

A Survey of Augmented Reality

Ronald T. Azuma
Hughes Research Laboratories
3011 Malibu Canyon Road, MS RL96
Malibu, CA 90265
azuma@isl.hrl.hac.com
<http://www.cs.unc.edu/~azuma>
W: (310) 317-5151
Fax: (310) 317-5695

Abstract

This paper surveys the field of Augmented Reality, in which 3-D virtual objects are integrated into a 3-D real environment in real time. It describes the medical, manufacturing, visualization, path planning, entertainment and military applications that have been explored. This paper describes the characteristics of Augmented Reality systems, including a detailed discussion of the tradeoffs between optical and video blending approaches. Registration and sensing errors are two of the biggest problems in building effective Augmented Reality systems, so this paper summarizes current efforts to overcome these problems. Future directions and areas requiring further research are discussed. This survey provides a starting point for anyone interested in researching or using Augmented Reality.

1. Introduction

1.1 Goals

This paper surveys the current state-of-the-art in Augmented Reality. It describes work performed at many different sites and explains the issues and problems encountered when building Augmented Reality systems. It summarizes the tradeoffs and approaches taken so far to overcome these problems and speculates on future directions that deserve exploration.

A survey paper does not present new research results. The contribution comes from consolidating existing information from many sources and publishing an extensive bibliography of papers in this field. While several other introductory papers have been written on this subject [Barfield95] [Bowskill95] [Caudell94] [Drascic93b] [Feiner94a] [Feiner94b] [Milgram94b] [Rolland94], this survey is more comprehensive and up-to-date. This survey provides a good beginning point for anyone interested in starting research in this area.

Section 1 describes what Augmented Reality is and the motivations for developing this technology. Six classes of potential applications that have been explored are described in Section 2. Then Section 3 discusses the issues involved in

building an Augmented Reality system. Currently, two of the biggest problems are in registration and sensing: the subjects of Sections 4 and 5. Finally, Section 6 describes some areas that require further work and research.

1.2 Definition

Augmented Reality (AR) is a variation of *Virtual Environments* (VE), or Virtual Reality as it is more commonly called. VE technologies completely immerse a user inside a synthetic environment. While immersed, the user cannot see the real world around him. In contrast, AR allows the user to see the real world, with virtual objects superimposed upon or composited with the real world. Therefore, AR supplements reality, rather than completely replacing it. Ideally, it would appear to the user that the virtual and real objects coexisted in the same space, similar to the effects achieved in the film "Who Framed Roger Rabbit?" Figure 1 shows an example of what this might look like. It shows a real desk with a real phone. Inside this room are also a virtual lamp and two virtual chairs. Note that the objects are combined in 3-D, so that the virtual lamp covers the real table, and the real table covers parts of the two virtual chairs. AR can be thought of as the "middle ground" between VE (completely synthetic) and telepresence (completely real) [Milgram94a] [Milgram94b].




Figure 1: Real desk with virtual lamp and two virtual chairs. (Courtesy ECRC)

Some researchers define AR in a way that requires the use of Head-Mounted Displays (HMDs). To avoid limiting AR to specific technologies, this survey defines AR as systems that have the following three characteristics:

- 1) Combines real and virtual
- 2) Interactive in real time
- 3) Registered in 3-D

This definition allows other technologies besides HMDs while retaining the essential components of AR. For example, it does not include film or 2-D overlays. Films like "Jurassic Park" feature photorealistic virtual objects seamlessly blended with a real environment in 3-D, but they are not interactive media. 2-D virtual overlays on top of live video can be done at interactive rates, but the overlays are not combined with the real world in 3-D. However, this definition does allow monitor-

based interfaces, monocular systems, see-through HMDs, and various other combining technologies. Potential system configurations are discussed further in Section 3.

1.3 Motivation

Why is Augmented Reality an interesting topic? Why is combining real and virtual objects in 3-D useful? Augmented Reality enhances a user's perception of and interaction with the real world. The virtual objects display information that the user cannot directly detect with his own senses. The information conveyed by the virtual objects helps a user perform real-world tasks. AR is a specific example of what Fred Brooks calls *Intelligence Amplification* (IA): using the computer as a tool to make a task easier for a human to perform [Brooks96].

At least six classes of potential AR applications have been explored: medical visualization, maintenance and repair, annotation, robot path planning, entertainment, and military aircraft navigation and targeting. The next section describes work that has been done in each area. While these do not cover every potential application area of this technology, they do cover the areas explored so far.

2. Applications

2.1 Medical

Doctors could use Augmented Reality as a visualization and training aid for surgery. It may be possible to collect 3-D datasets of a patient in real time, using non-invasive sensors like Magnetic Resonance Imaging (MRI), Computed Tomography scans (CT), or ultrasound imaging. These datasets could then be rendered and combined in real time with a view of the real patient. In effect, this would give a doctor "X-ray vision" inside a patient. This would be very useful during minimally-invasive surgery, which reduces the trauma of an operation by using small incisions or no incisions at all. A problem with minimally-invasive techniques is that they reduce the doctor's ability to see inside the patient, making surgery more difficult. AR technology could provide an internal view without the need for larger incisions.

AR might also be helpful for general medical visualization tasks in the surgical room. Surgeons can detect some features with the naked eye that they cannot see in MRI or CT scans, and vice-versa. AR would give surgeons access to both types of data simultaneously. This might also guide precision tasks, such as displaying where to drill a hole into the skull for brain surgery or where to perform a needle biopsy of a tiny tumor. The information from the non-invasive sensors would be directly displayed on the patient, showing exactly where to perform the operation.

AR might also be useful for training purposes [Kancherla95]. Virtual instructions could remind a novice surgeon of the required steps, without the need to look away from a patient to consult a manual. Virtual objects could also identify organs and specify locations to avoid disturbing [Durlach95].

Several projects are exploring this application area. At UNC Chapel Hill, a research group has conducted trial runs of scanning the womb of a pregnant woman with an ultrasound sensor, generating a 3-D representation of the fetus inside the womb and displaying that in a see-through HMD (Figure 2). The goal is to endow the doctor with the ability to see the moving, kicking fetus lying inside the womb, with the hope that this one day may become a "3-D stethoscope" [Bajura92] [State94]. More recent efforts have focused on a needle biopsy of a breast tumor. Figure 3 shows a mockup of a breast biopsy operation, where the virtual objects identify the location of the tumor and guide the needle to its target [State96b]. Other groups at the MIT AI Lab [Grimson94] [Grimson95] [Mellor95a] [Mellor95b], General Electric [Lorenzen93], and elsewhere [Betting95] [Edwards95] [Taubes94] are investigating displaying MRI or CT data, directly registered onto the patient.



Figure 2: Virtual fetus inside womb of pregnant patient. (Courtesy UNC Chapel Hill Dept. of Computer Science.)

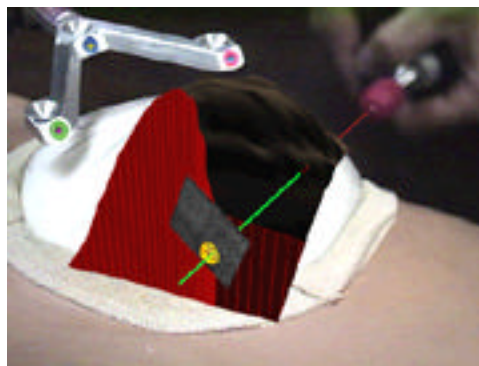


Figure 3: Mockup of breast tumor biopsy. 3-D graphics guide needle insertion. (Courtesy UNC Chapel Hill Dept. of Computer Science.)

2.2 Manufacturing and repair

Another category of Augmented Reality applications is the assembly, maintenance, and repair of complex machinery. Instructions might be easier to understand if they were available, not as manuals with text and pictures, but rather as 3-D drawings superimposed upon the actual equipment, showing step-by-step the tasks that need to be done and how to do them. These superimposed 3-D drawings can be animated, making the directions even more explicit.

Several research projects have demonstrated prototypes in this area. Steve Feiner's group at Columbia built a laser printer maintenance application [Feiner93a], shown in Figures 4 and 5. Figure 4 shows an external view, and Figure 5 shows the user's view, where the computer-generated wireframe is telling the user to remove the paper tray. A group at Boeing is developing AR technology to guide a technician in building a wiring harness that forms part of an airplane's electrical system. Storing these instructions in electronic form will save space and reduce costs. Currently, technicians use large physical layout boards to construct such harnesses, and Boeing requires several warehouses to store all these boards. Such space might be emptied for other use if this application proves successful [Caudell92] [Janin93] [Sims94]. Boeing is using a Technology Reinvestment Program (TRP) grant to investigate putting this technology onto the factory floor [BoeingTRP94]. Figure 6 shows an external view of Adam Janin using a prototype AR system to build a wire bundle. Eventually, AR might be used for any complicated machinery, such as automobile engines [Tuceryan95].



Figure 4: External view of Columbia printer maintenance application. Note that all objects must be tracked. (Courtesy Steve Feiner, Blair MacIntyre, and Dorée Seligmann, Columbia University.)




Figure 5: Prototype laser printer maintenance application, displaying how to remove the paper tray. (Courtesy Steve Feiner, Blair MacIntyre, and Dorée Seligmann, Columbia University.)



Figure 6: Adam Janin demonstrates Boeing's prototype wire bundle assembly application. (Courtesy David Mizell, Boeing)

2.3 Annotation and visualization

AR could be used to annotate objects and environments with public or private information. Applications using public information assume the availability of public databases to draw upon. For example, a hand-held display could provide information about the contents of library shelves as the user walks around the library [Fitzmaurice93] [Rekimoto95a] [Rekimoto95b]. At the European Computer-Industry Research Centre (ECRC), a user can point at parts of an engine model and the AR system displays the name of the part that is being pointed at [Rose95]. Figure 7 shows this, where the user points at the exhaust manifold on an engine model and the label "exhaust manifold" appears.



Figure 7: Engine model part labels appear as user points at them. (Courtesy ECRC)

Alternately, these annotations might be private notes attached to specific objects. Researchers at Columbia demonstrated this with the notion of attaching windows from a standard user interface onto specific locations in the world, or attached to specific objects as reminders [Feiner93b]. Figure 8 shows a window superimposed as a label upon a student. He wears a tracking device, so the computer knows his location. As the student moves around, the label follows his location, providing the AR user with a reminder of what he needs to talk to the student about.



Figure 8: Windows displayed on top of specific real-world objects. (Courtesy Steve Feiner, Blair MacIntyre, Marcus Haupt, and Eliot Solomon, Columbia University.)

AR might aid general visualization tasks as well. An architect with a see-through HMD might be able to look out a window and see how a proposed new skyscraper would change her view. If a database containing information about a building's structure was available, AR might give architects "X-ray vision" inside a building, showing where the pipes, electric lines, and structural supports are inside the walls [Feiner95]. Researchers at the University of Toronto have built a system called Augmented Reality through Graphic Overlays on Stereovideo (ARGOS) [Milgram95], which among other things is used to make images easier to understand during difficult viewing conditions [Drascic93a]. Figure 9 shows wireframe lines drawn on top of a space shuttle bay interior, while in orbit. The lines make it easier to see the geometry of the shuttle bay. Similarly, virtual lines and objects could aid navigation and scene understanding during poor visibility conditions, such as underwater or in fog.




Figure 9: Virtual lines help display geometry of shuttle bay, as seen in orbit. (Courtesy David Drascic and Paul Milgram, U. Toronto.)

2.4 Robot path planning

Teleoperation of a robot is often a difficult problem, especially when the robot is far away, with long delays in the communication link. Under this circumstance,

instead of controlling the robot directly, it may be preferable to instead control a virtual version of the robot. The user plans and specifies the robot's actions by manipulating the local virtual version, in real time. The results are directly displayed on the real world. Once the plan is tested and determined, then user tells the real robot to execute the specified plan. This avoids pilot-induced oscillations caused by the lengthy delays. The virtual versions can also predict the effects of manipulating the environment, thus serving as a planning and previewing tool to aid the user in performing the desired task. The ARGOS system has demonstrated that stereoscopic AR is an easier and more accurate way of doing robot path planning than traditional monoscopic interfaces [Drascic93b] [Milgram93]. Others have also used registered overlays with telepresence systems [Kim93] [Kim96] [Oyama93] [Tharp94] [Yoo93]. Figure 10 shows how a virtual outline can represent a future location of a robot arm.




Figure 10: Virtual lines show a planned motion of a robot arm (Courtesy David Drascic and Paul Milgram, U. Toronto.)

2.5 Entertainment

At SIGGRAPH '95, several exhibitors showed "Virtual Sets" that merge real actors with virtual backgrounds, in real time and in 3-D. The actors stand in front of a large blue screen, while a computer-controlled motion camera records the scene. Since the camera's location is tracked, and the actor's motions are scripted, it is possible to digitally composite the actor into a 3-D virtual background. For example, the actor might appear to stand inside a large virtual spinning ring, where the front part of the ring covers the actor while the rear part of the ring is covered by the actor. The entertainment industry sees this as a way to reduce production costs: creating and storing sets virtually is potentially cheaper than constantly building new physical sets from scratch. The ALIVE project from the MIT Media Lab goes one step further by

populating the environment with intelligent virtual creatures that respond to user actions [Maes95].

2.6 Military aircraft

For many years, military aircraft and helicopters have used Head-Up Displays (HUDs) and Helmet-Mounted Sights (HMS) to superimpose vector graphics upon the pilot's view of the real world. Besides providing basic navigation and flight information, these graphics are sometimes registered with targets in the environment, providing a way to aim the aircraft's weapons. For example, the chin turret in a helicopter gunship can be slaved to the pilot's HMS, so the pilot can aim the chin turret simply by looking at the target. Future generations of combat aircraft will be developed with an HMD built into the pilot's helmet [Wanstall89].

3. Characteristics

This section discusses the characteristics of AR systems and design issues encountered when building an AR system. Section 3.1 describes the basic characteristics of augmentation. There are two ways to accomplish this augmentation: optical or video technologies. Section 3.2 discusses their characteristics and relative strengths and weaknesses. Blending the real and virtual poses problems with focus and contrast (Section 3.3), and some applications require portable AR systems to be truly effective (Section 3.4). Finally, Section 3.5 summarizes the characteristics by comparing the requirements of AR against those for Virtual Environments.

3.1 Augmentation

Besides *adding* objects to a real environment, Augmented Reality also has the potential to *remove* them. Current work has focused on adding virtual objects to a real environment. However, graphic overlays might also be used to remove or hide parts of the real environment from a user. For example, to remove a desk in the real environment, draw a representation of the real walls and floors behind the desk and "paint" that over the real desk, effectively removing it from the user's sight. This has been done in feature films. Doing this interactively in an AR system will be much harder, but this removal may not need to be photorealistic to be effective.

Augmented Reality might apply to all senses, not just sight. So far, researchers have focused on blending real and virtual images and graphics. However, AR could be extended to include sound. The user would wear headphones equipped with microphones on the outside. The headphones would add synthetic, directional 3-D sound, while the external microphones would detect incoming sounds from the environment. This would give the system a chance to mask or cover up selected real sounds from the environment by generating a masking signal that exactly canceled

the incoming real sound [Durlach95]. While this would not be easy to do, it might be possible. Another example is haptics. Gloves with devices that provide tactile feedback might augment real forces in the environment. For example, a user might run his hand over the surface of a real desk. Simulating such a hard surface virtually is fairly difficult, but it is easy to do in reality. Then the tactile effectors in the glove can augment the feel of the desk, perhaps making it feel rough in certain spots. This capability might be useful in some applications, such as providing an additional cue that a virtual object is at a particular location on a real desk [Wellner93].

3.2 Optical vs. video

A basic design decision in building an AR system is how to accomplish the combining of real and virtual. Two basic choices are available: optical and video technologies. Each has particular advantages and disadvantages. This section compares the two and notes the tradeoffs. For additional discussion, see [Rolland94].

A see-through HMD is one device used to combine real and virtual. Standard *closed-view HMDs* do not allow any direct view of the real world. In contrast, a *see-through HMD* lets the user see the real world, with virtual objects superimposed by optical or video technologies.

Optical see-through HMDs work by placing optical combiners in front of the user's eyes. These combiners are partially transmissive, so that the user can look directly through them to see the real world. The combiners are also partially reflective, so that the user sees virtual images bounced off the combiners from head-mounted monitors. This approach is similar in nature to *Head-Up Displays* (HUDs) commonly used in military aircraft, except that the combiners are attached to the head. Thus, optical see-through HMDs have sometimes been described as a "HUD on a head" [Wanstall89]. Figure 11 shows a conceptual diagram of an optical see-through HMD. Figure 12 shows two optical see-through HMDs made by Hughes Electronics.

The optical combiners usually reduce the amount of light that the user sees from the real world. Since the combiners act like half-silvered mirrors, they only let in some of the light from the real world, so that they can reflect some of the light from the monitors into the user's eyes. For example, the HMD described in [Holmgren92] transmits about 30% of the incoming light from the real world. Choosing the level of blending is a design problem. More sophisticated combiners might vary the level of contributions based upon the wavelength of light. For example, such a combiner might be set to reflect all light of a certain wavelength and none at any other wavelengths. This would be ideal with a monochrome monitor. Virtually all the light from the monitor would be reflected into the user's eyes, while almost all the light from the real world (except at the particular wavelength) would reach the user's eyes. However, most existing optical see-through HMDs do reduce the amount of light from the real world, so they act like a pair of sunglasses when the power is cut off.




Figure 11: Optical see-through HMD conceptual diagram



Figure 12: Two optical see-through HMDs, made by Hughes Electronics

In contrast, video see-through HMDs work by combining a closed-view HMD with one or two head-mounted video cameras. The video cameras provide the user's view of the real world. Video from these cameras is combined with the graphic images created by the scene generator, blending the real and virtual. The result is sent to the monitors in front of the user's eyes in the closed-view HMD. Figure 13 shows a conceptual diagram of a video see-through HMD. Figure 14 shows an actual video see-through HMD, with two video cameras mounted on top of a Flight Helmet.




Figure 13: Video see-through HMD conceptual diagram



Figure 14: An actual video see-through HMD. (Courtesy Jannick Rolland, Frank Biocca, and UNC Chapel Hill Dept. of Computer Science. Photo by Alex Treml.)

Video composition can be done in more than one way. A simple way is to use chroma-keying: a technique used in many video special effects. The background of the computer graphic images is set to a specific color, say green, which none of the virtual objects use. Then the combining step replaces all green areas with the corresponding parts from the video of the real world. This has the effect of superimposing the virtual objects over the real world. A more sophisticated composition would use depth information. If the system had depth information at each pixel for the real world images, it could combine the real and virtual images by a pixel-by-pixel depth comparison. This would allow real objects to cover virtual objects and vice-versa.

AR systems can also be built using monitor-based configurations, instead of see-through HMDs. Figure 15 shows how a monitor-based system might be built. In this case, one or two video cameras view the environment. The cameras may be static or mobile. In the mobile case, the cameras might move around by being attached to a robot, with their locations tracked. The video of the real world and the graphic images generated by a scene generator are combined, just as in the video see-through HMD case, and displayed in a monitor in front of the user. The user does not wear the display device. Optionally, the images may be displayed in stereo on the monitor, which then requires the user to wear a pair of stereo glasses. Figure 16 shows an external view of the ARGOS system, which uses a monitor-based configuration.




Figure 15: Monitor-based AR conceptual diagram



Figure 16: External view of the ARGOS system, an example of monitor-based AR. (Courtesy David Drascic and Paul Milgram, U. Toronto.)

Finally, a monitor-based optical configuration is also possible. This is similar to Figure 11 except that the user does not wear the monitors or combiners on her head. Instead, the monitors and combiners are fixed in space, and the user positions her head to look through the combiners. This is typical of Head-Up Displays on military aircraft, and at least one such configuration has been proposed for a medical application [Peuchot95].

The rest of this section compares the relative advantages and disadvantages of optical and video approaches, starting with optical. An optical approach has the following advantages over a video approach:

1) *Simplicity*: Optical blending is simpler and cheaper than video blending. Optical approaches have only one "stream" of video to worry about: the graphic images. The real world is seen directly through the combiners, and that time delay is generally a few nanoseconds. Video blending, on the other hand, must deal with separate video streams for the real and virtual images. Both streams have inherent

delays in the tens of milliseconds. Digitizing video images usually adds at least one *frame time* of delay to the video stream, where a frame time is how long it takes to completely update an image. A monitor that completely refreshes the screen at 60 Hz has a frame time of 16.67 ms. The two streams of real and virtual images must be properly synchronized or temporal distortion results. Also, optical see-through HMDs with narrow field-of-view combiners offer views of the real world that have little distortion. Video cameras almost always have some amount of distortion that must be compensated for, along with any distortion from the optics in front of the display devices. Since video requires cameras and combiners that optical approaches do not need, video will probably be more expensive and complicated to build than optical-based systems.

2) *Resolution*: Video blending limits the resolution of what the user sees, both real and virtual, to the resolution of the display devices. With current displays, this resolution is far less than the resolving power of the fovea. Optical see-through also shows the graphic images at the resolution of the display device, but the user's view of the real world is not degraded. Thus, video reduces the resolution of the real world, while optical see-through does not.

3) *Safety*: Video see-through HMDs are essentially modified closed-view HMDs. If the power is cut off, the user is effectively blind. This is a safety concern in some applications. In contrast, when power is removed from an optical see-through HMD, the user still has a direct view of the real world. The HMD then becomes a pair of heavy sunglasses, but the user can still see.

4) *No eye offset*: With video see-through, the user's view of the real world is provided by the video cameras. In essence, this puts his "eyes" where the video cameras are. In most configurations, the cameras are not located exactly where the user's eyes are, creating an offset between the cameras and the real eyes. The distance separating the cameras may also not be exactly the same as the user's interpupillary distance (IPD). This difference between camera locations and eye locations introduces displacements from what the user sees compared to what he expects to see. For example, if the cameras are above the user's eyes, he will see the world from a vantage point slightly taller than he is used to. Video see-through can avoid the eye offset problem through the use of mirrors to create another set of optical paths that mimic the paths directly into the user's eyes. Using those paths, the cameras will see what the user's eyes would normally see without the HMD. However, this adds complexity to the HMD design. Offset is generally not a difficult design problem for optical see-through displays. While the user's eye can rotate with respect to the position of the HMD, the resulting errors are tiny. Using the eye's center of rotation as the viewpoint in the computer graphics model should eliminate any need for eye tracking in an optical see-through HMD [Holloway95].

Video blending offers the following advantages over optical blending:

1) *Flexibility in composition strategies*: A basic problem with optical see-through is that the virtual objects do not completely obscure the real world objects, because the optical combiners allow light from both virtual and real sources. Building an optical see-through HMD that can selectively shut out the light from the real world is difficult. In a normal optical system, the objects are designed to be in

focus at only one point in the optical path: the user's eye. Any filter that would selectively block out light must be placed in the optical path at a point where the image is in focus, which obviously cannot be the user's eye. Therefore, the optical system must have *two* places where the image is in focus: at the user's eye and the point of the hypothetical filter. This makes the optical design much more difficult and complex. No existing optical see-through HMD blocks incoming light in this fashion. Thus, the virtual objects appear ghost-like and semi-transparent. This damages the illusion of reality because occlusion is one of the strongest depth cues. In contrast, video see-through is far more flexible about how it merges the real and virtual images. Since both the real and virtual are available in digital form, video see-through compositors can, on a pixel-by-pixel basis, take the real, or the virtual, or some blend between the two to simulate transparency. Because of this flexibility, video see-through may ultimately produce more compelling environments than optical see-through approaches.

2) *Wide field-of-view*: Distortions in optical systems are a function of the radial distance away from the optical axis. The further one looks away from the center of the view, the larger the distortions get. A digitized image taken through a distorted optical system can be undistorted by applying image processing techniques to unwarpage the image, provided that the optical distortion is well characterized. This requires significant amounts of computation, but this constraint will be less important in the future as computers become faster. It is harder to build wide field-of-view displays with optical see-through techniques. Any distortions of the user's view of the real world must be corrected *optically*, rather than digitally, because the system has no digitized image of the real world to manipulate. Complex optics are expensive and add weight to the HMD. Wide field-of-view systems are an exception to the general trend of optical approaches being simpler and cheaper than video approaches.

3) *Real and virtual view delays can be matched*: Video offers an approach for reducing or avoiding problems caused by temporal mismatches between the real and virtual images. Optical see-through HMDs offer an almost instantaneous view of the real world but a delayed view of the virtual. This temporal mismatch can cause problems. With video approaches, it is possible to delay the video of the real world to match the delay from the virtual image stream. For details, see Section 4.3.

4) *Additional registration strategies*: In optical see-through, the only information the system has about the user's head location comes from the head tracker. Video blending provides another source of information: the digitized image of the real scene. This digitized image means that video approaches can employ additional registration strategies unavailable to optical approaches. Section 4.4 describes these in more detail.

5) *Easier to match the brightness of real and virtual objects*: This is discussed in Section 3.3.

Both optical and video technologies have their roles, and the choice of technology depends on the application requirements. Many of the mechanical assembly and repair prototypes use optical approaches, possibly because of the cost and safety issues. If successful, the equipment would have to be replicated in large numbers to equip workers on a factory floor. In contrast, most of the prototypes for

medical applications use video approaches, probably for the flexibility in blending real and virtual and for the additional registration strategies offered.

3.3 Focus and contrast

Focus can be a problem for both optical and video approaches. Ideally, the virtual should match the real. In a video-based system, the combined virtual and real image will be projected at the same distance by the monitor or HMD optics. However, depending on the video camera's depth-of-field and focus settings, parts of the real world may not be in focus. In typical graphics software, everything is rendered with a pinhole model, so all the graphic objects, regardless of distance, are in focus. To overcome this, the graphics could be rendered to simulate a limited depth-of-field, and the video camera might have an autofocus lens.

In the optical case, the virtual image is projected at some distance away from the user. This distance may be adjustable, although it is often fixed. Therefore, while the real objects are at varying distances from the user, the virtual objects are all projected to the same distance. If the virtual and real distances are not matched for the particular objects that the user is looking at, it may not be possible to clearly view both simultaneously.

Contrast is another issue because of the large dynamic range in real environments and in what the human eye can detect. Ideally, the brightness of the real and virtual objects should be appropriately matched. Unfortunately, in the worst case scenario, this means the system must match a very large range of brightness levels. The eye is a logarithmic detector, where the brightest light that it can handle is about eleven orders of magnitude greater than the smallest, including both dark-adapted and light-adapted eyes. In any one adaptation state, the eye can cover about six orders of magnitude. Most display devices cannot come close to this level of contrast. This is a particular problem with optical technologies, because the user has a direct view of the real world. If the real environment is too bright, it will wash out the virtual image. If the real environment is too dark, the virtual image will wash out the real world. Contrast problems are not as severe with video, because the video cameras themselves have limited dynamic response, and the view of both the real and virtual is generated by the monitor, so everything must be clipped or compressed into the monitor's dynamic range.

3.4 Portability

In almost all Virtual Environment systems, the user is not encouraged to walk around much. Instead, the user navigates by "flying" through the environment, walking on a treadmill, or driving some mockup of a vehicle. Whatever the technology, the result is that the user stays in one place in the real world.

Some AR applications, however, will need to support a user who will walk around a large environment. AR requires that the user actually be at the place where the task is to take place. "Flying," as performed in a VE system, is no longer an

option. If a mechanic needs to go to the other side of a jet engine, she must physically move herself and the display devices she wears. Therefore, AR systems will place a premium on portability, especially the ability to walk around outdoors, away from controlled environments. The scene generator, the HMD, and the tracking system must all be self-contained and capable of surviving exposure to the environment. If this capability is achieved, many more applications that have not been tried will become available. For example, the ability to annotate the surrounding environment could be useful to soldiers, hikers, or tourists in an unfamiliar new location.

3.5 Comparison against virtual environments

The overall requirements of AR can be summarized by comparing them against the requirements for Virtual Environments, for the three basic subsystems that they require.

1) *Scene generator*: Rendering is not currently one of the major problems in AR. VE systems have much higher requirements for realistic images because they completely replace the real world with the virtual environment. In AR, the virtual images only supplement the real world. Therefore, fewer virtual objects need to be drawn, and they do not necessarily have to be realistically rendered in order to serve the purposes of the application. For example, in the annotation applications, text and 3-D wireframe drawings might suffice. Ideally, photorealistic graphic objects would be seamlessly merged with the real environment (see Section 7), but more basic problems have to be solved first.

2) *Display device*: The display devices used in AR may have less stringent requirements than VE systems demand, again because AR does not replace the real world. For example, monochrome displays may be adequate for some AR applications, while virtually all VE systems today use full color. Optical see-through HMDs with a small field-of-view may be satisfactory because the user can still see the real world with his peripheral vision; the see-through HMD does not shut off the user's normal field-of-view. Furthermore, the resolution of the monitor in an optical see-through HMD might be lower than what a user would tolerate in a VE application, since the optical see-through HMD does not reduce the resolution of the real environment.

3) *Tracking and sensing*: While in the previous two cases AR had lower requirements than VE, that is not the case for tracking and sensing. In this area, the requirements for AR are much stricter than those for VE systems. A major reason for this is the registration problem, which is described in the next section. The other factors that make the tracking and sensing requirements higher are described in Section 5.

4. Registration

4.1 The registration problem

One of the most basic problems currently limiting Augmented Reality applications is the registration problem. The objects in the real and virtual worlds must be properly aligned with respect to each other, or the illusion that the two worlds coexist will be compromised. More seriously, many applications *demand* accurate registration. For example, recall the needle biopsy application. If the virtual object is not where the real tumor is, the surgeon will miss the tumor and the biopsy will fail. Without accurate registration, Augmented Reality will not be accepted in many applications.

Registration problems also exist in Virtual Environments, but they are not nearly as serious because they are harder to detect than in Augmented Reality. Since the user only sees virtual objects in VE applications, registration errors result in visual-kinesthetic and visual-proprioceptive conflicts. Such conflicts between different human senses may be a source of motion sickness [Pausch92]. Because the kinesthetic and proprioceptive systems are much less sensitive than the visual system, visual-kinesthetic and visual-proprioceptive conflicts are less noticeable than visual-visual conflicts. For example, a user wearing a closed-view HMD might hold up her real hand and see a virtual hand. This virtual hand should be displayed exactly where she would see her real hand, if she were not wearing an HMD. But if the virtual hand is wrong by five millimeters, she may not detect that unless actively looking for such errors. The same error is much more obvious in a see-through HMD, where the conflict is visual-visual.

Furthermore, a phenomenon known as *visual capture* [Welch78] makes it even more difficult to detect such registration errors. Visual capture is the tendency of the brain to believe what it sees rather than what it feels, hears, etc. That is, visual information tends to override all other senses. When watching a television program, a viewer believes the sounds come from the mouths of the actors on the screen, even though they actually come from a speaker in the TV. Ventriloquism works because of visual capture. Similarly, a user might believe that her hand is where the virtual hand is drawn, rather than where her real hand actually is, because of visual capture. This effect increases the amount of registration error users can tolerate in Virtual Environment systems. If the errors are systematic, users might even be able to adapt to the new environment, given a long exposure time of several hours or days [Welch78].

Augmented Reality demands much more accurate registration than Virtual Environments [Azuma93]. Imagine the same scenario of a user holding up her hand, but this time wearing a see-through HMD. Registration errors now result in visual-visual conflicts between the images of the virtual and real hands. Such conflicts are easy to detect because of the resolution of the human eye and the sensitivity of the human visual system to differences. Even tiny offsets in the images of the real and virtual hands are easy to detect.

What angular accuracy is needed for good registration in Augmented Reality? A simple demonstration will show the order of magnitude required. Take out a dime and hold it at arm's length, so that it looks like a circle. The diameter of the dime covers about 1.2 to 2.0 degrees of arc, depending on your arm length. In comparison, the width of a full moon is about 0.5 degrees of arc! Now imagine a virtual object superimposed on a real object, but offset by the diameter of the full moon. Such a difference would be easy to detect. Thus, the angular accuracy required is a small fraction of a degree. The lower limit is bounded by the resolving power of the human eye itself. The central part of the retina is called the *fovea*, which has the highest density of color-detecting cones, about 120 per degree of arc, corresponding to a spacing of half a minute of arc [Jain89]. Observers can differentiate between a dark and light bar grating when each bar subtends about one minute of arc, and under special circumstances they can detect even smaller differences [Doenges85]. However, existing HMD trackers and displays are not capable of providing one minute of arc in accuracy, so the present achievable accuracy is much worse than that ultimate lower bound. In practice, errors of a few pixels are detectable in modern HMDs.

Registration of real and virtual objects is not limited to AR. Special-effects artists seamlessly integrate computer-generated 3-D objects with live actors in film and video. The difference lies in the amount of control available. With film, a director can carefully plan each shot, and artists can spend hours per frame, adjusting each by hand if necessary, to achieve perfect registration. As an interactive medium, AR is far more difficult to work with. The AR system cannot control the motions of the HMD wearer. The user looks where she wants, and the system must respond within tens of milliseconds.

Registration errors are difficult to adequately control because of the high accuracy requirements and the numerous sources of error. These sources of error can be divided into two types: static and dynamic. *Static* errors are the ones that cause registration errors even when the user's viewpoint and the objects in the environment remain completely still. *Dynamic* errors are the ones that have no effect until either the viewpoint or the objects begin moving.

For current HMD-based systems, dynamic errors are by far the largest contributors to registration errors, but static errors cannot be ignored either. The next two sections discuss static and dynamic errors and what has been done to reduce them. See [Holloway95] for a thorough analysis of the sources and magnitudes of registration errors.

4.2 Static errors

The four main sources of static errors are:

- Optical distortion
- Errors in the tracking system
- Mechanical misalignments

- Incorrect viewing parameters (e.g., field of view, tracker-to-eye position and orientation, interpupillary distance)

1) *Distortion in the optics:* Optical distortions exist in most camera and lens systems, both in the cameras that record the real environment and in the optics used for the display. Because distortions are usually a function of the radial distance away from the optical axis, wide field-of-view displays can be especially vulnerable to this error. Near the center of the field-of-view, images are relatively undistorted, but far away from the center, image distortion can be large. For example, straight lines may appear curved. In a see-through HMD with narrow field-of-view displays, the optical combiners add virtually no distortion, so the user's view of the real world is not warped. However, the optics used to focus and magnify the graphic images from the display monitors can introduce distortion. This mapping of distorted virtual images on top of an undistorted view of the real world causes static registration errors. The cameras and displays may also have nonlinear distortions that cause errors [Deering92].

Optical distortions are usually systematic errors, so they can be mapped and compensated. This mapping may not be trivial, but it is often possible. For example, [Robinet92b] describes the distortion of one commonly-used set of HMD optics. The distortions might be compensated by additional optics. [Edwards93] describes such a design for a video see-through HMD. This can be a difficult design problem, though, and it will add weight, which is not desirable in HMDs. An alternate approach is to do the compensation digitally. This can be done by image warping techniques, both on the digitized video and the graphic images. Typically, this involves predistorting the images so that they will appear undistorted after being displayed [Watson95]. Another way to perform digital compensation on the graphics is to apply the predistortion functions on the vertices of the polygons, in screen space, before rendering [Rolland93]. This requires subdividing polygons that cover large areas in screen space. Both digital compensation methods can be computationally expensive, often requiring special hardware to accomplish in real time. Holloway determined that the additional system delay required by the distortion compensation adds more registration error than the distortion compensation removes, for typical head motion [Holloway95].

2) *Errors in the tracking system:* Errors in the reported outputs from the tracking and sensing systems are often the most serious type of static registration errors. These distortions are not easy to measure and eliminate, because that requires another "3-D ruler" that is more accurate than the tracker being tested. These errors are often non-systematic and difficult to fully characterize. Almost all commercially-available tracking systems are not accurate enough to satisfy the requirements of AR systems. Section 5 discusses this important topic further.

3) *Mechanical misalignments:* Mechanical misalignments are discrepancies between the model or specification of the hardware and the actual physical properties of the real system. For example, the combiners, optics, and monitors in an optical see-through HMD may not be at the expected distances or orientations with respect to each other. If the frame is not sufficiently rigid, the various component parts may change their relative positions as the user moves around, causing errors. Mechanical misalignments can cause subtle changes in the position and orientation of the

projected virtual images that are difficult to compensate. While some alignment errors can be calibrated, for many others it may be more effective to "build it right" initially.

4) *Incorrect viewing parameters*: Incorrect viewing parameters, the last major source of static registration errors, can be thought of as a special case of alignment errors where calibration techniques can be applied. Viewing parameters specify how to convert the reported head or camera locations into viewing matrices used by the scene generator to draw the graphic images. For an HMD-based system, these parameters include:

- Center of projection and viewport dimensions
- Offset, both in translation and orientation, between the location of the head tracker and the user's eyes
- Field of view

Incorrect viewing parameters cause systematic static errors. Take the example of a head tracker located above a user's eyes. If the vertical translation offsets between the tracker and the eyes are too small, all the virtual objects will appear lower than they should.

In some systems, the viewing parameters are estimated by manual adjustments, in a non-systematic fashion. Such approaches proceed as follows: place a real object in the environment and attempt to register a virtual object with that real object. While wearing the HMD or positioning the cameras, move to one viewpoint or a few selected viewpoints and manually adjust the location of the virtual object and the other viewing parameters until the registration "looks right." This may achieve satisfactory results if the environment and the viewpoint remain static. However, such approaches require a skilled user and generally do not achieve robust results for many viewpoints. Achieving good registration from a single viewpoint is much easier than registration from a wide variety of viewpoints using a single set of parameters. Usually what happens is satisfactory registration at one viewpoint, but when the user walks to a significantly different viewpoint, the registration is inaccurate because of incorrect viewing parameters or tracker distortions. This means many different sets of parameters must be used, which is a less than satisfactory solution.

Another approach is to directly measure the parameters, using various measuring tools and sensors. For example, a commonly-used optometrist's tool can measure the interpupillary distance. Rulers might measure the offsets between the tracker and eye positions. Cameras could be placed where the user's eyes would normally be in an optical see-through HMD. By recording what the camera sees, through the see-through HMD, of the real environment, one might be able to determine several viewing parameters. So far, direct measurement techniques have enjoyed limited success [Janin93].

View-based tasks are another approach to calibration. These ask the user to perform various tasks that set up geometric constraints. By performing several tasks, enough information is gathered to determine the viewing parameters. For example, [Azuma94] asked a user wearing an optical see-through HMD to look straight

through a narrow pipe mounted in the real environment. This sets up the constraint that the user's eye must be located along a line through the center of the pipe. Combining this with other tasks created enough constraints to measure all the viewing parameters. [Caudell92] used a different set of tasks, involving lining up two circles that specified a cone in the real environment. [Oishi96] moves virtual cursors to appear on top of beacons in the real environment. All view-based tasks rely upon the user accurately performing the specified task and assume the tracker is accurate. If the tracking and sensing equipment is not accurate, then multiple measurements must be taken and optimizers used to find the "best-fit" solution [Janin93].

For video-based systems, an extensive body of literature exists in the robotics and photogrammetry communities on camera calibration techniques; see the references in [Lenz88] for a start. Such techniques compute a camera's viewing parameters by taking several pictures of an object of fixed and sometimes unknown geometry. These pictures must be taken from different locations. Matching points in the 2-D images with corresponding 3-D points on the object sets up mathematical constraints. With enough pictures, these constraints determine the viewing parameters and the 3-D location of the calibration object. Alternately, they can serve to drive an optimization routine that will search for the best set of viewing parameters that fits the collected data. Several AR systems have used camera calibration techniques, including [ARGOS94] [Bajura93] [Drascic91] [Tuceryan95] [Whitaker95].

4.3 Dynamic errors

Dynamic errors occur because of system delays, or lags. The *end-to-end system delay* is defined as the time difference between the moment that the tracking system measures the position and orientation of the viewpoint to the moment when the generated images corresponding to that position and orientation appear in the displays. These delays exist because each component in an Augmented Reality system requires some time to do its job. The delays in the tracking subsystem, the communication delays, the time it takes the scene generator to draw the appropriate images in the frame buffers, and the scanout time from the frame buffer to the displays all contribute to end-to-end lag. End-to-end delays of 100 ms are fairly typical on existing systems. Simpler systems can have less delay, but other systems have more. Delays of 250 ms or more can exist on slow, heavily loaded, or networked systems.

End-to-end system delays cause registration errors only when motion occurs. Assume that the viewpoint and all objects remain still. Then the lag does not cause registration errors. No matter how long the delay is, the images generated are appropriate, since nothing has moved since the time the tracker measurement was taken. Compare this to the case with motion. For example, assume a user wears a see-through HMD and moves her head. The tracker measures the head at an initial time t . The images corresponding to time t will not appear until some future time t_2 , because of the end-to-end system delays. During this delay, the user's head remains in motion, so when the images computed at time t finally appear, the user sees them at a different location than the one they were computed for. Thus, the images are

incorrect for the time they are actually viewed. To the user, the virtual objects appear to "swim around" and "lag behind" the real objects. This was graphically demonstrated in a videotape of UNC's ultrasound experiment shown at SIGGRAPH '92 [Bajura92]. In Figure 17, the picture on the left shows what the registration looks like when everything stands still. The virtual gray trapezoidal region represents what the ultrasound wand is scanning. This virtual trapezoid should be attached to the tip of the real ultrasound wand. This is the case in the picture on the left, where the tip of the wand is visible at the bottom of the picture, to the left of the "UNC" letters. But when the head or the wand moves, large dynamic registration errors occur, as shown in the picture on the right. The tip of the wand is now far away from the virtual trapezoid. Also note the motion blur in the background, which is caused by the user's head motion.




Figure 17: Effect of motion and system delays on registration. Picture on the left is a static scene. Picture on the right shows motion. (Courtesy UNC Chapel Hill Dept. of Computer Science)

System delays seriously hurt the illusion that the real and virtual worlds coexist because they cause large registration errors. With a typical end-to-end lag of 100 ms and a moderate head rotation rate of 50 degrees per second, the angular dynamic error is 5 degrees. At a 68 cm arm length, this results in registration errors of almost 60 mm. System delay is the largest single source of registration error in existing AR systems, outweighing all others *combined* [Holloway95].

Methods used to reduce dynamic registration fall under four main categories:

- Reduce system lag
- Reduce apparent lag
- Match temporal streams (with video-based systems)
- Predict future locations

1) Reduce system lag: The most direct approach is simply to reduce, or ideally eliminate, the system delays. If there are no delays, there are no dynamic errors. Unfortunately, modern scene generators are usually built for throughput, not minimal latency [Foley90]. It is sometimes possible to reconfigure the software to sacrifice throughput to minimize latency. For example, the SLATS system completes rendering a pair of interlaced NTSC images in one field time (16.67 ms) on Pixel-Planes 5 [Olano95]. Being careful about synchronizing pipeline tasks can also reduce the end-to-end lag [Wloka95a].

System delays are not likely to completely disappear anytime soon. Some believe that the current course of technological development will automatically solve this problem. Unfortunately, it is difficult to reduce system delays to the point where they are no longer an issue. Recall that registration errors must be kept to a small fraction of a degree. At the moderate head rotation rate of 50 degrees per second, system lag must be 10 ms or less to keep angular errors below 0.5 degrees. Just scanning out a frame buffer to a display at 60 Hz requires 16.67 ms. It might be possible to build an HMD system with less than 10 ms of lag, but the drastic cut in throughput and the expense required to construct the system would make alternate solutions attractive. Minimizing system delay is important, but reducing delay to the point where it is no longer a source of registration error is not currently practical.

2) *Reduce apparent lag*: Image deflection is a clever technique for reducing the amount of apparent system delay for systems that only use head orientation [Burbidge89] [Regan94] [Riner92] [So92]. It is a way to incorporate more recent orientation measurements into the late stages of the rendering pipeline. Therefore, it is a feed-forward technique. The scene generator renders an image much larger than needed to fill the display. Then just before scanout, the system reads the most recent orientation report. The orientation value is used to select the fraction of the frame buffer to send to the display, since small orientation changes are equivalent to shifting the frame buffer output horizontally and vertically.

Image deflection does not work on translation, but image warping techniques might [Chen93] [McMillan95a] [McMillan95b]. After the scene generator renders the image based upon the head tracker reading, small adjustments in orientation and translation could be done after rendering by warping the image. These techniques assume knowledge of the depth at every pixel, and the warp must be done much more quickly than rerendering the entire image.

3) *Match temporal streams*: In video-based AR systems, the video camera and digitization hardware impose inherent delays on the user's view of the real world. This is potentially a blessing when reducing dynamic errors, because it allows the temporal streams of the real and virtual images to be matched. Additional delay is added to the video from the real world to match the scene generator delays in generating the virtual images. This additional delay to the video stream will probably not remain constant, since the scene generator delay will vary with the complexity of the rendered scene. Therefore, the system must dynamically synchronize the two streams.

Note that while this reduces conflicts between the real and virtual, now *both* the real and virtual objects are delayed in time. While this may not be bothersome for small delays, it is a major problem in the related area of telepresence systems and will not be easy to overcome. For long delays, this can produce negative effects such as pilot-induced oscillation.

4) *Predict*: The last method is to predict the future viewpoint and object locations. If the future locations are known, the scene can be rendered with these future locations, rather than the measured locations. Then when the scene finally appears, the viewpoints and objects have moved to the predicted locations, and the graphic images are correct at the time they are viewed. For short system delays

(under ~80 ms), prediction has been shown to reduce dynamic errors by up to an order of magnitude [Azuma94]. Accurate predictions require a system built for real-time measurements and computation. Using inertial sensors makes predictions more accurate by a factor of 2-3. Predictors have been developed for a few AR systems [Emura94] [Zikan94b], but the majority were implemented and evaluated with VE systems (see the reference list in [Azuma94]). More work needs to be done on ways of comparing the theoretical performance of various predictors [Azuma95a] [Azuma95b] and in developing prediction models that better match actual head motion [Wu95].

4.4 Vision-based techniques

Mike Bajura and Ulrich Neumann [Bajura95] point out that registration based solely on the information from the tracking system is like building an "open-loop" controller. The system has no feedback on how closely the real and virtual actually match. Without feedback, it is difficult to build a system that achieves perfect matches. However, video-based approaches can use image processing or computer vision techniques to aid registration. Since video-based AR systems have a digitized image of the real environment, it may be possible to detect features in the environment and use those to enforce registration. They call this a "closed-loop" approach, since the digitized image provides a mechanism for bringing feedback into the system.

This is not a trivial task. This detection and matching must run in real time and must be robust. This often requires special hardware and sensors. However, it is also not an "AI-complete" problem because this is simpler than the general computer vision problem.

For example, in some AR applications it is acceptable to place fiducials in the environment. These fiducials may be LEDs [Bajura95] or special markers [Mellor95a] [Mellor95b] [Neumann96]. Recent ultrasound experiments at UNC Chapel Hill have used colored dots as fiducials [State96a]. The locations or patterns of the fiducials are assumed to be known. Image processing detects the locations of the fiducials, then those are used to make corrections that enforce proper registration.

These routines assume that one or more fiducials are visible at all times; without them, the registration can fall apart. But when the fiducials are visible, the results can be accurate to one pixel, which is as about close as one can get with video techniques. Figure 18, taken from [Bajura95], shows a virtual arrow and a virtual chimney exactly aligned with their desired points on two real objects. The real objects each have an LED to aid the registration. Figures 19 through 21 show registration from [Mellor95a], which uses dots with a circular pattern as the fiducials. The registration is also nearly perfect. Figure 22 demonstrates merging virtual objects with the real environment, using colored dots as the fiducials in a video-based approach. In the picture on the left, the stack of cards in the center are real, but the ones on the right are virtual. Notice that they penetrate one of the blocks. In the image on the right, a virtual spiral object interpenetrates the real blocks and table and also casts virtual shadows upon the real objects [State96a].




Figure 18: A virtual arrow and virtual chimney aligned with two real objects. (Courtesy Mike Bajura, UNC Chapel Hill Dept. of Computer Science, and Ulrich Neumann, USC)



Figure 19: Real skull with five fiducials. (Courtesy J.P. Mellor, MIT AI Lab)



Figure 20: Virtual wireframe skull registered with real skull. (Courtesy J.P. Mellor, MIT AI Lab)



Figure 21: Virtual wireframe skull registered with real skull moved to a different position. (Courtesy J.P. Mellor, MIT AI Lab)




Figure 22: Virtual cards and spiral object merged with real blocks and table. (Courtesy Andrei State, UNC Chapel Hill Dept. of Computer Science.)

Instead of fiducials, [Uenohara95] uses *template matching* to achieve registration. Template images of the real object are taken from a variety of viewpoints. These are used to search the digitized image for the real object. Once that is found, a virtual wireframe can be superimposed on the real object.

Recent approaches in video-based matching avoid the need for any calibration. [Kutukalos96] represents virtual objects in a non-Euclidean, affine frame of reference that allows rendering without knowledge of camera parameters. [Iu96] extracts contours from the video of the real world, then uses an optimization technique to match the contours of the rendered 3-D virtual object with the contour extracted from the video. Note that calibration-free approaches may not recover all the information required to perform all potential AR tasks. For example, these two approaches do not recover true depth information, which is useful when compositing the real and the virtual.

Techniques that use fiducials as the sole tracking source determine the *relative* projective relationship between the objects in the environment and the video camera. While this is enough to ensure registration, it does not provide all the information one might need in some AR applications, such as the absolute (rather than relative) locations of the objects and the camera. Absolute locations are needed to include virtual and real objects that are not tracked by the video camera, such as a 3-D pointer or other virtual objects not directly tied to real objects in the scene.

Additional sensors besides video cameras can aid registration. Both [Mellor95a] [Mellor95b] and [Grimson94] [Grimson95] use a laser rangefinder to acquire an initial depth map of the real object in the environment. Given a matching virtual model, the system can match the depth maps from the real and virtual until they are properly aligned, and that provides the information needed for registration.

Another way to reduce the difficulty of the problem is to accept the fact that the system may not be robust and may not be able to perform all tasks automatically. Then it can ask the user to perform certain tasks. The system in [Sharma94] expects manual intervention when the vision algorithms fail to identify a part because the view is obscured. The calibration techniques in [Tuceryan95] are heavily based on computer vision techniques, but they ask the user to manually intervene by specifying correspondences when necessary.

4.5 Current status

The registration requirements for AR are difficult to satisfy, but a few systems have achieved good results. [Azuma94] is an open-loop system that shows registration typically within ± 5 millimeters from many viewpoints for an object at about arm's length. Closed-loop systems, however, have demonstrated nearly perfect registration, accurate to within a pixel [Bajura95] [Mellor95a] [Mellor95b] [Neumann96] [State96a].

The registration problem is far from solved. Many systems assume a static viewpoint, static objects, or even both. Even if the viewpoint or objects are allowed to move, they are often restricted in how far they can travel. Registration is shown under controlled circumstances, often with only a small number of real-world objects, or where the objects are already well-known to the system. For example, registration may only work on one object marked with fiducials, and not on any other objects in the scene. Much more work needs to be done to increase the domains in which registration is robust. Duplicating registration methods remains a nontrivial task, due to both the complexity of the methods and the additional hardware required. If simple yet effective solutions could be developed, that would speed the acceptance of AR systems.

5. Sensing

Accurate registration and positioning of virtual objects in the real environment requires accurate tracking of the user's head and sensing the locations of other objects in the environment. The biggest single obstacle to building effective Augmented Reality systems is the requirement of accurate, long-range sensors and trackers that report the locations of the user and the surrounding objects in the environment. For details of tracking technologies, see the surveys in [Ferrin91] [Meyer92] and Chapter 5 of [Durlach95]. Commercial trackers are aimed at the needs of Virtual Environments and motion capture applications. Compared to those two applications, Augmented Reality has much stricter accuracy requirements and demands larger working volumes. No tracker currently provides high accuracy at long ranges in real time. More work needs to be done to develop sensors and trackers that can meet these stringent requirements.

Specifically, AR demands more from trackers and sensors in three areas:

- Greater input variety and bandwidth
- Higher accuracy
- Longer range

5.1 Input variety and bandwidth

VE systems are primarily built to handle output bandwidth: the images displayed, sounds generated, etc. The input bandwidth is tiny: the locations of the user's head and hands, the outputs from the buttons and other control devices, etc. AR systems, however, will need a greater variety of input sensors and much more input bandwidth [Buxton93]. There are a greater variety of possible input sensors than output displays. Outputs are limited to the five human senses. Inputs can come from anything a sensor can detect. Robinett speculates that Augmented Reality may be useful in any application that requires displaying information not directly available or detectable by human senses by making that information visible (or audible, touchable, etc.) [Robinett92a]. Recall that the proposed medical applications in Section 2.1 use CT, MRI and ultrasound sensors as inputs. Other future applications might use sensors to extend the user's visual range into infrared or ultraviolet frequencies, and remote sensors would let users view objects hidden by walls or hills. Conceptually, anything not detectable by human senses but detectable by machines might be transduced into something that a user can sense in an AR system.

Range data is a particular input that is vital for many AR applications [Aliaga97] [Breen96]. The AR system knows the distance to the virtual objects, because that model is built into the system. But the AR system may not know where all the real objects are in the environment. The system might assume that the entire environment is measured at the beginning and remains static thereafter. However, some useful applications will require a dynamic environment, in which real objects move, so the objects must be tracked in real time. However, for some applications a depth map of the real environment would be sufficient. That would allow real objects to occlude virtual objects through a pixel-by-pixel depth value comparison. Acquiring this depth map in real time is not trivial. Sensors like laser rangefinders might be used. Many computer vision techniques for recovering shape through various strategies (e.g., "shape from stereo," or "shape from shading") have been tried. A recent work [Wloka95b] uses intensity-based matching from a pair of stereo images to do depth recovery. Recovering depth through existing vision techniques is difficult to do robustly in real time.

Finally, some annotation applications require access to a detailed database of the environment, which is a type of input to the system. For example, the architectural application of "seeing into the walls" assumes that the system has a database of where all the pipes, wires and other hidden objects are within the building. Such a database may not be readily available, and even if it is, it may not be in a format that is easily usable. For example, the data may not be grouped to segregate the parts of the model that represent wires from the parts that represent pipes. Thus, a significant modelling effort may be required and should be taken into consideration when building an AR application.

5.2 High accuracy

The accuracy requirements for the trackers and sensors are driven by the accuracies needed for visual registration, as described in Section 4. For many

approaches, the registration is only as accurate as the tracker. Therefore, the AR system needs trackers that are accurate to around a millimeter and a tiny fraction of a degree, across the entire working range of the tracker.

Few trackers can meet this specification, and every technology has weaknesses. Some mechanical trackers are accurate enough, although they tether the user to a limited working volume. Magnetic trackers are vulnerable to distortion by metal in the environment, which exists in many desired AR application environments. Ultrasonic trackers suffer from noise and are difficult to make accurate at long ranges because of variations in the ambient temperature. Optical technologies [Janin94] have distortion and calibration problems. Inertial trackers drift with time. Of the individual technologies, optical technologies show the most promise due to trends toward high-resolution digital cameras, real-time photogrammetric techniques, and structured light sources that result in more signal strength at long distances. Future tracking systems that can meet the stringent requirements of AR will probably be hybrid systems [Azuma93] [Durlach95] [Foxlin96] [Zikan94b], such as a combination of inertial and optical technologies. Using multiple technologies opens the possibility of covering for each technology's weaknesses by combining their strengths.

Attempts have been made to calibrate the distortions in commonly-used magnetic tracking systems [Bryson92] [Ghazisaedy95]. These have succeeded at removing much of the gross error from the tracker at long ranges, but not to the level required by AR systems [Holloway95]. For example, mean errors at long ranges can be reduced from several inches to around one inch.

The requirements for registering other sensor modes are not nearly as stringent. For example, the human auditory system is not very good at localizing deep bass sounds, which is why subwoofer placement is not critical in a home theater system.

5.3 Long range

Few trackers are built for accuracy at long ranges, since most VE applications do not require long ranges. *Motion capture* applications track an actor's body parts to control a computer-animated character or for the analysis of an actor's movements. This is fine for position recovery, but not for orientation. Orientation recovery is based upon the computed positions. Even tiny errors in those positions can cause orientation errors of a few degrees, which is too large for AR systems.

Two scalable tracking systems for HMDs have been described in the literature [Ward92] [Sowizral93]. A scalable system is one that can be expanded to cover any desired range, simply by adding more modular components to the system. This is done by building a cellular tracking system, where only nearby sources and sensors are used to track a user. As the user walks around, the set of sources and sensors changes, thus achieving large working volumes while avoiding long distances between the current working set of sources and sensors. While scalable trackers can

be effective, they are complex and by their very nature have many components, making them relatively expensive to construct.

The Global Positioning System (GPS) is used to track the locations of vehicles almost anywhere on the planet. It might be useful as one part of a long range tracker for AR systems. However, by itself it will not be sufficient. The best reported accuracy is approximately one centimeter, assuming that many measurements are integrated (so that accuracy is not generated in real time), when GPS is run in differential mode. That is not sufficiently accurate to recover orientation from a set of positions on a user.

Tracking an AR system outdoors in real time with the required accuracy has not been demonstrated and remains an open problem.

6. Future directions

This section identifies areas and approaches that require further research to produce improved AR systems.

Hybrid approaches: Future tracking systems may be hybrids, because combining approaches can cover weaknesses. The same may be true for other problems in AR. For example, current registration strategies generally focus on a single strategy. Future systems may be more robust if several techniques are combined. An example is combining vision-based techniques with prediction. If the fiducials are not available, the system switches to open-loop prediction to reduce the registration errors, rather than breaking down completely. The predicted viewpoints in turn produce a more accurate initial location estimate for the vision-based techniques.

Real-time systems and time-critical computing: Many VE systems are not truly run in real time. Instead, it is common to build the system, often on UNIX, and then see how fast it runs. This may be sufficient for some VE applications. Since everything is virtual, all the objects are automatically synchronized with each other. AR is a different story. Now the virtual and real must be synchronized, and the real world "runs" in real time. Therefore, effective AR systems must be built with real-time performance in mind. Accurate timestamps must be available. Operating systems must not arbitrarily swap out the AR software process at any time, for arbitrary durations. Systems must be built to guarantee completion within specified time budgets, rather than just "running as quickly as possible." These are characteristics of flight simulators and a few VE systems [Krueger92]. Constructing and debugging real-time systems is often painful and difficult, but the requirements for AR demand real-time performance.

Perceptual and psychophysical studies: Augmented Reality is an area ripe for psychophysical studies. How much lag can a user detect? How much registration error is detectable when the head is moving? Besides questions on perception, psychological experiments that explore *performance* issues are also needed. How

much does head-motion prediction improve user performance on a specific task? How much registration error is tolerable for a specific application before performance on that task degrades substantially? Is the allowable error larger while the user moves her head versus when she stands still? Furthermore, not much is known about potential optical illusions caused by errors or conflicts in the simultaneous display of real and virtual objects [Durlach95].

Few experiments in this area have been performed. Jannick Rolland, Frank Biocca and their students conducted a study of the effect caused by eye displacements in video see-through HMDs [Rolland95]. They found that users partially adapted to the eye displacement, but they also had negative aftereffects after removing the HMD. Steve Ellis' group at NASA Ames has conducted work on perceived depth in a see-through HMD [Ellis94] [Ellis95]. ATR has also conducted a study [Utsumi94].

Portability: Section 3.4 explained why some potential AR applications require giving the user the ability to walk around large environments, even outdoors. This requires making the equipment self-contained and portable. Existing tracking technology is not capable of tracking a user outdoors at the required accuracy.

Multimodal displays: Almost all work in AR has focused on the visual sense: virtual graphic objects and overlays. But Section 3.1 explained that augmentation might apply to all other senses as well. In particular, adding and removing 3-D sound is a capability that could be useful in some AR applications.

Social and political issues: Technological issues are not the only ones that need to be considered when building a real application. There are also social and political dimensions when getting new technologies into the hands of real users. Sometimes, perception is what counts, even if the technological reality is different. For example, if workers perceive lasers to be a health risk, they may refuse to use a system with lasers in the display or in the trackers, even if those lasers are eye safe. Ergonomics and ease of use are paramount considerations. Whether AR is truly a cost-effective solution in its proposed applications has yet to be determined. Another important factor is whether or not the technology is perceived as a threat to jobs, as a replacement for workers, especially with many corporations undergoing recent layoffs. AR may do well in this regard, because it is intended as a tool to make the user's job easier, rather than something that completely replaces the human worker. Although technology transfer is not normally a subject of academic papers, it is a real problem. Social and political concerns should not be ignored during attempts to move AR out of the research lab and into the hands of real users.

7. Conclusion

Augmented Reality is far behind Virtual Environments in maturity. Several commercial vendors sell complete, turnkey Virtual Environment systems. However, no commercial vendor currently sells an HMD-based Augmented Reality system. A few monitor-based "virtual set" systems are available, but today AR systems are primarily found in academic and industrial research laboratories.

The first deployed HMD-based AR systems will probably be in the application of aircraft manufacturing. Both Boeing [ARPA95] [BoeingTRP94] and McDonnell Douglas [Neumann96] are exploring this technology. The former uses optical approaches, while the latter is pursuing video approaches. Boeing has performed trial runs with workers using a prototype system but has not yet made any deployment decisions. Annotation and visualization applications in restricted, limited-range environments are deployable today, although much more work needs to be done to make them cost effective and flexible. Applications in medical visualization will take longer. Prototype visualization aids have been used on an experimental basis, but the stringent registration requirements and ramifications of mistakes will postpone common usage for many years. AR will probably be used for medical training before it is commonly used in surgery.

The next generation of combat aircraft will have Helmet-Mounted Sights with graphics registered to targets in the environment [Wanstall89]. These displays, combined with short-range steerable missiles that can shoot at targets off-boresight, give a tremendous combat advantage to pilots in dogfights. Instead of having to be directly behind his target in order to shoot at it, a pilot can now shoot at anything within a 60-90 degree cone of his aircraft's forward centerline. Russia and Israel currently have systems with this capability, and the U.S. is expected to field the AIM-9X missile with its associated Helmet-Mounted Sight in 2002 [Dornheim95a] [Dornheim95b]. Registration errors due to delays are a major problem in this application [Dornheim95c].

Augmented Reality is a relatively new field, where most of the research efforts have occurred in the past four years, as shown by the references listed at the end of this paper. The SIGGRAPH "Rediscovering Our Fire" report identified Augmented Reality as one of four areas where SIGGRAPH should encourage more submissions [Mair94]. Because of the numerous challenges and unexplored avenues in this area, AR will remain a vibrant area of research for at least the next several years.

One area where a breakthrough is required is tracking an HMD outdoors at the accuracy required by AR. If this is accomplished, several interesting applications will become possible. Two examples are described here: navigation maps and visualization of past and future environments.

The first application is a navigation aid to people walking outdoors. These individuals could be soldiers advancing upon their objective, hikers lost in the woods, or tourists seeking directions to their intended destination. Today, these individuals must pull out a physical map and associate what they see in the real environment around them with the markings on the 2-D map. If landmarks are not easily identifiable, this association can be difficult to perform, as anyone lost in the woods can attest. An AR system makes navigation easier by performing the association step automatically. If the user's position and orientation are known, and the AR system has access to a digital map of the area, then the AR system can draw the map in 3-D directly upon the user's view. The user looks at a nearby mountain and sees graphics directly overlaid on the real environment explaining the mountain's name, how tall it is, how far away it is, and where the trail is that leads to the top.

The second application is visualization of locations and events as they were in the past or as they will be after future changes are performed. Tourists that visit historical sites, such as a Civil War battlefield or the Acropolis in Athens, Greece, do not see these locations as they were in the past, due to changes over time. It is often difficult for a modern visitor to imagine what these sites really looked like in the past. To help, some historical sites stage "Living History" events where volunteers wear ancient clothes and reenact historical events. A tourist equipped with an outdoors AR system could see a computer-generated version of Living History. The HMD could cover up modern buildings and monuments in the background and show, directly on the grounds at Gettysburg, where the Union and Confederate troops were at the fateful moment of Pickett's charge. The gutted interior of the modern Parthenon would be filled in by computer-generated representations of what it looked like in 430 BC, including the long-vanished gold statue of Athena in the middle. Tourists and students walking around the grounds with such AR displays would gain a much better understanding of these historical sites and the important events that took place there. Similarly, AR displays could show what proposed architectural changes would look like before they are carried out. An urban designer could show clients and politicians what a new stadium would look like as they walked around the adjoining neighborhood, to better understand how the stadium project will affect nearby residents.

After the basic problems with AR are solved, the ultimate goal will be to generate virtual objects that are so realistic that they are virtually indistinguishable from the real environment [Fournier94]. Photorealism has been demonstrated in feature films, but accomplishing this in an interactive application will be much harder. Lighting conditions, surface reflections, and other properties must be measured automatically, in real time. More sophisticated lighting, texturing, and shading capabilities must run at interactive rates in future scene generators. Registration must be nearly perfect, without manual intervention or adjustments. While these are difficult problems, they are probably not insurmountable. It took about 25 years to progress from drawing stick figures on a screen to the photorealistic dinosaurs in "Jurassic Park." Within another 25 years, we should be able to wear a pair of AR glasses outdoors to see and interact with photorealistic dinosaurs eating a tree in our backyard.

Acknowledgements

This paper is an updated version of my course notes for an ACM SIGGRAPH '95 class called "Developing Advanced Virtual Reality Applications," given in Los Angeles, CA, in August 1995.

I thank the anonymous reviewers for their constructive criticism and suggestions for improving this paper. I also thank the following individuals and organizations for sending pictures to include with this paper:

- Mike Bajura, Andrei State, and Linda Houseman, University of North Carolina at Chapel Hill Department of Computer Science

- David Drascic and Paul Milgram, University of Toronto
- Steve Feiner and Blair MacIntyre, Columbia University
- Alessandro Giacalone, The European Computer-Industry Research Centre (ECRC) and David Breen, Caltech
- J.P. Mellor, MIT AI Laboratory
- David Mizell, Boeing
- Ulrich Neumann, University of Southern California
- Jannick Rolland, Center for Research and Engineering in Optics and Lasers (CREOL) at the University of Central Florida (rolland@creol.ucf.edu)

References and Bibliography

Note: some of these references are available electronically at the following sites on the World Wide Web:

- | | |
|-----------------|--|
| Columbia U. | http://www.cs.columbia.edu/graphics/ |
| ECRC | http://www.ecrc.de/ |
| MIT AI Lab | http://www.ai.mit.edu/ |
| UNC Chapel Hill | http://www.cs.unc.edu/ |
| U. Toronto | http://vered.rose.utoronto.ca/etc-lab.html/ |
| Ahlers95 | Ahlers, Klaus H., André Kramer, David E. Breen, Pierre-Yves Chevalier, Chris Crampton, Eric Rose, Mihran Tuceryan, Ross T. Whitaker and Douglas Greer. Distributed Augmented Reality for Collaborative Design Applications. <i>Proceedings of Eurographics '95</i> (Maastricht, The Netherlands, August 1995), 3-14. |
| Aliaga97 | Aliaga, Daniel G. Virtual Objects in the Real World. <i>Communications of the ACM</i> 40, 3 (March 1997), 49-54. |
| ARGOS94 | ARGOS Virtual Pointer Camera Calibration Procedure. WWW page = http://vered.rose.utoronto.ca/people/david_dir/POINTER/Calibration.html |
| ARPA95 | WWW page = http://molothrus.sysplan.com/ESTO/ |
| Azuma93 | Azuma, Ronald. Tracking Requirements for Augmented Reality. <i>Communications of the ACM</i> 36, 7 (July 1993), 50-51. |
| Azuma94 | Azuma, Ronald, and Gary Bishop. Improving Static and Dynamic Registration in a See-Through HMD. <i>Proceedings of SIGGRAPH</i> |

- '94 (Orlando, FL, 24-29 July 1994). In *Computer Graphics, Annual Conference Series*, 1994, 197-204.
- Azuma95a Azuma, Ronald T. Predictive Tracking for Augmented Reality. Ph.D. dissertation. UNC Chapel Hill Department of Computer Science technical report TR95-007 (February 1995).
- Azuma95b Azuma, Ronald, and Gary Bishop. A Frequency-Domain Analysis of Head-Motion Prediction. *Proceedings of SIGGRAPH '95* (Los Angeles, CA, 6-11 August 1995). In *Computer Graphics, Annual Conference Series*, 1995, 401-408.
- Barfield95 Barfield, Woodrow, Craig Rosenberg, and Wouter A. Lotens. Augmented-Reality Displays. In Barfield, Woodrow and Thomas A. Furness III (editors). *Virtual Environments and Advanced Interface Design*. Oxford University Press (1995), 542-575. ISBN 0-19-507555-2.
- Bajura92 Bajura, Mike, Henry Fuchs, and Ryutarou Ohbuchi. Merging Virtual Reality with the Real World: Seeing Ultrasound Imagery Within the Patient. *Proceedings of SIGGRAPH '92* (Chicago, IL, 26-31 July 1992). In *Computer Graphics* 26, 2 (July 1992), 203-210.
- Bajura93 Bajura, Mike. Camera Calibration for Video See-Through Head-Mounted Display. UNC Chapel Hill Department of Computer Science technical report TR93-048 (July 7, 1993), 6 pages.
- Bajura95 Bajura, Michael and Ulrich Neumann. Dynamic Registration Correction in Video-Based Augmented Reality Systems. *IEEE Computer Graphics and Applications* 15, 5 (September 1995), 52-60.
- Betting95 Betting, Fabienne, Jacques Feldmar, Nicholas Ayache, and Frédéric Devernay. A New Framework for Fusing Stereo Images with Volumetric Medical Images. *Proceedings of Computer Vision, Virtual Reality, and Robotics in Medicine '95 (CVRMed '95)* (Nice, France, 3-6 April 1995), 30-39.
- BoeingTRP94 WWW page = <http://esto.sysplan.com/ESTO/Displays/HMD-TDS/Factsheets/Boeing.html> (July 1994).
- Bowskill95 Bowskill, Jerry and John Downie. Extending the Capabilities of the Human Visual System: An Introduction to Enhanced Reality. *Computer Graphics* 29, 2 (May 1995), 61-65.
- Breen96 Breen, David E., Ross T. Whitaker, Eric Rose and Mihran Tuceryan. Interactive Occlusion and Automatic Object Placement for Augmented Reality. *Proceedings of Eurographics '96* (Futuroscope - Poitiers, France, 26-30 August 1996), 11-22.

- Brooks96 Brooks, Frederick P. Jr. The Computer Scientist as Toolsmith II. *CACM* 39, 3 (March 1996), 61-68.
- Bryson92 Bryson, Steve. Measurement and Calibration of Static Distortion of Position Data from 3D Trackers. *Proceedings of SPIE Vol. 1669: Stereoscopic Displays and Applications III* (San Jose, CA, 12-13 February 1992), 244-255.
- Burbidge89 Burbidge, Dick, and Paul M. Murray. Hardware Improvements to the Helmet-Mounted Projector on the Visual Display Research Tool (VDRT) at the Naval Training Systems Center. *SPIE Proceedings Vol. 1116 Head-Mounted Displays* (1989), 52-59.
- Buxton93 Buxton, Bill. Personal communication. MIT Workshop on Ubiquitous Computing and Augmented Reality (Cambridge, MA, 24-25 February 1993).
- Caudell92 Caudell, Thomas P. and David W. Mizell. Augmented Reality: An Application of Heads-Up Display Technology to Manual Manufacturing Processes. *Proceedings of Hawaii International Conference on System Sciences* (January 1992), 659-669.
- Caudell94 Caudell, Thomas P. Introduction to Augmented Reality. *SPIE Proceedings volume 2351: Telemanipulator and Telepresence Technologies* (Boston, MA, 31 October - 4 November 1994), 272-281.
- Chen93 Chen, Shenchang Eric, and Lance Williams. View Interpolation for Image Synthesis. *Proceedings of SIGGRAPH '93* (Anaheim, CA, 1-6 August 1993). In *Computer Graphics, Annual Conference Series, 1993*, 279-288.
- Deering92 Deering, Michael. High Resolution Virtual Reality. *Proceedings of SIGGRAPH '92* (Chicago, IL, 26-31 July 1992). In *Computer Graphics* 26, 2 (July 1992), 195-202.
- Doenges85 Doenges, Peter K. Overview of Computer Image Generation in Visual Simulation. *SIGGRAPH '85 Course Notes #14 on High Performance Image Generation Systems* (San Francisco, CA, 22 July 1985).
- Dornheim95a Dornheim, Michael A. and David Hughes. U.S. Intensifies Efforts to Meet Missile Threat. *Aviation Week and Space Technology* 143, 16 (16 October 1995), 36-39.
- Dornheim95b Dornheim, Michael A. U.S. Fighters to Get Helmet Displays After 2000. *Aviation Week and Space Technology* 143, 17 (23 October 1995), 46-48.

- Dornheim95c Dornheim, Michael A. Helmet-Mounted Sights Must Overcome Delays. *Aviation Week and Space Technology* 143, 17 (23 October 1995), 54.
- Drascic91 Drascic, David and Paul Milgram. Positioning Accuracy of a Virtual Stereographic Pointer in a Real Stereoscopic Video World. *SPIE Proceedings Volume 1457 - Stereoscopic Displays and Applications II* (San Jose, CA, February 1991), 302-313.
- Drascic93a Drascic, D., J.J. Grodski, P. Milgram, K. Ruffo, P. Wong, and S. Zhai. ARGOS: A Display System for Augmenting Reality. *Video Proceedings of INTERCHI '93: Human Factors in Computing Systems* (Amsterdam, the Netherlands, 24-29 April 1993). Also in ACM SIGGRAPH Technical Video Review, Volume 88. Extended abstract in *Proceedings of INTERCHI '93*, 521.
- Drascic93b Drascic, David. Stereoscopic Vision and Augmented Reality. *Scientific Computing & Automation* 9, 7 (June 1993), 31-34.
- Durlach95 Durlach, Nathaniel I. and Anne S. Mavor (editors). Virtual Reality: Scientific and Technological Challenges. (Report of the Committee on Virtual Reality Research and Development to the National Research Council) National Academy Press (1995). ISBN 0-309-05135-5.
- Edwards93 Edwards, Emily, Jannick Rolland, and Kurtis Keller. Video See-through Design for Merging of Real and Virtual Environments. *Proceedings of IEEE VRAIS '93* (Seattle, WA, 18-22 September 1993), 222-233.
- Edwards95 Edwards, P.J., D.L.G. Hill, D.J. Hawkes, R. Spink, A.C.F. Colchester, A. Strong, and M. Gleeson. Neurosurgical Guidance Using the Stereo Microscope. *Proceedings of Computer Vision, Virtual Reality, and Robotics in Medicine '95 (CVRMed '95)* (Nice, France, 3-6 April 1995), 555-564.
- Ellis94 Ellis, Stephen R. and Urs J. Bucher. Distance Perception of Stereoscopically Presented Virtual Objects Optically Superimposed on Physical Objects by a Head-Mounted See-Through Display. *Proceedings of 38th Annual Meeting of the Human Factors and Ergonomics Society* (Nashville, TN, 24-28 October 1994), 1300-1305.
- Ellis95 Ellis, Stephen R. and Brian M. Menges. Judged Distance to Virtual Objects in the Near Visual Field. *Proceedings of 39th Annual Meeting of the Human Factors and Ergonomics Society* (San Diego, CA, 1995), 1400-1404.
- Emura94 Emura, Satoru and Susumu Tachi. Compensation of Time Lag Between Actual and Virtual Spaces by Multi-Sensor Integration.

Proceedings of the 1994 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (Las Vegas, NV, 2-5 October 1994), 463-469.

- Feiner93a Feiner, Steven, Blair MacIntyre, and Dorée Seligmann. Knowledge-based Augmented Reality. *Communications of the ACM* 36, 7 (July 1993), 52-62.
- Feiner93b Feiner, Steven, Blair MacIntyre, Marcus Haupt, and Eliot Solomon. Windows on the World: 2D Windows for 3D Augmented Reality. *Proceedings of UIST '93* (Atlanta, GA, 3-5 November 1993), 145-155.
- Feiner94a Feiner, Steven. Augmented Reality. Course Notes, 2: ACM SIGGRAPH 1994, 7:1-7:11.
- Feiner94b Feiner, Steven. Redefining the User Interface: Augmented Reality. Course Notes, 2: ACM SIGGRAPH 1994, 18:1-18:7.
- Feiner95 Feiner, Steven K., Anthony C. Webster, Theodore E. Krueger III, Blair MacIntyre, and Edward J. Keller. Architectural Anatomy. *Presence: Teleoperators and Virtual Environments* 4, 3 (Summer 1995), 318-325.
- Ferrin91 Ferrin, Frank J. Survey of Helmet Tracking Technologies. *SPIE Proceedings Vol. 1456 Large-Screen Projection, Avionic, and Helmet-Mounted Displays* (1991), 86-94.
- Fitzmaurice93 Fitzmaurice, George. Situated Information Spaces: Spatially Aware Palmtop Computers. *CACM* 36, 7 (July 1993), 38-49.
- Foley90 Foley, James D., Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics: Principles and Practice* (2nd edition). Addison-Wesley (1990).
- Fournier94 Fournier, Alain. Illumination Problems in Computer Augmented Reality. *Journée INRIA, Analyse / Synthèse D'Images* (January 1994), 1-21.
- Foxlin96 Foxlin, Eric. Inertial Head-Tracker Sensor Fusion by a Complementary Separate-Bias Kalman Filter. *Proceedings of VRAIS '96* (Santa Clara, CA, 30 March - 3 April 1996), 185-194.
- Ghazisaedy95 Ghazisaedy, Morteza, David Adamczyk, Daniel J. Sandin, Robert V. Kenyon, and Thomas A. DeFanti. Ultrasonic Calibration of a Magnetic Tracker in a Virtual Reality Space. *Proceedings of VRAIS '95* (Research Triangle Park, NC, 11-15 March 1995), 179-188.

- Grimson94 Grimson, W., T. Lozano-Pérez, W. Wells, G. Ettinger, S. White and R. Kikinis. An Automatic Registration Method for Frameless Stereotaxy, Image Guided Surgery, and Enhanced Reality Visualization. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (Los Alamitos, CA, June 1994), 430-436.
- Grimson95 Grimson, W.E.L., G.J. Ettinger, S.J. White, P.L. Gleason, T. Lozano-Pérez, W.M. Wells III, and R. Kikinis. Evaluating and Validating an Automated Registration System for Enhanced Reality Visualization in Surgery. *Proceedings of Computer Vision, Virtual Reality, and Robotics in Medicine '95 (CVRMed '95)* (Nice, France, 3-6 April 1995), 3-12.
- Holloway95 Holloway, Richard. Registration Errors in Augmented Reality. Ph.D. dissertation. UNC Chapel Hill Department of Computer Science technical report TR95-016 (August 1995).
- Holmgren92 Holmgren, Douglas E. Design and Construction of a 30-Degree See-Through Head-Mounted Display. UNC Chapel Hill Department of Computer Science technical report TR 92-030 (July 1992), 4 pages.
- Iu96 Iu, Siu-Leong and Kevin W. Rogovin. Registering Perspective Contours with 3-D Objects Without Correspondence Using Orthogonal Polynomials. *Proceedings of VRAIS '96* (Santa Clara, CA, 30 March - 3 April 1996), 37-44.
- Jain89 Jain, Anil K. Fundamentals of Digital Image Processing. Prentice Hall (1989). ISBN 0-13-336165-9.
- Janin93 Janin, Adam L., David W. Mizell, and Thomas P. Caudell. Calibration of Head-Mounted Displays for Augmented Reality Applications. *Proceedings of IEEE VRAIS '93* (Seattle, WA, 18-22 September 1993), 246-255.
- Janin94 Janin, Adam, Karel Zikan, David Mizell, Mike Banner, and Henry Sowizral. A Videometric Head Tracker for Augmented Reality. *SPIE Proceedings volume 2351: Telemicroscopy and Telepresence Technologies* (Boston, MA, 31 October - 4 November 1994), 308-315.
- Kancherla95 Kancherla, Anantha R., Jannick P. Rolland, Donna L. Wright, and Grigore Burdea. A Novel Virtual Reality Tool for Teaching Dynamic 3D Anatomy. *Proceedings of Computer Vision, Virtual Reality, and Robotics in Medicine '95 (CVRMed '95)* (Nice, France, 3-6 April 1995), 163-169.
- Kim93 Kim, Won S. Advanced Teleoperation, Graphics Aids, and Application to Time Delay Environments. *Proceedings of the 1st*

Industrial Virtual Reality Show and Conference (IVR '93)
(Makuhari Meese, Japan, 23-25 June 1993), 202-207.

- Kim96 Kim, Won S. Virtual Reality Calibration and Preview / Predictive Displays for Telerobotics. *Presence: Teleoperators and Virtual Environments* 5, 2 (Spring 1996), 173-190.
- Krueger92 Krueger, Myron W. Simulation Versus Artificial Reality. *Proceedings of IMAGE VI Conference* (Scottsdale, AZ, 14-17 July 1992), 147-155.
- Kutulakos96 Kutulakos, Kiriakos N. and James Vallino. Affine Object Representations for Calibration-Free Augmented Reality. *Proceedings of VRAIS '96* (Santa Clara, CA, 30 March - 3 April 1996), 25-36.
- Lenz88 Lenz, Reimar K. and Roger Y. Tsai. Techniques for Calibration of the Scale Factor and Image Center for High Accuracy 3-D Machine Vision Metrology. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10, 5 (September 1988), 713-720.
- Lion93 Lion, Dav, Craig Rosenberg, and Woodrow Barfield. Overlaying Three-Dimensional Computer Graphics with Stereoscopic Live Motion Video: Applications for Virtual Environments. *Society for Information Display International Symposium Digest of Technical Papers* (Seattle, WA, 18-29 May 1993), 483-486.
- Lorensen93 Lorensen, William, Harvey Cline, Christopher Nafis, Ron Kikinis, David Altobelli, and Langham Gleason. Enhancing Reality in the Operating Room. *Proceedings of Visualization '93* (Los Alamitos, CA, October 1993), 410-415.
- MacKay93 MacKay, Wendy, Gilles Velay, Kathy Carter, Chaoying Ma, and Daniele Pagani. Augmenting Reality: Adding Computational Dimensions to Paper. *CACM* 36, 7 (July 1993), 96-97.
- Maes95 Maes, Pattie. Artificial Life Meets Entertainment: Lifelike Autonomous Agents. *CACM* 38, 11 (November 1995), 108-114.
- Madritsch96 Madritsch, F. and M. Gervautz. CCD-Camera Based Optical Beacon Tracking for Virtual and Augmented Reality. *Proceedings of Eurographics '96* (Futuroscope - Poitiers, France, 26-30 August 1996).
- Mair94 Mair, Susan G. Preliminary Report on SIGGRAPH in the 21st Century: Rediscovering Our Fire. *Computer Graphics* 28, 4 (November 1994), 288-296.

- McMillan95a McMillan, Leonard and Gary Bishop. Head-Trackd Stereoscopic Display Using Image Warping. *SPIE Proceedings 2409 Electronic Imaging Science and Technology* (San Jose, CA, 5-10 February 1995), 21-30.
- McMillan95b McMillan, Leonard and Gary Bishop. Plenoptic Modeling. *Proceedings of SIGGRAPH '95* (Los Angeles, CA, 6-11 August 1995). In *Computer Graphics, Annual Conference Series, 1995*, 39-46.
- Mellor95a Mellor, J. P. Enhanced Reality Visualization in a Surgical Environment. MS Thesis, Department of Electrical Engineering, MIT (13 January 1995).
- Mellor95b Mellor, J.P. Realtime Camera Calibration for Enhanced Reality Visualization. *Proceedings of Computer Vision, Virtual Reality, and Robotics in Medicine '95 (CVRMed '95)* (Nice, France, 3-6 April 1995), 471-475.
- Meyer92 Meyer, Kenneth, Hugh L. Applewhite, and Frank A. Biocca. A Survey of Position-Trackers. *Presence: Teleoperators and Virtual Environments 1, 2* (Spring 1992), 173-200.
- Milgram93 Milgram, Paul, Shumin Zhai, David Drascic, and Julius J. Grodski. Applications of Augmented Reality for Human-Robot Communication. *Proceedings of International Conference on Intelligent Robotics and Systems* (Yokohama, Japan, July 1993), 1467-1472.
- Milgram94a Milgram, Paul, and Fumio Kishino. A Taxonomy of Mixed Reality Virtual Displays. *IEICE Transactions on Information and Systems E77-D, 9* (September 1994), 1321-1329.
- Milgram94b Milgram, Paul, Haruo Takemura, Akira Utsumi, and Fumio Kishino. Augmented Reality: A Class of Displays on the Reality-Virtuality Continuum. *SPIE Proceedings volume 2351: Telem manipulator and Telepresence Technologies* (Boston, MA, 31 October - 4 November 1994), 282-292.
- Milgram95 Milgram, Paul, David Drascic, Julius J. Grodski, Anu Restogi, Shumin Zhai, and Chin Zhou. Merging Real and Virtual Worlds. *Proceedings of IMAGINA '95* (Monte Carlo, 1-3 February 1995), 218-230.
- Mine93 Mine, Mark R. Characterization of End-to-End Delays in Head-Mounted Display Systems. UNC Chapel Hill Department of Computer Science technical report TR 93-001 (March 1993), 11 pages.

- Neumann96 Neumann, Ulrich and Youngkwan Cho. A Self-Tracking Augmented Reality System. *Proceedings of VRST '96* (Hong Kong, 1-4 July 1996), 109-115.
- Oishi96 Oishi, Takashi and Susumu Tachi. Methods to Calibrate Projection Transformation Parameters for See-Through Head-Mounted Displays. *Presence: Teleoperators and Virtual Environments* 5, 1 (Winter 1996), 122-135.
- Olano95 Olano, Marc, Jon Cohen, Mark Mine, and Gary Bishop. Combating Graphics System Latency. *Proceedings of 1995 Symposium on Interactive 3D Graphics* (Monterey, CA, 9-12 April 1995), 19-24.
- Oyama93 Oyama, Eimei, Naoki Tsunemoto, Susumu Tachi, and Yasuyuki Inoue. Experimental Study on Remote Manipulation Using Virtual Reality. *Presence: Teleoperators and Virtual Environments* 2, 2 (Spring 1993), 112-124.
- Pausch92 Pausch, Randy, Thomas Crea, and Matthew Conway. A Literature Survey for Virtual Environments: Military Flight Simulator Visual Systems and Simulator Sickness. *Presence: Teleoperators and Virtual Environments* 1, 3 (Summer 1992), 344-363.
- Peuchot95 Peuchot, Bernard, Alain Tanguy, and Michel Eude. Virtual Reality as an Operative Tool During Scoliosis Surgery. *Proceedings of Computer Vision, Virtual Reality, and Robotics in Medicine '95 (CVRMED '95)* (Nice, France, 3-6 April 1995), 549-554.
- Regan94 Regan, Matthew, and Ronald Pose. Priority Rendering with a Virtual Reality Address Recalculation Pipeline. *Proceedings of SIGGRAPH '94* (Orlando, FL, 24-29 July 1994). In *Computer Graphics*, Annual Conference Series, 1994, 155-162.
- Rekimoto95a Rekimoto, Jun, and Katashi Nagao. The World Through the Computer: Computer Augmented Interaction with Real World Environments. *Proceedings of UIST '95* (Pittsburgh, PA, 14-17 November 1995), 29-36.
- Rekimoto95b Rekimoto, Jun. The Magnifying Glass Approach to Augmented Reality Systems. *Proceedings of ICAT '95* (Makuhari Messe, Chiba, Japan, 20-22 November 1995).
- Riner92 Riner, Bruce and Blair Browder. Design Guidelines for a Carrier-Based Training System. *Proceedings of IMAGE VI* (Scottsdale, AZ, July 14-17 1992), 65-73.

- Robinett92a Robinett, Warren. Synthetic Experience: A Proposed Taxonomy. *Presence: Teleoperators and Virtual Environments 1*, 2 (Spring 1992), 229-247.
- Robinett92b Robinett, Warren and Jannick Rolland. A Computational Model for the Stereoscopic Optics of a Head-Mounted Display. *Presence: Teleoperators and Virtual Environments 1*, 1 (Winter 1992), 45-62.
- Rolland93 Rolland, Jannick P., and Terry Hopkins. A Method of Computational Correction for Optical Distortion in Head-Mounted Displays. UNC Chapel Hill Department of Computer Science technical report TR93-045 (1993).
- Rolland94 Rolland, Jannick, Rich Holloway, and Henry Fuchs. A Comparison of Optical and Video See-Through Head-Mounted Displays. *SPIE Proceedings volume 2351: Telemicroscopy and Telepresence Technologies* (Boston, MA, 31 October - 4 November 1994), 293-307.
- Rolland95 Rolland, Jannick, Frank Biocca, Todd Barlow, and Anantha Kancherla. Quantification of Adaptation to Virtual-Eye Location in See-Thru Head-Mounted Displays. *Proceedings of IEEE VRAIS '95* (Research Triangle Park, NC, 11-15 March 1995), 56-66.
- Rose95 Rose, Eric, David Breen, Klaus Ahlers, Chris Crampton, Mihran Tuceryan, Ross Whitaker, and Douglas Greer. Annotating Real-World Objects Using Augmented Reality. *Proceedings of Computer Graphics International '95* (Leeds, UK, 25-30 June 1995), 357-370.
- Rosen96 Rosen, Joseph M., Hooman Soltanian, Richard J. Redett and Donald R. Laub. Evolution of Virtual Reality: From Planning to Performing Surgery. *IEEE Engineering in Medicine and Biology 15*, 2 (March / April 1996), 16-22.
- Sharma94 Sharma, Rajeev, and Jose Molineros. Role of Computer Vision in Augmented Virtual Reality. *SPIE Proceedings volume 2351: Telemicroscopy and Telepresence Technologies* (Boston, MA, 31 October - 4 November 1994), 220-231.
- Simon94 Simon, David A., Martial Hebert, and Takeo Kanade. Techniques for Fast and Accurate Intra-Surgical Registration. *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery (MRCAS)* (September 1994), 90-97.
- Sims94 Sims, Dave. New Realities in Aircraft Design and Manufacture. *IEEE Computer Graphics and Applications 14*, 2 (March 1994), 91.

- So92 So, Richard H. Y. and Michael J. Griffin. Compensating Lags in Head-Coupled Displays Using Head Position Prediction and Image Deflection. *Journal of Aircraft* 29, 6 (November - December 1992), 1064-1068.
- Sowizral93 Sowizral, Henry, and James Barnes. Tracking Position and Orientation in a Large Volume. *Proceedings of IEEE VRAIS '93* (Seattle, WA, 18-22 September 1993), 132-139.
- State94 State, Andrei, David T. Chen, Chris Tector, Andrew Brandt, Hong Chen, Ryutarou Ohbuchi, Mike Bajura and Henry Fuchs. Case Study: Observing a Volume Rendered Fetus within a Pregnant Patient. *Proceedings of IEEE Visualization '94* (Washington D.C., 17-21 October 1994), 364-368.
- State96a State, Andrei, Gentaro Hirota, David T. Chen, Bill Garrett, and Mark Livingston. Superior Augmented Reality Registration by Integrating Landmark Tracking and Magnetic Tracking. *Proceedings of SIGGRAPH '96* (New Orleans, LA, 4-9 August 1996), 429-438.
- State96b State, Andrei, Mark A. Livingston, Gentaro Hirota, William F. Garrett, Mary C. Whitton, Henry Fuchs and Etta D. Pisano. Techniques for Augmented-Reality Systems: Realizing Ultrasound-Guided Needle Biopsies. *Proceedings of SIGGRAPH '96* (New Orleans, LA, 4-9 August 1996), 439-446.
- Taubes94 Taubes, Gary. Surgery in Cyberspace. *Discover* 15, 12 (December 1994), 84-94.
- Tharp94 Tharp, Greg, Samad Hayati and Linh Phan. Virtual Window Telepresence System for Telerobotic Inspection. *SPIE Proceedings volume 2351: Telem manipulator and Telepresence Technologies* (Boston, MA, 31 October - 4 November 1994), 366-373.
- Torrance95 Torrance, Mark C. Advances in Human-Computer Interaction: The Intelligent Room. *CHI '95 Research Symposium* (Denver, CO, 6-7 May 1995).
- Tuceryan95 Tuceryan, Mihran, Douglas S. Greer, Ross T. Whitaker, David Breen, Chris Crampton, Eric Rose, and Klaus H. Ahlers. Calibration Requirements and Procedures for Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* 1, 3 (September 1995), 255-273.
- Uenohara95 Uenohara, Michihiro and Takeo Kanade. Vision-Based Object Registration for Real-Time Image Overlay. *Proceedings of Computer Vision, Virtual Reality, and Robotics in Medicine '95 (CVRMed '95)* (Nice, France, 3-6 April 1995), 13-22.

- Utsumi94 Utsumi, Akira, Paul Milgram, Haruo Takemura, and Fumio Kishino. Effects of Fuzziness in Perception of Stereoscopically Presented Virtual Object Location. *SPIE Proceedings volume 2351: Telemanipulator and Telepresence Technologies* (Boston, MA, 31 October - 4 November 1994), 337-344.
- Wanstall89 Wanstall, Brian. HUD on the Head for Combat Pilots. *Interavia 44* (April 1989), 334-338. [A89-39227].
- Ward92 Ward, Mark, Ronald Azuma, Robert Bennett, Stefan Gottschalk, and Henry Fuchs. A Demonstrated Optical Tracker with Scalable Work Area for Head-Mounted Display Systems. *Proceedings of 1992 Symposium on Interactive 3D Graphics* (Cambridge, MA, 29 March - 1 April 1992), 43-52.
- Watson95 Watson, Benjamin and Larry Hodges. Using Texture Maps to Correct for Optical Distortion in Head-Mounted Displays. *Proceedings of IEEE VRAIS '95* (Research Triangle Park, NC, 11-15 March 1995), 172-178.
- Welch78 Welch, Robert B. Perceptual Modification: Adapting to Altered Sensory Environments. Academic Press (1978). ISBN 0-12-741850-4.
- Wellner93 Wellner, Pierre. Interacting with Paper on the DigitalDesk. *CACM 36*, 7 (July 1993), 86-96.
- Whitaker95 Whitaker, Ross T., Chris Crampton, David E. Breen, Mihran Tuceryan and Eric Rose. Object Calibration for Augmented Reality. *Proceedings of Eurographics '95* (Maastricht, The Netherlands, August 1995), 15-27.
- Wloka95a Wloka, Matthias M. Lag in Multiprocessor Virtual Reality. *Presence: Teleoperators and Virtual Environments 4*, 1 (Winter 1995), 50-63.
- Wloka95b Wloka, Matthias M. and Brian G. Anderson. Resolving Occlusion in Augmented Reality. *Proceedings of 1995 Symposium on Interactive 3D Graphics* (Monterey, CA, 9-12 April 1995), 5-12.
- Wu95 Wu, Jiann-Rong and Ming Ouhyoung. A 3D Tracking Experiment on Latency and its Compensation Methods in Virtual Environments. *Proceedings of UIST '95* (Pittsburgh, PA, 14-17 November 1995), 41-49.
- Yoo93 Yoo, Terry S., and T. Marc Olano. Instant Hole™ (Windows Onto Reality). UNC Chapel Hill Department of Computer Science technical report TR93-027 (1993), 13 pages.

- Zikan94a Zikan, Karel, W. Dan Curtis, Henry Sowizral, and Adam Janin. Fusion of Absolute and Incremental Position and Orientation Sensors. *SPIE Proceedings volume 2351: Telem manipulator and Telepresence Technologies* (Boston, MA, 31 October - 4 November 1994), 316-327.
- Zikan94b Zikan, Karel, W. Dan Curtis, Henry A. Sowizral, and Adam L. Janin. A Note on Dynamics of Human Head Motions and on Predictive Filtering of Head-Set Orientations. *SPIE Proceedings volume 2351: Telem manipulator and Telepresence Technologies* (Boston, MA, 31 October - 4 November 1994), 328-336.