Fuchs 99

Chapter 1

Displays for Augmented Reality: Historical Remarks and Future Prospects

Henry Fuchs and Jeremy Ackerman

1.1 Introduction

An augmented or "mixed" reality system is one in which computer graphics objects are added into the user's surrounding scene in a meaningful way. Augmented reality includes applications ranging from "heads up" displays for fighter pilots to complex real-time visualizations of multiple data sets in medical applications. Other applications include architectural previewing and engineering design.

Augmented reality systems require the use of specialized display devices. Typically optical see-through or video see-through displays are used. It is important to note that the conception of augmented reality systems should not be limited to applications using head mounted displays. In the following sections the principles of several display technologies are presented and examples of their application to augmented reality are discussed.

1.2 Summary of Augmented Reality Display Technologies and Examples of Their Use

Display devices for augmented reality applications are most easily categorized as head mounted and non-head mounted devices.

1.2.1 Head Mounted Displays

Head mounted displays have been used for augmented reality applications since the 1960s [1], and have been the mainstay of many augmented reality systems to date.

Two kinds of head mounted displays have been used: optical see-through and video see-through. As their name implies, optical see-through systems combine the real and synthetic imagery via some optical merging mechanism, most commonly a "half-silvered" mirror. Video see-through systems combine synthetic images with "real" images of the user's surroundings by combining two video streams, one typically generated by computer, the other coming from a video camera mounted on (or near) the user's head.

Optical See-Through Head Mounted Displays

Perhaps the first augmented reality system was developed in the late 1960s by Ivan Sutherland [2], first at Harvard University and then at the University of Utah. This display is shown in figure 1.1.

The display demonstrates the key components and principles needed for optical see-through displays: There is a distinct display for each eye, since each eye's view of the world is from a slightly different vantage point; there is a see-through capability allowing each eye to view simultaneously the user's surroundings and the computer-generated imagery; a half-silvered prism sitting directly in front of each eye allows some of the light of the user's surroundings to pass through to the user's eyes. Also passing through to the user's eyes is some of the light that originates at the face of the CRTs mounted on the sides of the head. Each eye observes the combination of two images superimposed one on top of the other. If light levels in the environment are carefully adjusted, and both the environment and the synthetic imagery are sufficiently simple, both can often be comprehended simultaneously.

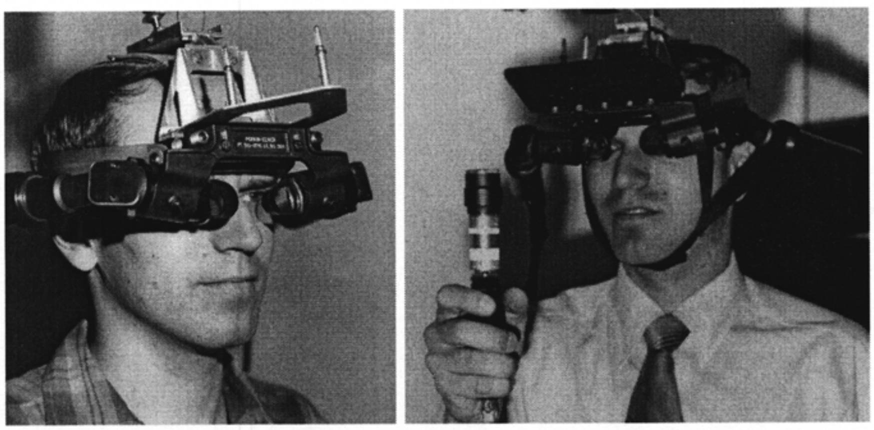

Figure 1.1: Two images of Ivan Sutherland's head-mounted display (circa 1970). One should note the CRTs on either side of the head and the optics in front of the CRTs to put the image in front of the eyes. Images are viewed on half silvered prisms, giving the user a view of both the synthetic imagery and the real world. (Courtesy, Dept. of Computer Science, University of Utah).

More recent implementations of optical see-through displays have utilized newer display technologies and novel optical systems to improve the performance of this type of head mounted display. Other display devices, such as small LCD panels, can be placed directly between the viewer's eyes and the environment, and significantly reduce the complexity of the optics and the weight and cost of the device. Several displays based on this principle are now commercially available.

Optical see-through displays have the very difficult problem of accurate registration of real and synthetic imagery in space and time. Changes in head position or orientation can occur very rapidly during one's normal interaction with the surrounding environment. The combined processes of tracking (orientation and position), and rendering and display of synthetic imagery can easily take close to 100 milleseconds. By the time an image corresponding to the tracking data gathered at the beginning of the rendering cycle can be displayed, the head has already moved considerably. The image now displayed is misregistered with the real world, which may lead to confusion, misinterpretation of the scene, and disorientation.

A second difficulty with optical see-through displays is light intensity. If the synthetic imagery is too bright relative to the ambient light, the user's environment will not be visible. If the reverse is true, the synthetic imagery will be either invisible or visible only as a vague shadow in the environment.

A final difficulty with optical see-through displays is that of occulusions. If a synthetic object should appear in front of a real world object, it will (in general) appear to be a semitransparent ghost floating in front of the object. If parts of the real world object are brightly lit, those portions of the synthetic object that are in front of the bright spots will appear more transparent or invisible.

An optical see-through display system manufactured by CAE is used in their simulators and trainers for military aviation. This head mounted display is expensive and heavy, making it impractical for many other applications. This particular head mounted display uses two light-valve type projectors for each eye. The central part of the user's visual field, the area with greatest acuity, is fed with one projector, while the larger surrounding, and lower acuity, region is fed by the second projector. An optical system is used to combine the light from each projector. This system provides greatest display resolution where it is needed most.

In the CAE system, use of a head mounted display reduces the space required for the simulator to the size of the simulated aircraft's cockpit, without the additional space needed for other display techniques used by other manufacturers. A see-through head mounted display is preferable to a virtual reality type display to enable the pilots to see the instruments and their own hands as they learn to operate an unfamiliar cockpit. This environment is ideal for the limitations of optical see-through devices discussed above. There is no reason why synthetic imagery (of the environment that the aircraft is flying through) should ever need to be overlaid on the real objects (inside the cockpit). The simulator designers can design the lighting systems inside the cockpit to insure that both the real world and the synthetic imagery are always visible to the user.

Optical see-through displays were also a part of the demonstration of augmented reality air hockey given by Mixed Reality Labs at SIGGRAPH'98. In this demonstration, both optical and video see-through displays were available to conference attendees to play air hockey. The puck and paddles were rendered (registered to a real table) in the head mounted displays of two players. The displays used were relatively light weight and comfortable [3]

Video See-Through HMDs

A video see-through display gives the user a view of the world through one or more video cameras mounted on the display. Synthetic imagery is combined with the image captured through the video cameras by the computer, and the combined signal is sent to the display.

No video see-through displays are commercially available at this time; it is, however, relatively simple to mount one or two small video cameras on a standard head mounted display designed for virtual reality (figure 1.2). Additional hardware and software are needed to combine real and synthetic imagery. One approach is to employ a frame grabber for each camera. Alternatively, techniques like chroma-keying could be employed to add the video signal into the data stream after the synthetic imagery is rendered.

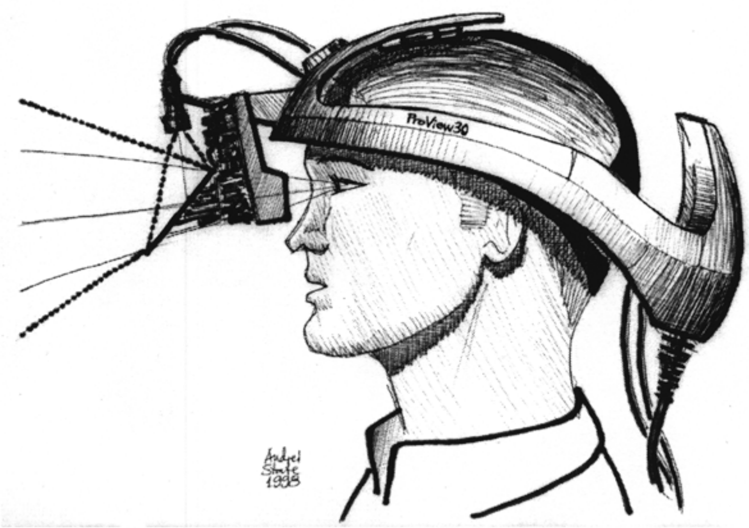

Figure 1.2: This drawing, by Andrei State of UNC-Chapel Hill, shows one method proposed to mount cameras on a Kaiser Electro Optics Proview 30 for use in the UNC augmented reality systems. The image of the world is reflected up to the camera so that the optical path length between world objects and the camera is very close to the path length to the user's eyes. The mirrors are moved with the displays when the displays adjusted for optimal viewing. ((c)1998, Dept. of Computer Science, UNC-Chapel Hill)

Video see-through also has a number of underlying difficulties. Among these are camera offset, limited field of view, temporal lag, fixed focus, and limited depth of field. When only a single camera is used, stereopsis is also lost. Some of these difficulties (camera offset and limited field of view) have been overcome in an experimental display built collaboratively by the University of North Carolina and the University of Utah (figure 1.3).

Without carefully designed optics, it is difficult to align the camera's view with the normal viewing axis of the eye. This means that the image sent to each of

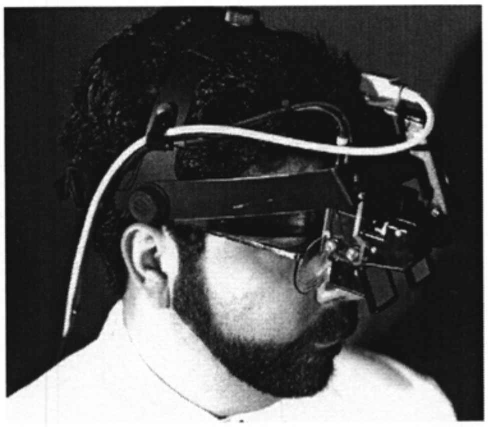

Figure 1.3: This photograph shows David Casalino, MD wearing the video see-through head mounted display, built in collaboration with the UNC-Chapel Hill and the University of Utah. This display features a relatively narrow field of view for the displays, but has an open design allowing unencumbered vision around the display. The displays and cameras are mounted on independently adjustable pods to optimize viewing performance. The camera view is aligned with the user's eyes by optics which fold the optical path. ((c)1998, Dept. of Computer Science, UNC-Chapel Hill)

the user's eyes is taken from a perspective other than that of the observer's eyes. A significant change in horizontal spacing between the cameras distorts the user's sense of depth because the stereo pairs have an effective interpupillary distance different from that to which the user is accustomed. Typically it is easiest to mount the camera or cameras above the user's eyes vertically. This displacement gives the user a false sense of height, particularly for relatively nearby objects.

Another problem is that of field of view. Care must be taken in the selection of cameras and lenses so that the field of view of the camera and lens combination matches the field of view of the displays. Without this precaution, the user's view of the world will appear noticeably distorted.

Like optical see-through displays, video see-through displays can suffer from temporal lag. However, it is possible to compensate for lag, such that the effects of lag are not noticeable in a given image being displayed by the HMD. This can be achieved by synchronizing video image capture with the report from the tracker. The synthetic imagery generated from that tracker data must be displayed with the video image captured at the same point in time. This mechanism of dealing with the lag inherent in the system solves the misregistration problem, but may introduce noticeable delays between when movements in the real world occur and when they are seen.

Fixed focus and depth of field are problems inherent in using video cameras. Objects outside the focal distance of the cameras will be out of focus, and typical cameras have a relatively narrow depth of field. While it may be possible to auto-focus type mechanisms to insure that objects seen by the camera are in focus, changing the focus may change the camera model significantly and lead to misregistration of synthetic imagery.

The problem of environmental lighting faced by optical see-through systems is not a concern in video see-through systems. Lighting must be considered when selecting cameras and camera settings, as cameras operate optimally on a much narrower range of illuminance than can the human eye.

Occlusion poses a challenging problem in video see-through head mounted displays. When a real object passes between the user and the location of a virtual object in space, the synthetic object will typically, and incorrectly, fully occlude the real object that it should appear to be behind. This confusion of depth cues can destroy the illusion that the virtual objects are truly existing in the user's environment. In some applications moving objects may be tracked with position sensors, or another method may be employed to determine the three-dimensional geometry of the real world. Without this type of three-dimensional information about real objects, this type of occlusion error is difficult, if not impossible, to solve.

At the University of North Carolina, video see-through head mounted displays have been used extensively for the surgical applications being developed [4]. Video see-through displays were chosen because of the ability to synchronize synthetic imagery with the video image from the real world, thereby reducing the likelihood that any individual image shown to the user would misrepresent the spatial relationship between real and virtual objects. Our system has used a number of displays, both that we were involved in designing (figure 1.3), and those made from commercially available displays (figure 1.2).

Our system has used several displays designed for use in virtual reality applications. We have learned that the wide field of view displays that these systems typical have, while helpful for a sense of immersion in virtual environments, do not provide adequate resolution for the surgical site at the center of the user's view. Our medical collaborators have also been unhappy with the inability to use peripheral vision and to see around the display easily. Our own design (figure 1.3) and our most recently acquired commercially available display (a Kaiser Electro Optics Proview 30, figure 1.2) solve these two problems by leaving the display open lateral to the display for each eye and by providing relatively narrow field of view displays. These devices also have the ability to flip the displays up and out of view if it is not needed.

1.2.2 Non-head Mounted Displays

There are some environments and situations where wearing a head mounted display is impractical or unfeasible. This may be true of situations requiring long term use, or extensive use of other resources in a non-augmented environment. The following subsections describe a few approaches to other types of displays.

Systems of this type provide the benefit of offering the higher resolution of a conventional monitor for augmented reality applications. The user also does not experience the fatigue and discomfort of wearing a heavy and unwieldy head mounted device.

These devices are limited in their usefulness by their limited working volume and by potential space constraints.

Small Displays

Displays the size of a small monitor (around 30cm by 20cm) can be employed to build an augmented reality system. One approach was proposed in 1975 by Knowlton [5]. This display was proposed as a way of creating a programmable keyboard. In his system, a real but unlabeled keypad was to be augmented with button labels drawn by computer graphics. In his approach, a half-silvered mirror was placed between the user and the keypad. The half-silvered mirror provides the user a view of the real objects behind the mirror, and also reflects the image from a conventional display (such as a CRT) to the user. This type of system can be thought of more generally as placing arbitrary real objects behind the mirror and adding whatever synthetic imagery is needed.

Devices of this type can be further enhanced in several ways. If the user's head position is tracked, synthetic imagery can be drawn from the correct perspective. This device can be used to reflect alternating left and right stereo pairs rather than single perspective images. A stereo illusion of dimensionality can be created if the user uses the display wearing shutter glasses [6].

Displays of this type face a problem with illumination similar to the one encountered in optical see-through head mounted devices. Real objects behind the mirror will wash out the synthetic image if they are more brightly lit than the illumination provided by the monitor. The optical see-through problem of not being able to fully occlude real surfaces with synthetic imagery is present with these devices.

These displays have a working volume limited to the space behind the mirror. Generally the mirror needs to be maintained at a fixed geometry relative to the conventional display and is, therefore, not easily moveable to a new position.

In situations where the user may want to move very close to the mirror, the display may fail in two ways. First, the user may obstruct the path of the light from the monitor to the mirror, thereby blocking the light path needed to view the synthetic imagery. Second, as the user approaches the display, each pixel subtends a greater portion of the user's field of view. As a result, the user can gain no better enhancement of their view of the synthetic objects by moving the head closer to where they perceive the object to be.

This type of display is currently being used in surgical simulators (so the surgeon's own hands are visible), and in some telesurgical devices being developed. Figure 1.4 shows how it was proposed to use these types of displays in UNC's surgical augmented reality applications.

In the field of neurosurgery another novel approach to augmented reality is used. Preoperatively acquired images, including three dimensional graphics created during surgical planning, may be overlayed onto the surgeon's view through the surgical microscope. This system is known as COMPAS and is rapidly gaining acceptance in the field of neurosurgery (Kelly).

Large Area Displays

Large displays have perhaps received less recognition than they deserve for augmented reality applications. These displays may be employed in many configurations: front projection, back projection, and even conventional CRTs. Use of larger

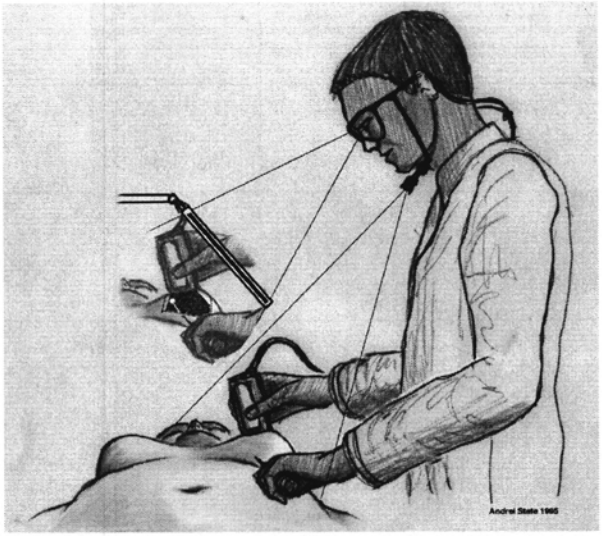

Figure 1.4: This drawing, by Andrei State of UNC-Chapel Hill, presents one vision of how a relatively small display might be employed for use in an augmented reality medical system. This particular drawing shows an alternate implementation in which a strategically placed camera enables use of video see-through. ((c)1998, Dept. of Computer Science, UNC-Chapel Hill)

displays in augmented reality is perhaps best explained by example, because they are not usually thought of as augmented reality display devices.

A display like an ImersaDesk or Virtual Workbench is usually a large monitor, set on either a horizontal or tilted work surface. A user wears shutter glasses so that the display can show three-dimensional objects in stereo. A variety of input devices can be employed to interact with the objects being displayed. These displays are limited to displaying objects near their surface. A virtual object can appear no "taller" than the plan from the far edge of the display to the user's eyes [8].

The large monitors and projector systems in civilian flight simulators may also be thought of as augmented reality devices. These displays augment the real environment of a large jet's cockpit with images of what the outside world should appear like.

Projects like UNC's Office of the Future (figure 1.5) take the previous two concepts a step further. In the Office of the Future, multiple projectors will be used to provide a virtual window into another Office. In addition to providing a compelling telepresent experience, workers in an Office will be able to collaborate on three dimensional models both locally and with remote participants. Large area displays, unlike the head mounted technologies, allow a user to easily switch back and forth between making use of the virtual objects and taking or reading notes, or using a conventional monitor.

Figure 1.5: This drawing, by Andrei State of UNC-Chapel Hill, shows a view of the Office of the Future. Large displays (provided by the projectors mounted in the ceiling), cover the walls and desk. The user is able to participate in a discussion with his colleagues who appear to him as sitting on the other side his desk. He is able to use his personal computer to make notes while working collaboratively on a design for a new head mounted display, which appears to be hovering over his desk. ((c)1998, Dept. of Computer Science, UNC-Chapel Hill)

1.3 Conclusion

Augmented reality systems have the advantage of enhancing a user's environment. The promise of this technology is very large in, among other field, engineering design, military, and medicine. The difficulties in building augmented reality systems are equally large, even at the level of the displays. For some of the problems with current displays there are potential solutions that may soon be available, but whether the problems can be solved for others is not clear.

While the information provided in this paper is not definitive, we hope that it will be useful in gaining a better understanding of the abilities of current display devices. It is important to begin developing augmented reality applications that are amenable to the current limitations, while also keeping an open mind to new ideas, which will, in time, reduce these limitations.

1.4 Acknowledgements

This work was supported by (1) "Enabling Technologies and Application Demonstrations for Synthetic Environments" (DARPA contract DABT63-93-C-0048), Frederick P. Brooks, Jr. and Henry Fuchs (UNC-CH), Principal Investigators; (2) the "National Tele-Immersion Initiative" (Jaron Lanier, Chief Scientist) of Advanced

Network and Services, (Al Weis, President and Chief Executive Officer); (3) "Science and Technology Center for Computer Graphics and Scientific Visualization" (National Science Foundation Cooperative Agreement no. EIA-8920219), Center Director Richard Riesenfeld (University of Utah). Andy van Dam (Brown University), Al Barr (Caltech), Donald Greenberg (Cornell University), and Henry Fuchs (UNC-Chapel Hill), Principal Investigators; (4) "Three-Dimensional Medical Image Display" (National Institutes of Health contract P01 CA47982-09), Stephen Pizer, Program Director. Elizabeth Bullitt MD, Edward Chaney MD, Henry Fuchs, Principal Investigators (UNC-Chapel Hill).

We gratefully acknowledge discussions over many years with close colleagues: Michael Bajura, Gary Bishop, Frederick P. Brooks Jr., David Casalino MD, Vern Chi, D'Nardo Colucci, Kurtis Keller, Mark Livingston, Leonard McMillan, Anthony A. Meyer MD, Ryutarou Ohbuchi, Etta Pisano MD, Stephen Pizer, Jannick Rolland, Michael Rosenthal, Andrei State, Greg Welch. We thank Andrei State in particular for the drawings and Andrew Ade for editorial and production assistance.

Bibliography

[1] Sutherland I. (1965) The ultimate display. In *Proc. IFIP Congress* vol. 2, pp. 506-8

[2] Sutherland I. (1968) A head-mounted three dimensional display. In *Proc. Fal Joint Computer Conference, AFIPS Conf. Proc.* vol. 33, 757-64

[3] T. Ohshima, K. (1998) AR2Hockey: A case study of collaborative augmented reality. *Proc. VRAIS'98*, pp.268-295.

[4] State, A. Technologies for Augmented-Reality Systems: realizing Ultrasound-Guided Needle Biopsies. *SIGGRAPH 96 Conference Proceedings* (New Orleans, LA, August 4-9, 1996). In *Computer Graphics Proceedings, Annual Conference Series* 1996, ACM SIGGRAPH, pgs. 439-446.

[5] Knowlton, K. (1975) Virtual pushbuttons as a means of person-machine interaction. In *Proc. Conf. Computer Graphics, Pattern Recognition and Data Structure*, 14-16 May 1975

[6] Knowlton, K.C. (1977). Computer displays optically superimposed on input devices. *Bell System Technical Journal*, 56(3), 367-83

[7] Raskar, R., G. Welch, M. Cutts, A. Lake, L. Stesin, H. Fuchs (1998) The Office of the Future : A Unified Approach to Image-Based Modeling and Spatially Immersive Displays. *SIGGRAPH 98 Conference Proceedings* (Orlando, FL, July 19-24, 1998). In *Computer Graphics Proceedings, Annual Conference Series* 1998, ACM SIGGRAPH, pgs. 179-188.

[8] Kalawsky, R.S. (1993) *The Science of Virtual Reality and Virtual Environments.* Addison-Wesley