

Combining Head-Mounted and Projector-Based Displays for Surgical Training

Kok-Lim Low Adrian Ilie Greg Welch Anselmo Lastra
Department of Computer Science
University of North Carolina at Chapel Hill
{lowk, adyilie, welch, lastra}@cs.unc.edu

Abstract

We introduce and present preliminary results for a hybrid display system combining head-mounted and projector-based displays. Our work is motivated by a surgical training application, where it is necessary to simultaneously provide both a high-fidelity view of a central close-up task (the surgery) and visual awareness of objects and events in the surrounding environment. In particular, for trauma surgeons it would be valuable to learn to work in an environment that is realistically filled with both necessary and distracting objects and events.

In this paper, we motivate the use of a hybrid display system, discuss previous work, describe a prototype along with methods for geometric calibration, and present results from a controlled human subject experiment.

1. Motivation

Today, the pace of surgical innovations has increased dramatically, as have the societal demands for safe and effective practices. The mechanisms for training and re-training suffer from inflexible timing, extended time commitments, and limited content. Video instruction has long been available to help surgeons learn new procedures, but it is generally viewed as marginally effective at best for a number of reasons, such as the fixed point of view that is integral to the narration, lack of depth perception and interactivity, and missing information [1]. In short, the experience of watching a video is not sufficiently close to being there and seeing the procedure.

A paradigm that uses immersive Virtual Reality could be a more effective approach to allow surgeons to witness and explore a past surgical procedure as if they were there. We are indeed pursuing such an immersive paradigm together with our medical collaborators at the UNC-Chapel Hill School of Medicine (Dr. Bruce Cairns and Dr. Anthony Meyer), and our computer graphics collaborators at Brown University (Andy van Dam et al). This paradigm demands methods to record the procedure and to reconstruct the original time-varying events to create an immersive 3D virtual environment of the real




Figure 1. Different views of a surgical operation.

scene. A more complete solution should also allow relevant instructions and information, such as vocal narration, 3D annotations and illustrations, to be added by the original surgeon or other instructors.

Besides the recording and the reconstruction, providing an effective way to *display* a 3D virtual environment to the user is also a major challenge. In this paper, we introduce a hybrid approach to address this challenge.

During a typical use of the training system, the trainee would usually stand beside the patient paying close attention to the surgery. She might even stand in the position of a surgeon and observe the procedure from his




Figure 2. A user using our prototype system based on our hybrid display approach that combines a HMD and a projector-based display.

point of view. At the same time, the trainee is also required to be aware of the surrounding events that could affect the surgeons' actions. Such surrounding events include the actions of other surgeons and technicians, changes in monitoring and life-support devices, and overall awareness of the patient's dynamic condition. Figure 1(a) shows a close-up view of a real surgical operation in progress, and Figure 1(b) shows a snapshot of the many events happening in the operation room.

The visual needs of the trainee can be divided into two main parts. The first part requires high-quality stereo view of the objects and events that the trainee is paying direct attention to, such as the main surgical procedure. High-quality and high-resolution views are needed to discern the great intricacy of the surgery, and stereovision is needed for better spatial understanding. The second part of a trainee's visual needs is the peripheral view of her surroundings. This is needed by the trainee to maintain visual awareness of the surrounding events. Our medical collaborators, and others in the field, feel that visual awareness of the entire patient and the surroundings is a critical component of surgical training. In particular, with trauma surgery there is typically a lot of relevant activity in the operating room. It has been found that in the human visual system, resolution in the periphery is less dense than in the fovea [2], therefore peripheral view need not be high-resolution and high-quality.

Traditionally, head-mounted displays (also called head-worn displays) have been used to provide high-quality stereo visualization of 3D virtual environments. However, most HMDs offer limited fields of view, often only 40° to 60° horizontally and 30° to 45° vertically. Wide-FOV HMDs have been manufactured, but they are rare, expensive and heavy to wear. We are aware of no HMD that can fully cover the human field of view of approximately 200° horizontally and 135° vertically [3]. Although HMDs are good at providing high-quality stereo views, the generally narrow FOV has rendered them less than ideal for providing peripheral views.

The common alternatives to HMDs for immersive visualization of 3D virtual environments are immersive projector-based displays, such as the CAVE™ [4]. Most immersive projector-based displays are capable of providing very wide-field-of-view visualization, and like CAVE™, some of them are even capable of fully covering the human field of view. Because of the relatively large display surfaces and the fact that the user may move close to them, the image quality and resolution of such projector-based systems may be insufficient for applications that require the display of fine details.

1.1. Our Approach

HMDs or projector-based displays alone cannot easily meet the visual needs of a surgical trainee. However, we

noticed potential advantages that could be exploited in each display system, especially the high-resolution stereo view in HMDs and the wide-FOV view in projector-based displays. The two systems seemed to complement each other, so the next logical step was to combine them.

In this paper, we propose and investigate a hybrid display that combines a HMD and a projector-based display. By providing high-resolution stereo view through the HMD and lower-resolution peripheral view using a set of projectors, we hope that they would effectively complement each other to satisfy the visual needs of a trainee.

As a proof of concept and a means of investigation, we have implemented a simple prototype of the hybrid display system (see Figure 2). To understand its effectiveness in combining the useful attributes of the HMD and the projector-based display, we have also conducted a human subject experiment. The preliminary results obtained from the experiment support our belief that the hybrid display is suitable for applications like 3D immersive surgical training, which involves visual tasks that require both up-close detail and peripheral awareness.

1.2. Contributions

The main contribution of this work is conceiving, articulating, and prototyping a hybrid display system that leverages the complementary merits of head-mounted and projector-based displays. A second contribution is our controlled human subject experiment, which offered some surprising (to us) and valuable support to the idea. Finally, we also introduce a new approach to calibrating opaque HMDs that incorporates peripheral vision.

1.3. Paper Organization

In the next section, we look at some previous work related to our hybrid display. In Section 3, we offer a detailed description of the hybrid display and our simple prototype, and include discussion about the advantages, drawbacks and implementation issues for the hybrid display. The later part of Section 3 also describes our new HMD calibration method. In Section 4, we describe our human subject experiment and present its results and analyses. Finally, in Section 5 we share some closing thoughts.

2. Related Work

While there are many comparable systems, the CAVE™ [4] is probably the most commonly cited example of general-purpose projector-based display system for immersive virtual reality and scientific visualization. A CAVE™ usually consists of a cube made of rear-projection screens for the walls and the ceiling, and a down-projection screen for the floor. It is capable of

providing stereo visualization to the user, usually with the use of light polarization and special projection screens, or the use of a time-multiplexing approach to alternate images for the left and the right eyes.

Other approaches for immersive projector-based display have also emerged, such as the work by Raskar et al that aims at reconfigurability, color uniformity and seamlessness [6]. Their approach allows non-planar display surfaces to be casually set up and cameras are then used to quickly calibrate the projectors and display surfaces to allow correct display of view-dependent imagery during subsequent uses.

Our previous work on life-sized projector-based dioramas [5] is a more general effort aimed at creating life-sized immersive virtual environments of real or imagined places. The basic idea is to size and arrange physical projection surfaces to closely approximate the real geometry in the virtual environment, providing both visual and spatial realism.

HMDs have been used as means to display immersive virtual environments for the last three decades. Research efforts in this field were generally concentrated on improving the resolution and the optics and minimizing negative effects such as cybersickness.

Research more closely related to our work includes attempts to incorporate peripheral vision into HMDs. Whitestone and Robinette concluded that practically this means increasing the FOV of HMDs, which translates into the increase of the size and resolution of the display screen, the increase of the diameter of the lens, or the decrease of the focal length [7].

Slater and Usoh [8] simulated a wider FOV using distortion of the displayed images in a way that compresses more of the scene into the outer regions of the display area. The foveal region of the display is projected normally, while in the periphery, a larger portion of the field of view than normal is imaged by adding a shear to the projection.

A similar approach is to display high-resolution images in a small area-of-interest and low-resolution images elsewhere. Toto and Kundel [9], and Yoshida et al [10] constructed area-of-interest HMDs with eye-tracking systems that move the high-resolution inset to where the eye is looking at a given time. Watson et al. [11] argued that there is no significant decrease in performance if the inset is fixed, since typical eye movements do not go farther than 15 degrees from the fovea.

The Kaiser “full immersion head-mounted display” [12] achieves wide FOV by tiling multiple screens to form the display. The HMD contains 12 liquid crystal displays placed in a 3 by 2 arrangement in each eye. It provides a total FOV of approximately 176° horizontal by 47° vertical. Unfortunately, as far as we know, this system was a one-time prototype, and is unavailable for purchase.

Others have suggested adding low-fidelity displays to a conventional HMD to provide peripheral information. The Ambient HMD project at Monterey Technologies has shown that adding peripheral LCD displays improves performance in a helicopter simulator [13].

Kevin Arthur provides in his Ph.D. dissertation [14] a much more detailed overview of numerous attempts to incorporate peripheral vision into HMDs.

3. The Hybrid Display

Our hybrid display combines an HMD and a projector-based display. The idea is to address the visual needs of a user of the 3D immersive surgical training system by leveraging the complementary merits and shortcomings of the HMDs and projector-based displays. We have implemented a simple prototype as a proof of the concept. In this prototype, we use a Kaiser ProView™ 30 stereo HMD, which does not have baffling material around the displays that blocks the wearer’s peripheral view of the real physical world (see Figure 3). This allows a user of the hybrid display system to also see peripheral imagery projected by projectors on display surfaces around her.

Ideally, to completely cover the peripheral vision of the user, the projector-based display should surround the user, even above and below. Instead, for our prototype and user study, we have built only a partial surround-display, which is shown in Figure 4. The display surfaces were constructed using white stackable Styrofoam blocks and a white foam-reinforced cardboard. When in use, four



Figure 3. A user wearing a Kaiser ProView™ 30 HMD. Her peripheral vision is unobstructed.



Figure 4. Our prototype's projector display surfaces.

projectors are used to front-project view-dependent imagery on top of the cardboard and on the vertical Styrofoam wall. In the spirit of life-sized projector-based dioramas [5], the cardboard actually corresponds to a surgical table in the virtual operation room. Besides providing the user with a stronger sense of spatial connection between the real and virtual environments, the main purpose is to increase the user's vertical FOV when the user's head is directly above the table.

When a user is standing at the center approximately 65 cm from the Styrofoam wall (which corresponds to a location a user would usually stand in the virtual environment), the projector-based display is able to partially cover the user's visual field approximately 130° horizontally and 120° vertically. In comparison to the FOV of the Kaiser HMD, which is only 24° horizontally and 18° vertically, we believed the projector-based display was wide enough to evaluate the potential improvement over the HMD alone.

For the first prototype, we created a simplified synthetic model of an operation room with a virtual patient lying on the surgical table and three virtual surgeons/nurses moving around near the table (see Figure 2). Sometimes, one of the surgeons/nurses would extend her arm towards the center of the table, mimicking a real nurse handling a surgical tool to a surgeon during a real surgery. When we tested the prototype ourselves, we would attempt to act in a way a surgical trainee would when using an actual training system. We tried to concentrate on the virtual patient's abdominal area while attempting to maintain awareness of the movements of the virtual surgeons/patients around us. For comparison, each of us had opportunities to try two different setups, one with the use of both the HMD and projectors (the hybrid display setup), and another with the use of the HMD alone. For the HMD-only setup, the projectors were switched off and the laboratory was kept dark.

After having experienced the two setups, every one of us felt that with the hybrid display setup, it was much easier to maintain awareness of the movements of the virtual surgeons/patients while concentrating on the patient's abdominal area. On the other hand, with the HMD-only setup, it was almost impossible to know the positions of the virtual surgeons/nurses without frequently turning our heads to scan the surroundings. Based on these informal preliminary experiments we decided to undertake a more objective, formal human subject experiment. We describe this experiment in Section 4.

Besides serving as a means to investigate and demonstrate the effectiveness of the hybrid approach, another objective of the prototype is to identify potential limitations and implementation issues. These limitations and implementation issues, as well as the advantages of the approach, are discussed in the following subsections.

3.1. Advantages

Through the HMD, a hybrid display can achieve high resolution that is hard to match by any immersive projector-based display. The user of the hybrid display can selectively increase the image details of any part of the virtual environment by moving closer to it. This directly corresponds to what a person would naturally do in the real world. Projector-based tiled-displays have long been used to display high-resolution images. However, their use for immersive virtual reality would be a relatively inefficient solution, since the user does not need high resolution everywhere around her at the same time.

As mentioned before, by using a traditional immersive projector-based surround display, the hybrid display is able to cover the entire visual field of the user. Since the user is seeing the projected imagery in her peripheral vision, there is no need to provide high resolution in the projected imagery. This means fewer projectors are needed. Moreover, since the responsibility to provide stereovision is already taken over by the HMD, there could be further savings in the number of projectors and rendering resources.

With a hybrid display, it is no longer desirable to use wide-FOV HMD, which is usually very expensive and heavy to wear. Narrower FOV in the HMD also allows some rendering accelerations, such as view-frustum culling, to be more effective.

In summary, we believe a hybrid display can be a more versatile, more cost effective, more efficient, and more functionally effective approach for immersive VR.

3.2. Drawbacks, Limitations and Issues

One drawback of the hybrid display system is that the user still has to wear a HMD on her head, which sometimes can be encumbering. However, because our HMD does not need to be wide-FOV and does not have baffling material around the display, it can be more easily made smaller and lighter. We also note that even users of immersive projector-based display systems have to wear *something* on their heads—stereo glasses and tracking devices.

A negative artifact that we have noticed in practice is the dark border (opaque frame) around each of the HMD's displays. As with any Augmented Reality application, such borders break the continuity between the central (HMD) and peripheral (projected) views. Our sense is that for the available commercial systems, these borders tend to be relatively thick. For our HMD in particular, we estimate that each vertical segment of the border is approximately 7° wide, and each horizontal segment is approximately 12° tall. We did not perform any experiments to investigate how these borders affect a user's performance or sense of presence.

The overlapped region between the FOVs of the left and right eyes forms the human stereo visual field. This stereo visual field spans almost the entire vertical FOVs of the eyes. In a hybrid display system, much of the user's stereo visual field is outside the HMD, and is served by non-stereo projected imagery from the projectors. The impact of this on performance and the sense of presence has not been understood, and is an important issue that needs further investigation.

A further limitation of a hybrid system using an opaque HMD is that the user cannot see her own hands when she is manipulating virtual objects seen through the HMD. Research has been done to successfully incorporate real objects, such as the user's hands, into the virtual environments. One of such methods is described by Benjamin Lok in his Ph.D. dissertation [15]. He also presented a comprehensive survey of the other work in the area. Fortunately, in the peripheral view, the user of a hybrid display system can still see her body and legs. We believe this helps a lot when the user is navigating in the virtual environment. For the future, we are considering the use of a see-through HMD in the hybrid display.

There are also other issues that are inherent in projector-based displays, such as intensity blending at overlapping regions, color uniformity across different projectors, and back-projection versus front-projection. We refer the reader to [6] and [5] for further details.

3.3. HMD Calibration

To implement a hybrid display system, one of the issues that needs to be addressed is the geometric calibration of the HMD. This is necessary to properly align the images seen in the HMD and the projected imagery from the projectors. While it is relatively easy to calibrate an optical see-through HMD [16], we are aware of no published method to calibrate an opaque HMD. Here, we present a simple and novel approach that makes use of the user's peripheral vision.

We divide the calibration of each HMD's display into two stages. In the first stage, we find the position of the eye in the tracker target's coordinate frame. (The tracker *target* is the part of the tracker that is rigidly attached to the user's head.) This can be done by finding at least two lines in the tracker target's coordinate frame that intersect at the location of the eye. Details of the procedure to determine such lines are described in [17].

In the second stage, we display two horizontal and two vertical 2D lines in the HMD (see Figure 5). These lines intersect at four known screen positions. In the physical real world, single horizontal and vertical 3D lines are marked on a wall. These lines should be long enough to extend beyond the FOV of the HMD. The 3D positions of the two physical lines are known in the tracker coordinate frame, which is fixed with respect to the physical world.




Figure 5. The HMD is calibrated using the peripheral vision to line up some known 2D lines in the HMD display with some known 3D physical lines.

One at a time, the user uses her peripheral vision in both eyes to line up each 2D horizontal (vertical) line in the HMD with the 3D horizontal (vertical) line in the real world. This procedure allows us to determine four planes in the tracker target's coordinate frame that intersect to form four rays. Using a variant of the camera calibration technique presented in [18], we solve for the view parameters of the HMD display and use them to construct OpenGL viewing matrices during run-time rendering.

This procedure has worked well for us. Our experience is that a user can typically align the displayed 2D and real 3D lines with less than 0.5° error, which for our setup results in relatively good registration between the HMD and projected imagery.

4. Experiment and Results

For an objective evaluation of the hybrid display, we conducted a formal human subject experiment involving 25 subjects randomly recruited from the UNC-Chapel Hill student population. The applicants were screened for (among other things) normal vision and prior experience with Virtual Reality systems (we wanted none).

While we could not test the subjects on visual tasks that could *only* be performed in an immersive surgical training system, we have designed simpler tasks that we believe are similar or analogous to the actual tasks. In keeping with the visual training requirements described in Section 1, our primary criterion for the experimental tasks was that they should force the subject to concentrate on a central up-close task while simultaneously maintaining awareness of peripheral events.

4.1. Experiment Description

The overall purpose of the experiment was to study the effect of two different displays on a user's performance in tasks that simultaneously require

1. the ability to visually and mentally concentrate on a central static virtual object, and
2. the ability to be visually aware of changes in his/her surroundings in the virtual environment.

The two different displays considered were

1. a hybrid display (both HMD and projectors), and
2. an HMD alone.

For the hybrid display, we used the same physical setup as our prototype, and for the HMD-only setup, we just switched off the projectors and kept our laboratory dark.

We decided to not test a projector-only setup for a few reasons. Given the available equipment, we were unable to produce stereo projected images of sufficient resolution. In addition, we wanted to keep the user study manageable in scope.

The Tasks

We created a virtual environment that consisted of a simple room with a wooden table in the middle (see Figure 6). On the table, there was a virtual screen used to display puzzles to the test subject. These puzzles were used to force the subject to visually and mentally concentrate on a static object. We used two different sets of puzzles. The first set consisted of 10 word-search puzzles (one is shown in Figure 7). For each word-search puzzle, the subject was shown 10 words in the right column of the virtual screen, and asked to find any 8 of them in the grid of letters. When a word was found, the subject was required to speak the word and describe its orientation (top-down, bottom-up, left-to-right or right-to-left), so that one of the experimenters could cross the word out on the virtual screen by pressing a key on a keyboard.

The second set of puzzles consisted of 20 colorful and detailed pictures from some "Where's Waldo?" books (author Martin Handford). For each Waldo puzzle, the subject was required to look for the fictional cartoon character Waldo in the picture. Once Waldo was found, the subject was asked to speak the coordinates of Waldo's location. (We overlaid an annotated grid on each picture.)

In each trial in the experiment, the subject was given 10 minutes to solve either the set of word-search puzzles or the set of Waldo puzzles. The choice was made by us in advance, with a uniform random distribution.

To test whether the subject was visually aware of the virtual surroundings, occasionally (about four times every



Figure 6. The virtual environment used in the experiment.



Figure 7. A word-search puzzle and a penalty token.

minute), a *token* (a spinning cube) appeared in the far left or right side of the subject, and floated toward the virtual screen in the center. When near the virtual screen, it would reverse its course and move back to where it came from, eventually disappearing from the room. The lifetime of each token was 8 seconds. There were two types of tokens: *bonus* tokens and *penalty* tokens. The bonus tokens had a single number (a point value) on all sides, while a penalty token had the letter "X" on all sides. Once a bonus token appeared, the digit shown was decremented every second. The subject was asked to press a button as soon as he/she was aware of a *bonus* token, and do nothing if he/she noticed a *penalty* token. The users were not actually penalized for choosing a *penalty* token, but we wanted them to think they would be. The idea was to discourage them from continually pushing the button without actually looking at the tokens.

Each subject was told that their objective was to solve as many puzzles and correctly identify as many bonus tokens as he/she could within 10 minutes. To each subject, we specifically emphasized that the two tasks were equally important. The subject was given about one minute to practice before the actual trial began.

The Trials

During a session, each subject was required to do two different trials. There was a 10-minute break between the end of the first trial and the start of the second. Each trial tested one of the following four combinations:

1. HMD-only and Waldo puzzles,
2. hybrid display and Waldo puzzles,
3. HMD-only and word-search puzzles,
4. hybrid display and word-search puzzles.

A test subject was randomly assigned Combinations 1 & 4, 4 & 1, 2 & 3 or 3 & 2, for the session. Figure 8 shows a subject doing a Combination-2 trial.

4.2. Results and Analyses

Our experimental tasks were designed to force a user to concentrate on a central static virtual object while



Figure 8. Using the hybrid display, a test subject is trying to “capture” a bonus token while finding Waldo in a Waldo puzzle.

simultaneously monitoring their virtual surroundings. We developed two hypotheses for these conditions: (**H1**) hybrid display users are more visually aware of changes in their virtual surroundings, and (**H2**) users of the hybrid display can visually and mentally concentrate better on the central static virtual object. Hypothesis **H1** would be supported if the number of bonus tokens identified by the hybrid display users was significantly higher than those identified by the HMD-only users, whereas, Hypothesis **H2** would be supported if the number of puzzles solved and the number of words found by the hybrid display users was significantly higher.

Table 1 and 2 show a summary of some of the data we obtained from the experiment. “Bonus Tokens %” is the percentage of bonus tokens that were successfully identified during the duration of a trail. The large differences between the mean bonus tokens % of the hybrid display and the mean bonus tokens % of the HMD clearly supports Hypothesis **H1**.

However if we look at the “Puzzles / min” and the “Words / min” rows, we do not see any significant

Table 1.

WALDO PUZZLES	Combination 1 HMD Only		Combination 2 Hybrid Display	
	mean	stddev	mean	stddev
Bonus Tokens %	54.4	16.3	99.1	2.1
Puzzles / min	2.0	0.5	2.1	0.7
Ave Angular Speed (°/s)	23.8	5.6	9.5	3.1
Head Turns / min	9.0	2.2	0.5	0.9

Table 2.

WORD-SEARCH PUZZLES	Combination 3 HMD Only		Combination 4 Hybrid Display	
	mean	stddev	mean	stddev
Bonus Tokens %	55.2	21.4	97.7	8.1
Words / min	4.4	1.4	5.4	2.3
Ave Angular Speed (°/s)	18.4	5.8	5.3	2.3
Head Turns / min	9.1	3.1	0.4	0.9

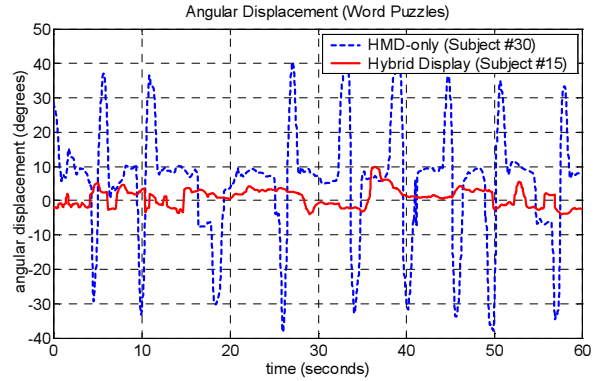


Figure 9. An angular-displacement plot comparing the amount and frequency of head rotations made by a HMD user and a hybrid display user

differences between the hybrid display and the HMD. Therefore, Hypothesis **H2** is not supported by our experimental data. This was quite unexpected because during the experiment, we saw the HMD users making large and frequent head turns to scan for tokens. We assumed this would have been very distracting and would have largely reduced the users’ ability to solve puzzles. The frequent head turns by the HMD users can be observed in the last rows of Table 1 and 2. Figure 9 is an angular-displacement plot that compares the amount and frequency of head rotations made by a HMD user and a hybrid display user.

The failure of the experimental data to support **H2** may lie in the inappropriate choice of tasks for our experiment. The ability to solve word-search or Waldo puzzles might vary significantly across the general population. In retrospect, we realize that a simpler central task, such as monitoring some changing digits for a “magic number,” would likely have less variance.

We also performed correlation analyses on the experimental data, and discovered some surprises in the results of Combination 3 (see Table 3). While we expected the number of bonus tokens identified and the frequency of head turns to be correlated (row 1), we did *not* expect to see the positive correlations between words found and head turn frequency (row 2), and bonus tokens and words found (row 3). In fact it appears that words found, bonus tokens, and head turns were mutually correlated. One possible explanation, which seems to be

Table 3. Correlation analysis of Combination 3.

	Correlation Coefficient	P-value
Bonus Tokens % vs. Head Turns / min	0.85	0.0005 < P < 0.0025
Words / min vs. Head Turn / min	0.56	0.025 < P < 0.05
Bonus Tokens % vs. Words / min	0.63	0.02 < P < 0.015

supported by the complete data, is that subjects who had difficulty with the word-search puzzles may have had less leftover mental capacity to maintain peripheral awareness.

5. Conclusions

We began this work with basically only intuition about the benefits of a hybrid approach, believing that it should be possible to provide and realize the combined benefits of head-tracked HMD and projector-based displays. At this point, we have a concrete prototype, and some objective assessment of the usability and benefits of the approach. While the analysis of our user study does not completely support our original hypotheses, we are encouraged by the results, enough to invest significant further effort in exploring larger prototypes, issues related to mono-stereo transitions, and variations such as the use of a see-through HMD.

While we would not claim that our hybrid approach is appropriate for general-purpose visualization, we believe it could extend to other applications. In particular, HMD-based applications where surrounding imagery would provide useful context or awareness, or projector-based applications where high-fidelity close-up 3D imagery would improve the user's insight into the data or task.

6. Acknowledgements

We thank Professor Henry Fuchs for his support and insightful suggestions, Tammy Pike for help in managing experiment participants, Haolong Ma for both use of his 3D operating room models and operational help with the experiment, and Hsin-Yee Teh for help in the statistical analyses. Ginnette Gill was our model in the video and some of the photographs that appear in this paper. We thank the members in the UNC EVE group, as well as Jim Mahaney, Herman Towles, David Harrison, and Kurtis Keller for the technical support. This research was supported by NSF ITR Grant no. 0121657: "Electronic Books for the Tele-Immersion Age."

7. References

- [1] S. Tachakra, P. Jaye, J. Bak, J. Hayes, and A. Sivakumar. Supervising trauma life support by telemedicine. *Journal of Telemedicine and Telecare*, 6(Suppl 1):S1:7–11, 2000.
- [2] Hugh Davson. *Physiology of the Eye*. (Fifth edition). New York: Pergamon Press, 1990.
- [3] Elliot B. Werner. *Manual of visual fields*. New York: Churchill Livingstone, 1991.
- [4] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. *Proceedings of SIGGRAPH 93*, July 1993.
- [5] Kok-Lim Low, Greg Welch, Anselmo Lastra, and Henry Fuchs. Life-Sized Projector-Based Dioramas. *Proceedings of the ACM Symposium on Virtual Reality and Software Technology (VRST 2001)*, pp. 93–101, November 2001.
- [6] Ramesh Raskar, Michael Brown, Ruigang Yang, Wei-Chao Chen, Greg Welch, Herman Towles, Brent Seales, and Henry Fuchs. Multi-Projector Displays Using Camera-Based registration. *Proceedings of the IEEE Conference on Visualization 99*, pp. 161–168, October 1999.
- [7] Jennifer J. Whitestone and Kathleen M. Robinette. Fitting to Maximize Performance of HMD Systems. In James E. Melzer & Kirk Moffitt (Eds.), *Head-Mounted Displays: Designing For the User* (pp. 175–206). New York: McGraw-Hill, 1997.
- [8] Mel Slater and Martin Usoh. Simulating peripheral vision in immersive virtual environments. *Computers & Graphics*, 17(6), pp. 643–653, 1993.
- [9] Lawrence Toto and Harold Kundel. An eye motion dependent interactive display. In J. K. O'Regan & A. Levy-Schoen (Eds.), *Eye Movements: From Physiology to Cognition* (pp. 643–649). Amsterdam: Elsevier Science Publishers, 1987.
- [10] Akitoshi Yoshida, Jannick P. Rolland, and John H. Reif. Design and applications of a high-resolution insert head-mounted display. *Proceedings of IEEE Virtual Reality Annual International Symposium 95*, pp. 84–93, March 1995.
- [11] Benjamin Watson, Neff Walker, Larry F. Hodges, and Martin Reddy. An Evaluation of Level of Detail Degradation in Head-Mounted Display Peripheries. *Presence*, 6(6), pp. 630–637, 1997.
- [12] Kaiser Electro-Optics. HiDef Family of Head-Mounted Displays. Available at <http://www.keo.com>, 2000.
- [13] Monterey Technologies. Presentation of Aircraft State Information in an Ambient Head-Mounted Display. Available at <http://www.montereytechnologies.com/hmd.htm>, 1999.
- [14] Kevin Arthur. Effects of Field of View on Performance with Head-mounted Displays. Ph.D. Dissertation, Department of Computer Science, University of North Carolina at Chapel Hill, April 2000.
- [15] Benjamin Lok. Interacting with Dynamic Real Objects in Virtual Environments. Ph.D. Dissertation, Department of Computer Science, University of North Carolina at Chapel Hill, April 2002.
- [16] Mihran Tuceryan and Nassir Navab. Single point active alignment method (SPAAM) for optical see-through HMD calibration for AR. *Proceedings of the IEEE and ACM International Symposium on Augmented Reality*, 2000.
- [17] Kok-Lim Low. Viewpoint Calibration With Respect to a Tracker. Technical Report TR02-019, Department of Computer Science, University of North Carolina at Chapel Hill, April 2002. Available at <ftp://ftp.cs.unc.edu/pub/publications/techreports/02-019.pdf>.
- [18] Olivier Faugeras. *Three-Dimensional Computer Vision*. MIT Press, 1993.