3D Stereo Interactive Medical Visualization

Damien Maupu
Ecole Supérieure de Chimie Physique Electronique de Lyon

Mark H. Van Horn, Susan Weeks, and Elizabeth Bullitt
University of North Carolina at Chapel Hill

Endovascular operations involve guidance of a catheter through a vascular network with the goal of delivering therapy to a specific target. This form of minimally invasive surgery has undergone rapid growth over the last 10 years and is now the procedure of choice for many patients who could only be treated by more invasive, open operations.

A primary difficulty with endovascular procedures is that they're typically guided only by flat-projection images. To image the vasculature, the interventionalist injects small amounts of a contrast agent and obtains fluoroscopic images. These images generally show only the connected portions of the vasculature downstream from the catheter tip. Thus, the interventionalist can't visualize any structures disconnected from the vascular anatomy shown in any particular image. Moreover, all 3D information is collapsed into a projection view, making 3D perception of the vascular network difficult.

We describe a new method of guiding endovascular procedures using interactive 3D stereo visualizations. We use as an example the Transjugular Intrahepatic Portosystemic Shunt (TIPS) procedure.[1] Our goal is to increase the speed and safety of endovascular procedures by providing the interventionalist with 3D information as the operation proceeds. The visualization methods employed aren't novel, but they haven't been used for the guidance of endovascular surgery and could change how clinicians perform these procedures.

Why is the TIPS procedure difficult?

TIPS is an endovascular procedure used to treat the effects of severe liver disease. The goal of the operation is to create a direct conduit between the portal vein (the venous system that supplies the liver) and the hepatic vein (the venous system that drains the liver). This connection allows decompression of elevated portal vein pressures and subsequent relief of many symptoms.

Needle insertion during the TIPS procedure has traditionally been guided with biplane fluoroscopic images acquired close to 90 degrees apart. A particular problem with TIPS, however, is that a needle must be inserted through the liver into the portal venous system, and we can only visualize these target vessels transiently by injecting CO_2. We then apply this static image to the real-time fluoroscopic image as an overlay, to provide an estimate of the location of the portal vein during needle passage.

During needle passage, however, the target veins can't be seen directly, and both respiratory motion and needle pressure may displace the liver and its portal veins. The operation therefore involves insertion of a needle toward a target we can't see and whose position changes as a result of the intervention. Figure 1a shows a typical pair of fluoroscopic images during needle insertion where the target portal vein has been outlined in black on each image. An arrow indicates the projection of the needle on each view.

Our goal is to provide 3D image guidance of the TIPS procedure so that the interventionalist can readily adjust the needle position and trajectory to reach the target on the first pass (see Figure 1b). We propose a 3D stereo display of the interventionalist's needle and target vessels. We also add interactivity via head tracking so that the interventionalist gains a better 3D sense of

> Our interactive, 3D stereo display helps guide clinicians during endovascular procedures, such as intraoperative needle insertion and stent placement relative to the target organs.
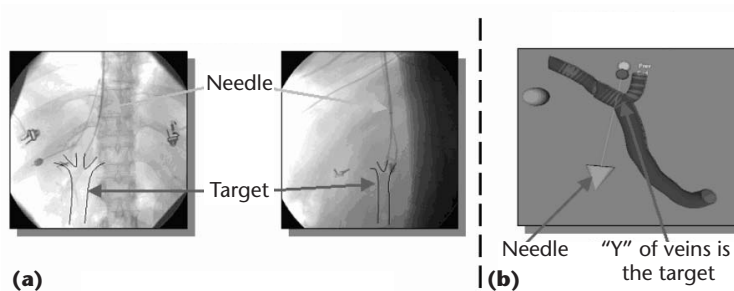

1 Comparison between (a) the normal procedure with 3D information on two biplane images and (b) the TIPS project with 3D stereo interactive images.

the relationship between the target vessels and the needle during needle advancement.

Segmentation and registration

Our approach requires the preoperative creation of a 3D model of each patient's portal venous system. We extract the images of these vessels from a preoperative contrast-enhanced computed tomographic scan (CTA) using the segmentation method of Aylward.[2] We then postprocess the segmented vessels to provide a set of connected vessel trees.[3]

2 Overview of the TIPS procedure. (a) Preoperative computed tomographic scan, (b) calibration and registration, (c) balloon and needle tracking, and (d) 3D reconstruction and visualization. (AP stands for anterior–posterior and LAT stands for lateral.)
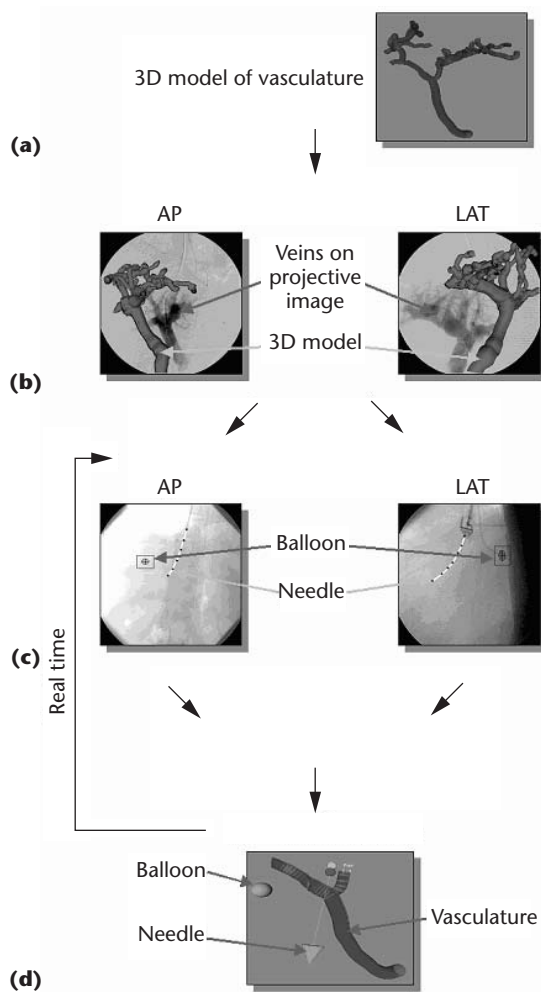
At the beginning of each operation, the two fluoroscopic views are calibrated so that the relative positions and orientations of the two views are known.[4] An 11-mm, contrast-filled balloon is also introduced endovascularly and is firmly wedged into a branch of the liver's hepatic venous system. This balloon can be seen on each fluoroscopic image (see Figure 2) and serves as a reference point that clinicians can use later to indicate rigid-body motion of the liver (and its enclosed portal veins) during respiration, heartbeats, and needle insertion.[5]

Prior to needle insertion, we can visualize the portal venous system transiently by CO_2 injection. This injection provides a snapshot on each of the two fluoroscopic images of both the balloon and of the target veins. We can then register the 3D vessel model with this pair of fluoroscopic views, using an iterative closest-point method that minimizes distance between the projections of 3D vessel skeleton curves and the skeleton curves of the 2D curves seen on the fluoroscopic images.[6] Although we can't see the portal venous system thereafter on the fluoroscopic images, we can estimate the displacement of the 3D portal veins during needle insertion by incorporating the changes in the position of the balloon, which we can track via fluoroscopy.

The interventionalist's needle is tracked on both fluoroscopic views throughout the procedure[7] and is reconstructed into 3D[8] using the principles of epipolar geometry.[9] The interventionalist's needle can then be shown in 3D and relative to the 3D portal veins throughout the procedure.

3D stereo display creation

We propose providing the clinician with a 3D stereo display of both the target portal veins and of the needle. Our goal is to provide a lightweight, inexpensive, easy-to-use system that will allow effective 3D visualization without impeding the clinician's ability to manipulate a catheter or to observe the actual fluoroscopic images.

Image creation

The stereo display provides each eye with a slightly different image at a high rate. To generate the 3D images, we use OpenGL and its ability to draw into two buffers: in our case, a right and a left buffer. The 3D information displayed includes the 3D vasculature and the clinician's needle reconstructed into 3D. A separate view is generated for each eye using the two OpenGL functions gluLookAt and glFrustum. The gluLookAt method sets the position of the camera. The glFrustum method sets the field of view of the camera. We successively display 3D objects in the left and right buffers, with synchronization between the left buffer and the stereo hardware controlled by the graphics card.

The function glFrustum requires six settings: left, right, bottom, top, near, and far (see Figure 3). Parameters near and far specify the distances to the near and far clipping planes.
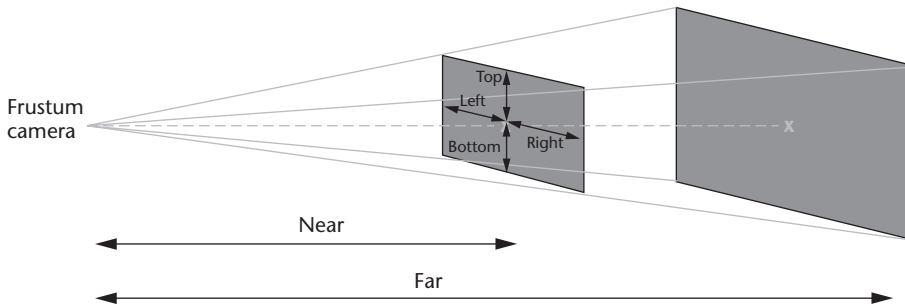
3 OpenGL frustum camera settings: glFrustum (left, right, bottom, top, near, and far).

Both distances must be positive. Parameters left and right specify the coordinates for the left and right vertical near-clipping planes. The parameters bottom and top specify the coordinates for the bottom and top horizontal near-clipping planes. The function glFrustum produces a perspective projection matrix. The points (left, bottom, and near) and (right, top, and near) specify the points on the near clipping plane that map to the lower-left and upper-right corners of the window.

The correct way to set the cameras for the left and right eyes is to use parallel axis, asymmetric frustum-perspective projection,[10] in which the view vectors for each eye remain parallel (see Figure 4). From a virtual camera placed between the two eyes, we can set the left and right eyes' glFrustum. The only information we need is the focal distance, the eye separation, the aspect ratio (width:height—we chose 4:3 to match the screen dimensions), and the field of view, plus near and far clipping planes. Then we can derive the settings of the gluLookAt and glFrustum using similar triangles.

For purposes of clarity, we provide more detail about the left eye's camera. The parameters near and far are the same as in our virtual camera. The parameters top and bottom are also the same:

$$top = near * \tan(FieldOfView/2)$$
$$bottom = top$$

The only difficulty is to compute the left and right parameters, which we label in Figure 4 as d_{left} and d_{right}. Here is how we process the parameters. From the camera setting, we can find distance a:

$$a = ratio * \tan(FieldOfView/2) * FocalLength$$

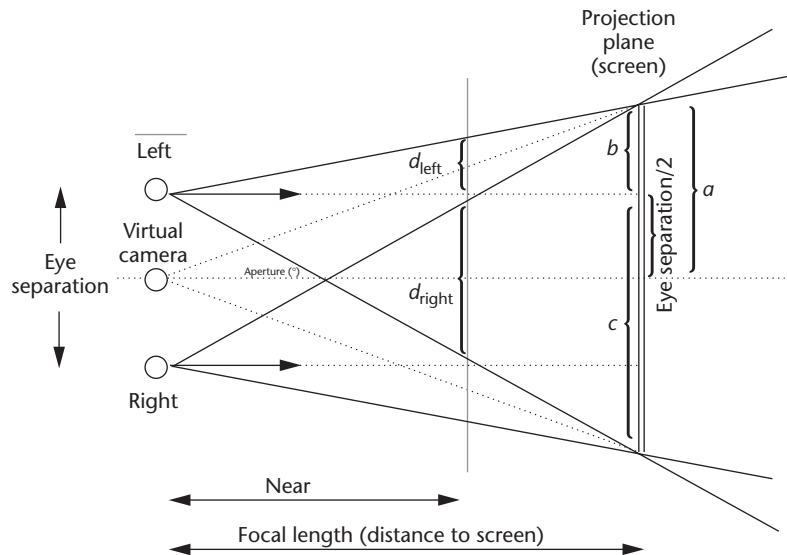From this we derive distances b and c:

$$b = a - EyeSep/2$$
$$c = a + EyeSep/2$$

The link between b, c, and respectively d_{left} and d_{right} is given by similar triangles:

$$\frac{d_{left}}{b} = \frac{d_{right}}{c} = \frac{near}{FocalLength}$$

We now have the left and right parameters for the left eye. We then make similar computations for the right eye.

Stereo hardware

The hardware used to create the 3D stereo display includes a StereoGraphics ZScreen 2000i polarized screen with passive glasses, a ViewSonic P95F 19-inch CRT monitor, and an Nvidia Quadro 4 graphics card. The ZScreen is a polarized screen, which is sold together with a lightweight pair of passive polarized glasses. The screen can be polarized in two directions. Each direction corresponds to the polarization of one of the eyes of the glasses. Each eye of the glasses alternately becomes opaque and allows visual input along the correct polarization. By synchronizing the ZScreen polarization with the image displayed at a high refresh rate (120 Hz for the screen, 60 images/second for each eye), the viewer perceives a 3D stereo image. We mount the ZScreen on the front of the monitor. Synchronization is provided by the Nvidia graphics card.

We selected this particular set of hardware for the guidance of endovascular surgery for several reasons. First, the system is wireless, so the physician isn't tethered by wiring, which might hamper movement. Second, the glasses are lightweight and comfortable. Because of the low optical density of the lenses, the physician's vision of the operating field and fluoroscopic images is relatively unaffected. Finally, we chose to use relatively inexpensive, commercial off-the-shelf components. The total cost of the system, not including the PC, is less than $5,500.

Adding interactivity

It's important to add interactivity to the software so that the interventionalist can make changes to the display. The software allows the user to rotate the vasculature in any direction within a fixed coordinate system. This is useful because the clinician's intuitive sense of object movement is preserved. Vessel segments can also be displayed in different ways to improve visibility.

Vessel display

We're able to display the vessel either as a surface, as a wire frame, or as a skeleton (see Figure 5, next page). This can provide helpful information during 2D–3D vessel registration when the doctor either needs to see through some vessels to perceive the underlying fluoro-

4 Parallel axis, asymmetric frustum-perspective projection.

5 Our system can display 3D vessels either as (a) a surface, (b) wire frame, or (c) skeleton.

(a)   (b)   (c)
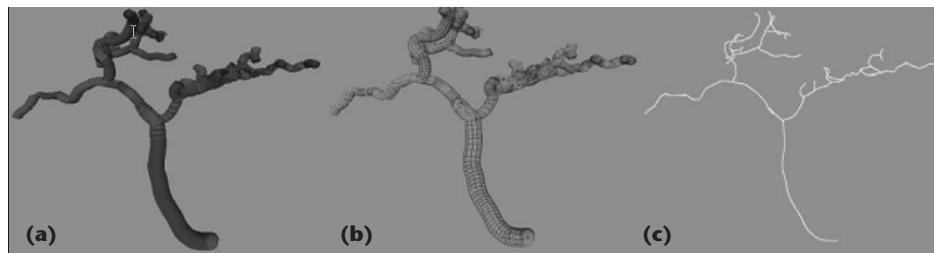
6 Rendered vessel with selected segment.
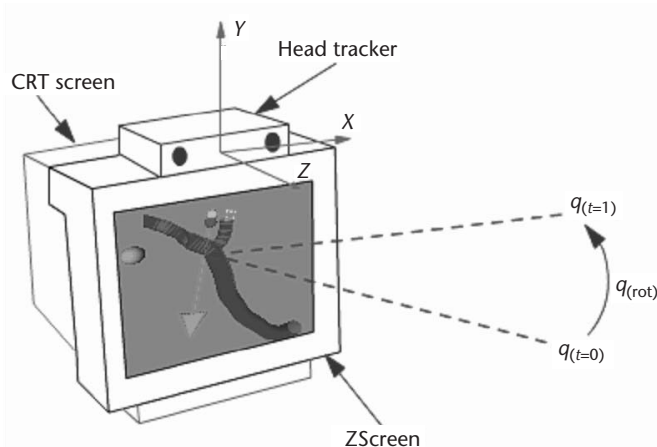
7 Display arrangement with CRT, polarizing screen, and head tracker.

scopic view, or to observe the curves of the vasculature, which is easy to do by displaying the skeleton.

The software also allows user-selected culling of the displayed 3D vasculature. The user can select an individual vessel and then either recolor or delete it from the display (see Figure 6). Similarly, since our data structures order each parent vessel's children according to the direction of flow in the parent vessel, users can delete all vessels downstream of a particular vessel or a selected vessel point. This capacity is particularly useful for guidance of endovascular surgery, since fluoroscopic images obtained during the early angiographic fill phase won't contain small terminal branches.

Head tracking

Another enhancement provided by our display system is tracking the position of the user's head. As the doctor's head moves, the system redraws images from the new point of view. With such tracking, the doctor can move around the 3D model, making it easier to visu-

alize 3D relationships.

To track head movement, we selected the Origin Instruments DynaSight infrared tracker (see Figure 7). The program queries the tracker 60 times per second and determines the target's position relative to the display. The target is a lightweight, 0.7-cm diameter passive reflective sticker. We put such a marker between the eyes of the glasses we use for 3D stereo display. This tracking solution is inexpensive and sufficiently accurate for our purposes. In addition, it's unaffected by radio of electromagnetic interference in the operating room, and since it's wireless, it doesn't restrict the physician's movement.

The view refreshes using the information provided by the tracker. Rather than moving the camera or eye position while the head moves around the 3D object, the 3D object is rotated on the screen. The scene can be rotated in software as well as by head motion. When the tracker reports a new head position, the camera information is updated with gluLookAt and glFrustum. These changes occur if the user moves closer to or further from the screen. Rotations are computed using quaternions. For example, at time t_0, the tracker reports a position, which can be represented as a normalized quaternion, q_0. At time t_1, we receive another position that we express again as a normalized quaternion, q_1. The relationship between these is also a quaternion that represents rotation by an angle θ around an axis:

$$q_1 = q_{rot} * q_0$$
$$q_{rot} = q_1 * q_0^{-1}$$

The quaternion q_{rot} gives us the head movement as a rotation that we then apply to the scene. Given q_{rot}, we can compute the axis and angle of rotation.[11,12]

Future work: Simulation

We've implemented the visualization methods we described, and initial reactions from the clinicians involved are positive. However, before using the proposed approach to guide clinical procedures directly, we require simulation tests to ensure that our method will be useful, accurate, and easy to use, and to formally assess which visualizations will be most useful (such as an external view of the vessels and needle, as currently implemented, or a needle eye's view).

We've already performed accuracy tests on many of the components involved in the overall system (vessel segmentation,[2] 3D curve reconstruction,[9] 2D-to-3D registration,[6] and balloon tracking).[5] We're planning formal clinical trials involving several radiologists to

compare the speed and accuracy of the traditional methods with the new visualization scheme. However, this requires modifying the system to allow direct input from the fluoroscopic units and increasing the speed of the current image-processing algorithms to allow real-time processing of the image streams. We estimate that we'll begin such trials in approximately one year. We believe that the proposed visualizations will prove useful not only for the TIPS procedure, but also for a wide range of endovascular operations. ∎

Acknowledgment

References

1. J.M. LaBerge et al., *Portal Hypertension: Options for Diagnosis and Treatment*, Soc. Cardiovascular and Interventional Radiology, 1995.
2. S. Aylward and E. Bullitt, "Initialization, Noise, Singularities and Scale in Height Ridge Traversal for Tubular Object Centerline Extraction," *IEEE Trans. Medical Imaging*, vol. 21, no. 2, 2002, pp. 61-75.
3. E. Bullitt et al., "Symbolic Description of Intracerebral Vessels Segmented from MRA and Evaluation by Comparison with X-Ray Angiograms," *Medical Image Analysis*, vol. 5, no. 2, 2001, pp. 157-169.
4. O. Faugeras, *Three-Dimensional Computer Vision: A Geometric Viewpoint*, MIT Press, 1996.
5. V. Venkatraman et al., "Liver Motion Due to Needle Pressure, Cardiac and Respiratory Motion During the TIPS Procedure," *Proc. Medical Image Computing and Computer-Assisted Intervention* (MICCAI 2004), LNCS 3217, Springer-Verlag, 2004, pp. 66-72.
6. A. Liu, E. Bullitt, and S.M. Pizer, "3D/2D Registration Using Tubular Anatomical Structures as a Basis," *Proc. Medical Image Computing and Computer-Assisted Intervention* (MICCAI 1998), LNCS 1496, Springer-Verlag, 1998, pp. 952-963.
7. B. Jolly et al., "Needle Tracking and Detection in the TIPS Endovascular Procedure," *Proc. Medical Image Computing and Computer-Assisted Intervention* (MICCAI 2003), LNCS 2878, Springer-Verlag, 2003, pp. 953-954.
8. R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge Univ. Press, 2000.
9. E. Bullitt, A. Liu, and S.M. Pizer, "Three-Dimensional Reconstruction of Curves from Projection Views, I. Algorithms," *Medical Physics*, vol. 24, no. 11, 1997, pp. 1671-78.
10. P. Bourke, "3D Stereo Rendering Using OpenGL (and GLUT)," 1999; http://astronomy.swin.edu.au/~pbourke/opengl/stereogl/.
11. A.S. Glassner, *Graphics Gems*, Academic Press, 1990.
12. K. Shoemake, "Animating Rotations with Quaternion Curves," *Proc. 12th Ann. Conf. Computer Graphics and Interactive Techniques*, ACM Press, 1985, pp. 245-254.

Damien Maupu *is finishing his diplôme d'ingénieur from the Ecole Supérieure de Chimie Physique Electronique de Lyon, France. He's currently attending the Swiss Federal Institute of Technology Lausanne, Switzerland. His interests include computer graphics and computer vision. Contact him at damien.maupu@cpe.fr.*

Mark H. Van Horn *is a research assistant professor in the Department of Surgery at the University of North Carolina (UNC) at Chapel Hill. His current research interests include medical image analysis and image-guided surgery. Van Horn received a PhD from North Carolina State University. He's a member of the IEEE, Eta Kappa Nu, and Tau Beta Pi. Contact him at mvanhorn@med.unc.edu.*

Susan Weeks *currently practices at UNC as an associate professor. Her professional interests include image-guided surgery, local tumor therapy, and abdominal transplantation. Weeks received an MD and completed a residency in radiology and a fellowship in vascular and interventional radiology from UNC. Contact her at sue_weeks@med.unc.edu.*

Elizabeth Bullitt *is a board-certified neurosurgeon and professor of neurosurgery at UNC. She directs the multidisciplinary CASILab (see http://casilab.med.unc.edu/), which focuses on image-guided surgery and computer-assisted diagnosis. Her research has centered upon vascular applications. Bullitt received a BA from Harvard University, an MD from the University of Colorado at Boulder, and she completed her residency at Duke University. She's an affiliate member of the IEEE. Contact her at bullitt@med.unc.edu.*

For further information on this or any other computing topic, please visit our Digital Library at http://www.computer.org/publications/dlib.