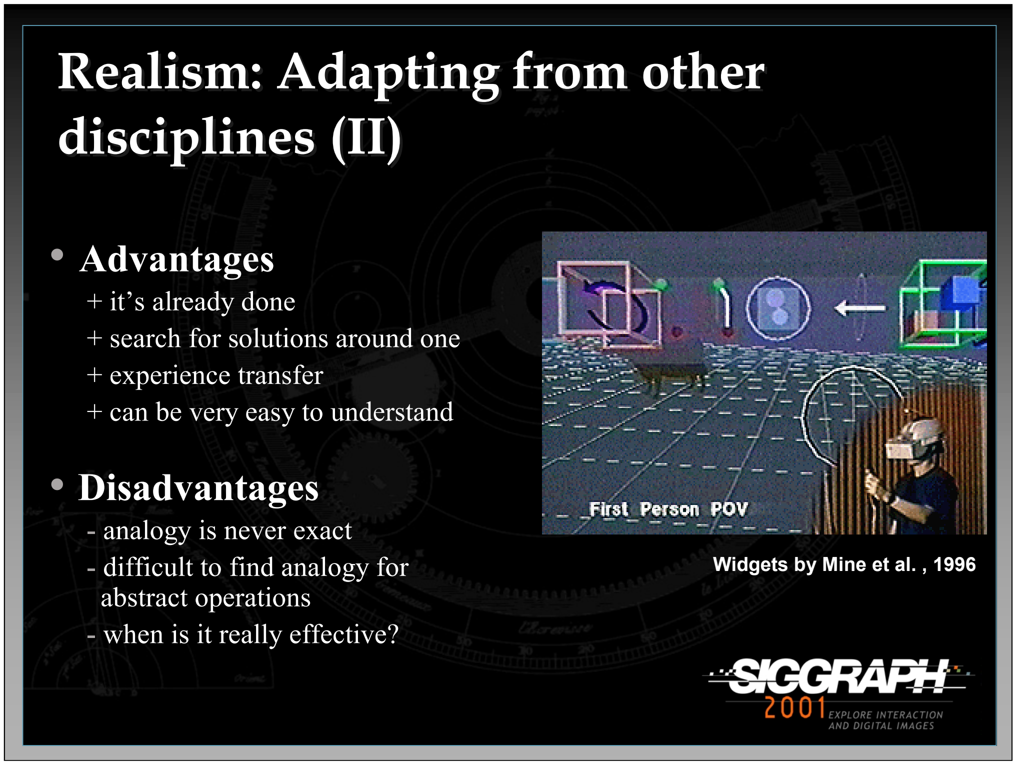

The advantages of this approach are that there is a large number of objects that we can adopt for VR interaction in almost any task, the user can transfer his or her own real world experience to virtual worlds, and analogies are usually easy to understand.

The disadvantages are that any analogy is never complete and it is usually difficult to find good analogies for abstract operations. For example in the dVise system from Division an egg widget was used to create new objects, a metaphor that is not so transparent. Finally, it is not clear if the adaptation is effective or not unless we conduct extensive experimental studies.

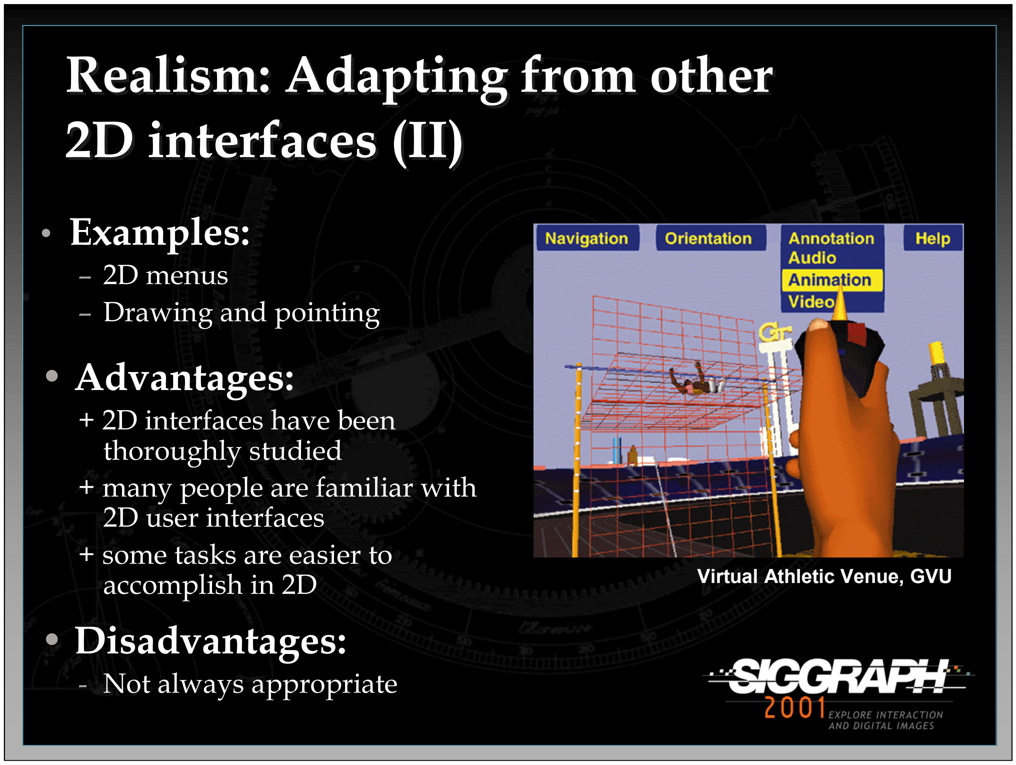

When we cannot borrow from the real world, why don't we borrow from 2D user interfaces? There have been quite a few attempts to do this, usually for system control tasks, menus and symbol control (e.g. Bolter, 1995). The major advantages are that 2D user interfaces have been thoroughly studied and today's users are quite familiar with 2D interfaces. The problem is that the 2D interfaces are not always appropriate for 3D interaction tasks simply because they have not been designed for 3D interfaces.

One of the approach of magic in 3D use interfaces would be to extend the user ability or change the geometrical properties of the real world. Two examples considered here are the Go-Go techniques (Poupyrev, et al. 1996) and World-in-Miniture (Stoackley, et al. 1995). The **Go-Go** technique, flexibly extends the virtual hand technique reaching distance by using a non-linear mapping function applied to the user's real hand extension. The space around the user is split into two concentric regions. While the user's the real hand is within the first closest region around the user, that is, the distance to the hand is smaller then some threshold distance D, the mapping is one-to-one and the movements of the virtual hand correspond to the real hand movements (see figure below). However, as the user extends her hand further than D, the mapping becomes non-linear and the virtual arm "grows" allowing the user to reach and manipulate distant objects.

-continued on the next page

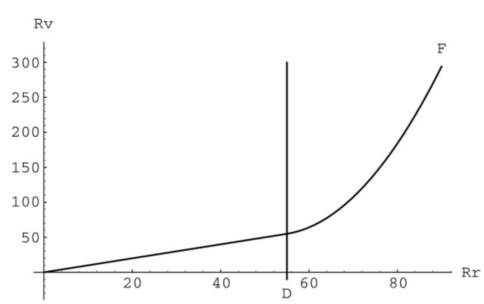

Mapping function for the Go-Go
(Poupyrev, et al. 1996)

The **World-In-Miniature** (WIM) technique (Stoakley et al. 1995) provides the user with a miniature hand-held model of the VE, which is scaled down using some constant coefficient (see figure below). The user can then indirectly manipulate virtual objects by interacting with their representations in the WIM.

The WIM technique is a powerful technique allowing easy object manipulation both within and outside of the area of user reach. It also can combine navigation with manipulation since the user can easily move his or her own representation on the WIM. The downside of the technique is that scaling large environment results in very small representations of objects in the WIM, so accurately manipulating small objects might be difficult. A technique that can choose the part of the environment within the WIM might overcome this problem.

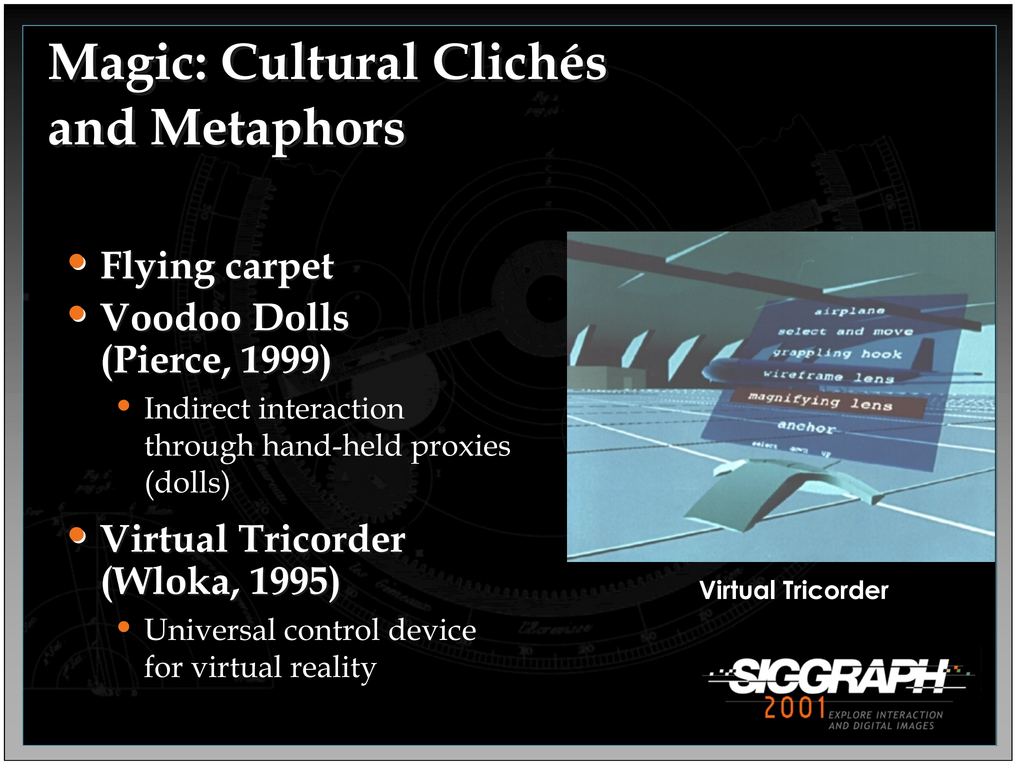

Cultural Clichés and metaphors, such as flying carpet, can also suggest an interesting approaches in designing 3D user interfaces. For example a **Voodoo Dolls** technique (Pierce, et al. 1999) is a two-handed interaction technique for manipulating objects at a distance in immersive virtual environments. The technique combines and builds upon a number of other techniques, such as Image Plane (Pierce et al., 1997) and WIM (Stoakley et al., 1995). Voodoo Dolls uses a couple of pinch gloves to allow the user to switch seamlessly between different modes of manipulation. It aims to provide an easy method of interacting with objects of widely varying sizes and at different distances. The technique is based on several ideas. First, to start object manipulation the user dynamically creates dolls: temporary, miniature, hand-held copies of objects. Similar to the WIM technique, the user can interact with objects in the environment by manipulating these dolls instead of directly manipulating the objects so that manipulated virtual objects can be at any distance, size and state of occlusion.

Second, the technique allows the user to explicitly and interactively specify a frame of reference for manipulation. The doll that the user holds in the non-dominant hand represents a stationary frame of reference, and the corresponding virtual object does not move when the the user moves this doll.

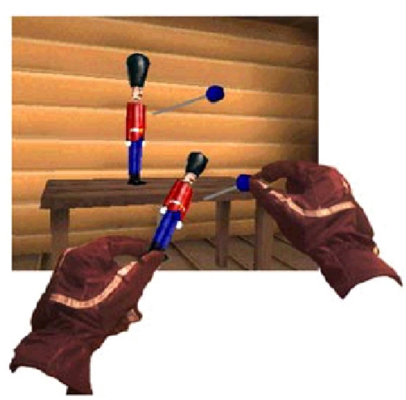

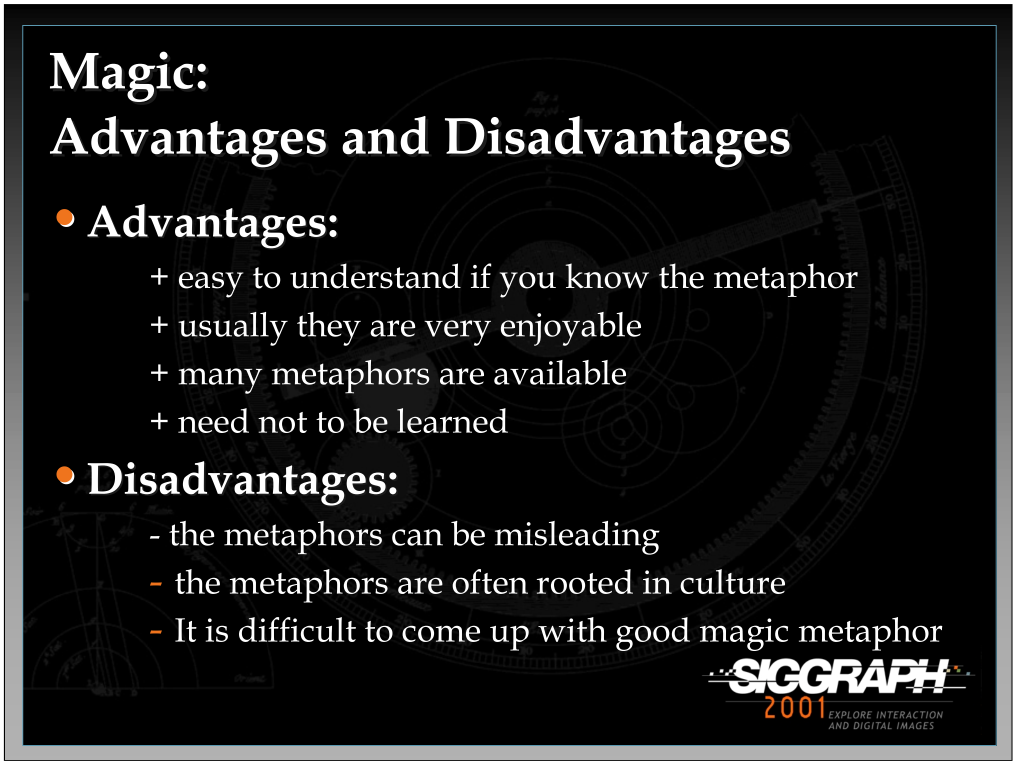

Magical approach does not try to incorporate properties of the physical world into the virtual environment, but rather extends them by inventing "magical" interfaces . All of these techniques are based on certain "magical" metaphors and are very easy to understand if one knows the metaphor. Many metaphors are available that can be used and they do not needed to be learned if the user already knows about them. Thus the resulting interface can be very easy to learn and used right away.

A problem with this approach is that metaphors are never complete, and they are often misleading, especially magical ones. Metaphors are rooted in culture: indeed if one has never heard about a flying carpet then the metaphor might not work. Finally, it is not that easy to find effective and compelling magical metaphors. However, if one is found it can provide a very enjoyable user interface.

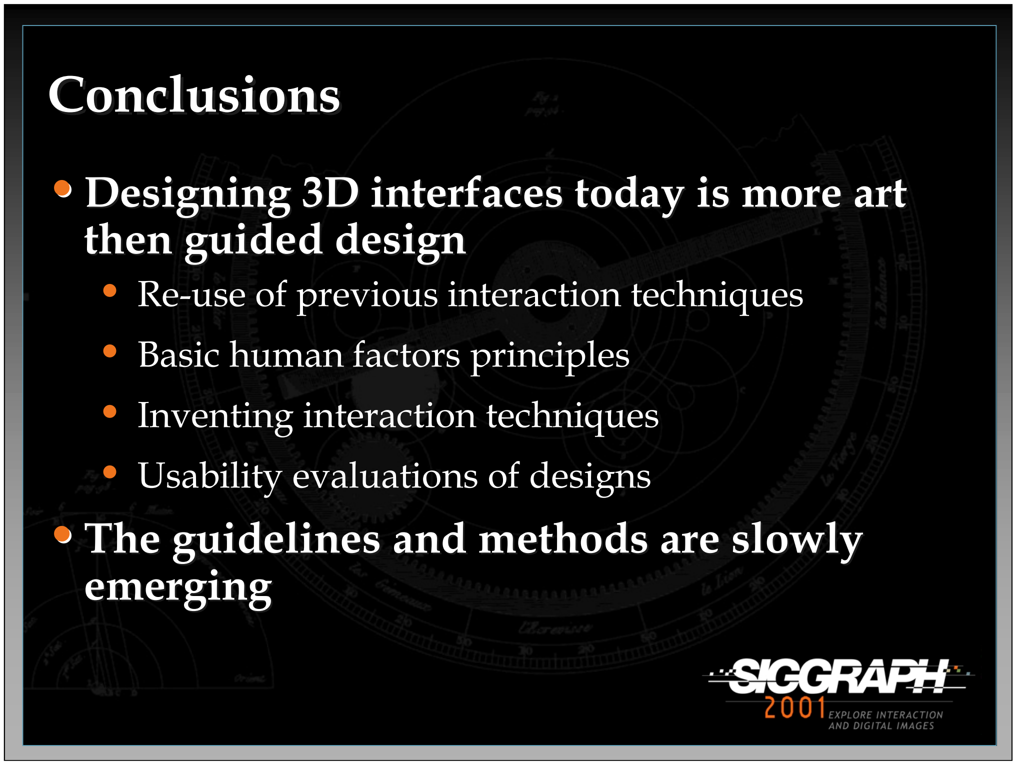

Bibliography

For most of the references see reference section of the notes. Below are only those references that are not in the main bibliography.

Apple Computer, I., *Macintosh Human Interface Guidelines*. 1992: Addison-Wesley Publishing Company. pp. 384.

Carlin, A., Hoffman, H., Weghorst, S., Virtual reality and tactile augmentation in the treatment of spider phobia: a case report. *Behavior Research and Therapy*, 1997. 35(2): pp. 153-159.

Buxton, W., Myers, B., A Study in Two-Handed Input. *Proceedings of CHI*. 1986. ACM. pp. 321-326.

Fitts, P., Jones, R., Compatibility: spatial characteristics of stimulus and response codes. *Journal of Experimental Psychology*, 1953 (46).

Shneiderman, B., *Designing the user interface: strategies for effective human-computer interaction*. 3 ed. 1998: Addison-Wesley. pp. 638.

Smith, T., Smith, K., Feedback-control mechanisms of human behavior. In *Handbook of human factors*, G. Salvendy, Editor. 1987, John Wiley and Sons. pp. 251-293.