Military Applications of Augmented Reality

Mark A. Livingston, Lawrence J. Rosenblum, Dennis G. Brown, Gregory S. Schmidt, Simon J. Julier, Yohan Baillot, J. Edward Swan II, Zhuming Ai, and Paul Maassel

1 Introduction

This chapter reviews military benefits and requirements that have led to a series of research efforts in augmented reality (AR) and related systems over the past few decades, beginning with the DARPA-funded research of Ivan Sutherland that initiated the field of interactive computer graphics. We will briefly highlight a few of the research projects that have advanced the field over the past five decades. We will then examine in detail the Battlefield Augmented Reality System at the Naval Research Laboratory, which was the first system developed to meet the needs of the dismounted warfighter. Developing this system has required advances, in particular in the user interface (UI) and human factors. We summarize our research and place it in the context of the field.

Military operations are becoming increasingly diverse in their nature. To cope with new and more demanding tasks, the military has researched new tools for use during operations and during training for these operations. There have been numerous goals driving this research over the past several decades. Many of the military requirements and capabilities have specifically driven development of AR systems. Thus we begin this chapter by discussing some military needs and challenges for which AR has been proposed to help. The following sections review some specific military applications of AR and examine some of the critical issues limiting the incorporation of AR in military applications. We conclude with a discussion of implications for the field of AR.

Situation Awareness

The environments in which military operations occur have always been complex, and modern operations have only served to increase this complexity. Dynamic sce-

Naval Research Laboratory, Washington, DC, e-mail: mark.livingston@nrl.navy.mil

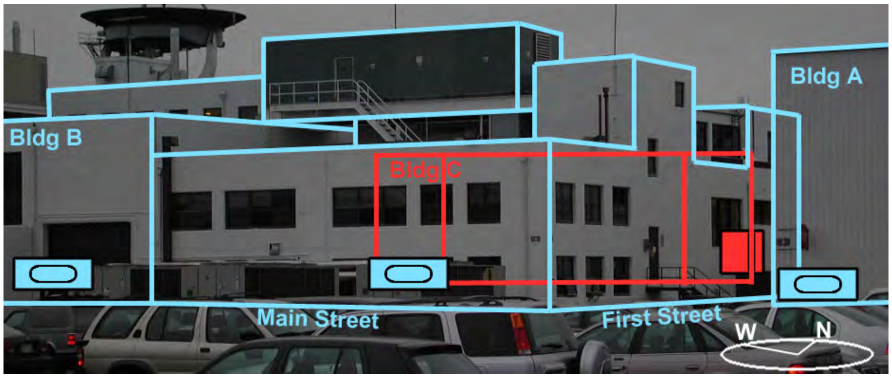

Fig. 1 This concept sketch shows information important for military personnel to establish and maintain SA: building and street labels, friendly (light rectangles) and enemy (dark square) icons, and a compass.

narios help create the "fog of war," according to the oft-quoted phrase. It is difficult to keep track of the many friendly and opposing forces operating in an environment. Keeping track of the past, present, and future during such a military operation has been termed *situation awareness* (SA) [Bolstad and Endsley(2002)]. The time scale considered to be part of SA varies, but the three categorical times remain. Even keeping track of basic information such as the locations of friendly forces, building and street names or identifiers, and orientation with respect to a global coordinate system become challenging, but critical, tasks. Early designs in our human-centered research process attempted to show multiple layers of geometric and human terrain that might be of interest to dismounted personnel (Fig. 1).

The Marine Corps Combat Development Command Concepts Division once described the issue as follows

> Units moving in or between zones must be able to navigate effectively, and to coordinate their activities with units in other zones, as well as with units moving outside the city. This navigation and coordination capability must be resident at the very-small-unit level, perhaps even with the individual Marine [Van Riper(1997)].

On top of this, recent trends towards asymmetric conflicts have witnessed civilians getting caught in the midst of battles – or worse, purposefully used as human shields by terrorists who do not operate under conventional rules of engagement. These asymmetric battles have become much more common in recent conflicts, and this trend is expected to continue. Increasingly, such battles are fought in dense urban environments, which are inherently more challenging to understand. The nature of battles in 3D urban structures and involving combined air-ground forces further stresses the cognitive load of the individual infantryman, pilot, or sailor, whether in command of some portion of the forces or reporting up the chain of command. With the ability of AR to augment one's view without obscuring that environment, AR became a natural paradigm in which to present military information. Head-up 3D visualization within urban structures was considered a key benefit over 2D map visualizations. This is similar in spirit to the insertion of the first-down line in broadcasts of (American) football; seeing the line as play unfolds gives viewers much greater awareness of the meaning of the play.

Information Overload

The counter-point to having myriad bits of information that give one a complete picture of the past history, current status, and potential consequences of actions in the environment is having too much information to process. Military commanders often compare information processing in the battlespace to attempting to sip water from a fire hose. The condition of *information overload* occurs when one is unable to process the information presented into coherent SA. With the rapidly expanding ability to collect data in (near) real-time about many locations and provide data abstractions to the warfighter at levels from the command center to individual field personnel, the danger of information overload has grown significantly.

The nature of AR is (generally) to add information to the user's view of an environment; clearly, the issue of information overload requires that this be done in a careful manner so as not to impede the user's ability to achieve or maintain SA. One corollary to this requirement is that the information presented to each user must be appropriate for that user's role in the team's mission. A commander may need to understand the global situation and how the various teams are expected to move through an environment, whereas a private on patrol may only be concerned with a very limited area of the environment. Similarly, a medic may need health records and a route to an injured soldier, whereas a forward observer may need a few days' worth of reconnaisance information in order to detect unusual or unexpected enemy actions. Ideally, an AR system (or any information delivery system) would be aware of these various tasks, the mission plans (including any contingencies), and the current roles any particular user may be fulfilling at a given time.

It should also be evident at this point that an AR system for military applications bridges two somewhat disparate fields. SA implies the introduction of visual representations of data. This type of data abstraction is in itself a major sub-field within the field of computer graphics. Overlaying information is a fundamental characteristic of AR, and this sensory integration can both limit the types of abstractions that make sense for a given application and push the application designer to create new methods of understanding perceptual or cognitive cues that go beyond typical human sensory experiences.

Training

When conceiving of virtual training, most people immediately think of immersive virtual environment systems, rather than AR and its overlaying of information on the real world. One research thrust that is gaining interest is the use of wearable virtual reality systems for embedded training. For example, a warfighter en route to a deployment may wear a system like the Land Warrior system [Army(2001)] containing a wearable computer and head-mounted display designed for the display of SA information. But the system could load a virtual training application to better use this travel time. Systems of this type include VICTER [Barham et al(2002)], DAGGERS and ExpeditionDI®[Quantum3D(2010)], Virtual Warrior [GD(2010)], Nett Warrior [Gould(2010)], and COMBATREDI®[Cubic(2011)].

AR offers some practical advantages over virtual environments. Embedding virtual training applications in existing live-action training facilities can reduce modeling (and rendering) requirements and other infrastructure costs. Modeling an accurate virtual environment and the unknown fidelity requirements of such a model make this an expensive need for immersive virtual environments. Furthermore, this AR facility would maintain the natural haptic cues one gets from walls and other real objects. Virtual environments often require unnatural (e.g. joystick-based) navigation methods; AR eliminates this and allows the user to walk normally, albeit by requiring a large tracking range. Given that AR may one day be an operational tool, using it for training follows the goal for the military to "train as you fight." AR allows for more realistic interaction among multiple trainees, since they see each other through their natural vision (as opposed to an avatar representing a particular person). Finally, instead of using personnel resources to take the roles of potential adversaries or having trainees learn against empty space, a warfighter could train against avatars.

A projection-display based version of mixed reality (MR) training was implemented in the Future Immersive Training Environment Joint Capability Technology Demonstration [Muller(2010)]. In the first implementation, avatars appear on projection screens within a real training environment, technology that is often known as spatial AR. This limits flexibility in the location of avatars, but still supports effective training. The ability to reduce the number of personnel required for effective training (by substituting avatars for real actors to play opposing forces) translates into cost savings. One advantage of the use of AR for this training is that the amount of infrastructure that must be changed from a live-action training facility is small compared to that required by an immersive virtual environment training facility. An improved version of the system used video-based, head-worn AR displays that incorporated a computer vision-based tracking system to reduce the errors in registration of the avatars to the real environment. Registration error could cause avatars to be improperly occluded by the real environment or appear to float above the ground. One limitation of this system is that it currently allows only small units to train together, either a fire team (four people) or squad (thirteen). Another limitation is the size, weight, and power requirements of the head-worn apparatus.

In general, the disadvantages of AR for training are that, like virtual environments, AR systems are not easy to implement, and the technology has struggled to meet some of the minimal requirements in order to be a useful system. Several ongoing research projects are aimed at improving the displays and tracking systems that are critical components for an AR system. Some aspects of AR technology, such as the display, have more stringent requirements than immersive training simulations.

So while AR clearly has powerful potential as a training tool, whether it is best for a particular application is not so clear.

A related concept to the planning or rehearsal of a mission is the analysis of a completed mission for future training. In the military, such an analysis is known as *after-action review* (AAR). Both virtual environments and AR generate data that may be used for this type of training. In the same way that AR could reduce the modeling costs associated with virtual training, AR might help reduce the expense of setting up a formal AAR.

Another possible use in training is for specific skills that are basic to numerous military roles. Patrols use particular search patterns to maintain awareness of potential threats. While actors could be trained to approach from specific directions, it can be a more repeatable and cost-effective system to implement virtual avatars for training such a fundamental skill. In this way, basic requirements in military training can be met in an individual instructional phase, allowing each trainee to progress at his or her own pace. It also affords the instructor the ability to test specific difficulties a trainee has had in the past in a repeatable fashion.

Quick-reaction Forces

Another increasing emphasis for military operations is the faster pace at which decisions must be made, while the cost of poor decisions can be catastrophically high. If an AR system can present exactly the right pieces of information, better and faster decisions can be made and turned into correct actions. This operates at multiple levels: an individual on an operation might make a better decision about whether an approaching vehicle is a threat, a squad might be able to come to the aid of another unit that is under fire, or a battalion can take advantage of a quickly-configurable AR training facility and be ready to respond to an opportunity that is available for only a brief time. Without such information, perhaps the advantage of pro-active manuevers will be lost, an operation would be too high a risk to undertake, or decisions have a lower probability of positive outcomes (e.g. a higher rate of losses). Because AR can theoretically present training scenarios with low configuration cost in terms of the scenario (if not the AR infrastructure with current technology), it offers hope for the future of quick-reaction forces.

2 AR Projects for the Military

The concept of a display system indistinguishable from reality was introduced by Ivan Sutherland [Sutherland(1965)]; a preliminary realization of this "Ultimate Display" for the visual sense was described subsequently [Sutherland(1968)]. The system included not only the head-worn display (HWD), but also an image generation subsystem and a tracking subsystem for the user's head and one of the user's hands. Thus it marked the starting point for both virtual environments and AR research.

It is interesting to note that this first HWD was an AR display, not a completely immersive display suitable for immersive virtual environments.

The system required other novel hardware, notably a “clipping divider” that could properly render perspective views (well before commodity graphics cards became standard) and two position tracking systems (mechanical arm and ultrasonic). One important difficulty noted in early tests of the system was the ambiguous nature of the 3D images. Users visualized a cyclo-hexane molecule; those familiar with the shape had no trouble recognizing it, but other users misinterpreted the shape. This foundational work foreshadowed the difficulties faced by later systems being applied to military applications.

2.1 The “Super Cockpit”

The first specific application of AR technology was for fighter pilots. The Super Cockpit was the forerunner of the modern head-up display still used now by fighter pilots and available in some passenger cars. The original implementations used both virtual environment and see-through display metaphors, to enable the pilot to use the system at night. The system was developed at Wright-Patterson Air Force Base beginning in the late 1960s [Furness(1969)].

Visibility out of a cockpit is limited, and airborne tasks such as low-altitude navigation, target acquisition, and weapons delivery require pilots to reference landmarks on the terrain. However, sensors mounted on the aircraft can create visibility in areas that are occluded by the aircraft structure, or in conditions such as low light that prevent the pilot from seeing the real world. The system superimposed flight and target data into the pilot’s visual field and provided sound cues to assist localization.

The key feature of this system was providing spatial awareness for the pilot to understand and incorporate into his course of action a variety of incoming data streams. The horizon became visible through the cockpit window, rather than being conveyed on an indicator on the instrument panel. Targets, navigation waypoints, and threats could similarly be registered to their 3D locations. The concept was that such a view would improve upon using a dashboard display, leaving the pilot to mentally merge the virtual map with his visual field. This is not an easy task and would have required the pilot to take his eyes off the real environment many times in order to align the virtual information. Another feature provided a rear-view mirror, similar to the standard mechanism in a car.

Early versions of the system pointed out the need for study of the human factors of such systems. Spatial reasoning is a complex task, even more so under the physical duress of flight and the emotional intensity of combat. An intuitive interface could take advantage of the natural abilities of many people to reason in three dimensions, rather than have them reason in two dimensions and try to apply that to the real environment. This 3D spatial reasoning is also not a trivial task, but pilots are screened for high spatial reasoning ability, so it seems natural to supply them with an inherently 3D view.

2.2 Aspen Movie Map

One long-standing goal of military training is for forces to know the environment in which an operation will take place, enabling them to navigate and make decisions much faster than if they had to focus on a possibly inaccurate mental map and consider the choices available to them. The interactive movie map was an early attempt to provide this "mechanism for pre-experiencing an unfamiliar locale that allows the acquisition of spatial knowledge to take place in a meaningful, natural, and accurate way." [Mohl(1981)] The Aspen Movie Map [Naimark(1979)] was the first of these systems, building on the newly-available optical video disc technology of the 1970s to enable interactive computer control of the display of video frames. In this regard, the movie map shares many characteristics with video-based AR systems; the major difference being the spatial and temporal separation of the user from the displayed environment. However, as this was intended to investigate training applications for the military, it sparked much research in virtual environments and AR, including some of the systems discussed below.

The goal of this system was to convey a sense of being in Aspen, CO to the user, such that the user would know how to navigate in the town without ever having been there. Thus the UI and the controls offered were important components of the system. Most relevant to our discussion of AR and MR was the overview map mode, in which the user could trace the route taken or specify a route for the system to follow; these graphics were merged with aerial overviews. Speed and direction were controlled through a scrollbar-like widget on the touch-screen display. The user could zoom in on the map view, change the season through a toggle switch, or engage a slide show about a specific building from the map.

In an informal user evaluation, subjects who experienced Aspen through the interactive movie map were found to sketch similar maps of Aspen as residents made. Similarity was judged in terms of the number of errors, degree of uncertainty, and level of detail (especially along the main street). It appeared that the movie map users composed their maps based more on linear paths than on areas. This was perhaps a consequence of the linear, grid-like street system of the town, which gave the only available travel routes in the movie map. These users were more likely to be unsure about distances from the movie map, and even when using the system, noted uncertainty about distance traveled along the real terrain and the number of degrees turned, even to the point of some being unsure whether they had turned 90° or 180°. One user sketched an extremely accurate map from the movie map, so at least the potential for learning the space seemed to be a feature.

2.3 Battlefield Augmented Reality System

The overall goal of the Battlefield Augmented Reality System™(BARS) was to do for the dismounted warfighter what the Super Cockpit and its successors had done for the pilot. Initial funding came from the Office of Naval Research. The challenges associated with urban environments were a particular concern: complex 3D environment, dynamic situation, and loss of line-of-sight contact of team members [Livingston et al(2002)]. Unambiguously referencing landmarks in the terrain and integrating unmanned systems into an operation can also be difficult for distributed users[Livingston et al(2006a), Livingston et al(2006b)]. All of these examples show the impairment of SA in urban operations [Van Riper(1997)]. The belief was that the equivalent of a head-up display (such as in the Super Cockpit) would help solve these. By networking the mobile users together and with a command center, BARS could assist in establishing collaborative SA by a dispersed team.

This raises numerous issues in system configuration. BARS includes an information database, which can be updated by any user. Sharing information across the area of operation is a critical component of team SA. We designed an information distribution system [Brown et al(2004b)] so that updates would be sent across the network. We enabled BARS to communicate with semi-automated forces (SAF) software [Brown et al(2004a)] to address the training issues discussed above. We chose to use commercially-available hardware components so that we could easily upgrade BARS as improved hardware became available. We built UI components so that routes could be drawn on the terrain in the command center application and assigned to mobile users, or drawn by mobile users and suggested to the commander or directly to other mobile users. Typical AR system issues like calibration [Baillot et al(2003)] were investigated.

The BARS research program helped spawn ongoing efforts in both operations [DARPA(2010)] and training [ONR(2010)]. Land Warrior [Cox(2008)] shares some capabilities of team SA and has been praised for eliminating confusion in time-sensitive target missions and stopping fratricides in terrain that prevents squad elements from seeing each other. Specific research efforts within BARS for the UI and human factors aspects are the focus of the next section.

2.4 C-130 Loadmaster Training

Flight simulators have become standard tools for training pilots, but other aspects of military flights were also considered to have the potential to benefit from virtual training systems. The USAF Air Education and Training Command conducted a project to determine if AR was an effective tool for training C-130 loadmaster normal and emergency procedures. A C-130 is a military transport aircraft typically used in cargo and resupply operations by numerous countries. The loadmaster is a cargo handler and rigging expert that delivers cargo by airdropping equipment and personnel out of the back of the C-130 [Gardley et al(2008)]. Current training for this job uses ground-based fuselages mounted in a hangar without wings or tail, but with the appropriate working interior. However, the number of training devices and the fidelity of existing devices was judged insufficient to meet the training objectives.

A system was built from commercially-available hardware components and custom software. Head-worn AR displays were mounted to standard-issue flight helmets for students of the training. After a calibration phase, 15 students were trained in procedures for engine start-up and another 15 students in procedures for cargo airdrop. After this training, the students were asked to evaluate the effectiveness of the AR system as a training tool. Six instructors were also asked to evaluate the AR system. A third scenario, simulation of smoke and fumes in the cargo compartment, was shown to students but not included in the evaluation. Questionnaires were used to gather evaluation data.

The overall reaction to the AR system was positive [Gardley et al(2008)]; 63% of the students said that AR was an improvement over the current instructional training. The visualization capability in the AR system was judged to be "far superior" to "checklists and class discussions." However, only 68% thought the AR display was comfortable to wear. Although only six students had used AR prior to this study, 21 recommended it for future use in the training. Among the six instructors surveyed, there was unanimous agreement that the AR system enhanced student training and gave a realistic portrayal of events in the aircraft. Three instructors felt that the training they received on the use of the AR device was less than adequate. Only two of the instructors had never used AR before. Students criticized the AR display as hampering their vision and being too dim. The helmet and display assembly were also judged to be uncomfortable. Finally, students commented that the reaction time of the software was too slow. One missing piece of data that the evaluators were not able to collect were the grades received in the training course to determine whether the AR training had helped students achieve better grades. Lessons learned from these surveys were used to upgrade the AR system. New software and hardware (including display goggles) will be tested in 2011 on a larger group of students. Future evaluations will use surveys, interviews, and reviews of student training records to judge the efficiency and effectiveness of the AR training for loadmaster procedures.

2.5 Summary

In addition to the applications described here, military tasks encompass a wide range of roles that are filled in other professions as well. Industrial and manufacturing (Chapter 30) applications are discussed elsewhere in this volume. Maintenance and repair of military vehicles is a critical element of mission requirements. Many military vehicles feature a complex, special-purpose set of equipment that makes repairs especially challenging. One particular difficulty can be the densely-packed interior of the vehicles. An evaluation of a prototype AR system for repair of an armored personnel carrier turret [Henderson and Feiner(2011)] found that six recent graduates of a US Marine Corps mechanic's course exhibited faster task location times with AR guidance versus an electronic manual running on a notebook computer and versus the use of a head-worn display with head-fixed (i.e. non-registered) graphics. The study also found that head motion (translation and rotation) was lowered with the AR system versus the electronic manual. The head-worn display may have induced the users to move less than they would have otherwise, so the cause-and-effect relationship must be investigated further. Medical applications (Chapters 27 and 28) are another field in which military interest in AR technology has a long history. The Ultrasound Augmented Reality project [State et al(1996)] is one noteworthy example of a project supported by military research funding.

The remainder of this chapter will describe the multi-faceted BARS program's points of emphasis. The next section discusses case studies of how BARS has been or could be instantiated. The following section describes a number of research efforts in the UI and human factors aspects of the BARS system instantiations. These were motivated by difficulties encountered in implementing the prototype applications and were viewed as critical problems to solve in order to transition these applications out of the lab and into the hands of real users.

3 BARS Case Studies

BARS was originally envisioned as a system for operations, but it became clear that there was great benefit in taking some of the concepts into the training domain as well. These are two vast areas of military applications. As noted above, the military has long supported AR efforts in medicine, maintenance and repair, and other areas that are of interest but not exclusive to the military. In this section, we present four examples of applications built from the BARS software and infrastructure. These discussions will discuss hardware as appropriate, mostly from the point of view of the requirements or benefits of certain technologies. Since some of these applications were pursued for a finite period of time which has long since passed, some specific hardware would now seem to be an anachronism and will thus be omitted from the discussion.

3.1 Dismounted Warfighter Operations

Military operations in urban terrain (MOUT) present many unique and challenging conditions for the warfighter. The environment is extremely complex and inherently 3D. Above street level, buildings serve varying purposes (such as hospitals or communication stations). They can harbor many risks, such as snipers or explosives, which can be located on different floors. Below street level, there can be an elaborate network of sewers and tunnels. The environment can be cluttered and dynamic. Narrow streets restrict line of sight and make it difficult to plan and coordinate group activities. Threats, such as snipers, can continuously move and the structure of the environment itself can change. For example, a damaged building can fill a street with rubble, making a once-safe route impassable. Such difficulties are compounded by the need to minimize civilian casualties and damage to civilian targets.

Fig. 2 We built this early prototype of our wearable AR system from commercial components. This and similar realizations were used in the research described later in this chapter.

A number of research programs have explored the means by which navigation and coordinated information can be delivered to dismounted warfighters. Many of these approaches are based on handheld maps or opaque head-mounted displays (HMDs). For example, Land Warrior introduced a head-worn display that combined a map and a "rolling compass" [Gumm et al(1998)]. Unfortunately, these methods have a number of limitations. They obscure the user's field of view and do not truly represent the 3D nature of the environment. Moreover they require the user to integrate the graphical display within the environment to make sense of it. This work is sometimes difficult and distracting from the current task. We believe a mobile AR system best meets the needs of the dismounted warfighter [Livingston et al(2002)]; we began assembling hardware and writing software to build prototype wearable systems (Fig. 2).

Through the ability to present direct information overlays, integrated into the user's environment, AR has the potential to provide significant benefits in many application areas. Many of these benefits arise from the fact that the virtual cues presented by an AR system can go beyond what is physically visible. Visuals include textual annotations, directions, instructions, or "X-ray vision," which shows objects that are physically present, but occluded from view. One important design consideration was not to limit the perceptual capabilities of the warfighter, so we chose optical see-through displays and worked with display manufacturers to reduce the loss of peripheral vision around the display. However, optical systems require significant range of brightness in order to be effective in the wide range of outdoor lighting conditions (day or night).

3.2 Mounted Warfighter Operations

Military vehicles are also increasingly operated in complex urban environments, forcing vehicle operators to face many of the same SA challenges as dismounted troops. Vehicle commanders are trained to maintain SA by cross-referencing two-dimensional map products, both digital and paper, to the live tactical environment outside the vehicle. Even for experienced operators, this process can be time-consuming and error-prone in the urban environment. AR systems are designed to merge the relevant aspects of the spatial data in the digital map environment into a view of the live tactical environment. A well-designed AR system will display spatial data intuitively with the real world. In a military vehicle, AR systems can enhance the SA of commander, driver, or gunners.

Vehicles are particularly well-suited for AR systems. Typically, the limitations of power, size, and weight that constrain wearable systems are less critical in the vehicle-mounted AR system. In addition, advanced military vehicles may already provide key components such as high-performance GPS and inertial navigation systems, external imaging sensors, digital computers, and video display screens. An early prototype of our vehicle-borne AR implementation is shown in Fig. 3.

Optical see-through AR systems must overcome some unique challenges to be effective in vehicle-mounted systems. The most obvious solution is a "head-up display" as previously discussed in reference to military aircraft. The display device could be mounted in the windshield in front of a vehicle commander. The primary issue with this approach is maintaining alignment of the symbols with the real world, since motion of the operator's head will create significant errors in registration caused by parallax. Fighter aircraft overcome this issue by constraining the pilot's position using an adjustable seat and a harness or placing the display in the fighter's helmet, which is in turn tightly worn on the head. Registration can be accomplished dynamically by tracking the operator's head with a tracking system and combining that information with the vehicle position and attitude to create the visual overlay. This approach adds complexity to the overall system and limits the operator view to a narrow window. Another solution involves having the user wear optical see-through displays inside the vehicle and then tracking the operator's head movements. This approach is currently being utilized by several military aircraft platforms, including the USMC AH-1Z and the Joint Strike Fighter. The approach is highly effective for providing real-time SA but the high costs to develop, purchase, and maintain such systems prevent it from being used in most land vehicles.

Video AR applications can provide similar functionality with less cost and complexity. The Meissa project at the Naval Research Laboratory looked to develop a simple AR system that could be installed in USMC and US Army HMMWV

Fig. 3 An early prototype for a vehicle-mounted system included a four-antenna GPS unit on the roof for position and orientation and mobile computers inside for computation and display.

vehicles to enhance the vehicle commander's SA. The system comprises a wide field-of-view camera installed behind the windshield, a commercially-available attitude GPS system, a rugged vehicle PC, and an LCD touch screen display. Custom-developed software merges symbols representing static and dynamic geo-spatial data onto the live video feed from the camera (Fig. 4). The end result is effectively a Head-Down-Display for the commander, which provides a SA alternative between two-dimensional map products and optical see-through options.

3.3 Embedded Training

MOUT training requires that trainees operate in urban structures or against other live trainees. Often the training uses simulated small-arms munitions and pits instructors against students in several scenarios. Many military bases have "towns" for training that consist of concrete block buildings with multiple levels and architectural configurations. AR and MR can enhance this training by providing synthetic opposing forces and non-combatants. MR MOUT [Hughes et al(2002)] provided virtual targets on a realistic set using MR (similar to the FITE project, described above). This does not require the trainee to wear an AR display but nearly eliminates

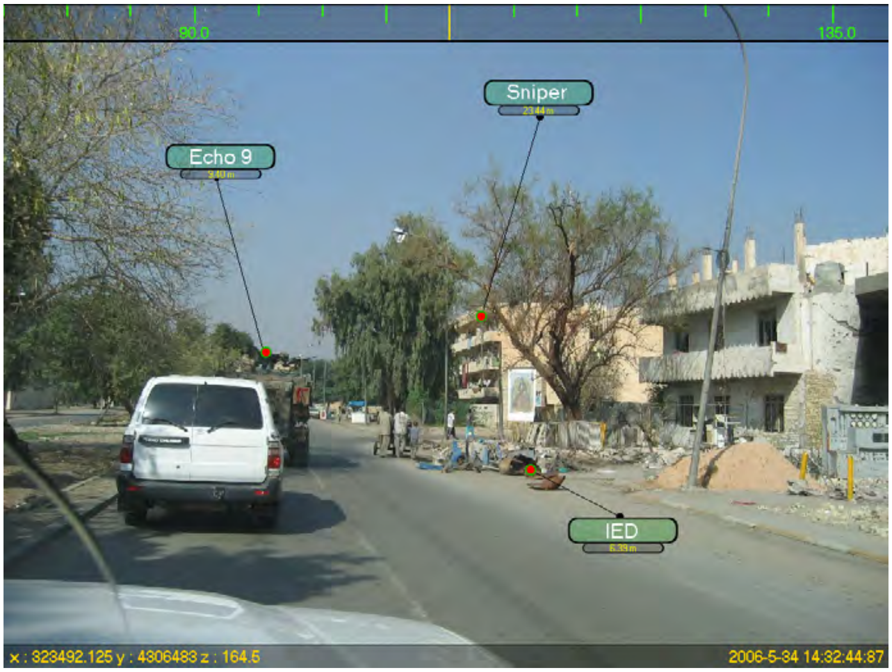

Fig. 4 The Meissa software displays cues for SA and threat potential.

portability. Training transfer work [Schmorrow et al(2004)] found that a "sensory-multiplexing" approach, in which the system takes advantage of the user's ability to gather information from multiple senses simulataneously – assuming the sensory input makes sense – increased the effectiveness of the training in the environment. We theorized that an AR trainer, getting many of these sensory inputs for "free," could be especially effective.

Using AR for MOUT training is a difficult undertaking. Once one has cleared acceptance and logistics issues, there are many technical challenges to face. Many of these challenges are the same as those as described earlier when AR is used for operations – wearable form factor, accurate tracking indoors and outdoors, and so on. One unique challenge to using AR for training operations is that the simulated forces need to give the illusion that they exist in the real world (Fig. 3.3).

Model fidelity is limited by the modeler's time and the rendering capability of the training computer. For mobile devices, it is only recently that sufficient capabilities have existed to display more than a few models with geometric and textural details. However, since only the virtual forces need to be drawn in an AR training application, rather than the entire environment, these resources can be focused on virtual forces. Lighting of the forces should match the environment, which requires measurement of the natural light, rendering capability to reproduce it, and a display algorithm that can make the rendered image appear as desired on the display. This last issue is problematic for optical see-through displays. Finally, the virtual

Fig. 5 *Above*: The wearable prototype for our embedded training uses a similar set of equipment as our earlier backpacks. *Right*: Virtual forces in the embedded training system must appear to exist in the real world and be subject to gravity and other laws of nature.

objects must properly occlude and be occluded by real objects. Optical see-through displays again present difficulties here. Whereas in operational contexts, seeing objects through walls might be an advantage, in training, this would break the illusion. Modeling the environment can provide the ability to occlude graphics with real objects by simply not rendering graphics where they are computed to be hidden from view; this is a standard property of the depth buffer in graphics hardware. Only a few research prototype optical see-through displays can fully occlude the real environment, by using a second display surface to become occlusive where desired. Video AR systems can provide this capability, at a cost of limiting the user to the geometric and color resolution of the camera that captures the real environment. We opted for the latter choice in our embedded training prototypes.

3.4 Forward Observer Training

The USMC's Fire Support Team training begins with small-scale (1:40) pneumatic mortars on a field. The purpose of this training is to hone the communication skills between the forward observer and the Fire Direction Center (FDC). In the current training plan, a forward observer visually locates targets, identifies and determines grid coordinates using binoculars and a map, and recommends a call for fire to the FDC. Once the shots are fired, the training instructor (not a part of the operational fire support team) determines the accuracy of the shots and the effect on the target: catastrophic hit, mobility hit, or no effect. The calls for fire are adjusted until the

Fig. 6 The forward observer training application shows virtual targets and buildings.

team has the desired effect on the target. Before the introduction of the AR system, the team fired upon static and unrealistic proxy targets.

One system [Brown et al(2005)] was demonstrated at Quantico Marine Corps Base in October 2004. It provided a head-mounted display for the forward observer and a touch screen for the instructor, each showing virtual targets on the real range. Fig. 6 shows the observer's view of virtual targets and buildings on the range. The observer can have the computer simulate a magnified view (including a reticle) that binoculars provide, to determine target identity and grid coordinates. The targets move along preset routes and are started and stopped by the instructor through a simple interface. As before, the forward observer calls for fire on the targets and a round is fired. The instructor sees where the round lands in the augmented touch screen view and designates the effect on the target. Through the dynamic shared database the forward observer sees that effect and revises the call for fire. Inserting AR into the training plan resulted in no significant changes to the duties and actions of the participants, but it enabled them to fire on moving targets.

The virtual targets for training were received well by trainees and instructors at Quantico; however, rigorous studies and measurements of effectiveness are yet to be done. The system can also insert virtual terrain and control measures into the display, and both capabilities were preliminarily tested at Quantico.

4 Human-Centered Research

As with the Super Cockpit [Furness(1986)] and Aspen Movie Map [Mohl(1981)], we found that BARS could not succeed in the applications described above without innovation in the UI and human factors of the hardware and software. For BARS research, we followed a user-centered approach [Livingston et al(2004)]; thus, we conducted a domain analysis [Livingston et al(2006b)]. In this section, we summarize a number of research thrusts in these arenas, with references to papers for readers interested in further details. These efforts were motivated by difficulties faced in the applications, as noted in the discussions below.

4.1 Depth Perception and Occlusion Representation

Among the things our initial domain analysis [Gabbard et al(2002)] indicated as a potential advantage for AR for dismounted troops was the ability to show where distributed troops were in an urban area of operations. Later, client interest included the ability to communicate points of interest in the environment to distributed team members (without the benefit of line-of-sight contact between team members). Both of these goals require the AR system to identify objects that are occluded from the user. This became a central focus of the BARS research program.

The metaphor of "Superman's X-ray vision" has long been applied to the capability of AR to depict a graphical object that is occluded by real objects [Stix(1992)]. There are three aspects to the problem of displaying cues that correspond to occluded virtual objects. First, the alignment or *registration* of the graphics on the display must be accurate. This is a defining aspect of AR [Azuma(1997)]. Second, the ordinal depth between the real and virtual objects must be conveyed correctly to the user. Because we selected optical see-through HWD for operational reasons, we needed to replace the natural occlusion cue for depth ordering. Third, the metric distance of the virtual object must be understood to within a sufficient accuracy that the user can accomplish the task. This requires the cues that are provided to be sufficiently accurate to estimate distance. Further, each successive aspect depends on the previous ones.

We began our investigation with a study that identified graphical cues that helped convey the ordinal depth of graphical objects [Livingston et al(2003)]. We found that changing the drawing style, decreasing the opacity with increasing distance, and decreasing intensity with increasing distance helped users properly order graph-

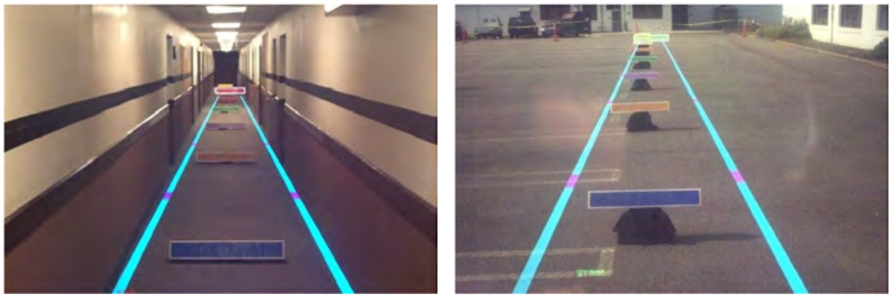

Fig. 7 *Left*: The indoor portion of our most recent AR depth perception featured strong (real and virtual) linear perspective cues. *Right*: The outdoor portion of the experiment tested the utility of only the virtual perspective cues, with limited success for distant targets.

ical depictions of buildings that corresponded to occluded buildings on our campus. This task was detached from the real world, however, so our next experiment used a depth-matching task with the same graphical cues. This enabled us to measure metric distance judgments and forced our users to compare the depth of graphical objects to the real environment. We verified that users were behaving similarly with virtual objects as with real objects [Livingston et al(2005)], and then studied distance estimation [Swan II et al(2006)]. We found an underestimation of distance at medium-field distances of 5-23 meters (similar to results for virtual environments), but overestimation beyond these distances. We found that the presence of an occluder increased the absolute distance error about 8 cm for every meter of distance from the user versus when no object occluded the region surrounding the virtual object's location relative to the real environment.

We next moved our experiment to the outdoor scenario for which BARS was intended and that this series of experiments was designed to solve. Because we were seeing such strong linear perspective cues in our indoor (hallway) environment, we decided to test whether replicating this cue would assist users in estimating the distance (Fig. 7). Our results [Livingston et al(2009b)] showed that users underestimated distance indoors, but overestimated distance outdoors. This differed from our previous indoor data, which showed a switch from underestimation in the medium-field but overestimation at far-field distances.

An analysis showed that the difference between the two data sets overall was significant. A closer analysis comparing the distances revealed that only for a reference object at approximately 33 meters was the difference for a particular distance significant. Thus we considered experimental conditions that were changed to suggest reasons for this difference in user behavior with respect to under- or over-estimation. The most obvious difference was the orientation of the reference objects. For the earlier experiment, we oriented the reference objects vertically, and the near edge of the virtual object (also oriented vertically) was a few inches away from the real reference objects. When the experiment was replicated for an outdoor environment, it became difficult to keep the reference objects upright, so we oriented both the real reference objects and the virtual targets horizontally in both indoor and outdoor environments. This separated the real and virtual objects by a couple of feet (as seen in Fig. 7). It is also true that for the second experiment, we compressed the distances slightly to fit into a smaller experimental space. One of these changes appeared to cause the difference in the indoor data.

The more important result from this experiment, however, was that the linear perspective cues caused users to reduce their estimation of the distance of the virtual object for only the most distant reference objects. In the outdoor environment, this improved the performance, since users were overestimating the distance. However, for the indoor environment, this increased the error, since users were consistently under-estimating the distance already. At distances of under 25 m, the linear perspective cues seemed to make no significant difference in the users' performance.

To bring this series of experiments full circle, we needed to return to the ecologically valid task of the MOUT scenario. In our most recent experiment, we made one more important change in the experimental design [Livingston et al(2011)]. We used military standard map icons [DISA(2008)], and applied the drawing styles discovered early in our sequence of experiments to these icons. We compared this to several other techniques for displaying occluded information that had appeared in the AR literature. The opacity and drawing style techniques were not as effective as newer techniques (Fig. 8). A *virtual tunnel* [Bane and Höllerer(2004)] built by drawing virtual holes in known occluding infrastructure led to the lowest error in interpreting the ordinal depth of a virtual squad icon amongst real buildings. The next best technique was one we devised for this study, a *virtual wall* metaphor with the number of edges increasing with ordinal depth. However, both of these techniques led users to perceive the icons as closer than they were intended to be. A *ground grid* technique which drew concentric circles on the ground plane (of the nearly flat experimental environment with the 18 cm deviation) resulted in the signed error that was closest to zero, even though users made more errors in this condition.

We observed other significant effects (not published elsewhere). The type of icon (friendly or enemy) had a significant effect on signed error – $F(1,13)=43.935$, $p=0.000$ – but no significant effect on unsigned error – $F(1,13)=0.584$, $p=0.458$. These results were echoed in our finding of a significant interaction between icon type and occlusion metaphor for signed error – $F(6,78)=7.651$, $p=0.000$ – but not for unsigned error – $F(6,78)=1.536$, $p=0.178$. The cyan friendly icon was interpreted to be closer than it really was, whereas the red enemy icon was interpreted to be slightly farther than it really was. The colors were matched to the MilStd2525c document, but not corrected for the AR display, which pushes the cyan (used for friendly icons) toward the purple while pushing the red (used for enemy) slightly more toward the orange [Livingston et al(2009c)]. This may have caused some change of apparent brightness, which in turn affected the perceived distance from the user.

We asked further for users to estimate (in meters, feet, or yards) the difference in depth between the two icons. A significant main effect of occlusion metaphor on unsigned distance error – $F(6,78)=13.426$, $p=0.000$ – was found, but not for signed error – $F(6,78)=0.555$, $p=0.765$, which raises concern that users got confused about which order implied a negative number should be entered. Users were most accurate with the ground grid, which explicitly shows the absolute egocentric distances. We

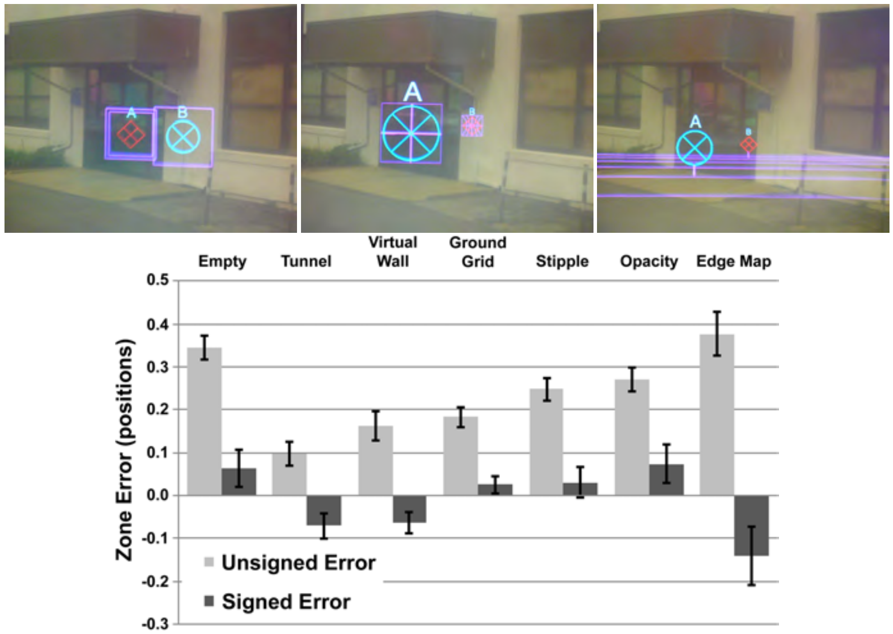

Fig. 8 The new techniques for displaying personnel locations occluded from line-of-sight contact performed well. The top row shows the virtual tunnel (*left*), virtual wall (*center*), and ground grid (*right*) techniques. The graph shows the performance by users in our study. Negative signed error indicates users thought the objects were closer than they were.

also saw a main effect of occlusion metaphor on response time – F(6,78)=2.660, p=0.021. Users were fastest with the empty design (mean response of 2.70 sec), closely followed by the virtual tunnel (2.78 sec). The ground grid (3.71 sec) and virtual wall (3.73 sec) were slower than all methods except an overlaid edge map (which inspired the virtual wall). We did not sub-divide response time into the subtasks of depth for each icon and the distance estimation task, and it is possible that users conceived all three answers before entering any responses for a trial. Finally, we noted a standard practice effect: users were faster with successive blocks of trials (where blocks had a constant occlusion metaphor). We recorded subjective workload responses, but did not find a significant main effect. We did see some slight evidence – F(6,78)=1.832, p=0.104 – for users to feel that (in this order) the virtual tunnel, virtual wall, and ground plane had the lowest workload (measured by NASA TLX [Hart and Staveland(1988)]).

To summarize this discussion, we found good designs for display the information that can help dismounted personnel meet some of their SA needs. This line of research represents how an interesting ecological problem suggested by subject matter experts can spark an interesting scientific question, which can be pursued independently or in the context of that specific application. While even the most recent test

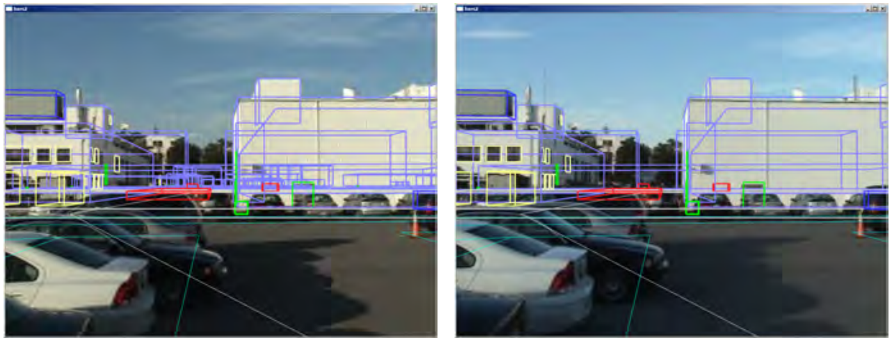

Fig. 9 *Left*: Showing all information and labels in the database can overwhelm the user and prevent a Marine from achieving SA. *Right*: Our information filter uses semantic keys and the concept of area of operations to limit the information shown, which enables the user in this case to discern the location of enemy tanks – spotted by another user – much more easily.

did not provide a final answer, it gave us great insight into directions for how the scientific and ecological questions inform the design of military AR applications.

4.2 Information Filtering

The issue of information overload, as noted above, can become a primary difficulty in MOUT. The physical environment is complex, 3D, and dynamic, with people and vehicles moving throughout. In addition, these entities may be hostile, neutral, or friendly to troops – and even these relationships may change depending on the recent course of events. Thus one may think that more information would be of obvious assistance to the military personnel engaged in such operations. But the amount of information can become too much to process in the dynamic pace of military operations, to the point where it inhibits the ability of personnel to complete their assigned tasks. We have thus developed algorithms for restricting information that is displayed to users.

Based on interviews with subject matter experts over the extended course of the project, our filtering algorithm evolved from a region-based filter [Julier et al(2002)] to a hybrid of the spatial model of interaction [Benford and Fahlén(1993)], rule-based filtering, and the military concept of an area of operations. The resulting algorithm [Livingston et al(2011)] enables sufficient flexibility to update the area of operations, the user's area of interest, the area in which a threat has potential impact, and of course the user's position and orientation in the environment. Objects are displayed when their impact can be felt within the user's area of interest (a spatial calculation) or the rule-based filter determines that the information is vital (Fig. 9).

Compounding the difficulty of having too much information is the issue of how well registered the annotating graphics are to their proper location. As noted above, this underlies the presentation of depth, but it also supports the filtering operation. If graphics can be properly aligned, then the cognitive load to understand the graphics' meaning or information content is reduced, enabling the user to understand and integrate the merged real and virtual image. Thus more information can be shown (assuming a consistent level of registration). We studied techniques to compensate for improper registration [MacIntyre et al(2002)] and how much mis-registration might be acceptable in certain military tasks [Livingston and Ai(2008)].

In the UI architecture, we had to determine how to merge the sometimes competing directives from the components. We began with a simple architecture, proposed a complicated mediation architecture [Julier et al(2003)], but then settled back to a simple pipeline architecture, assisted by incorporating the occlusion representation into the information filter [Livingston et al(2011)].

4.3 Object Selection

In order to query, manipulate, or act upon objects, the user must first select these objects. BARS allows a user to select objects by combining gaze direction (using tracking of the head) with relative pointing within the field of view using a 2D or 3D mouse or eye tracker. The complex nature of the selection operation makes it susceptible to equipment error, scene ambiguities, and user error. Equipment error includes tracking noise, drift, latency, and insufficient resolution for the desired precision. Scene ambiguities arise from the geometric environment, such as when objects overlap in their projections to the current viewpoint. In BARS, with the "X-ray vision" paradigm, these occlusion relationships complicate matters more than many applications. Human error includes imprecision due to lack of experience, poor motor control skills needed for fine-grain selections, and fatigue developed during a session. All these errors lead to selections that were not intended.

To mitigate these errors, we designed a multimodal (speech and gesture) probabilistic selection algorithm [Schmidt et al(2006)]. This algorithm incorporates an object hierarchy (e.g., in the BARS object database, a door is a child of a wall, which is a child of a building, and so on), several gaze and pointing algorithms, and speech recognition. The pointing algorithms rank the objects by distance to the pointing vector, a weighting scheme combining size and distance to the pointing vector, and a weighting scheme combining distance to the view plane and pointing vector. For each pointing selection, the user issues a voice command including the type of object to be selected. The speech recognizer returns a ranked list of candidate objects that it interpreted as what the user intends to select. The object hierarchy, in conjunction with the voice commands, reduces the search space for which object is being selected. The algorithms are combined using a weighted voting scheme to disambiguate any remaining selection discrepancies. We estimated best weighting assignments by evaluating the algorithms through a series of user experiments.

4.4 Collaboration Techniques

From its inception, BARS was envisioned to include a command-and-control (C2) application that would enable a commander to maintain SA of the mobile users in the field; this is also a critical aspect of a military scenario. Commanders naturally have more information about the entire area of operations and situation that subordinates in the field do not have. A C2 application can enable them to get a bird's-eye view, which can be useful for understanding positioning or routes on the ground plane, or they may see the environment from a particular mobile user's vantage. They can direct personnel out of harm's way or to come to the assistance of personnel who encounter a crisis. They can see conflicts between units and perhaps through this capability, reduce friendly fire incidents. We implemented a route-drawing feature in the C2 application, enabling the commander to place waypoints on the ground plane simply by clicking the mouse.

Virtual globe applications can provide an excellent platform for this type of C2 application; we found Google Earth to be suitable due to the 3D building layer and the API that enabled rapid prototyping of environments and an application [Ai and Livingston(2009)]. We simulated having sensors in the environment by merging in live camera views onto this simple 3D terrain. We computed the projection of the camera's image onto known geometry to approximate a live view of the environment (Fig. 10).

Fig. 10 A command-and-control (C2) application might show icons for forces and live sensor (including camera) data over a mixture of satellite imagery and virtual terrain.

Mobile users also need to collaborate directly with each other. From our interviews with domain experts, we learned how military personnel often draw maps in the dirt in order to coordinate a new plan of action. We decided to extend the map view mode of our mobile AR system to incorporate this paradigm. The map view has always been a popular feature of our and others' AR software [Feiner et al(1997)]. Technically, this breaks the first-person paradigm that is often considered a fundamental aspect of AR; we give an essentially 2D view of the world by raising the virtual viewpoint high, directly over the user's position. We then extended the route-drawing feature of the C2 application into the mobile application's map mode, and gave users the ability to communicate these objects through the data distribution system. As with all objects, routes are subject to the filter parameters and rules. We then found that our domain experts reacted more positively to viewing the filter results in map mode than in the head-up AR view. This became a good way to preview the filter results.

4.5 Evaluation of Vehicle AR for IED Awareness

Several important results were identified during the development, testing, and demonstration of the Meissa system. Testing showed that in a moving vehicle environment, the AR display should be considered an enhancement, not a replacement, for the map display. The two-dimensional map provided excellent SA at greater distances. The AR display was most effective at closer distances where the camera could see clearly. Meissa was designed to allow the user to toggle between using the map or AR display as a primary display with a smaller window visible for the alternate view. Also, a significant effort was made to ensure that symbols were consistent between the two views, to allow the operator to switch quickly between views without losing SA.

A somewhat surprising result identified during the operational demonstrations was the enthusiasm the operators had for the Meissa mission recording capability. Operators appreciated the real-time SA, but they were often more interested in the Meissa system's ability to record video, geospatial data, events, and audio annotations. These features allowed for much more effective AAR using the AR technology to overlay aspects of the recorded mission. The geo-registered video also enabled much more accurate extraction of valuable intelligence from missions than the hand-written notes and human recollection used in most military patrols. The mission recording capability of Meissa was not an initial focus of the research and development effort, but based on user feedback this technology became a focus of subsequent development efforts.

4.6 Urban Skills Evaluations

To determine the suitability for mobile AR for urban skills training, we conducted two evaluations. The first evaluation considered the training transfer to teams of novices for the task of clearing multiple rooms [Brown et al(2006)]. The second, expert evaluation, designed using lessons learned from the first, looked at skills improvement of experienced warfighters, Corporals and Lance Corporals with combat experience. Room clearing [USMC(1998)] is a task in which a four-person fire team methodically enters a room and engages any threats in that room. Each member of the team enters the room a certain way (cross, hook, etc) and has a particular area of responsibility within the room.

The novice evaluation used two-person teams; scenarios were designed such that the missing team members would have had no responsiblities in the small rooms in the experiment. The primary independent variable was the training mode: with AR or without AR. Eight teams participated, receiving their training through a video, demonstration of techniques by a subject matter expert, and then practice in their assigned training mode. All participants wore the AR backpack, but only teams assigned to AR training practiced against virtual forces (Fig. 11); the other teams practiced against empty rooms. Teams were encouraged to perform several repetitions of the task in the fifteen minutes allocated to practice. After the instructional period ended, the subjects moved to another part of the test site to be evaluated. Here, participants performed six room-clearing scenarios against real people. Each scenario had enemy and neutral forces in different positions. The subjects and the people playing the enemy and neutral forces traded fire using "laser-tag-style" weapons. This weapon system counts the number of hits on the subjects and on the enemy and neutral forces. The participants wore the AR backpacks solely for tracking and logging the users' actions.

While there were no main effects on objective or subjective performance measures designed by our subject matter expert, we did find an interesting interaction between training method and the number of trials performed. The dependent measure that showed this interaction was visual sweep. Visual sweep is a composite of both speed and effectiveness in the participant's initial entry into the room, and was created based upon recommendations from our SME. It is the angular rotation (based upon the head tracking data) of the participants during their initial 3.5 seconds in each room. The brevity of this time period was enough to capture the subjects' first sweep of the room, but eliminated all motions that they made after finishing their initial sweep, when the subjects would turn and discuss the situation, or turn completely around and leave the room. A test of simple main effects revealed that on trial 1, the AR group had a significantly smaller room sweep compared to the non-AR group. However, on the last trial, the AR group's room sweep was significantly larger than that of the non-AR group. In essence, novice subjects who trained against AR-generated forces were learning to look at more of the room to be cleared than novice subjects who trained against empty rooms.

We simplified the expert evaluation to allow us to collect more detailed data and to make practical the task of performing the evaluation off-site. Subjects were

Fig. 11 Virtual enemy forces populated the training rooms for the urban skills evaluation teams randomly selected to train in the AR condition.

evaluated individually, and the scenarios were set up such that the threat was always in the subject's area of responsibility. This evaluation had three training conditions: with AR, without AR using a live enemy, and without AR using static targets. The purpose of these studies was to measure the usefulness of AR at the application level and to set the stage for future work. This evaluation used a pre-test, a training period, and a post-test for each subject. The test periods ran the subject through six scenarios each. The training period contained 24 scenarios to be completed regardless of time needed rather than a fixed time period as in the first evaluation. In this evaluation, the subjects were experienced warfighters, mainly Corporals and Lance Corporals. This subject pool was more homogeneous than that of the novice evaluation, and due to their experience and expert knowledge in room clearing, we could focus on skills improvement rather than training transfer. The pre- and post-tests were always against live threats regardless of the training condition. As before, the scenarios were designed such that the single user would only face threats in that user's area of responsibility, retaining the spirit of the room clearing doctrine.

We ran 24 subjects through our three training conditions: AR (seven subjects), live (eight), and static targets (nine). We did not see any significant results, and the trends noted in the data seemed to indicate the difficulties our Marine subjects had with the AR system rather than measurements of the effectiveness of the application as a concept. One benefit of testing on enlisted Marines was getting feedback from people who would use a system like this if it were fielded. The most frequent comment was that subjects had difficulty sighting with the AR-based weapon. This feedback reflected the difficulty of achieving precise registration. Other comments focused on the poor visibility through the video-based AR display (which prompted the basic perception experiments described in Section 4.7) and on ergonomic issues. But subjects liked the concept, and the manpower savings that could be realized by replacing human actors with virtual enemy forces in training could be significant. As noted above, this aspect of the BARS project has been carried forward into other programs, currently sponsored by the Office of Naval Research.

4.7 Basic Perception

One of the problems encountered by users in our urban skills evaluation was an extreme difficulty in seeing clearly through the video AR display we selected in order to overcome the occlusion issue noted in Section 3.3. This prompted an investigation into exactly how well users could see through the AR displays. We began to consider several aspects of basic perception in AR displays: contrast sensitivity, color perception, and stereo perception. (Users also noted problems with depth perception, an issue addressed in Section 4.1.) It should be noted that these issues occurred in a controlled laboratory setting, whereas military operations may occur at any time of day or night and will likely involve both indoor and outdoor settings. This increases the range of background lighting conditions under which good visibility of both the real and virtual portions of the AR scene must be maintained.

Visual acuity is perhaps the most basic visual capability that can be measured. An optometrist determines this through a series of recognizable targets (often letters or simple shapes) at a standard contrast value. Contrast sensitivity accounts for the varying size requirements for different levels of contrast. But such a *contrast sensitivity function* for AR has two forms: one can measure the ability of the user to see the graphics presented on the AR display or measure the ability to see the real environment through the AR display. Full details of our analysis to date [Livingston(2006), Livingston et al(2009c)] may summarized as follows (Fig. 12). Some optical see-through displays inhibited the user's ability to see the real world. Some graphical presentations were interpreted at the visual acuity one would expect from the geometric resolution of the display device. The video AR display (which was used in the urban skills evaluation) did not fare well in either measure, owing to the camera resolution and display quality in commercial products of the time. However, even with moderate contrast levels, users were not able to achieve performance that would have corresponded to the maximum visual acuity score for which the stimuli could test; this acuity was in turn below that implied by the pixel resolution of the displays. Thus it is fair to speculate whether poor visual quality of the display devices could be blamed for difficulties in any of the applications or evaluations we conducted.

Color perception can also be a key display property for military applications and a particularly novel hazard for optical see-through displays. Black in a rendering buffer becomes transparent on an optical see-through display, allowing the real world to be seen. So one can easily imagine that dark colors will be perceived improperly. But even bright colors do not fully occlude the real-world background, and

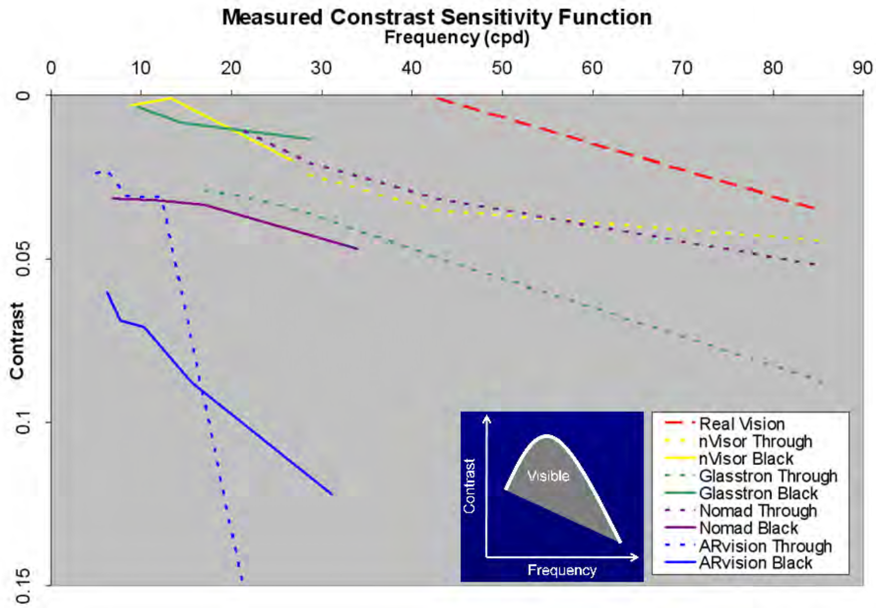

Fig. 12 The measured contrast sensitivity function (CSF) shows that AR displays severely reduce users' visual capabilities relative to their natural vision. The inset next to the legend shows the canonical form of the CSF we have sampled.

thus they too are subject to incorrect perception depending on the color (or mix of colors) that appears behind them in the real environment. We measured the color distortion seen in AR displays, two optical see-through and one video. We found that all three displays distorted colors that were seen on white or black backgrounds, and that this occurred with both graphics presented on the displays and real-world targets seen through the displays. With our particular devices, the video AR display and one of the optical see-through displays were particularly limiting of the color brightness, making all colors appear closer to gray than they were intended to be. Additionally, the video AR display caused high variance in user responses to our color matching task (Fig. 13).

Stereo presentation of graphical images has often been considered a requirement for AR displays. The belief is that in order for the user to perceive graphics as representing 3D objects existing in the surrounding 3D environment, the graphics must be in stereo. One limiting factor in whether a user is able to fuse two images for the left and right eye is vertical alignment. Using nonius lines, we detected errors in alignment ranging from a few hundreths of a degree (well within the tolerance of the human visual system) to four tenths of a degree (an amount that would likely cause eye fatigue or headaches if the user were to force the images to fuse). Simple geometric corrections applied to one eye were sufficient to alleviate these errors [Livingston et al(2006c)]. We then were able to measure the stereo acuity that

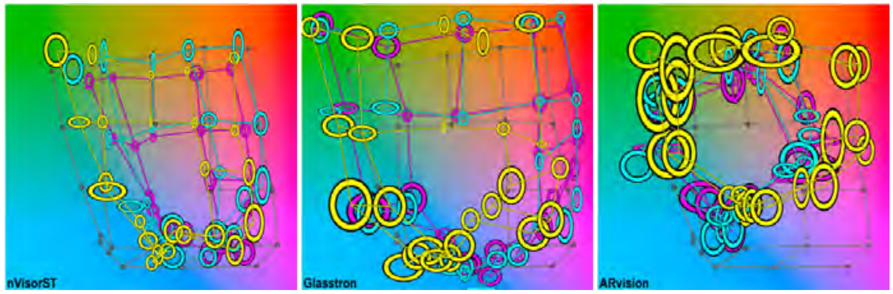

Fig. 13 The perceived color gamut for three AR displays shows the distortion that optical see-through displays nVisorST (*left*) and Glasstron (*center*) as well as video-overlay display ARvision (*right*) cause in the perception of colors, both of the graphical overlays and of real-world objects.

users experience with AR displays, again finding that the differences in depth that a user could detect between real and graphical imagery were well above thresholds in normal human vision of real objects [Livingston et al(2009a)]. This gave us further evidence of the limitations relative to normal human vision caused by a commercial AR display.

5 Challenges of Designing Military AR Applications

We begin our summary of applications of AR to the military with a discussion of the aspects of military operations and personnel that make designing AR applications a challenging problem.

5.1 Mobility

Among the greatest technical challenges for BARS was mobility; this also makes BARS somewhat unique from several other military applications of AR. A pilot or boat captain moves or turns his or her head in a constrained domain. Dismounted and vehicle-mounted personnel have a much freer range of movement relative to fixed infrastructure that is central to their tasks. This requires the tracking system to be usable nearly anywhere on the globe, in theory. To achieve the kind of accuracy in tracking that will lead to proper registration over this range is of course a heavy requirement. In the BARS program, we tested technologies including GPS, magnetic systems, inertial systems, video-metric systems, and hybrids of these without finding a solution that we felt was robust and accurate enough to achieve the kind of results we could get in a laboratory setting; these in turn were barely sufficient to perform controlled versions of the tasks we envisioned as being of potential value to military users of AR. Note that pilots are in particular known for having a high dynamic range in their head orientation; we expect dismounted personnel on a patrol could exhibit similarly sudden movements, further stressing the tracking system.

Another crucial difficulty created by the need for mobility was finding a suitable display device. We tested numerous head-worn displays. As noted above, optical see-through devices have the obvious advantage of permitting the user to take advantage of peripheral vision around the display and natural vision where the graphics are not present. Video overlay displays were judged to be advantageous for training systems because we could completely control the occlusion relationships between real and virtual entities. But to use either type of display outdoors requires that the brightness of the display be sufficiently high to be visible in bright sunlight. At the same time, a display must be usable at night without being visible to a third-party observer who may be hostile to the user; any light emitted by the display becomes a danger to the user. And, as we found, these displays reduce human visual capabilities in several ways.

Another consideration of the mobility requirement is the size, weight, and power requirements of the AR system components. While these metrics have been rapidly improving thanks to the efforts of hardware designers and manufacturers, reducing the amounts of each of these required to drive the AR system will always be a worthy goal for mobile AR. Within the BARS research program, we focused our efforts on other aspects than the hardware requirements, preferring to communicate our requirements to hardware vendors who were interested.

5.2 User Profiles

One challenge that heavily affected our research – especially on the UI for BARS – was the array of tasks and experience that BARS users were expected to have. Even when focusing on a dismounted task of maintaining SA, the type of role the user plays in the organization may differ significantly enough to warrant very different information being presented. A fire team leader (with three subordinates) may need different information about neighboring units than a company commander (typically responsible for approximately 125 subordinates). A medic, supply specialist, forward observer, or interpreter might need a very different set of information about the surrounding environment, people encountered, or the plans for the immediate future (the third aspect of SA).

Another consideration is the background or previous experience with computer graphics that the user is likely to have. There is a significant difference between designing for a dismounted Marine private (typically a male of age 18-20 years) and a field commander with 15 or more years of experience in the military. The private is likely to have grown up with computers and be a frequent player of video games, especially the first-person shooter games that are often similar to training tools for basic strategy and techniques. He is more likely to be comfortable with novel computer interfaces and will be familiar with the standard voice and gesture commands that are used to communicate with friendly forces in the battlefield. Field commanders are more likely to be an experienced leaders and military strategists. Most served as field officers but (for the time being) are less likely to have grown up using computers and mobile electronic devices on a regular basis. They integrate reports from forces in the field, and they resolve conflicts between pieces of information (which can be reasonably easy) and between evaluations of the situation (which were shown to be harder to resolve for pilots and air traffic controllers [Orasanu(2005)]). They may be less comfortable using even traditional computer interfaces but will have familiarity with standard protocols for voice communication with field officers and with superior officers in the C2 center.

5.3 Task Context

One overriding concern for operational applications for military AR systems is that the users will be working in a high-stress environment. This argues for having a simple UI that does not distract the user's attention from the task at hand. By this, we mean that the physical actions required to perform functions in the UI and the ease of understanding the presentation of information must be as simple and intuitive as possible. For the functions, we focused on using the metaphors and commands (voice and gesture) that military personnel currently use to build the UI for BARS. For the understanding, we eventually settled on building on the standard military icons in our most recent implementations, and conducting human factors evaluations as an integral part of the research and development process.

One issue in the UI was that dismounted military personnel generally need their hands free to perform existing military tasks or protocols. While gestural commands may be incorporated as commands to the AR system, interacting with the AR system should intrude on military protocols as little as possible. However, the higher the level of a commander, the less likely that his hands will be occupied with such tasks. Thus, it becomes reasonable that the UI could include more commands that require hands-on operation of the AR system for such personnel.

5.4 Collaboration between Distributed Users

As noted above, one of the most important tasks that was identified as a potential area for BARS to improve the team's SA was to highlight the locations of users who were not within line-of-sight contact. This is a frequently-occuring challenge for MOUT, which are themselves a more frequently-occurring aspect of military operations. Thus we constructed an information database and distribution system and focused the UI visualizations on troop locations and routes. We also built a UI for reporting new or changing information, such as the sighting of an enemy or conflicts between the urban terrain database and the real environment. While this action required a more complicated UI than would be advisable for certain tasks, it is appropriate for forward observers and other personnel.

This dynamic aspect leads to another type of collaboration, that of synchronizing movements and routes. This is also an important aspect of MOUT. Routes and lines of fire for one unit should not intersect those for friendly units without careful consideration of the timing and danger that this implies for both units. During MOUT, available routes can change dramatically and quickly, increasing the difficulty of keeping movements synchronized. Since the dynamic nature of military operations means that it is nearly inevitable that the initial plan will need to be adapted to the situation on the ground, keeping SA of troop locations and movements is a challenging task. AR has the potential to assist dismounted or mounted personnel with this aspect of SA. An important, connected issue to this is length of routes. In one study [Bowman et al(2006)], routes and directions communicated during an operation were generally short, with close contact between units encouraged. When teams must refrain from communications, the duration and distance potentially grow, complicating the task of avoiding conflicts between movements. Control measures such as phase lines, and the progress of other units towards these points of synchronization along respective routes, are another potentially useful visualization that adds to the collaboration between distributed users.

6 Summary

The field of AR has a long and rich history of applications specific to military tasks, dating back to the beginning of the field. In addition, the military has helped to push the technology forward in numerous application areas from medicine to maintenance and repair. AR technology has been demonstrated to have the potential to benefit military applications. However, the military problems are difficult and, for mobile AR systems, more complex than civilian applications. Hard AR research and systems research challenges remain to produce useful prototypes.

It is true that – as with many other applications of AR – military applications have often been limited by the hardware available to system designers and builders. But just as critical to the success of these applications has been the user interface and human subject evaluations (which often reveal limitations of the hardware). BARS focused much effort on these two aspects of AR design and implementation, while trying to shape the interests of hardware vendors to coincide with the demands of the envisioned military applications. One can see similar threads in the other military-specific AR applications discussed here. Indeed, one of the most famous AR applications (at least within the research community) failed not for primarily technical reasons, but for lack of user acceptance [Curtis et al(1998)]. The physical and psychological stress inherent to military tasks increases the importance of these aspects of the AR system.

With all the efforts underway, there is reason to believe that AR will one day find a place in the standard gear of military personnel. It remains to be seen whether AR equipment can make itself critical enough to the individual soldier or Marine to be issued to all personnel, or whether only a unit leader (fire team, squad, company, or battalion) needs to see the augmented common operating picture through the lens of augmented reality.

Acknowledgments

A complex system such as BARS cannot be constructed by even a modest-sized team such as the author list of this text. Numerous employees, interns, and university collaborators have been an integral part of the research over the past decade: Marco Lanzagorta, Doug Maxwell, Steven Feiner, Blaine Bell, Deborah Hix, Joseph L. Gabbard, Tobias Höllerer, Blair MacIntyre, Enylton Coelho, Ulrich Neumann, Suya You, Reinhold Behringer, Catherine Zanbaka, Aaron Kotranza, Robert Dickerson, Jane Barrow, Evan Suma, Brent Daugherty, Dennis Lin, Bryan Hurley, Elliot Cooper-Balis, Adam Lederer, Jason Jerald, Erik Tomlin, Eric Burns, Donald Charity, Joshua Eliason, Jesus Arango, and Scott Frees. In addition, the authors would like to thank Henry Fuchs, Pete Muller, Steven Feiner, and Randy Mayberry for assistance in the preparation of this manuscript.

Current affiliations for the authors no longer affiliated with the Naval Research Laboratory are: Lawrence J. Rosenblum, National Science Foundation; Simon J. Julier, University College London; Dennis G. Brown, USAF Materiel Command, Warner Robins Air Logistics Center; Gregory S. Schmidt, SPADAC, Inc.; Yohan Baillot, Simulation3D, LLC; J. Edward Swan II, Mississippi State University; Paul Maassel, Reallaer, LLC. Interested readers may contact Mark A. Livingston for further information.

References

[Ai and Livingston(2009)] Ai Z, Livingston MA (2009) Integration of georegistered information on a virtual globe. In: IEEE International Symposium on Mixed and Augmented Reality 2009 (Poster Session)

[Azuma(1997)] Azuma RT (1997) A survey of augmented reality. Presence: Teleoperators and Virtual Environments 6(4):355–385

[Baillot et al(2003)] Baillot Y, Julier SJ, Brown D, Livingston MA (2003) A tracker alignment framework for augmented reality. In: IEEE International Symposium on Mixed and Augmented Reality, pp 142–150

[Bane and Höllerer(2004)] Bane R, Höllerer T (2004) Interactive tools for virtual x-ray vision in mobile augmented reality. In: International Symposium on Mixed and Augmented Reality, pp 231–239

[Barham et al(2002)] Barham P, Plamondon B, Dumanoir P, Garrity P (2002) VICTER: An embedded virtual simulation system for land warrior. In: 23rd Army Science Conference

[Benford and Fahlén(1993)] Benford S, Fahlén L (1993) A spatial model of interaction in large virtual environments. In: Third European Conference on Computer-Supported Cooperative Work, pp 109–124

[Bolstad and Endsley(2002)] Bolstad CA, Endsley MR (2002) Tools for supporting team sa and collaboration in army operations. Tech. rep., SA Technologies

[Bowman et al(2006)] Bowman E, Cosenzo K, Hill SG, Grynovicki J, Branscome TA, Brelsford M, Savage-Knepshield P (2006) C4ISR on-the-move testbed 2005 experiment: Human factors analysis. Tech. rep., Army Research Laboratory, Human Research and Engineering Directorate

[Brown et al(2004a)] Brown DG, Baillot Y, Julier SJ, Maassel P, Armoza D, Livingston MA, Rosenblum LJ (2004a) Building a mobile augmented reality system for embedded training: Lessons learned. In: Interservice/Industry Training, Simulation, and Education Conference

[Brown et al(2004b)] Brown DG, Julier SJ, Baillot Y, Livingston MA, Rosenblum LJ (2004b) Event-based data distribution for mobile augmented reality and virtual environments. Presence: Teleoperators and Virtual Environments 13(2):211–221

[Brown et al(2005)] Brown DG, Baillot Y, Bailey MP, Pfluger KC, Maassel P, Thomas J, Julier SJ (2005) Using augmented reality to enhance fire support team training. In: Interservice/Industry Training, Simulation, and Education Conference

[Brown et al(2006)] Brown DG, Stripling R, Coyne JT (2006) Augmented reality for urban skills training. In: IEEE Virtual Reality Conference

[Cox(2008)] Cox M (2008) Land Warrior: Now or later? Army Times Retrieved from `http://www.armytimes.com/news/2008/10/army_landwarrior_101308w/`, 22 December 2010

[Cubic(2011)] Cubic Corp (2011) COMBATREDI®. Retrieved 12 January 2011 from `http://www.cubic.com/cda1/prod_&_serv/defense/training_products/virtual_training/combatredi.html`

[Curtis et al(1998)] Curtis D, Mizell D, Gruenbaum P, Janin A (1998) Several devils in the details: Making and AR application work in the airplane factory. In: Augmented Reality: Placing Artificial Objects in Real Scenes (Proceedings of the First International Workshop on Augmented Reality

[DARPA(2010)] Defense Advanced Research Projects Agency (2010) Urban leader tactical response, awareness and visualization. `http://www.darpa.mil/i2o/programs/uvis/uvis.asp`, retrieved 22 December 2010

[Army(2001)] Department of the Army (2001) Operational requirements document (revised) for land warrior

[DISA(2008)] DISA Standards Management Branch (2008) Department of Defense interface standard: Common warfighting symbology. http://assist.daps.dla.mil/quicksearch/basic_profile.cfm?ident_number=114934

[USMC(1998)] Marine Corps Doctrine Division (1998) Military operations on urbanized terrain (MOUT). Marine Corps Warfighting Publication No. MCWP 3-35.3

[Feiner et al(1997)] Feiner S, MacIntyre B, Höllerer T, Webster A (1997) A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment. In: First International Symposium on Wearable Computing (ISWC), pp 74–81

[Furness(1969)] Furness LT TA (1969) The application of head-mounted displays to airborne reconnaissance and weapon delivery. Tech. Rep. TR-69-241, U.S. Air Force Avionics Laboratory, Wright-Patterson AFB

[Furness(1986)] Furness TA (1986) The super cockpit and its human factors challenges. In: Proceedings of the Human Factors Society 30th Annual Meeting, Human Factors and Ergonomics Society, pp 48–52

[Gabbard et al(2002)] Gabbard JL, Swan II JE, Hix D, Lanzagorta M, Livingston MA, Brown D, Julier S (2002) Usability engineering: Domain analysis activities for augmented reality systems. In: Woods AJ, Bolas MT, Merritt JO, Benton SA (eds) Proceedings of SPIE: Stereoscopic Displays and Virtual Reality Systems IX, Vol. 4660, pp 445–457

[Gardley et al(2008)] Gardley M, Bishop L, McCollon B, Clemons L, Mayberry R, Kohler T (2008) Loadmaster training utilizing augmented reality. Tech. rep., Air Education and Training Command, Randolph AFB

[GD(2010)] General Dynamics Corp (2010) Virtual warrior. Retrieved 02 December 2010 from http://www.gdc4s.com/content/detail.cfm?item=d3ae855e-74ab-4af0-8c56-abad8be8ea0f&page=1&ParentPageID=177d49b8-73c6-479f-94e8-f79826e1d422

[Gould(2010)] Gould J (2010) Infantry battalion tests nett warrior. Army Times Retrieved from http://www.armytimes.com/news/2010/10/infantry-battalion-tests-nett-warrior-103010w/, 22 December 2010

[Gumm et al(1998)] Gumm M, Marshakm W, Branscome T, McWesler M, Patton D, Mullins L (1998) A comparison of solider performance using current land navigation equipment with information integrated on a helmet-mounted display. Tech. Rep. ARL-TR-1604, Army Research Laboratory, Aberdeen Proving Ground

[Hart and Staveland(1988)] Hart SG, Staveland LE (1988) Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research, North Holland Press, pp 239–250

[Henderson and Feiner(2011)] Henderson SJ, Feiner S (2011) Exploring the benefits of augmented reality documentation for maintenance and repair. *To appear in* IEEE Transactions on Visualization and Computer Graphics DOI \url{http://doi.ieeecomputersociety.org/10.1109/TVCG.2010.245}

[Hughes et al(2002)] Hughes CE, Stapleton CB, Moshell JM, Micikevicius P (2002) Challenges and opportunities simulating future combat systems via mixed reality. In: 23rd Army Science Conference

[Julier et al(2002)] Julier S, Lanzagorta M, Baillot Y, Brown D (2002) Projects in VR: Information filtering for mobile augmented reality. IEEE Computer Graphics & Applications 22(5):12–15

[Julier et al(2003)] Julier S, Livingston MA, Swan II JE, Baillot Y, Brown D (2003) Adaptive user interfaces in augmented reality. In: Proceedings of Workshop on Software Technology for Augmented Reality Systems at IEEE International Symposium on Mixed and Augmented Reality

[Livingston(2006)] Livingston MA (2006) Quantification of visual capabilities using augmented reality displays. In: IEEE International Symposium on Mixed and Augmented Reality, pp 3–12

[Livingston and Ai(2008)] Livingston MA, Ai Z (2008) The effect of registration error on tracking distant augmented objects. In: International Symposium on Mixed and Augmented Reality

[Livingston et al(2002)] Livingston MA, Rosenblum LJ, Julier SJ, Brown D, Baillot Y, Swan II JE, Gabbard JL, Hix D (2002) An augmented reality system for military operations in urban terrain. In: Interservice/Industry Training, Simulation, and Education Conference

[Livingston et al(2003)] Livingston MA, Swan II JE, Gabbard JL, Höllerer TH, Hix D, Julier SJ, Baillot Y, Brown D (2003) Resolving multiple occluded layers in augmented reality. In: IEEE International Symposium on Mixed and Augmented Reality, pp 56–65

[Livingston et al(2004)] Livingston MA, Swan II JE, Julier SJ, Baillot Y, Brown D, Rosenblum LJ, Gabbard JL, Höllerer TH, Hix D (2004) Evaluating system capabilities and user performance in the battlefield augmented reality system. In: Performance Metrics for Intelligent Systems Workshop (PerMIS), National Institute of Standards and Technology

[Livingston et al(2005)] Livingston MA, Zanbaka C, Swan II JE, Smallman HS (2005) Objective measures for the effectiveness of augmented reality. In: IEEE Virtual Reality, pp 287–288

[Livingston et al(2006a)] Livingston MA, Brown D, Julier SJ, Schmidt GS (2006a) Mobile augmented reality: Applications and human factors evaluations. In: Virtual Media for Military Applications (RTO Meeting Proceedings No. RTO-MP-HFM-136), Research and Technology Organisation, Neuilly-sur-Seine, France, pp 25–1—25–16

[Livingston et al(2006b)] Livingston MA, Julier SJ, Brown D (2006b) Situation awareness for teams of dismounted warfighters and unmanned vehicles. In: Verly JG, Guell JJ (eds) Proceedings of SPIE: Enhanced and Synthetic Vision, Vol. 6266, pp 62,260F.1–62,660F.12

[Livingston et al(2006c)] Livingston MA, Lederer A, Ellis SR, White SM, Feiner SK (2006c) Vertical vergence calibration for augmented reality displays. In: IEEE Virtual Reality (Poster Session)

[Livingston et al(2009a)] Livingston MA, Ai Z, Decker J (2009a) A user study towards understanding stereo perception in head-worn augmented reality displays. In: IEEE International Symposium on Mixed and Augmented Reality, pp 53–56

[Livingston et al(2009b)] Livingston MA, Ai Z, Swan II JE, Smallman HS (2009b) Indoor vs. outdoor depth perception for mobile augmented reality. In: IEEE Virtual Reality

[Livingston et al(2009c)] Livingston MA, Barrow JH, Sibley CM (2009c) Quantification of contrast sensitivity and color perception using head-worn augmented reality displays. In: IEEE Virtual Reality

[Livingston et al(2011)] Livingston MA, Ai Z, Karsch K, Gibson, MSC, USN LT GO (2011) User interface design for military ar applications. Virtual Reality (UK)

[MacIntyre et al(2002)] MacIntyre B, Coelho EM, Julier S (2002) Estimating and adapting to registration errors in augmented reality systems. In: IEEE Virtual Reality, pp 73–80

[Mohl(1981)] Mohl R (1981) Cognitive space in the interactive movie map: An investigation of spatial learning in virtual environments. PhD thesis, Massachusetts Institute of Technology

[Muller(2010)] Muller P (2010) The future immersive training environment (FITE) JCTD: Improving readiness through innovation. In: Interservice/Industry Training, Simulation, and Education Conference

[Naimark(1979)] Naimark M (1979) Spatial correspondence – a study in environmental media. Master's thesis, Massachusetts Institute of Technology

[ONR(2010)] Office of Naval Research (2010) Human performance training and education thrust. http://www.onr.navy.mil/Science-Technology/Departments/Code-30/All-Programs/Human-Performance-Training.aspx, retrieved 22 December 2010

[Orasanu(2005)] Orasanu J (2005) What neeeds to be shared in shared situation models? In: DoD Human Factors Engineering Technical Advisory Group Meeting 54

[Quantum3D(2010)] Quantum3D, Inc (2010) Expeditiondi®. Retrieved 02 December 2010 from http://www.quantum3d.com/solutions/immersive/expedition_di.html

[Schmidt et al(2006)] Schmidt GS, Brown DG, Tomlin EB, Swan II JE, Baillot Y (2006) Probabilistic algorithms, integration, and empirical evaluation for disambiguating multiple selections in frustum-based pointing. Journal of Multimedia 1(3):1–12

[Schmorrow et al(2004)] Schmorrow C, Cohn L, Stripling R, Kruse A (2004) Enhancing virtual environments using sensory-multiplexing. In: Interservice/Industry Training, Simulation, and Education Conference

[State et al(1996)] State A, Livingston MA, Hirota G, Garrett WF, Whitton MC, Pisano ED, Fuchs H (1996) Technologies for augmented reality systems: Realizing ultrasound-guided needle biopsies. In: SIGGRAPH '96: Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, pp 439–446

[Stix(1992)] Stix G (1992) See-through view: Virtual reality may guide physician's hands. Scientific American 267(3):166

[Sutherland(1965)] Sutherland IE (1965) The ultimate display. In: Proceedings of the IFIP Conference, pp 506–508

[Sutherland(1968)] Sutherland IE (1968) A head-mounted three-dimensional display. In: AFIPS'68: Fall Joint Computer Conference, pp 757–764

[Swan II et al(2006)] Swan II JE, Livingston MA, Smallman HS, Brown D, Baillot Y, Gabbard JL, Hix D (2006) A perceptual matching technique for depth judgments in optical, see-through augmented reality. In: IEEE Virtual Reality, pp 19–26

[Van Riper(1997)] Van Riper PK (1997) A concept for future military operations on urbanized terrain. Available from http://www.globalsecurity.org/military/library/policy/usmc/mout.pdf