Optical Merger of Direct Vision with Virtual Images for Scaled Teleoperation

Samuel T. Clanton, David C. Wang, Vikram S. Chib, Yoky Matsuoka, *Member*, *IEEE*, and George D. Stetten, *Member*, *IEEE*

Abstract—Scaled teleoperation is increasingly prevalent in medicine, as well as in other applications of robotics. Visual feedback in such systems is essential and should make maximal use of natural hand-eye coordination. This paper describes a new method of visual feedback for scaled teleoperation in which the operator manipulates the handle of a remote tool in the presence of a registered virtual image of the target in real time. The method adapts a concept already used successfully in a new medical device called the Sonic Flashlight, which permits direct *in situ* visualization of ultrasound during invasive procedures. The Sonic Flashlight uses a flat-panel monitor and a half-silvered mirror to merge the visual outer surface of a patient with a simultaneous ultrasound scan of the patient's interior. Adapting the concept to scaled teleoperation involves removing the imaging device and the target to a remote location and adding a master-slave control device. This permits the operator to see his hands, along with what appears to be the tool, and the target, merged in a workspace that preserves natural hand-eye coordination. Three functioning prototypes are described, one based on ultrasound and two on light microscopy. The limitations and potential of the new approach are discussed.

Index Terms—Artificial, augmented, and virtual realities, image display, medical information systems, real time.

1 Introduction

Robotic assistants are currently being introduced into surgery because they hold the promise of aiding or enhancing the capabilities of surgeons to perform more effectively in certain circumstances. One class of surgical assistant is designed to transfer the motions of a surgeon to a different location and scale. These are used to establish operator telepresence for a surgeon at a remote location, to allow procedures to be conducted less invasively, or to otherwise enhance surgical performance. The purpose of our research is to create a new human interface for such systems, one that allows an operator to interact more naturally with a workspace located at a distance and of arbitrary size. Our intent is, as much as possible, to make operating at a different location and scale as easy and natural as performing more traditional local surgery.

The work described in this paper builds on research we have conducted in the real-time superimposition of medical images with a natural view of the patient. The interface that we employ to create the illusion of telepresence is based on the Sonic Flashlight, a device developed in our laboratory that enhances the visualization of ultrasound data. We have previously published descriptions of this device [1], [2], [3], [4], [5] and will only briefly cover it here before describing its extension to scaled teleoperation in the present paper.

The Sonic Flashlight combines an ultrasound transducer, a flat-panel display, and a half-silvered mirror to merge an ultrasound image of the interior of the patient with a natural view of the patient's exterior. The ultrasound image is reflected by the half-silvered mirror in such a way as to be overlaid on the operator's direct view of the patient. Normal stereoscopic vision applies and the merger is correct, regardless of the viewpoint of the observer.

Many approaches to merging medical images with natural sight rely on tracking the patient and/or the observer in order to display the merged medical image at an appropriate angle and location. By strategically placing the mirror, transducer, and display, however, the need for tracking the patient or observer is eliminated. The image of the ultrasound slice, displayed at the correct size, can be reflected such that the virtual image is at the correct location within the patient. The ultrasound data appears to emanate from its actual location.

In the present research, we extend the general approach of the Sonic Flashlight to create a system by which an operator can employ direct hand-eye coordination to interact with a remote environment at a different scale. In the Sonic Flashlight, an ultrasound image is registered with a direct view of the surface of the patient. In the new system, a remote effector is located in the operating field of a patient or other workspace. An image of that remote workspace, displayed at an arbitrary level of magnification, is merged with the direct view of a master instrument held by the operator and linked to the motion of the actual slave effector. The master effector is an appropriately scaled version of a manipulator handle for the slave effector, designed for optimal use in the hand of the operator. The

- *S.T. Clanton is with the School of Medicine, University of Pittsburgh, 443 S. Atlantic Ave, Pittsburgh, PA 15214. E-mail: sclanton@oeic.net.*
- *D.C. Wang is with the School of Medicine, University of Pittsburgh, 4625 Fifth Ave, Apt 405, Pittsburgh PA 15213. E-mail: david@wangmd.com.*
- *V.S. Chib is with the Biomedical Engineering Department, Northwestern University, 345 East Superior Street, Suite 1406, Chicago, IL 60611. E-mail: v-chib@northwestern.edu.*
- *Y. Matsuoka is with the Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213. E-mail: yoky@cs.cmu.edu.*
- *G.D. Stetten is with the University of Pittsburgh, Benedum Hall Room 749, Pittsburgh, PA 15261. E-mail: george@stetten.com.*

Manuscript received 7 Dec. 2004; revised 3 June 2005; accepted 27 June 2005; published online 10 Jan. 2005.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCGSI-0220-1204.

master effector is electromechanically or otherwise linked to the slave effector such that the motion of the master will cause equivalent, scaled motion of the slave effector in the remote workspace. An image of the target from the remote workspace is merged with the operator's view of the master effector in his hand, and it appears to the operator that he is interacting directly with the remote, scaled environment.

In the following section, we review some of the other approaches used to merge visual and medical images, as well as related methods, leading to the Sonic Flashlight and its adaptation to teleoperation. Then, we describe the prototypes we have built to demonstrate the adaptation to scaled teleoperation. We end with a discussion of the potential uses and limitations of this new technique.

2 Augmented Reality Approaches

The innovation described in this paper derives from an extensive body of prior work whose goal has been to look directly into the human body in a natural way. From the discovery of X-rays more than a century ago, clinicians have been presented with a broad assortment of imaging modalities capable of yielding maps of localized structure and function within the human body. Great advances continue to be made in magnetic resonance imaging (MRI), computerized tomography (CT), positron emission tomography (PET), single photon emission computerized tomography (SPECT), ultrasound, confocal microscopy, and optical coherence tomography (OCT). Each of these is a *tomographic* imaging modality, meaning that the data is localized into voxels rather than projected along lines of sight as are conventional X-ray images. Tomographic images, with their unambiguous voxel locations, are essential for our present method of merger with natural vision.

New techniques to display tomographic images directly within the patient have lagged behind the development of the imaging modalities themselves. In the practice of medicine, the standard method of viewing an image is still to examine a photographic film or an electronic screen rather than to look directly into the patient. Previous experimental approaches to fuse images with direct vision have not met with widespread acceptance, in part, because of their complexity. Our approach is simpler and, thus, we hope, more likely to find its way into clinical practice. If so, the approach could have a broad impact on the use of imaging in interventional diagnosis and the treatment of disease.

2.1 Tracking and Head Mounted Displays

Previous methods to fuse images with direct vision have generally relied on tracking devices and, on an apparatus borrowed from the virtual reality community, the head mounted display (HMD). State et al. have developed a HMD for ultrasound, combining a direct view of the patient with ultrasound images using miniature video cameras in the HMD and displaying the video and ultrasound images merged on miniature monitors in the HMD [6], [7]. The approach permits a graphically controlled merge, although it also introduces significant reduction in visual resolution. The HMD and the ultrasound transducer must be tracked so that the appropriate perspective can be computed for the ultrasound image at all times. Sauer et al. at the Siemens Corporation has developed an HMD-based ultrasound system along similar lines, but eliminating the room-based tracking used by State et al. in favor of a head-mounted tracking device. This has resulted in faster and smoother tracking [8]. Head-mounted displays, in general, restrict the operator's peripheral vision and freedom of motion, and they isolate the wearer from others in the room. They do, however, permit extensive use of computer vision and graphics techniques to analyze and enhance the video images, not just with the imaging data itself, but also with graphical overlays (cf. Nicolau et al. [9]), in ways that are not possible with the mirror-based systems described in this paper.

A number of researchers have pursued optical merger of images and graphics with direct vision using a half-silvered mirror, instead of the HMD approach. Mirror-based systems generate a virtual image from a display monitor that floats at a fixed location beyond the mirror, visually superimposed with what is actually seen through the mirror. Mirror-based systems for merging physical input devices with a virtual image were proposed as early as 1977 by Knowlton [10]. A subsequent version of this concept by Schmandt included stereo shutter glasses to permit the overlaid information to be perceived as out of plane from the virtual image [11].

In related work, DiGioia et al. have merged real-world images with CT data using a mirror to achieve a reduction in the total apparatus that the operator must wear, compared to the HMD [12], [13]. In their system, called *image overlay*, a large half-silvered mirror is mounted just above the patient with a flat panel monitor fixed above the mirror. Images of CT data on the monitor are reflected by the mirror and superimposed on the view of the patient through the mirror. The operator only needs to wear a small head-tracking optical transmitter so that the three-dimensional CT data can be rendered from his/her particular perspective. Special shutter glasses are needed only if stereoscopic visualization is desired. A second tracking device must be attached to the patient to achieve proper registration between the rendered CT data and the patient. A similar system, using a half-silvered mirror, has been developed by Albrecht et al. [14].

2.2 Real-Time Tomographic Reflection

Hofstein proposed a simpler system for *in situ* visualization in 1980 [15]. He displayed an ultrasound slice in real time, strategically positioning an ultrasound transducer, a half-silvered mirror, and a display such that the virtual image produced by the mirror was registered in space with the ultrasound scan. This eliminated the need for tracking either the observer or the patient. The Sonic Flashlight is an independent rediscovery of this idea, applied to the guidance of interventional procedures.

The lack of tracking with this approach is possible because of the nature of virtual images. The word "virtual" is used here in its classical sense: The reflected image is optically indistinguishable from an actual slice suspended in space. Ultrasound produces a tomographic slice within the patient representing a set of 3D locations that lie in a plane. The image

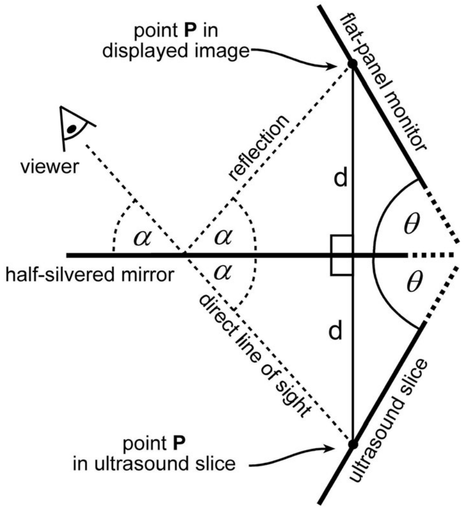

Fig. 1. Configuration of the Sonic Flashlight: A half-silvered mirror bisects the angle between the ultrasound slice (within the target) and the flat-panel monitor. Point P in the ultrasound slice and its corresponding location on the monitor are equidistant from the mirror along a line perpendicular to the mirror (distance = d). Because the angle of incidence equals the angle of reflectance (angle = α), the viewer (shown as an eye) sees each point in the reflection precisely at its corresponding physical 3D location, independent of viewer location.

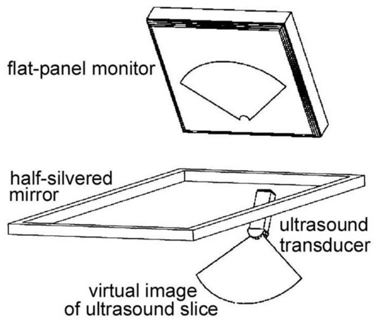

Fig. 2. Schematic representation of the Sonic Flashlight apparatus. A flat-panel monitor and an ultrasound transducer are placed on opposite sides of a half-silvered mirror such that the mirror bisects the angle between them.

of that tomographic slice, displayed at its correct size on a flat panel display, may be reflected to occupy the same physical space as the actual slice within the patient. If a half-silvered mirror is used, the patient may be viewed through the mirror with the reflected image of the slice superimposed, independent of viewer location. The reflected image is truly occupying its correct location within the patient and does not require any particular perspective to be rendered correctly. We have adopted the term *Real-Time Tomographic Reflection* (RTTR) to convey this concept.

To accomplish RTTR, certain geometric relationships must exist between the slice being scanned, the monitor displaying the ultrasound image, and the mirror. As shown in Fig. 1, the mirror must bisect the angle between the slice and the monitor. On the monitor, the image must be correctly translated and rotated so that each point in the image is paired with a corresponding point in the slice to define a line segment perpendicular to, and bisected by, the mirror. By fundamental laws of optics, the ultrasound image will thus appear at its physical location, independent of viewer position. The actual apparatus we have constructed is depicted in Fig. 2.

In Fig. 3, a human hand is seen with the transducer pressed against the soft tissue between the thumb and the index finger. While not a common target for clinical ultrasound, the hand was chosen because it clearly demonstrates successful alignment. The ultrasound image is consistent with the external landmarks of the hand. The photograph cannot convey the strong sense, derived from stereoscopic vision, that the reflected image is located within the hand. This perception is intensified with head motion because the image remains properly aligned from different viewpoints. To one experiencing the technique in person, anatomical targets within the hand visible in the ultrasound would clearly be accessible to direct percutaneous injection, biopsy, or excision.

Superimposing ultrasound images on human vision using RTTR may improve an operator's ability to find targets while avoiding damage to neighboring structures and, generally, facilitating interpretation of ultrasound images by relating them spatially to external anatomy. As such, it holds promise for increasing accuracy, ease, and safety during percutaneous biopsy of suspected tumors, amniocentesis, fetal surgery, brain surgery, insertion of catheters, and many other interventional procedures. We have tested the Sonic Flashlight on phantoms and have recently conducted our first clinical trial on patients to place vascular catheters.

Masamune et al. have demonstrated RTTR on CT data [16]. The application to CT was independently proposed by Stetten [1]. By properly mounting a flat-panel display and a half-silvered mirror above the gantry of a CT scanner, a slice displayed on a flat panel monitor can be reflected by the half-silvered mirror to its correct location within the patient. Assuming the patient remains motionless between

Fig. 3. Photograph, from the viewpoint of the operator, showing a scan of a hand using the apparatus in Fig. 2. The reflected ultrasound image is merged with the direct visual image.

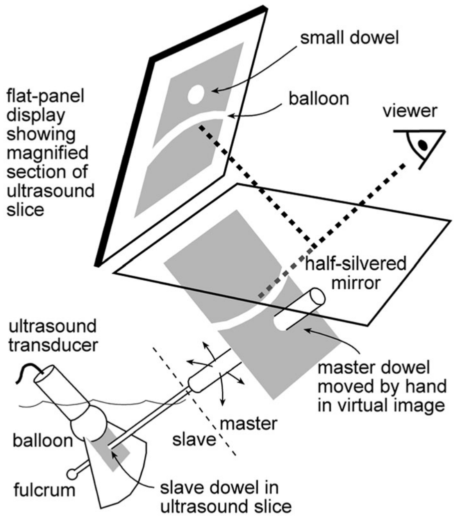

Fig. 4. Apparatus demonstrating magnified remote RTTR, using ultrasound to image a water-filled balloon and a lever to link a master controller to a remote effector at a reduced scale. Moving the dowel indents the balloon, as shown in Fig. 5 and Fig. 6.

the time of the CT scan and the viewing, no tracking is required. However, without repeated scans, the CT image will not be correctly updated during any invasive procedure that changes anatomical structures. The practicality of providing sufficiently continual updates during a procedure is questionable, given the presence of ionizing radiation. Ultrasound does not pose this problem.

2.3 Application of RTTR to Scaled Teleoperation

The application of RTTR to remote scaled procedures was first described by Stetten in 2000 [1]. We have implemented three prototypes, which we describe in the following sections of this paper. The unifying concept is this: The actual target is removed from the operator's immediate workspace, along with the imaging device and the interventional tool. The imaging device still produces a tomographic image of the target and this is displayed on a flat panel monitor, properly scaled, so that its reflection from a half-silvered mirror is registered with a master controller linked to the remote interventional tool. The master-slave system is thus provided with direct visual feedback, allowing the operator to see his or her hand controlling what appears to be an interaction of the tool with the virtual image properly aligned in 3D space.

Using this concept, we intend to develop systems that provide hand-eye coordination and even force-feedback for interventional procedures on patients, animals, tissue samples, and individual cells at mesoscopic and microscopic scales. Interventional procedures could be carried out under a microscope or at the end of a catheter using a robotic linkage. A number of other researchers are presently involved in this pursuit [17], [18], [19], [20], [21], but none has yet, to our knowledge, used an RTTR display. In particular, these systems use a HMD or a real image rather than a virtual image.

3 Magnified Ultrasound Prototype

Our first working demonstration of remote, scaled RTTR uses ultrasound, magnified by a factor of 4, and a simple mechanical master-slave linkage with 2 degrees of freedom, to indent a "remote" water-filled balloon.

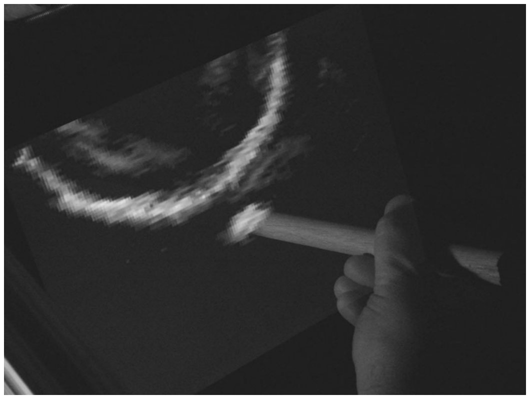

Fig. 5. Master controller (3/4" wooden dowel) interacting with the virtual image of a magnified ultrasound scan of the balloon, seen through the half-silvered mirror.

The apparatus is shown in Fig. 4. Unlike the original sonic flashlight (Fig. 1 and Fig. 2) the ultrasound transducer and the target are no longer in the operator's field of view. The target, instead of being a patient, is now a small water-filled balloon placed before the transducer in a water tank, beyond the direct view of the operator. A lever forms a simple master-slave system, with two degrees of freedom. The lever is a wooden rod formed by attaching a thick (3/4") and thin (3/16") section of wooden dowel, end to end. The thin dowel is attached to the wall of the water tank to create a fulcrum. The operator moves the thick dowel (master controller) through the virtual image, pressing the thin dowel (remote effector) into the balloon, thereby visibly indenting the balloon in the ultrasound image. The fulcrum is four times as far from the virtual image as it is from the actual ultrasound slice, resulting in a mechanical magnification of 4. This matches the ratio between the diameters of the thick dowel (3/4") and the thin dowel (3/16"). A section of the ultrasound slice is magnified by a factor of 4 and displayed on the flat-panel monitor so that the virtual image is reflected to merge visually with the thick dowel (master controller).

Fig. 5 and Fig. 6 show the operator moving the thick dowel to control the thin dowel remotely, producing an

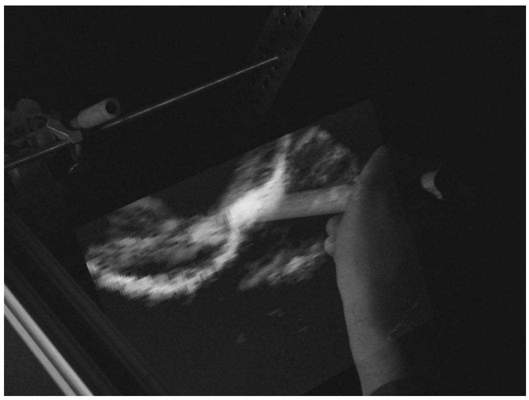

Fig. 6. Result of slave effector (3/16" dowel) pressing into the balloon as visualized by merging the master controller (3/4" dowel) with the virtual image of the magnified ultrasound slice.

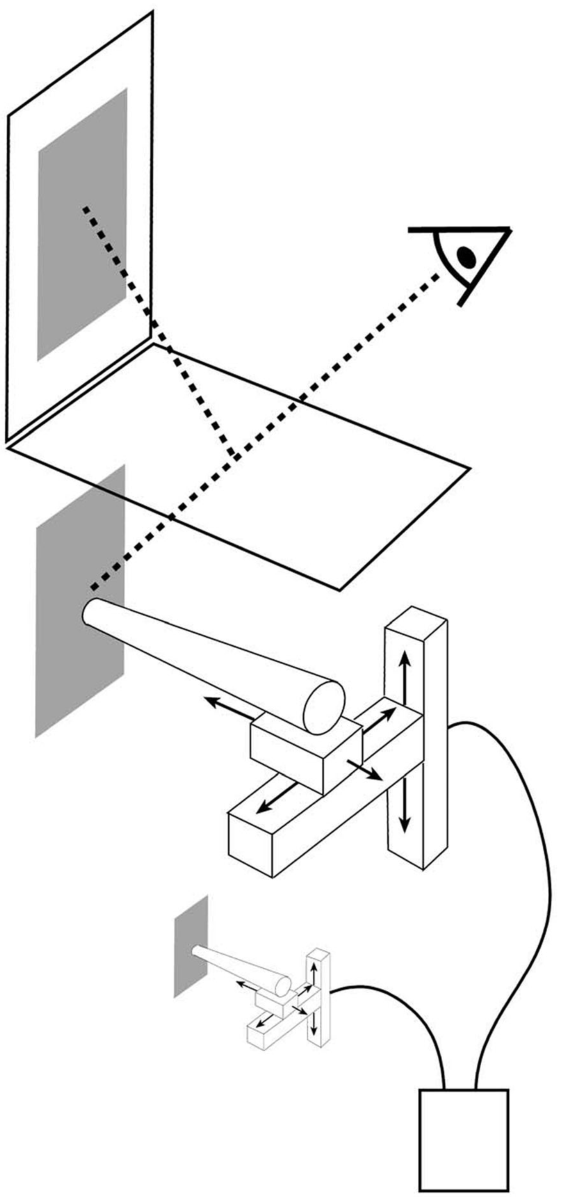

Fig. 7. Abstract illustration of an electromechanically linked system for remote scaled RTTR. The box represents electronic servo link.

indentation in the balloon visible by ultrasound. The pictures are captured with a camera from the point of view of the operator looking through the half-silvered mirror. The operator's hand is shown holding the thick dowel (master controller). A cross section of the slave effector (the thin dowel being scanned in the water tank) is magnified to 3/4″ in the virtual image and accurately tracks the master controller as it appears to cause the indentation in the magnified virtual image of the balloon. The extension of the thin dowel into the water bath is hidden from view by selective lighting.

4 Electromechanical Linkage

The prototype in the previous section demonstrates remote RTTR using a wooden dowel to mechanically link the master controller and the slave effector. Clearly, mechanical linkages have severe limitations for real microscopic manipulation. To create a more practical system, we need to develop electro-mechanical linkages that work on a similar principle, as shown in Fig. 7. A small slave effector (probe) is shown interacting with a tomographic slice (neither the imaging device nor the actual target is shown). A larger master controller is electromechanically linked (box) to the slave. The master and slave are scaled versions of each other and both capable of 3 degrees of translational freedom in this illustration, although rotations could also be incorporated. A semitransparent mirror visually merges a magnified image of the tomographic slice with the master controller using RTTR. The master acts as a servo controller so that the operator manually controls it using hand-eye coordination and the actual slave effector moves accordingly. We have implemented this system in two stages, as described in the following sections (Sections 5 and 6).

5 Simple Light Microscope Prototype

We have implemented a system based on light microscopy that features the basic desired image merging characteristics, although without any electromechanical linkage. The system produces the correct visual illusion of interaction with an environment at 40x magnification. Fig. 8 shows the

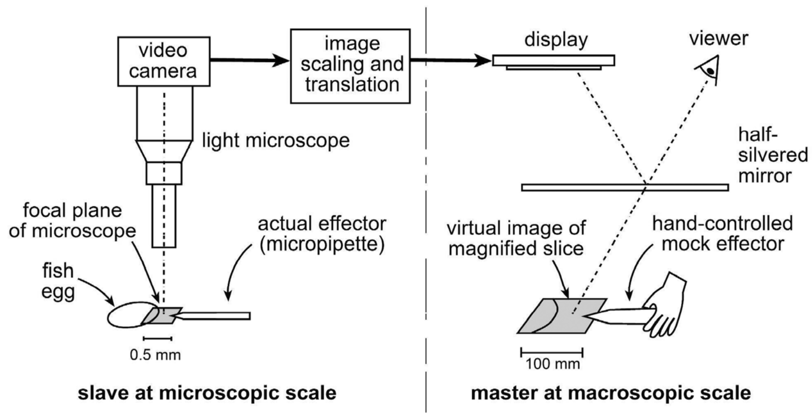

Fig. 8. Diagram of the simple light microscope prototype, demonstrating the visual merger of a master "mock effector" with a magnified image of the target (fish egg). The "mock effector" in this case is just a scaled-up version of the actual effector (a micropipette) and does not really control a master-slave linkage.

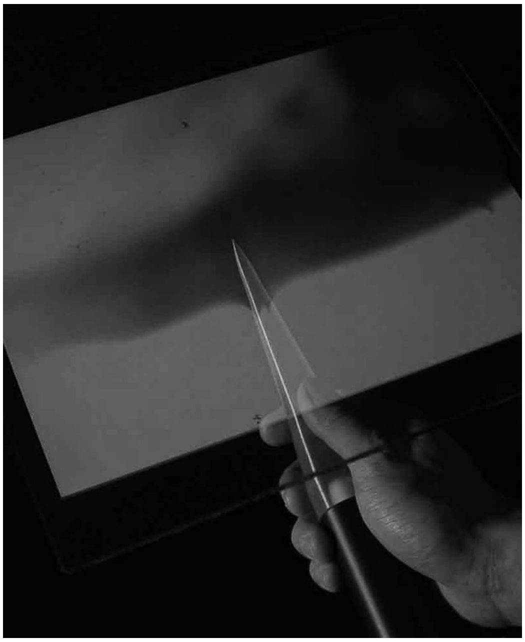

Fig. 9. View through the mirror of the apparatus depicted in Fig. 8. The master controller is a mockup of a micropipette. Although not an actual master-slave system in this example, moving the target (fish egg) into the actual slave effector (micropipette) gives the illusion of piercing the egg with the hand-held device.

apparatus. A fish egg (black caviar) is placed on a microscope slide adjacent to a pulled glass micropipette. The micropipette is fixed to the microscope frame, while the egg can be moved manually with the microscope stage. The video output of the microscope at 40x magnification is displayed at an even greater actual scale on a flat panel monitor. The virtual image is seen below the half-silvered mirror registered with a "mock effector," a scaled up version of the micropipette held in the operator's hand. When the microscope stage is moved toward the micropipette, the fish egg appears to be pierced by the mock effector (see Fig. 9). In a fully functional system, as described next, the mock effector is actually a master controller, linked to a slave effector by a servo.

6 Master Slave "Telepainter" Prototype

Our third prototype system, dubbed the "Telepainter," has been created as an implementation of remote RTTR, again using light microscopy, but this time with an electromechanical master-slave controller. We chose to demonstrate the basic image merge and motion transfer capabilities by implementing a system with which we could paint very small pictures remotely. Although not a clinical application, painting was chosen to demonstrate remote RTTR because it is tolerant of a wide range of forces, while permitting complex hand-eye tasks to be performed.

In this system (see Fig. 10), the workspace of a small robotic arm is viewed as a video image through a surgical microscope (VDI IR Pro video camera attached to a Zeiss OPMI-1 microscope). This image is visually superimposed on the natural workspace of the operator via a half-silvered mirror (34 $\times$ 23 cm) mounted 38 cm above a piece of black paper. A master-slave system is implemented using two SensAble Technologies Phantom haptic interface devices as the master and slave devices. The slave robot arm, a Premium 1.0 model Phantom haptic interface operating with 3 active degrees of freedom, holds a small paintbrush. The master controller is a SensAble 1.5 Premium model Phantom operating passively, with 3 degrees-of-freedom joint-angle encoding, holding a paintbrush handle. The master and slave robot arms are linked such that manual movement of the paintbrush handle (master controller) by the operator produces corresponding movement by the paintbrush (slave effector), scaled down by a factor of 10. A second piece of black paper is placed within the reachable extent of the brush and a small blob of tempura paint is placed on the paper. Photographs of the system are seen in Fig. 11 and Fig. 12.

The system was used to perform Chinese calligraphy with a paintbrush (Fig. 13), enabling the user to paint very small characters (roughly 2 cm square), among other things, while giving the impression of painting much larger characters (roughly 20 cm square). Note the relative size of the penny to the drawing in Fig. 12. To the operator, it seemed that his hand and the paintbrush were connected

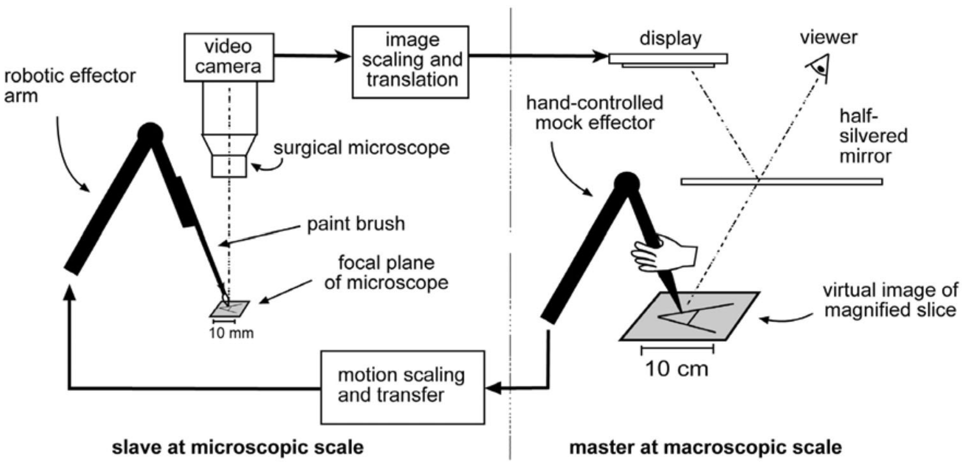

Fig. 10. Apparatus for the "Telepainter" prototype. A master controller paintbrush handle ("mock effector") is linked to a slave effector paintbrush at 1/10 scale. Video images are magnified and registered with the operator's workspace.

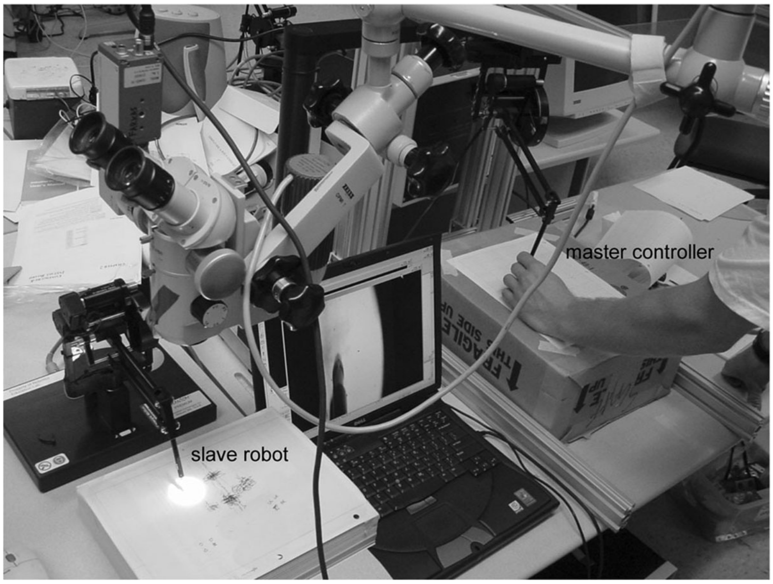

Fig. 11. Telepainter apparatus showing master and slave robots. The operator is manipulating the paintbrush handle held by the master (passive) robot while the slave robot is moving the paintbrush. The paper is white in this photo, though black paper was used during actual operation. The half-silvered mirror and flat-panel monitor over master controller are not shown.

and interacting with the paint and the paper in the remote environment.

It is interesting to note that the SensAble Technologies Phantom slave robot in the system is normally used as a haptic interface device rather than as an effector robot. To implement the scaled motion transfer feature of the system with the Phantom, a Proportional-Integral-Derivative (PID) controller was implemented to control the slave Phantom. Periodic procedures monitored the position of the input and output instruments, and a third periodic procedure used the PID controller to adjust a force on the output Phantom such that it would move to the correct scaled position. The PID parameters were adjusted so that the slave would quickly and accurately track the master. The slave Phantom consistently achieved a position within 0.5 mm of the correct scaled-down position of the input Phantom in the plane of drawing.

Since only 3 degrees of freedom were available for manipulation of the robot, only information about the tip location, without the tool orientation, could be transferred. The input and output devices were kinematically different and working at different scales, so the orientation of the tools between robots was skewed to some degree as the tools moved to the extents of their drawing planes. Using a 7 degree-of-freedom slave robot could overcome this limitation. In addition, the image merge of the position of the tool tips in the plane of drawing was correct from any

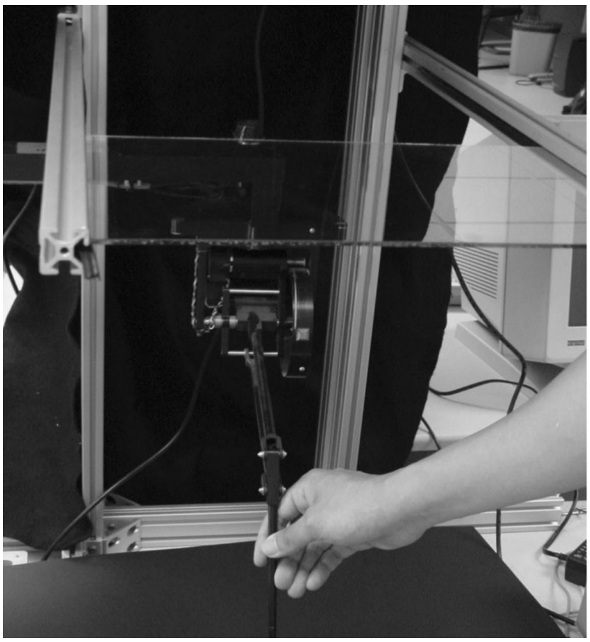

Fig. 12. The master controller is seen with its paintbrush handle beneath the half slivered mirror. Also shown is the black paper in the operator's workspace (no actual paint is placed there). The flat panel monitor (not shown) is mounted above the mirror a distance equal to that between the mirror and the black paper.

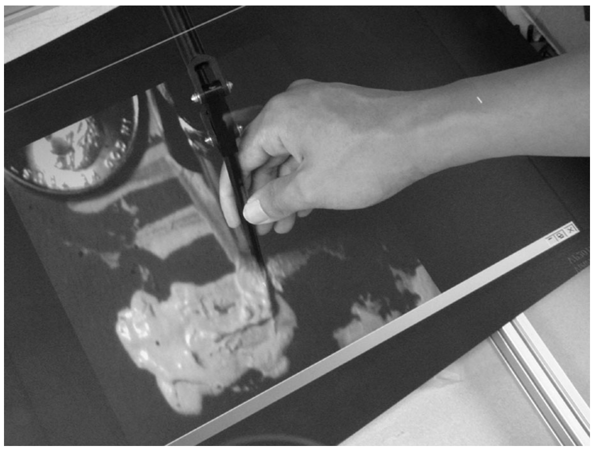

Fig. 13. The Micropainter system as viewed through the half-silvered mirror by the operator, showing the master handle registered with the remote paintbrush and paint. The remote environment is 10 times as small (notice the scale of the penny). Author David Wang is painting his name in Chinese.

viewpoint, but the out-of-plane location of the tools was skewed at different viewpoints. This is a problem inherent in remote RTTR due to the 2D nature of the display, which may, or may not, be counterbalanced by the advantages offered by RTTR over other methods of visualizing remotely controlled procedures.

7 Conclusions and Discussion

We have demonstrated the concept of remote RTTR as an effective method for superimposing visual feedback in real time on the natural workspace of the operator. By merging natural stereoscopic vision with a normal view of one's hands holding the tool, natural hand-eye coordination can be effectively exploited in a remote environment. The lack of a head-mounted display is a further attraction.

The system has possible applications in many areas of medicine, microbiology, and engineering. One can imagine a version in which forceps and needle holder motions are transferred to perform microsurgery, where an operator could manipulate individual cells with a robotically controlled micropipette, or where a machinist could perform microscopic fabrication in an engineering context. An important limitation of the current system for light microscopy is that the visual merge is only viewpoint independent in the plane of the painting. However, for 2D tomographic imaging modalities such as ultrasound or OCT, the visual merge with the master controller would remain accurate throughout the 3D workspace of the operator. Catheter-based procedures and *in vitro* microscopic procedures are particularly appealing candidates for this technology in clinical medicine and biomedical research.

An exciting extension of this approach, currently underway in our laboratory, involves the development of a holographic version of RTTR. Replacing the half-silvered mirror with a holographic optical element would enable greater diversity in the configuration of possible virtual images [22].

Another possible extension involves haptics. The integration of haptic feedback into the instrument linkage would further enhance the immersive environment for performing remote interventional procedures, allowing the operator to use the integrated senses of sight, touch, and proprioception to perform the remote procedures in a natural way.

Acknowledgments

This work was supported by a grant from the US National Institutes of Health R01-EB00860-1, the US National Science Foundation 0308096, and by a grant from the Whitaker Foundation.

References

[1] G.D. Stetten, "System and Method for Location-Merging of Real-Time Tomographic Slice Images with Human Vision," US Patent no. 6,599,247, July 2003.

[2] G.D. Stetten, V.S. Chib, and R.J. Tamburo, "Tomographic Reflection to Merge Ultra-Sound Images with Direct Vision," *Proc. Applied Imagery Pattern Recognition Workshop,* pp. 200-205, 2000.

[3] G. Stetten, V. Chib, D. Hildebrand, and J. Bursee, "Real Time Tomographic Reflection: Phantoms for Calibration and Biopsy," *Proc. IEEE and ACM Int'l Symp. Augmented Reality,* pp. 11-18, Oct. 2001.

[4] G.D. Stetten and V.S. Chib, "Overlaying Ultrasound Images on Direct Vision," *J. Ultrasound in Medicine,* vol. 20, pp. 235-240, 2001.

[5] W. Chang, G.D. Stetten, L. Lobes, and D. Shelton, "Guidance of Retrobulbar Injection with Real Time Tomographic Reflection," *J. Ultrasound in Medicine,* vol. 21, pp. 1131-1135, 2002.

[6] A. State, M. Livingston, W. Garret, G. Hirota, M. Whitton, E. Pisano, and H. Fuchs, "Technologies for Augmented Reality Systems: Realizing Ultrasound-Guided Needle Biopsies," *Proc. ACM SIGGRAPH,* pp. 439-446, 1996.

[7] M. Rosenthal, A. State, J. Lee, G. Hirota, J. Ackerman, K. Keller, E. Pisano, M. Jiroutek, K. Muller, and H. Fuchs, "Augmented Reality Guidance for Needle Biopsies: A Randomized, Controlled Trial in Phantoms," *Medical Image Analysis,* vol. 6, no. 3, pp. 313-320, 2002.

[8] F. Sauer, A. Khamene, B. Bascle, L. Schimmang, F. Wenzel, and S. Vogt, "Augmented Reality Visualization of Ultrasound Images: System Description, Calibration, and Features," *Proc. Int'l Symp. Augmented Reality 2001,* pp 30-39, 2001.

[9] S. Nicolau, A. Garcia, X. Pennec, L. Soler, and N. Ayache, "Augmented Reality Guided Radio-Frequency Tumor Ablation," *Computer Animation and Virtual World,* vol. 16, no. 1, 2005.

[10] K. Knowlton, "Computer Displays Optically Superimposed on Input Devices," *The Bell System Technical J.,* vol. 56, no. 3, pp. 367-383, Mar. 1977.

[11] C. Schmandt, "Spatial Input/Display Correspondence in a Stereoscopic Computer Graphic Work Station," *Proc. ACM SIGGRAPH,* vol. 17, no. 3, pp. 253-261, July 1983.

[12] A. DiGioia, B. Colgan, and N. Koerbel, *Cybersurgery,* R. Satava, ed., pp. 121-139. Wiley, 1998.

[13] M. Blackwell, F. Morgan, and A. DiGioia, "Augmented Reality and Its Future in Orthopaedics," *Clinical Orthopaedic and Related Research,* vol. 345, pp. 111-122, 1998.

[14] K. Albrecht, M. Braun, S. Conrad, G. Goebbels, A. Grab, N. Hanssen, F. Hasenbrink, A. Ivanovic, E. Keeve, Z. Krol, R. Sader, and H.-F. Zeilhofer, "ARSyS-Tricorder: Development of an Augmented Reality System for Intra-Operative Navigation with Respect to the Requirements of Individual Transplant Design in Maxillo-Facial Surgery," Abstract D-0020, European Congress of Radiology, Vienna, Mar. 2002.

[15] S. Hofstein, "Ultrasonic Scope," US Patent no. 4,200,885, Apr. 1980.

[16] K. Masamune, G. Fichtinger, A. Deguet, D. Matsuka, and R. Taylor, "An Image Overlay System with Enhanced Reality for Percutaneous Therapy Performed inside CT Scanner," *Proc. Conf. Medical Image Computing and Computer-Assisted Intervention 2002,* pp. 77-84, 2002.

[17] W. Birkfellner, M. Figl, K. Huber, F. Watzinger, F. Wanschitz, R. Hanel, A. Wagner, D. Rafolt, R. Ewers, and H. Bergmann, "The Varioscope AR—A Head-Mounted Operating Microscope for Augmented Reality," *Proc. Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI),* pp. 869-877, 2000.

[18] P.J. Edwards, D.J. Hawkes, and D.L. Hill, "Augmentation of Reality Using an Operating Microscope for Otolaryngology and Neurosurgical Guidance," *J. Image Guided Surgery,* vol. 1, pp. 172-178, 1995.

[19] B.J. Nelson and B. Vikramaditya, "Visually Servoed Micropositioning for Robotic Micromanipulation," *Microcomputer Applications,* vol. 18, pp. 23-31, 1999.

[20] R. Taylor, W. Robinett, V. Chi, F. Brooks, W. Wright, R. Williams, and E. Snyder, "The Nanomanipulator: A Virtual-Reality Interface for a Scanning Tunneling Microscope," *Proc. SIGGRAPH '93,* pp. 127-134, 1993.

[21] R. Superfine, M. Falvo, R. Taylor, and S. Washburn, "Nanomanipulation: Buckling, Transport and Rolling at the Nanoscale," *CRC Handbook of Nanoscience, Eng., and Technology,* S. Lyshevski, D. Brenner, J. Lafrate, and W. Goddard, eds., CRC Press LLC, year?

[22] A. Nowatzyk, D. Shelton, J. Galeotti, and G. Stetten, "Extending the Sonic Flashlight to Real Time Tomographic Holography," *Proc. AMI-ARCS 2004, Workshop Augmented Environments for Medical Imaging Including Augmented Reality in Computer-Aided Surgery,* Sept. 2004.

Samuel T. Clanton received the BS degree from Johns Hopkins University in biomedical engineering and has been a consultant for the NASA Ames Research Laboratory and Barrett Technologies, Inc. He is an MD-PhD student at the University of Pittsburgh School of Medicine, conducting research at the University of Pittsburgh Department of Bioengineering and the Carnegie Mellon University Robotics Institute.

David C. Wang received the BS and MS degrees in computer science from Cornell University in 2000 and 2001, respectively, and has worked as a software engineer at Oracle, Inc. He is an MD-PhD student at the University of Pittsburgh School of Medicine, conducting research at the Carnegie Mellon Department of Biomedical Engineering.

Vikram S. Chib received the BS degree in bioengineering from the University of Pittsburgh and the MS degree in bioengineering from Northwestern University, where he is currently a doctoral candidate in bioengineering. He has been awarded a US National Institutes of Health NRSA Predoctoral Fellowship.

Yoky Matsuoka received the BS degree in electrical engineering and computer science from the University of California at Berkeley and the MS/PhD degrees in electrical engineering and computer science from the Massachusetts Institute of Technology. She was a postdoctoral fellow in the Biorobotics Laboratory at Harvard University and is currently an assistant professor in the Robotics Institute at Carnegie Mellon University. In 2004, she received a US National Science Foundation Presidential Early CAREER Award. She is a member of the IEEE.

George D. Stetten received the MS degree in biology from New York University in 1986, the MD degree from the State University of New York at Syracuse in 1991, and the PhD degree in biomedical engineering from the University of North Carolina at Chapel Hill in 1999. He was a research assistant professor in biomedical engineering at Duke University and is currently an associate professor in bioengineering at the University of Pittsburgh and a research scientist in robotics at Carnegie Mellon University. He is a member of the IEEE.