

Case Study: Observing a Volume Rendered Fetus within a Pregnant Patient

Andrei State, David T. Chen, Chris Tector, Andrew Brandt, Hong Chen,
Ryutarou Ohbuchi, Mike Bajura and Henry Fuchs

University of North Carolina

Abstract

Augmented reality systems with see-through head-mounted displays have been used primarily for applications that are possible with today's computational capabilities. We explore possibilities for a particular application—in-place, real-time 3D ultrasound visualization—without concern for such limitations. The question is not “How well could we currently visualize the fetus in real time,” but “How well could we see the fetus if we had sufficient compute power?”

Our video sequence shows a 3D fetus within a pregnant woman's abdomen—the way this would look to a HMD user. Technical problems in making the sequence are discussed. This experience exposed limitations of current augmented reality systems; it may help define the capabilities of future systems needed for applications as demanding as real-time medical visualization.

1 Introduction

Interpreting 3D radiological data is difficult for non-experts because understanding spatial relationships between patient and data requires mental fusion of the two. Volume rendering has been useful for visualization but has typically been viewed separately from the patient.

Ideally one would like to see directly inside a patient. Ultrasound echography allows dynamic live scanning and patient-doctor interaction; an augmented reality system displaying live ultrasound data in real time and properly registered in 3D space within a scanned subject would be a powerful and intuitive tool; it could be used for needle-guided biopsies, obstetrics, cardiology, etc.

2 Previous work

Many researchers have attempted visualization of 3D echography data [1, 2, 3, 4]; some have volume visualized data sets that were acquired as a series of hand-guided 2D echography slices with 6 DOF [5, 6]. Compared to

echography imaging by current state-of-the-art 2D scanners, such volume visualizations promise to reduce the difficulty of mentally combining 2D echography slices into a coherent 3D volume.

However, almost all of these systems have used conventional stationary video monitors for presentation, so that a user must still mentally fuse 3D volume images on the monitor with the 3D volume of the patient.

One system [7] tried to visualize live 2D echography images in-place within the patient using a see-through HMD system. While that system demonstrated the initial concept of “augmented reality,” it could show only a few image slices (no 3D volume) at a relatively low frame rate. Those few ultrasound images appeared to be pasted *in front* of the patient's body rather than fixed within it.

3 Near-real-time visualization system

In January 1993 we attempted to improve upon [7] with a system designed to perform real-time, in-place volume visualization of a live human subject. It contained two major real-time features: a continuously updated and rendered volume, and an image compositor. The volume was updated from a series of 2D echography images acquired by a tracked ultrasound transducer. The image compositor combined each frame of the volume rendering with a live HMD video image of the patient (Plate 1).

Unfortunately, due to the requirements of real-time operation, the performance was seriously inadequate. The system's major shortcomings were:

- the ultrasound acquisition rate was only 3 ultrasound frames per second, causing both temporal *and spatial* undersampling of the volume,
- the reconstructed volume was crudely sampled (to 100 x 100 x 100), and could not be updated at more than 1 ultrasound slice per second,
- the volume was rendered at low resolution (65 x 81 rays cast) and interpolated to the display resolution of (512 x 640) to achieve 10 fps,
- the tracking system resolution and accuracy were poor, with significant lag and noise in the data.

As a result of these problems, the (interactive) volume

renderings displayed by this system were unrecognizable. The scans of a nearly full-term fetus did not reveal human-like forms to any member of the research team other than the M.D. ultrasonographer.

4 Hybrid real-time / off-line system

To answer the question “How much better would the visualization be if the real-time computational demands were met by the available resources,” this system was designed so that the most computationally expensive tasks (volume reconstruction and rendering) are done off-line.

Figure 1 shows the steps involved in generating a video sequence that combines volume rendered echography data with HMD camera visuals. The figure shows the dependency of various tasks on one another. The tasks can be grouped roughly into calibration tasks performed prior to scanning a subject, data acquisition tasks performed during the scan, and image generation tasks which are a post process.

4.1 Calibration

Calibration procedures compute several sets of parameters. One such set relates echography pixels to the transducer tracker origin; together with the transducer tracker position, it determines the location of the echography pixels in 3D world space. Similarly, the camera position and orientation relative to the HMD tracker origin must be determined; the optical distortion of the camera lens must be modeled so that the CG imagery can be made to match the camera images.

Echography pixel to tracker calibration. This function (which is not necessarily linear in pixel space) is measured by imaging a point target (a 4 mm bead obtained from GE medical systems and suspended at the tip of a pin in a water tank) at a known location relative to the transducer. Plate 2 shows the calibration setup for the ultrasound transducer. The transducer is attached to a precision translation stage which moves under computer control to chart out the point spread functions of the echography pixels (this function grows with distance from the transducer tip). The 2D transducer slice was measured to be rotated 2.3 degrees from the transducer's axis.

Camera calibration. Position and orientation of the HMD camera relative to the HMD tracking origin are determined by an iterative semi-automatic method. The optical distortion of the lens is determined by imaging a grid pattern (Plate 2, inset). A circularly symmetric model based on a 5th degree polynomial is used.

Our optical tracking system is described in [8]; the calibration methods used for it are described in [9].

4.2 Real-time acquisition

Unburdened by visualization processing needs, this phase takes place in *true* real-time. Both the ultrasound images and the HMD camera images are recorded at 30 fps on Sony D2 digital tape recorders. At the same time, tracking data for the ultrasound transducer and for the HMD is saved on a UNIX™ workstation which controls the D2 recorders. With each tracker record the time code of the corresponding video frame is stored for later synchronization, thus establishing correspondence between the tracking and video data streams.

To create an illusion of the visualized volume

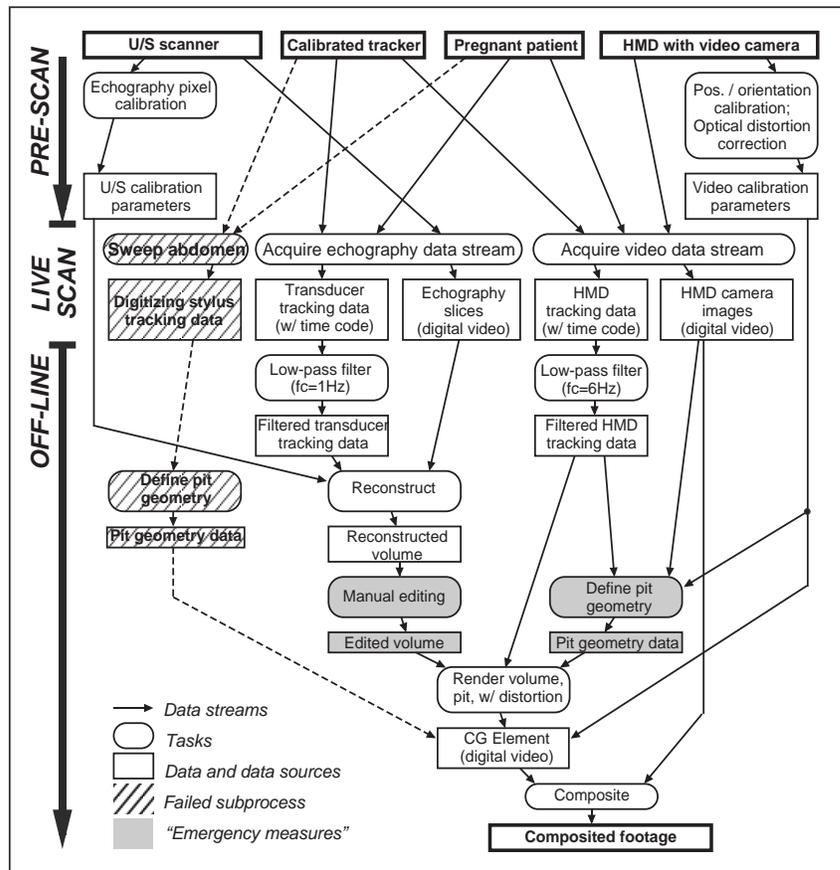


Figure 1: Flow diagram for hybrid experiment combining real-time acquisition with off-line visualization. The top row shows the “basic ingredients” of our system.

residing inside the abdomen, a (polygonal) model of a "pit" must be created; the pit must conform to the shape of the abdomen and be placed at the correct position. To achieve this, the geometry of the abdomen is acquired by making a zig-zag sweep of the abdomen. The tip of the tracked transducer is used as a 3D digitizing stylus.

4.3 Off-line image generation

In this phase volume visualized echography images and the images captured by the HMD-mounted camera are combined into composite HMD viewpoint imagery. The major steps in generating the composite are:

Tracking noise filtering. The tracking data we acquire fluctuates even if the tracked target remains stationary. Such noise causes misregistration between video and CG imagery. To reduce this effect, a non-causal low-pass filter without phase shift is used [10], with cut-off frequencies of 1 Hz for the transducer tracker and 6 Hz for the HMD tracker.

Reconstruction. The echography pixels are positioned and resampled into a regularly gridded volume, using a simple approximation algorithm based on a linear combination of Gaussian weighting functions which are translated and scaled to minimize artifacts [6, 11]. Size and shape of an echography pixel in world (or tracker) space are approximated by a point spread function which falls off away from the world space position as a non-spherical Gaussian. Since an image frame and its tracking information are related by the frame's time code, digitization of echography frames from video tape and volume reconstruction can take place automatically on a workstation under program control.

Visualization. A volume renderer *vol2* [12] running on a graphics multicomputer *Pixel-Planes 5* is used to render the reconstructed volume(s) from viewpoints matching those of the HMD-mounted camera. By modulating the direction of rays, the images are distorted based on the model described above; the polygonal pit, built using the data from the abdomen geometry sweep, is included in the rendering. The images are recorded in single-frame mode onto digital video tape.

Compositing. Camera and CG images are mixed by chroma-keying on a Sony video mixer. The mixer replaces blue background in the CG frames by HMD frames; the time codes recorded during HMD acquisition are used to ensure synchronization of the 2 elements.

5 Live subject experiment and results

In January 1994 we scanned two pregnant subjects with the hybrid system. Since the current tracking setup allows only one target at a time, the abdomen sweep data, the ultrasound data (Plate 3, left) and the head camera

data (Plate 4) had to be recorded in 3 successive passes. The patients had to remain motionless throughout the acquisition phase.

From the acquired ultrasound imagery, we selected and digitized a short sequence of about 15 seconds, during which the ultrasonographer had made a continuous sweep (455 slices) of the fetus from the middle of the skull down to the bottom of the hip (Plate 3, right). The slices were reconstructed into a 165 x 165 x 150 volume with a resolution of 8.2 voxels/cm (or a voxel size of .122 cm³), which is better than the highest resolution of the ultrasound machine/transducer combination.

After reconstruction, objects such as uterus and placenta were edited out manually using an editing mask with Gaussian fall-off to avoid introducing artifacts in the volume. Finally a small 3D Gaussian filter (standard deviation 2 voxels) was applied to the volume.

The abdomen geometry sweep had failed due to a minor technical problem during the live scan; we derived abdomen geometry data by triangulating a number of small structures visible on the skin of the abdomen in frames videographed from different viewpoints and extracted from the HMD video sequence.

The reconstructed and edited volume (Plate 5) and the pit were rendered by *vol2* with optical distortion (Plate 5, inset); the CG sequence was combined with the HMD camera sequence (Plate 6).

6 Conclusion and future directions

Many aspects of our experiment suffered from lack of immediate, real-time 3D feedback. During HMD video acquisition, the HMD wearer could not really see inside the patient; thus we unfortunately ended up looking at the patient from the "wrong" side and thus viewing the fetus from behind in the resulting composite video sequence. During echography acquisition, we were unable to gather enough echography slices to reconstruct a complete fetus due to lack of real-time 3D feedback on the geometry of the scanned anatomy. Still, in separately generated images from viewpoints chosen more advantageously than those of the HMD camera during the HMD video acquisition, one can recognize more anatomical features (Plate 5).

What resources would be required to present an on-line visualization of comparable quality in an augmented-reality system? We expect advances in volume rendering software and hardware to soon provide high-speed stereoscopic rendering of volumetric data sets (see for example [13]). As for reconstruction, our volume contained nearly 4 times as many voxels as the one used in the 1993 experiment. Since the latter was being reconstructed at a rate of about 1 Hz, we need an increase of 2 orders of magnitude in computational speed for the

reconstruction subsystem.

We learned from the January 1993 experiment and others that, besides the image generation frame rate, short lag in both volume reconstruction and visualization is very important. Our hybrid system has avoided this problem through off-line processing. For an on-line real-time system, however, we need to design and implement hardware and algorithms that provide not only high throughput but also short lag. In addition, we need fast, minimal-lag, high-precision tracking.

In the area of our application—visualizing ultrasound as a “flashlight” into the body—we conclude that a step forward has been achieved. The sequence showing a fetus registered inside the pregnant subject provides a “brass standard” (if not a gold one) as a target for our next real-time efforts. One way to acquire the high amounts of computing resources needed for such efforts is through the current research on high-bandwidth links between powerful computing stations, which hints at computational capabilities that might be available in the next few years and are within the desired range of power.

In general, complex visualizations presented within augmented vision systems make greater application demands than either virtual environments or scientific visualization individually. Any closed virtual environment or scientific visualization system lacks the error-emphasizing cues that a combined system provides. An augmented reality application provides sufficient information to enable the user to easily notice registration errors, tracker lag, computational errors, calibration errors and real time delays which destroy the attempted illusion. Augmented systems or any virtual reality system dealing directly with the real world will not be easy to create. Simple applications with little computational demand, such as overlays for wiring guides or informational pointers, will be able to get by, but applications with complex visualization goals will be heavily burdened by the demands. Researchers should be sensitive to these issues and attempt to evaluate carefully their impact in the intended application.

7 Acknowledgments

Jack Goldfeather provided mathematical models for the echography calibration procedure. Gary Bishop wrote filtering software for the tracking data. Nancy Chescheir, M.D. and Vern Katz, M.D. were our ultrasonographers. Kathryn Y. Tesh and Eddie Saxe provided experimental assistance. Scott Shauf created an early version of figure 1. We thank our anonymous subject for her patience.

Funding was provided by ARPA (ISTO DABT 63-92-C-0048 and ISTO DAEA 18-90-C-0044) and by CNRI / ARPA / NSF / BellSouth / GTE (NCR-8919038).

8 References

1. McCann, H.A., Sharp, J.S., Kinter, T.M., McEwan, C.N., Barillot, C., and Greenleaf, J.F. “Multidimensional Ultrasonic Imaging for Cardiology.” *Proc. IEEE* 76.9 (1988): 1063-1073.
2. Lalouche, R.C., Bickmore, D., Tessler, F., Mankovich, H. K., and Kangaraloo, H. “Three-dimensional reconstruction of ultrasound images.” *SPIE’89, Medical Imaging*. SPIE, 1989. 59-66.
3. Pini, R., Monnini, E., Masotti, L., Novins, K. L., Greenberg, D. P., Greppi, B., Cerofolini, M., and Devereux, R. B. “Echocardiographic Three-Dimensional Visualization of the Heart.” *3D Imaging in Medicine*. Ed. Fuchs, H. Hühne, K. H., and Pizer, S. M. *NATO ASI Series*. Travemünde, Germany: Springer-Verlag, 1990. F 60: 263-274.
4. Tomographic Technologies, GMBH. *4D Tomographic Ultrasound, A clinical study*. 1993.
5. Ganapathy, U., and Kaufman, A. “3D acquisition and visualization of ultrasound data.” *Visualization in Biomedical Computing 1992*. Chapel Hill, NC: SPIE, 1992. 1808: 535-545.
6. Ohbuchi, R., Chen, D., and Fuchs, H. “Incremental volume reconstruction and rendering for 3D ultrasound imaging.” *Visualization in Biomedical Computing 1992*. Chapel Hill, NC: SPIE, 1992. 1808: 312-323.
7. Bajura, M., Fuchs, H., and Ohbuchi, R. “Merging Virtual Objects with the Real World: Seeing Ultrasound Imagery within the Patient.” *Computer Graphics (Proceedings of SIGGRAPH’92)* 26.2 (1992): 203-210.
8. Wang, J., Chi, V., and Fuchs, H. “A Real-Time Optical 3D Tracker for Head-Mounted Display System.” *Computer Graphics (Proceedings of 1990 Symposium on Interactive 3D Graphics)* 24.2 (1990): 205-215.
9. Gottschalk, S. “Autocalibration for Virtual Environment Tracking Hardware.” *Computer Graphics (Proceedings of SIGGRAPH’93)* 27 (1993): 65-72.
10. “Digital Low-Pass Filter Without Phase Shift.” *NASA Tech Briefs KSC-11471*. John F. Kennedy Space Center, Florida.
11. Ohbuchi, R. “Incremental Acquisition and Visualization of 3D Ultrasound Images.” Doctoral dissertation. University of North Carolina at Chapel Hill, Computer Science Department, 1994.
12. Neumann, U., State, A., Chen, H., Fuchs, H., Cullip, T. J., Fang, Q., Lavoie, M., and Rhoades, J. “Interactive Multimodal Volume Visualization for a Distributed Radiation-Treatment Planning Simulator.” *Technical Report TR94-040*. University of North Carolina at Chapel Hill, Computer Science Department, 1994.
13. Cullip, T. J., and Neumann, U. “Accelerating Volume Reconstruction With 3D Texture Hardware.” *Technical Report TR93-027*. of North Carolina at Chapel Hill, Computer Science Department, University 1994.

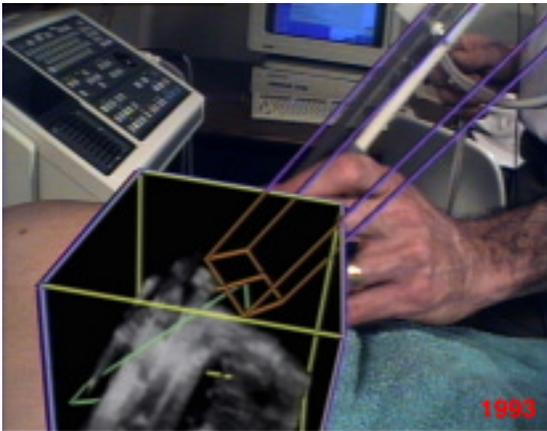


Plate 1: HMD view in near-real-time system with on-line concurrent processing.

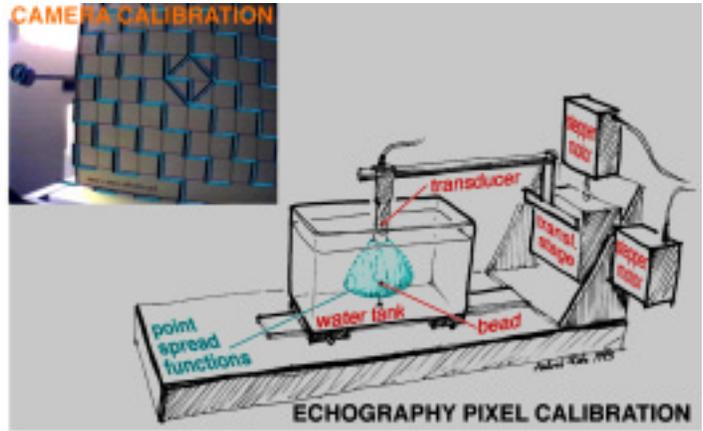


Plate 2: Ultrasound transducer calibration setup. Inset—HMD view during camera calibration run.

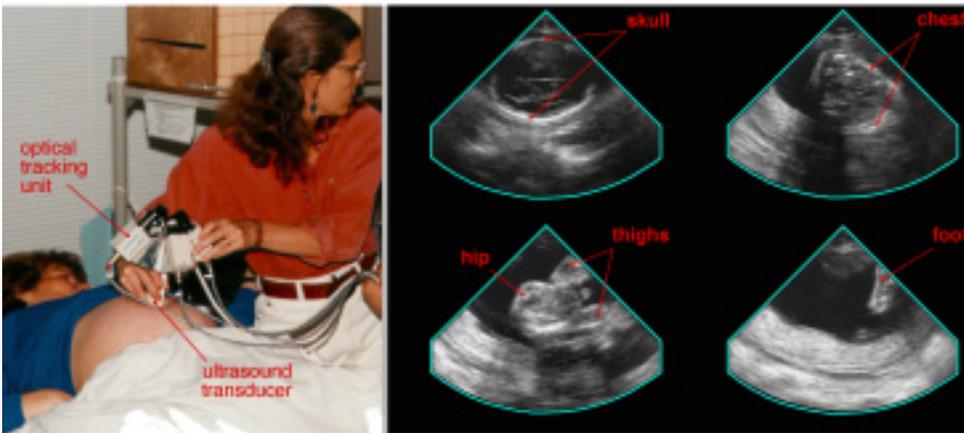


Plate 3: Left—ultrasonographer acquiring echography slices. Right—4 of the 455 slices used for volume reconstruction.



Plate 4: Acquiring HMD camera video.

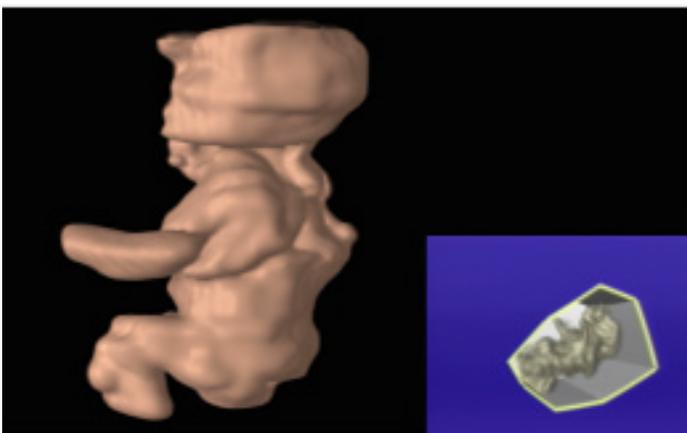


Plate 5: Reconstructed fetus data set. Inset—CG frame matching the HMD camera view in Plate 4.



Plate 6: View of volume rendered fetus within pregnant subject.