

An Image Overlay System for Medical Data Visualization

Mike Blackwell¹, Constantinos Nikou¹, Anthony M. DiGioia^{1,2}, and
Takeo Kanade¹

¹ Center for Medical Robotics and Computer Assisted Surgery,
Carnegie Mellon University, Pittsburgh, PA

{mkb, costa, tk}@cs.cmu.edu <http://www.mrcas.ri.cmu.edu>

² Center for Orthopaedic Research, UPMC Shadyside Hospital, Pittsburgh, PA
digioia@cor.ssh.edu <http://www.cor.ssh.edu>

1 Introduction

Image Overlay is a computer display technique which superimposes computer images over the viewer's direct view of the real world. The positions of the viewer's head, objects in the environment, and components of the display system are all tracked in space. These positions are used to transform the images so they appear to be an integral part of the real world environment. By utilizing semi-transparent display devices, the images can appear to the viewer to be inside of real objects. For example, a 3D image of a bone, reconstructed from CT data, can be displayed to a surgeon inside the patient's anatomy at exactly the location of the real bone, regardless of the position of either the surgeon or the patient. In effect, Image Overlay can provide the viewer "X-ray vision" in a wide variety of applications.

Image Overlay is a form of "augmented reality" in that it merges computer generated information with real world images. In a typical augmented reality system the real world images come from a video camera, which are enhanced with properly aligned computer images and displayed on a computer monitor. The neurosurgical system from Brigham and Women's [6] is an excellent example of this type of system. There are some limitations with typical augmented reality systems, however, which are addressed by Image Overlay. With Image Overlay the surgeon does not need to look away from the surgical field to view the image – it appears within the patient. Also, the real world information is not limited by the resolution or field of view of a camera, since the surgeon views the patient directly. Image Overlay is also distinct from "virtual reality" in that the images are based on real data and are merged at the proper scale and location into the viewer's direct experience of the environment, instead of replacing the viewer's senses with purely virtual data. With Image Overlay the surgeon does not need to wear bulky head gear which would limit his visual acuity – at most the system might require lightweight clear polarized glasses.

We have built two prototype Image Overlay systems to evaluate the technology for applications in the medical domain, especially in areas where improved

visualization can reduce the surgical exposure. Target applications include intraoperative tool localization and guidance, as well as surgical education and training. This paper describes the prototype systems and initial experimental results.

2 System Description

An Image Overlay system consists of four basic components:

- Computer graphic workstation,
- Semi-transparent display,
- Position tracking system, and
- Software to correlate positions and transform images.

2.1 Workstation

In order to present a convincing illusion that the computer generated images are part of the environment, the workstation must be able to transform and re-display images in real-time, at least 30 updates per second. We are using an Indigo-II R10K (Silicon Graphics Inc., Mountain View, CA), which handles complex 3D geometric images with ease.

2.2 Semi-Transparent Display

There are several techniques for producing a semi-transparent display. One is to use a standard flat panel liquid crystal display (LCD) with the backlight removed. The result is a display in which one can see through the image (some computer overhead projectors use this technique). The problem with this technique is that the image is typically very dim and the transparency is not high, because LCD panels are not optimized for this purpose. The technique we use is based on a standard computer display monitor coupled with a half silvered mirror or beam splitter glass. The viewer looks directly at the environment through the glass, while simultaneously seeing a reflection of the computer display from the glass, as shown in Figure 1.

The result of the display/mirror combination is a “virtual” image which appears to float below the mirror exactly the same distance as the display sits above the mirror. It is important that the virtual image be positioned such that it appears to lie within the patient, and not too far above or below the region in which the surgeon is working. This is because in order to see the patient clearly the surgeon must focus his eyes on the patient. If the virtual image is located at a different focal distance, the surgeon’s eyes must re-focus to see it clearly. If this focal disparity is too great it will lead to eye strain and fatigue. This is the primary cause of “VR sickness” that some people experience with head-mounted displays – the viewer’s eyes must focus on the display, which is very close to the eye, but artificial visual cues fool the brain in to thinking that the displayed

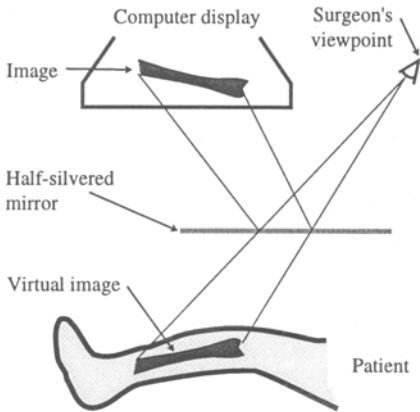


Fig. 1. Using a half-silvered mirror to create a semi-transparent “virtual” display.

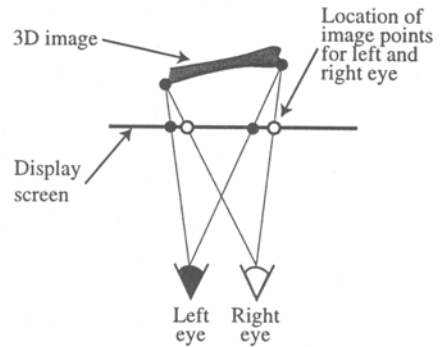


Fig. 2. Location of image points on display for each eye to produce stereoscopic effect.

objects are much farther away, leading to a discrepancy. Image Overlay avoids this problem by placing the virtual image within the patient.

The virtual image produced by the display/mirror system is a flat two-dimensional image, exactly the same as the surface of the display. An enhancement to the system is to utilize a stereoscopic technique to produce three-dimensional images [5], [9]. This is done by presenting a slightly different image to the left and right eyes – the different images the eyes would see if looking at an actual 3D scene due to their slightly different positions in space, as shown in Figure 2. There are a number of ways to produce this stereoscopic effect. One is with the Crystal Eyes system (StereoGraphics Corp., San Rafael, CA), in which the left and right eye images are alternated very rapidly on the computer monitor. A pair of “shutter glasses” is synchronized to this alternation, blacking out one lens or the other to ensure the each eye sees the appropriate image. By alternating the images fast enough, no flicker will be detected and the stereoscopic effect will be achieved. Another system (VRex, Elmsford, NY) uses a special polarized film applied to the computer monitor. This film polarizes every other scan line of the monitor in the opposite direction. Coupled with glasses with oppositely polarized lenses, one eye will see the even scan lines, and the other eye the odd scan lines. By simultaneously drawing the left and right eye images to the alternate scan lines, the stereoscopic effect will again be achieved.

2.3 Tracking System

The primary requirements for the position tracking system are that it must be at least as accurate as the spatial resolution of the display system and that it be as fast as the image update time. One such system is the OptoTrak 3020 (Northern Digital, Toronto, ON), which uses three cameras to triangulate the

positions of infrared light emitting diodes (LEDs) attached to the objects to be tracked. Other systems include different optical tracking systems (Northern Digital Polaris, Image Guided Technologies FlashPoint), magnetic field trackers (Ascension Flock of Birds, Polhemus FastTrak), mechanical linkages (Faro Arm), or some combination. Whichever system is used, it must be capable of tracking the viewer's eye positions, the display itself, and the patient. For some applications it may also be necessary to track additional objects, such as surgical instruments or structures within the patient.

2.4 Registration and Display Software

The remaining components of an Image Overlay system are a means of properly registering image data to the external environment and software to transform the images so they appear to be part of the environment. Registration is the process of sensing objects in the environment and then corresponding those objects to their counterparts in the image data. There are many ways to achieve registration, but most rely on sensing the three-dimensional shape of objects, using, for example, multiple camera views, digitizing probes, or ultrasound probes. These 3D shapes can then be matched to geometric features in the image data, producing a registration transformation [10].

Once the objects to be displayed have all been transformed in to the same coordinate system via the calibration and registration transformations, they must be displayed so that they appear to the eye exactly as if they existed in their virtual spatial locations. We use the OpenGL graphical libraries to render three-dimensional objects on the two-dimensional display, so this process requires correctly computing the projection matrix. The projection matrix is defined by a viewing frustum with the apex positioned at the eye and the edges passing through the corners of the virtual image, as shown in Figure 3. The viewing frustum differs from that typically used in computer graphics in two ways. First, the frustum is skewed depending on the eye position, since the viewing axis is not necessarily normal to the screen. Second, objects which lie exactly on the virtual image plane will be display in fixed locations regardless of the eye position. For example, if the eye gets closer to the image the object appears to grow larger because it is filling more of the field of view even though it is still drawn the same size. Stereoscopic displays are produced by computing a different projection matrix for each of the two eye positions.

An additional subtlety in the display rendering is that the mirror flips one axis, yielding a left-handed coordinate system. Depending on how the data is represented, this can be accommodated by scaling one of the axes by -1 at some point in the display pipeline.

3 Image Overlay Prototype Systems

In our first Image Overlay prototype the display was built around a 21-inch CRT monitor coupled with a large (1-meter square) half-silvered mirror (Figure 4).

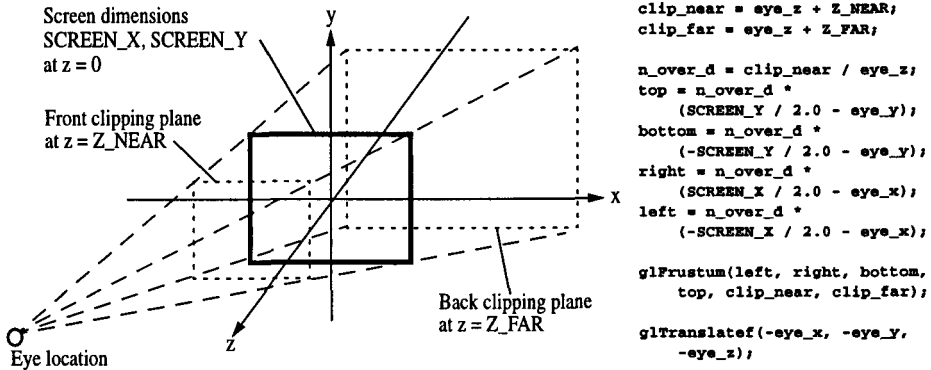


Fig. 3. Viewing frustum for one eye. The physical screen (or virtual image) lies in the XY plane. OpenGL pseudo-code to compute the projection matrix is shown on the right.

The Crystal Eyes stereoscopic system was used to produce 3D images. The OptoTrak system was used to track the display, patient, and viewer's head (with a target attached to the Crystal Eyes glasses). This prototype convinced us that the Image Overlay concept worked, and was the basis for some initial accuracy studies [2]. The prototype also revealed a number of limitations. Foremost was its size – it was extremely large and cumbersome, difficult to move, hard to calibrate, and impractical for any sort of clinical setting. Another problem was its accuracy. For computational efficiency, the display software assumed that the display was perfectly flat and the pixels evenly spaced. The CRT monitor was in fact fairly curved. The thick glass on the face of the CRT also caused diffraction effects. The curvature and diffraction lead to significant inaccuracies in the system, especially near the edges of the workspace (these effects are analyzed in [4]).

Based on our experiences with the first prototype a new Image Overlay system was recently constructed, pictured in Figure 5. This system is much smaller and easier to move. The display is a high resolution flat LCD panel (Silicon Graphics Presenter 1280). The geometry of the mirror with respect to the display and user's head was carefully chosen to minimize the size of the mirror while still allowing a large field of view. The display/mirror combination is attached to an articulating arm so the system can be easily moved in and out of the surgical field as appropriate. OptoTrak targets are again used to track the head, display, and patient. The system is initially two-dimensional, but we are beginning experimentation with a VRex style stereoscopic display.

The new Image Overlay system is much more suitable for a clinical environment. Its small size and articulation allow it to be easily moved about the OR, so the OptoTrak camera can be positioned to maintain the required line of sight to the tracking targets. Sterilization is addressed by “bagging” the system in sterile plastic draping (similar to that used for surgical microscopes) and at-




Fig. 4. Configuration of original prototype Image Overlay system.




Fig. 5. The new Image Overlay prototype in a pelvic screw fixation simulation.

taching a sterile disposable OR lamp handle to the display so the surgeon can move it. The mirror is currently glass, but will be replaced by Lexan to prevent breakage if hit by a tool handle (the illusion of the virtual image is good enough that you quickly forget that the mirror is there!).

4 Applications

There are many potential medical applications of the Image Overlay technology. We are focusing on two specific areas, intraoperative guidance and surgical education.

4.1 Intraoperative Guidance

The first clinical application is assisting the surgeon in the proper positioning of the acetabular and femoral prosthetic implant components during a total hip replacement surgery. The Image Overlay system is used in conjunction with the Hipnav system [11], which we are currently using in a pre-clinical trial with a simpler guidance feedback technique. In the Hipnav system the proper locations and orientations of the implant components are determined preoperatively from kinematic and biomechanical simulations based on CT data. During the surgery an Optotrak system is used to track the patient and the surgical tools. A digitizing probe is used to collect points on the pelvis, and these points are matched to the CT data to determine the correct registration transformation. Once the registration is known the Hipnav system can guide the surgeon to properly place

the tools, and thus the implants. Currently, the Hipnav system uses a conventional two-dimensional computer display, mounted on the OR wall, to provide guidance to the surgeon. Although simple, this has the disadvantage that the surgeon must look away from the patient to receive this guidance. By utilizing the Image Overlay system to display the proper implant positions and graphical indications such as arrows on which way to move the tools, we can potentially increase the accuracy and decrease the time of the surgery.

Closed reduction of fractures is another application of Image Overlay which we are experimenting with. The Image Overlay system allows the surgeon to visualize the positions and orientations of bone fragments (generated from CT images) and provides guidance for the proper placement of screws or intramedullary rods. Orthopaedic applications in general lend themselves well to Image Overlay for two reasons. First, the applications typically involve rigid bony anatomy, so models built from preoperative data tend to remain valid during the surgery. In contrast soft tissues will deform during surgery, so Image Overlay would need a method of intraoperative sensing to accurately display their shapes and positions. Second, techniques already exist for registering and tracking rigid anatomy using systems such as Optotrak [10] or ultrasound [12].

We are also experimenting with neurosurgical procedures, in particular tumor resection based on MRI data. Using Image Overlay to display the three-dimensional locations of tumors, vessels, and critical areas, the surgeon can plan the craniotomy and approach to minimize the surgical exposure while avoiding the critical areas. Registration is performed via a stereotactic frame which is attached to the patient's head during the scan and in the same position during the procedure. Conceptually, Image Overlay could also be used to display the positions of tools in the brain with respect to the tumor, but questions of overall system accuracy and the deformation of the brain following the craniotomy [3] must be resolved before this can be considered safe or appropriate.

A system very similar to Image Overlay is already being tested clinically in a neurosurgical procedure. The VolumeGraph system [7] from the Tokyo Womens Medical College displays a static three-dimensional image produced from preoperative MRI data, and reflects the image off a half-silvered mirror so the 3D image appears aligned inside of the patient's head. While this system has the same accuracy constraints as Image Overlay, it illustrates that the concept can find clinical utility.

4.2 Surgical Education

Our other main area of research with Image Overlay is in simulations for surgical education. The use of the system is similar to its use in the clinical applications, except a cadaver specimen or plastic phantom is used in place of the patient. Virtually any surgical procedure can be enhanced with the Image Overlay system, allowing the surgeon to directly observe the results of every action. The constraints on the system are greatly reduced for simulations, primarily because the accuracy requirements are reduced. Preoperative data can be collected at as high a resolution as necessary without regard to radiation exposure, or generic

preoperative data can be substituted (from the Visible Human data set, for example). Motions of anatomical structures can be estimated or internally instrumented. Registration can be performed by attaching fiducial landmarks to the phantom patient prior to scanning and then recording their positions prior to the procedure. The first procedure we are simulating with Image Overlay is the insertion of an intramedullary rod to reduce a femur fracture, followed by screw insertion to lock the distal end of the rod. In this simulation the surgeon can see and manipulate the two femoral bone fragments, and well as the rod inside of the bone. During distal locking, the position of the distal hole and the trajectories for the screw are displayed in their proper locations.

5 Conclusion

Our experience with the two prototypes has shown that Image Overlay produces a very convincing illusion that the user is really seeing objects “inside” of the patient. Anecdotally, users tell us that the system becomes “transparent” and they quickly forget that it is there (and are surprised when they accidentally smash a tool in to the mirror – in the first prototype we resorted to covering the tool tops with foam). While this shows that Image Overlay achieves its desired effect, it also points out a large danger, that the user will tend to believe what he sees, even if it is incorrect. This can be especially problematic when using Image Overlay to observe anatomical structures that aren’t otherwise visible – it’s main strength – since there may be no other cues to its accuracy.

Inaccuracies can arise from a number of sources: the registration can be incorrect, the patient’s anatomy can change from the preoperative data, or the display device can be improperly calibrated. Some of these errors are endemic to all image guided surgical procedures, and are being addressed by several research groups, including ours. Some inaccuracy is more subjective – some users are better able to use the system than others, perhaps due to better stereoscopic perception. We are continuing usability and accuracy studies by having users perform different positioning tasks. The time to complete the task and errors in the final position are measured to characterize both the subjective and objective errors of the system in different configurations [2].

One source of error in our current prototype is the location of the user’s eyes. The system tracks the location of glasses that the user is wearing, and then makes assumptions on where the eyeballs lie behind the glasses and that the eyes are looking straight forward. These assumptions will invariably be approximations. While the error is minimized by distance to the display and patient, it still adds to an overall system error. One possible solution that we are exploring is attaching eye pieces, or perhaps just a nose bridge, to the display/mirror device. This would constrain the user’s eyes to one position. It would also eliminate the need for the user to wear head gear to track the head, and also allow a reduction in the size of the mirror. The system would then be used like a surgical microscope (a “macroscope” in this case), pulled in to the field of view when required, and

pushed away when not needed. This reduces some of the flexibility of the system, but it also simplifies it and improves the accuracy.

Overall, our research has shown Image Overlay to be a promising new data visualization technique with a wide variety of medical applications. Surgical education will probably be one of the first applications of Image Overlay due to the reduced accuracy requirements. Clinical applications will follow, as enhancements to image guided surgical techniques, as the accuracy issues are better understood and addressed.

References

1. W. Bauer, H.J. Bullinger, A. Hinkenjann, O. Riedel: "Planning of Orthopaedic Surgery in Virtual Environments by the Example of Osteotomy Operations," *Interactive Technology and the New Paradigm for Healthcare*, edited by Morgan et al., IOS Press, 1995.
2. M. Blackwell, R.V. O'Toole, F. Morgan, L. Gregor: "Performance and Accuracy Experiments with 3D and 2D Image Overlay Systems," *Proc. Medical Robotics and Computer Assisted Surgery*, 1995.
3. R.D. Bucholz, D.D. Yeh, J. Trobaugh, et al.: "The Correction of Stereotactic Inaccuracy Caused by Brain Shift Using an Intraoperative Ultrasound Device," *CVRMed-MRCAS 97*, pp. 459-466, 1997.
4. M. Deering: "High Resolution Virtual Reality," *Computer Graphics* 26, no. 2, July 1992.
5. D. Drascic, P. Milgram: "Positioning accuracy of a virtual stereographic pointer in a real stereoscopic video world," *SPIE Stereoscopic Displays and Applications II*, Vol. 1457, 1991.
6. W.E.L. Grimson, T. Lozano-Perez, W.M. Wells, et al.: "An Automated Registration Method for Frameless Stereotaxy, Image Guided Surgery, and Enhanced Reality Visualization," *Trans. Medical Imaging*, 1996.
7. H. Iseki, Y. Masutani, M. Iwahara, et al.: "Volumegraph (three-dimensional image-guided navigation), Clinical Application of Augmented Reality in Neurosurgery," *VSMM'96*, 1996.
8. R.V. O'Toole, M. Blackwell, F.M. Morgan, et al.: "Image Overlay for Surgical Enhancement and Telemedicine," *Interactive Technology and the New Paradigm for Healthcare*, vol. 18, pp. 271-273. IOS Press, January 1995.
9. C. Schmandt: "Spatial input/display correspondence in a stereoscopic computer graphics work station," *Computer Graphics*, 17, no. 3, pp. 253-261, July 1983.
10. D. Simon, M. Hebert, T. Kanade: "Techniques for Fast and Accurate Intrasurgical Registration," *Journal of Image Guided Surgery* 1, no. 1, pp. 17-29, 1995.
11. D. Simon, B. Jaramaz, M. Blackwell, et al.: "Development and Validation of a Navigational Guidance System for Acetabular Implant Placement," *CVRMed-MRCAS 97*, pp. 583-592, 1997.
12. J. Tonetti, L. Carrat, S. Lavallée, et al.: "Ultrasound-based registration for percutaneous computer assisted pelvis surgery: Application to iliosacral screwing of pelvis ring fractures," *CAR97*, 1997.