

Mobile Augmented Reality: Applications and Human Factors Evaluations

Mark A. Livingston, Ph.D., Dennis G. Brown, Simon J. Julier, Ph.D., and Greg S. Schmidt, Ph.D.

Advanced Information Technology Code 5580

Naval Research Laboratory

4555 Overlook Ave SW

Washington, DC 20375

mark.livingston@nrl.navy.mil

ABSTRACT

Recent trends in military operations (quick-reaction forces, putting fewer warfighters at risk, and increasing the use of unmanned vehicles) have increased the difficulty in acquiring and maintaining situation awareness (SA). Augmented reality (AR) has the potential to meet some of these new challenges. AR systems integrate computer-generated graphics (or annotations) with the user's view of the real world. These annotations can be cues to establish and maintain SA, or they can provide virtual opposing forces (OPFOR) for training scenarios. However, the design of the user interface of a mobile AR system presents a unique set of technical challenges. The interface must be capable of automatically deciding what annotations need to be shown. Furthermore, it must select the characteristics of those annotations (including appearance, size, and drawing style) to ensure the display is intuitive and unambiguous. In the training applications, the virtual OPFOR must appear and behave realistically. We discuss the development of our augmented reality system and the human factors testing we have performed. We apply the system to two military needs: situation awareness during operations and training.

1.0 INTRODUCTION

The trends in military operations towards quick-reaction forces, putting fewer warfighters at risk, and increasing the use of unmanned vehicles combine to increase the information requirements for an individual in the battlespace. As units become more dispersed and specialized, acquiring and maintaining situation awareness (SA) becomes harder. The predictive aspect of SA becomes especially difficult in urban operations, where line of sight contact with even friendly forces is unlikely to be maintained for long periods of time. In principle, some of these difficulties can be overcome through the use of a display that can automatically organize and present information to the user. One promising approach is augmented reality (AR) [1]. An AR system mixes computer-generated graphics (or annotations) with the real world. The annotations can provide information aimed at establishing situation awareness or to provide realistic training for such scenarios. The design of the user interface of a mobile AR system presents a unique set of technical challenges. An AR display must be capable of automatically deciding what annotations need to be shown. Furthermore, the system must select characteristics of those annotations (such as appearance, size and drawing style) that ensure the display is intuitive and unambiguous. In the training applications, the virtual opposing forces (OPFOR) must appear and behave realistically.

Providing such a display is the goal of the Battlefield Augmented Reality System (BARS) project [8] at the U.S. Naval Research Laboratory (NRL). We have developed methods to filter the important virtual cues, represent objects or forces hidden by the 3D terrain, and draw cues in semantically meaningful ways. We

have incorporated a network interface to allow fully distributed, multi-user operations. Section 2 describes the general system implementation. In a separate application, we bridge our system to a semi-automated forces (SAF) system and create a virtual training tool. With this implementation, described in Section 3, the user can train with or against SAF agents.

One important aspect of ensuring the utility of such systems is evaluation through user studies. We have conducted a series of user studies that examined detailed perceptual effects of representations as well as performance on military tasks. We began our exploration of the human factors issues by talking to domain experts in one expected operational use of AR, military operations in urban terrain (MOUT). We summarize results from a variety of user studies inspired by this problem domain in Section 4. We conclude with a discussion of future directions and important open issues.

2.0 SYSTEM DESCRIPTION

BARS is a *mobile augmented reality* system [5], consisting of a computer, a tracking system, and a see-through wearable display (Figure 1). The system tracks the position and orientation of the user's head and superimposes graphics and annotations that are aligned with real objects in the user's field of view. With this approach, complex 3D spatial information can be directly aligned with the environment. For example, the name of a building could appear as a "virtual sign post" attached directly to the side of the building. BARS networks multiple outdoor, mobile users together with a command center.



Figure 1: A prototype implementation of BARS using commercially available components. The wearable computer produces graphics seen in the display. Audio and wireless hand-held devices are used to interact with the system.

2.1 Hardware Implementations

Built from commercial, off-the-shelf (COTS) products, the mobile prototypes for BARS are composed of a computer with an advanced graphics processor, a see-through display, and a number of interaction devices. Our current implementations use either a Sony Glasstron LDI-D100B personal display or a Microvision

Nomad display. Both of these displays are optical see through devices; the user always sees the real world directly, even if the system is switched off. The Glasstron permits color and bi-ocular imagery presented with disparity between the eyes. The Nomad displays a monocular image, but its retinal scanning laser display offers better brightness against the real background, which is useful in bright outdoor environments. Our early prototype used a PC104-based microcomputer with a configurable graphics processor, enabling us to upgrade as new products became available. While upgrades are still frequent, we have switched to a more robust Quantum 3D Thermanite computer, which has the graphics processor embedded.

In order for the rendering system to draw the graphics with the proper perspective, the system must track the user's position and orientation in the world. We currently use an Ashtech GG24-Surveyor with real-time kinematic and differential GPS for tracking the position and an Intersense InertiaCube2 for tracking the user's orientation. We have tested experimental software for videometric tracking of landmarks in the environment but, for robustness reasons, we currently do not use these implementations.

We have a variety of methods to interact with the system. One method is through voice commands over standard audio hardware connected to the PC. Another is with mouse devices; we have used touch-pad mice and a Gyro-Mouse, which measures tilt on two axes to control the two linear dimensions on the screen.

In addition to the hardware, BARS encompasses a number of software systems to perform a variety of functions in presenting information to the user or interpreting user commands. The following subsections describe these components.

2.2 Information Filter

A BARS system contains a 3D model of the environment in which an operation is to occur. Such a model might be obtained from any of a number of intelligence sources. We envision BARS will also have mission plans, such as objectives, landing and extraction zones, proposed routes, or tactical information such as enemy locations or patterns that might prove useful. Currently, we draw routes in real-time from a command center application. Enemy locations may be highlighted from the command center or by BARS users. Since any object in the database may be shared, this information can instantly be passed to all users who need to know. The database also enables semantic tags such as relevance to a task, threat level, or timeliness of the data.



Figure 2: Two augmented views of a building: the first shows all information available, and the second shows only a route and an item of interest.

The shared database contains much information about the local environment. Showing all of this information can lead to a cluttered and confusing display. We use an information filter (Figure 2) to add objects to or remove objects from the user's display [6]. We use a *spatial filter* to show only those objects that lie in a

certain zone around the user. This zone can be visualized as a cylinder whose main axis is perpendicular to the ground plane. Objects within the cylinder's walls are shown, and the user can vary the inner and outer diameters of the cylinder walls. We also use *semantic filters* based on the user's task or orders from a commander. For example, a route associated with a task will be shown regardless of the user's spatial filter settings, and threats will be shown at all times.

2.3 Representations of Depth

One important problem in urban operations is that of *troop location*, knowing where friendly forces are within the environment. Since the urban environment often breaks line-of-sight contact and maintaining radio silence is often required, it can be difficult to always know where friendly forces are. This prompted us to develop a set of representations of depth information [9]. Drawing inspiration from methods used in technical illustrations, we use graphical parameters, such as stipple effects (dashed or dotted lines or filled shapes) or opacity to vary representations based on the distance to those objects. Figure 3 illustrates some examples of showing building locations. In each case, the colored building lies *behind* the visible buildings and cannot be directly seen.



Figure 3: Candidate representations of occluded terrain and forces in urban environments.

These candidate representations show ordinal depth information. In Section 3, we will discuss human factors experiments that we extended to include metric depth matching. When such information is presented in AR, this creates a metaphor of “x-ray vision” to allow users to see spatial information that may be occluded by real or graphical objects. This is an unnatural percept and has proven difficult to provide in an intuitive manner. It also leads to difficulties in interaction.

2.4 Interaction Methods

We expect that a BARS user will want to specify objects in the environment for such purposes as identifying landmarks for other users (Figure 4), retrieving more detailed information, or modifying the database to reflect changes in the environment. While there are many ways to specify objects or locations, pointing is a common and natural method. Pointing may be performed using a range of devices: a hand-held mouse or head orientation tracker indicating the position in the field of view, a 3D tracking device encircling an object, or an eye tracker measuring gaze direction. Selections may also be performed by sketching over or circling an object, and then using the object which has the largest intersection as the choice. Authoring new objects or annotations uses similar methods of specifying features, or may use a menu of pre-defined objects.

We assert that all pointing-based selection or drawing operations are susceptible to error. Human error comes from lack of experience, poor motor control during fine-grained pointing, or fatigue developed during a

session. Equipment error could be noise, drift, or latency in a position and orientation tracking system, or insufficient resolution on a wheel-based device to perform fine selections. Finally, there are ambiguities associated with the scene itself, such as when the user tries to select one object occluded by another object. In mobile AR, this arises from the “X-ray vision” metaphor. These errors can lead to selections that are incorrect or to imperfections in the shape or placement of primitives authored into the environment.

For BARS, we designed a pointing-based probabilistic selection algorithm that alleviates some of the error in user pointing-based selections [13]. The algorithm generates lists of candidate objects the user may have meant to select and probability estimates of how likely it is the user meant to select each object. The algorithm combines three low-level intersection algorithms and the hierarchical structure of the dataset (e.g., a door is in a wall, which is part of a building, and so on), and then integrates the resulting candidates. The three low-level intersection algorithms have differing utility depending on the user's preferences for making selections, on what type of object the user is trying to select, and on its relationship to other objects in the scene. The preferences for the three algorithms are: (1) select the item nearest the central pointing ray; (2) select the largest item in the viewing frustum; and (3) select using a filtering approach that weights the objects by applying a Gaussian function based on how far away they are from the center of the viewing frustum. These algorithms are run in parallel and their probabilistic outputs are fused using several weighting schemes. The combined selection algorithm works effectively at disambiguating multiple selections.



Figure 4: The result of a user drawing in the world. If another user is to interpret these locations, the drawing mechanism must be sufficiently precise to make the annotations unambiguous.

2.5 Collaboration Mechanism

The BARS collaboration system [2] shares relevant parts of the database with each networked machine. Figure 5 shows how this functionality is key to providing multiple mobile users a common set of information, as one user can see another user's position and current path, updated in real time. The fundamental design is an abstraction of the IP multicast standard. Some implementations do use IP multicast, however, other networking methods are used to transport. Information is deemed relevant to a particular user based on the information filter described previously. Based on the importance of the data, communications use reliable and unreliable transport mechanisms in order to keep network traffic low. For example, under optimal conditions,

user positions are updated in real time (at least 30 Hz) using unreliable transport, but with a frequency of around 5 Hz, user positions are sent reliably so that those with overloaded connections will at least get positions at a usable rate.




Figure 5: A remote BARS user is highlighted in the user interface. His route and an occluded building are also depicted. Text in the bottom center shows position and orientation data, while text in the bottom left shows status information.

A *channel* contains a class of objects and distributes information about those objects to members of the channel. Some channels are based on physical areas, and as the user moves through the environment or modifies parameters of his spatial filter, the system automatically joins or leaves those channels. Other channels are based on semantic information, such as route information only applicable to one set of users, or phase lines only applicable to another set of users. In this case, the user voluntarily joins the channel containing that information, or a commander can join that user to the channel. Figure 6 shows how multiple units share a single common database on the left, and on the right, shows how the system was extended to support multiple channels of data.

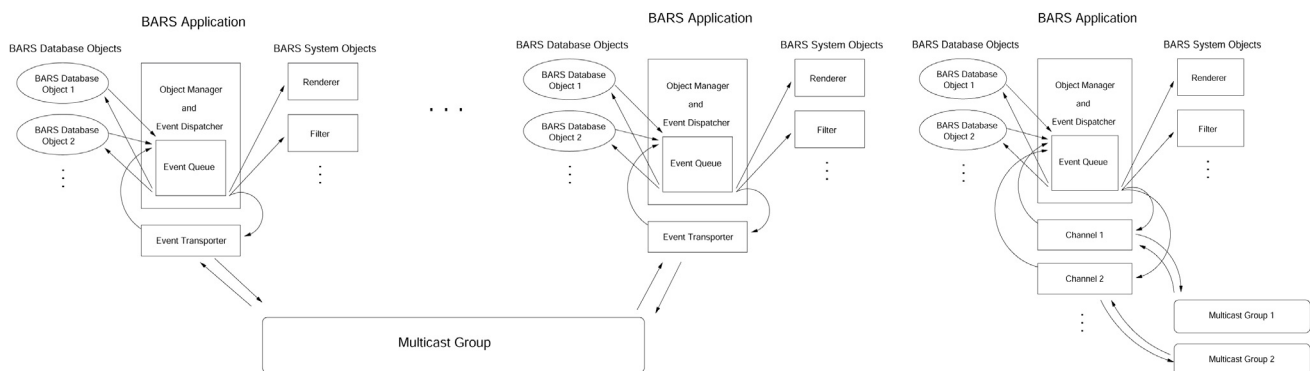


Figure 6: The BARS distribution system. The left image shows how multiple applications share a common database over a single multicast group. The right image shows how the system has been extended to support several channels of data.

3.0 TRAINING APPLICATIONS

Although BARS was originally designed for providing situation awareness during operations, its components can be reused for training in real environments by augmenting the real world with simulated forces and other factors [3]. BARS works for embedded MOUT training as follows:

1. Simulated forces are rendered on the display, so as the user looks around the real MOUT facility, forces appear to exist in the real world (within current graphics limitations) even though they do not truly exist. At the same time, fellow real trainees remain visible.
2. Spatialized audio is sent through the headphones to replicate the aural cues that the simulated forces would make if they were real. These sounds include footsteps, shouting, helicopters, and so on. Since the sound is spatialized, the user can determine the location of the simulated force by listening, like in the real world.
3. Interaction with the simulated forces is very limited at this time. Real and virtual forces can shoot at each other.
4. Simulated forces are controlled through various means and are distributed to the trainees using the BARS distribution system.

There are several technical challenges to this task, even with all of the work already completed for BARS.

3.1 Interaction Methods

The simulated forces need to appear on the user's display to give the illusion that they exist in the real world (Figure 7). There are several inherent problems: model fidelity, lighting to match the real environment, and occlusion by real objects.




Figure 7. Two simulated forces in an office environment.

Model fidelity is controlled by the modeler and is limited by the power of the machine running the application. Although models that can be rendered in real time still look computer generated, just like in VR-based simulations, the limited AR model representation capabilities are adequately realistic for embedded simulation and training. AR actually has an advantage over VR with respect to rendering: the AR graphics

system does not need to draw an entire virtual world, only the augmented forces, so they could potentially be more detailed than those in VR-based simulations.

Lighting the rendered forces is a problem our team has not approached yet. This task would require knowing the lighting conditions of the real environment in which the model would appear, and changing the renderer's light model to match. Another limitation is the display itself, as it is very sensitive to outside light, and even if the image is rendered with perfect lighting, it still might not appear correctly on the display.

The problem of occlusion of simulated objects by real objects, more than lighting or model complexity, is the one that would most likely ruin the immersion of training using AR. Imagine using an AR training system and seeing a simulated force, which is supposed to be behind a building, rendered in front of the building. This property is actually a feature of BARS—it gives the user a way to see through walls. However, today's dismounted warriors cannot see through walls, and so in the AR-based trainer, they should not see simulated forces that should be occluded by real objects.

Solving the occlusion problem first requires creating a model of the training environment [7]. In the AR system for operations, it is known where the user is looking and the system can draw an augmenting model of buildings and features superimposed on the real features. In AR for training, this same model is rendered in flat black. Since the graphics processor compares the depth values, these black features will occlude the parts of the simulated forces the user should not see. However, since black is the “see through” color on the AR display, the user will still see the real world, along with the correct non-occluded parts of the simulated forces. This solution was introduced for indoor applications [14] and applied to outdoor models [12] for use in outdoor AR gaming. Figure 8 shows a sequence of images demonstrating this technique. Figure 8A shows the real-world scene with no augmentation. In Figure 8B, the same scene is shown but with simulated forces simply drawn over the scene at their locations in the world—there is no occlusion. It is hard to tell if all of the forces are intended to be in front of the building, or if they are just drawn there due to limitations of the system. Figure 8C shows the simulated forces occluded by a gray model, however, the model also occludes some of the real world. Finally, Figure 8D shows the scene rendered using a black model, which occludes the simulated forces properly and allows the user to see the real world.




Figure 8. Stages in the development of AR models for embedded training—see explanation in text.

3.2 Inserting Aural Cues

Since the system already has a 3D world model, and the locations of the user and the simulated forces are known, existing 3D sound libraries are used to provide spatialized audio. Sound streams are simply attached to simulated forces and the audio library is updated with the positions of those forces and with the user’s listening attitude. Open-air headphones naturally mix the sounds of the real world with the computer-generated sounds.

3.3 Interacting With Simulated Forces

The simulated forces can be controlled in several ways including simple animation scripts. However, the animations are not reactive and tend to create a simple “shooting gallery” type of simulation. They can also be controlled by users of immersive VR simulations that participate on the same network as the AR user. Finally, they can be controlled through Semi-Automated Forces (SAF) systems.

BARS communicates with outside information systems using bridge applications, as described in the previous section. By creating a bridge application between BARS and a SAF system, the years of work already put into simulating forces for both non-immersive and immersive VR-based training can be leveraged, and the user interact with those forces in a real training environment.

Figure 9 shows a set of BARS applications for an embedded training scenario: two trainees using wearable systems, a trainee using an immersive VR system, an observer using a VR system, and a bridge synchronizing the entities in BARS and a connected SAF system. The bridge converts SAF entities into BARS entities and vice-versa. It keeps those entities updated on each side of the bridge as they change by converting BARS events into DIS or HLA packets and vice-versa. The bridge is not a simple filter for converting these events; it must maintain internal state information in order to convert the events and packets properly. In addition to sharing entity information, the system allows BARS users to engage the simulated forces and allows the simulated forces to retaliate.

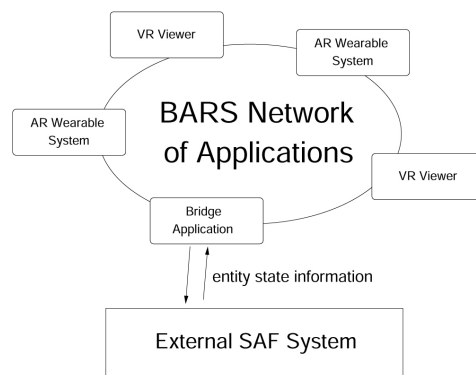


Figure 9. BARS and an external SAF system share information using a bridge application.

4.0 HUMAN FACTORS TESTS

We have adopted a layered concept for our human factors testing. The most basic layer is the *perceptual layer*, in which tasks are abstract and not connected to a particular military task. The next layer up consists of

basic cognitive functions such as prediction of events or decision-making. The highest layer is that of tasks in which we expect our system to assist a user; this is essentially a field-test. We do not follow a strict order for human factors tests, but perform tests as the need for understanding arises in our evaluations. The first test we conducted was a cognitive test. The results of this test indicated a need for perceptual tests. We have conducted a number of studies at the perceptual level. When the system is deemed sufficiently mature for a field test on a particular task, we will conduct such an evaluation. We have not done so yet. Note that even such tests may result in the need for further tests at lower levels and may not give insight into the system's performance on other tasks.

4.1 Situation Awareness Evaluation

Among the important difficult problems in MOUT that experts identified was that of *troop location*, knowing where friendly forces are at all times during an operation. In complex urban environments, people are easily hidden within or behind buildings, and tunnels hide infrastructure such as subways or electrical conduits.

We applied user interface principles to create visual representations of occluded objects, focusing on vehicles and small teams of people. (Candidate representations appear in Figure 3.) We then designed an evaluation of these representations using questions that tested users' SA. Questions included identifying which of several objects – people, vehicles, or buildings hidden within the urban canyon – was closer to the subject's location, relative distances between remote objects, absolute distance to a remote object (using a legend in the display), or the heading of a moving object. We used two classes of subjects, user interface experts and active-duty Marines (Figure 10). The overall result of this test was that the representations were successful, with subjects answering approximately 85% of the questions correctly. More importantly, we found that subjects were generally able to interpret the representations as they were intended.



Figure 10: A subject in our first study of situation awareness gestures while answering questions.

4.2 Depth Perception Studies

The first test revealed which questions subjects struggled to answer; many of these questions required subjects to understand depth relationships between real and graphical objects in the field of view. While an overhead map view can provide clear answers to this type of question in most situations, it is not ideal for a 3D environment and requires the user to switch context from the real world. We hoped to provide visualizations

that would enable the user to understand the relative depth of real and graphical objects. Our initial designs took advantage of graphical parameters, as described in Section 2.3. We wanted to study the relative importance of the various parameters and see what constituted appropriate values in different situations. Thus our second test focused on the issue of relative depth among graphical objects, all of which were hidden behind a real building [9]. Users were asked to identify whether a red target building was in front, between, or behind two blue buildings. All buildings were at virtual distances that were behind a real building, and the users were told that the graphical buildings were behind this real building.

The test found that wire-frame outlines of buildings were not as effective at conveying depth as filled shapes with wire-frame outlines. This was an expected result, but important to quantify since we had up to that point been primarily relying on wire-frame outlines to show objects. We found that the opacity parameter offered in the color specification on modern graphics processors was effective at conveying the depth of an object. This corresponds to the atmospheric effect for human vision, in which colors fade with increasing distance. Most interestingly, we found that the combination of the drawing style (filled shapes with wire-frame outlines) and approximated atmospheric effect was statistically equivalent at conveying depth as the perspective constraint provided by a flat ground plane.

In a follow-up study [15], we had subjects place a graphical object at the depth of a real object. We gave the user control of the virtual distance with a trackball and placed several real targets, differentiated by color, in a 50-meter hallway (Figure 11). This task was structured such that subjects had to attend to both the real and graphical objects simultaneously, a flaw in our first experiment. We found that although the task appears to be solvable through the use of only two-dimensional cues such as relative size, subjects appear to experience depth in a manner consistent with 3D depth perception. We thus hope to use this design to build towards future studies of depth perception between real and graphical objects. A variation on this experiment tested users' depth matching with a graphical target against their matching with a real target (Figure 12). We found that users had approximately equivalent abilities, further validating our belief that users are able to understand the graphical objects as if they were present in the environment.




Figure 11: Subjects in the depth matching experiment were asked to place the graphical target at the same distance as one of the rectangular referents. The task was done with the target both higher and lower in the field of view than the referents.



Figure 12: The real target used for comparison with placing the graphical target at the depth of real referents.

4.3 Basic Perception Studies

We performed two tests of basic perception with the system, in order to verify that subjects could reasonably be expected to resolve objects and properly verge the image pairs. These two functions are so fundamental to looking through the see-through display that nothing else in the system works perceptually if either of these two tasks can not be performed by the user. To test the first, we encoded a standard Snellen eye chart in the display and had subjects read both real and graphical eye charts, just as an optometrist would have a patient read an eye chart. The optics were severely hampering users' ability to read the letters [10]. We assume that this effect was due to the lowered contrast one experiences when looking through the display. The users had no more trouble reading the graphical eye chart than they did reading a real eye chart through the display.

In a test of the users' ability to properly verge the two images (presented to each eye), we presented both real and virtual cross-hairs to subjects. We asked them to indicate when the graphical cross-hairs seemed to verge simultaneously with the real cross-hairs. We found that with some of our display units, this was automatic; no adjustment was necessary. Some displays, however, required significant adjustment [11]. Whether this was due to manufacturing defects, damage through extended use, or some other cause, we can not say.

4.4 Urban Skills Training Evaluation

As part of a project entitled Augmented Reality for Urban Skills Training (ARUST) [4], we ran a pilot study to evaluate the usefulness of wearable AR in teaching urban skills to teams, specifically, team room clearing. Participants, in teams of two, were briefed on room-clearing techniques, then allowed to practice these techniques with or without the AR system, and finally evaluated in a simulated room-clearing task, without AR, against real people acting as opposing forces. The evaluation testbed assembled for this project consisted of two wearable AR systems, wide-area indoor tracking, the Army's OneSAF to drive the computer-generated forces, and wireless networking to tie the systems together.

This purpose of this pilot study was to measure the usefulness of AR at the application level and to set the stage for future work. Two conditions were evaluated: training with AR and without AR. Eight individuals grouped into four teams were tested for each condition, for a total of sixteen individuals in eight teams. Each trial contained an instructional period and an evaluation period. During the instructional period, the team learned basic room clearing techniques. Part of this period included donning the AR backpacks and practicing room clearing techniques for fifteen minutes in the practice area. Subjects in both the AR and non-AR conditions were free to practice as they saw fit, but they were encouraged to perform several repetitions of clearing all of the rooms. In the AR condition, as a team started each new repetition, we would load a new SAF scenario, placing stationary but reactive enemy and neutral forces in the environment.

After the instructional period ended, the subjects were moved to another part of the test site to be evaluated. Here, participants performed in six room-clearing scenarios against real people. Each scenario had enemy forces and non-combatants in different positions. As in the training period, these forces were stationary and occupied a particular corner of a room. The subjects and the people playing the enemy forces traded fire using "laser-tag-style" weapons. This weapon system counts the number of hits on the subjects and on the enemy forces and non-combatants.

The subjects were evaluated using objective and subjective measures. The objective measures were numerical scores based on the number of team members who survived each trial, number of enemies killed, and number of neutrals left alone. The subjective measures, as observed by our SME, included aggressiveness, movement,

security, communication between teammates, and coordination between teammates.

We found no significant difference between the performances of subjects using AR and those not using AR, judging by the results of a repeated measures Analysis of Variance (ANOVA) calculation. We did find a significant learning effect in both conditions as the trials progressed—in other words, the subjects learned more performing the trials than they did in either training condition. There are several factors that we think caused these results. First, the subject pool consisted of scientists with varying levels of experience in weaponry, gaming, and so on. Second, the training time with the AR system was very short, and no feedback was provided during that time. Finally, the weapons used in the trials were inaccurate and resulted in unintentional friendly fire, among other problems. We are setting up another user study to measure the effectiveness of AR for training in which we will address these and other issues.

5.0 CONCLUSIONS

Although a broad definition of AR systems encompasses head-up displays (HUDs) for pilots, man-portable AR systems are clearly not as mature as HUDs. Historically, the hardware has been the limiting factor in development of AR systems. We believe that the advantage for a mobile warrior, during operations or training, can be analogous to the advantage for a pilot with a HUD.

Our current implementations of BARS enable laboratory studies, but are not yet ready for operational use as man-portable systems. Development continues on many aspects of both types of applications. Notably, the situation awareness application will be extended to more explicitly benefit collaboration between different users, who may have different roles and different information about the environment. By running user studies, we expect to learn which factors limit performance of the user in various situations.

We believe that the training application we describe here and similar applications will be the first military use of AR for dismounted warriors. This is because training scenarios are conducted under somewhat controlled circumstances. Thus, we can instrument the environment with systems that accurately track the users' movements, which remains a difficult technical obstacle for usable AR systems. Also, the displays, which are currently rather cumbersome, are less problematic in such environments. Significant technical advances are needed to mitigate these limitations before the systems are sufficiently unobtrusive as to be practical for operational use. Mobile computers have made this leap already, and there are displays under development that are coming close to the requirements. Alternate displays, such as integration with binoculars or other hand-held displays, offer another possible avenue for improvement. We see some hardware manufacturers taking an interest in personal systems that allow unencumbered movement and believe that when the right applications are in place, whether they will be for the military or perhaps for the computer gaming markets, the manufacturers will be able to provide suitable hardware platforms.

Ultimately, we must evaluate the effectiveness of the system. We believe that this will not rest solely on the hardware with which a system is implemented, but rather be determined by the capability of the software to provide information to the warrior at the right time. We see two directions for our future research. First, we need to determine what the right information is for a warrior on a particular task. We must also continue to improve the visualizations provided by our system and the methods with which warriors may interact with the data. We envision a continued series of user studies, graduating to field tests of prototype systems, in order to answer these questions.

The system has progressed significantly in its ability to filter out less important data, represent complex or

hidden urban terrain, allow the user to interact with the data, and communicate with the envisioned network-centric battlespace. All of these advances will help push the system towards usable and useful situation awareness information or training scenarios for the warrior.

6.0 REFERENCES

- [1] R. Azuma, Y. Baillet, R. Behringer, S. Feiner, S.J. Julier, and B. MacIntyre (2001). "Recent Advances in Augmented Reality," *IEEE Computer Graphics & Applications*, Nov/Dec 2001, 21(6):34-47.
- [2] D.G. Brown, S.J. Julier, Y. Baillet, M.A. Livingston, & L.J. Rosenblum (2004). "Event-Based Data Distribution for Mobile Augmented Reality and Virtual Environments," *Presence: Teleoperators and Virtual Environments*, 13(2):211-221.
- [3] D.G. Brown, Y. Baillet, S.J. Julier, P. Maassel, D. Armoza, M.A. Livingston, and L.J. Rosenblum (2004). "Building a Mobile Augmented Reality System for Embedded Training: Lessons Learned," *Proceedings of the Interservice / Industry Training, Simulation, & Education Conference (IITSEC '04)*.
- [4] D.G. Brown, R. Stripling, and J.T. Coyne (2006). "Augmented Reality for Urban Skills Training," *Proceedings of IEEE Virtual Reality 2006 Conference*.
- [5] S. Feiner, B. MacIntyre, T. Höllerer, and A. Webster (1997). "A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment". *Proceedings of International Symposium on Wearable Computers*.
- [6] S.J. Julier, M. Lanzagorta, Y. Baillet, L.J. Rosenblum, S. Feiner, T. Höllerer, and S. Sestito (2000). "Information Filtering for Mobile Augmented Reality", *Proceedings of IEEE International Symposium on Augmented Reality*.
- [7] S.J. Julier, Y. Baillet, M. Lanzagorta, L.J. Rosenblum, and D.G. Brown (2001). "Urban Terrain Modeling For Augmented Reality Applications." In *3D Synthetic Environment Reconstruction*, Kluwer Academic Publishers.
- [8] M.A. Livingston, L.J. Rosenblum, S.J. Julier, D.G. Brown, Y. Baillet, J.E. Swan II, J.L. Gabbard, and D. Hix (2002). "An Augmented Reality System for Military Operations in Urban Terrain." *Proceedings of Interservice / Industry Training, Simulation, & Education Conference*.
- [9] M.A. Livingston, J.E. Swan II, J.L. Gabbard, T. Höllerer, D. Hix, S.J. Julier, Y. Baillet, and D.G. Brown (2003). "Resolving Multiple Occluded Layers in Augmented Reality." *Proceedings of International Symposium on Mixed and Augmented Reality*.
- [10] M.A. Livingston, C. Zambaka, J.E. Swan II, and H.S. Smallman (2005). "Objective Measures for the Effectiveness of Augmented Reality." *Proceedings of IEEE Virtual Reality (Poster Session)*.
- [11] M.A. Livingston, A. Lederer, S.R. Ellis, S.M. White, and S.K. Feiner (2006). "Vertical Vergence Calibration for Augmented Reality Displays." *Proceedings of IEEE Virtual Reality (Poster Session)*.
- [12] W. Piekarski, and B.H. Thomas (2002). "ARQuake: The Outdoor Augmented Reality Gaming System,"

Communications of the ACM, 45(1):36-38.

- [13] G.S. Schmidt, D.G. Brown, E.B. Tomlin, J.E. Swan II, and Y. Baillet (2006). "Toward Disambiguating Multiple Selections for Frustum-Based Pointing," *IEEE 3D User Interface Symposium*.
- [14] A. State, M.A. Livingston, G. Hirota, W.F. Garrett, M.C. Whitton, E.D. Pisano, and H. Fuchs (1996). "Technologies for Augmented Reality Systems: Realizing Ultrasound-Guided Needle Biopsies," *SIGGRAPH 96 Conference Proceedings*.
- [15] J.E. Swan II, M.A. Livingston, H.S. Smallman, D.G. Brown, Y. Baillet, J.L. Gabbard, and D. Hix (2006). "A Perceptual Matching Technique for Depth Judgments in Optical, See-Through Augmented Reality." *Proceedings of IEEE Virtual Reality*.