EG UK Theory and Practice of Computer Graphics (2006)
M. McDerby, L. Lever (Editors)

A Mixed Reality Anatomy Teaching Tool

Rhys G. Thomas, Nigel W. John, Ik Soo Lim

School of Informatics, University of Wales, Bangor, LL57 1UT, UNITED KINGDOM

Abstract

In this paper we present an inexpensive Mixed Reality software tool for training medical students in anatomy. The software integrates the ARToolkit and Visualization Toolkit (VTK) to create a novel interactive environment in which the user can manipulate the position and orientation of the volume rendering using a plastic model of the organ to be observed. The volume rendering can then be clipped relative to an arbitrary plane to reveal data from its interior, using a second prop.

1. Introduction

The teaching of human anatomy makes use of many resources. Comprehensive and detailed text books are available, such as the classic Gray's Anatomy of the Human Body [Gra18], now in its 20th edition. Problem-based learning scenarios, prosections (pre-dissected specimens), and anatomy models are also commonly used. Performing a dissection of a human cadaver, however, has traditionally been considered as the optimum method for students to gain an excellent spatial understanding that is difficult to glean from a text book alone. Yet dissection has become less common today due to financial and ethical reasons, and some medical schools have taken the decision not to use cadavers in the teaching of anatomy in its undergraduate curriculum. Although, it should be noted that there is some indication of a resurgence in the use of hands on dissection as it is difficult to use other methods to study fascial planes and the pathological processes that can occur, such as fascitis, compartmental syndrome, and carpal tunnel syndrome.

A potential substitute for dissection is to make use of computer graphics technologies delivered using an appropriate Virtual Environment or Augmented Reality (AR) application. Anatomy can thus be explored in three dimensions (3-D), providing new educational tools ranging from interactive anatomical atlases to systems for surgery rehearsal [KHHT]. Previous work in this area, together with other training examples of AR applications is presented in the next section. Section 3 details the methods and techniques that we have used to build a novel anatomy teaching tool. Initial results are presented in Section 4, and the paper ends with conclusions and a discussion of future work.

Although there are excellent examples of AR being used in medical applications, there is little work published on using this technology to aid in the teaching of anatomy. To our knowledge, the tool that is described in this paper that links AR with volume rendering is currently unique.

2. Previous Work

The Visible Human Project [oMTvhpw] has provided a catalyst for the use of high resolution volume rendered medical data for anatomy teaching, for example [VSS]. Web based anatomy teaching tools have been developed based on this data set [BT02], and commercially available products include the VOXEL-MAN family of applications [vox]. The Visible Human male is around 15 GB and the female 45 GB in size. Our tool is also able to make use of these data sets, as well as any patient-specific medical scan. In [Azu97], an Augmented Reality system is defined as being one which has the following characteristics:

1. Combines real and virtual scenes.
2. Interactive in real time.
3. Registered in 3D.

There are several examples of AR being used as a training tool for both medical and non-medical applications. One non-medical purpose is for training drivers of motor vehicles to handle unexpected situations [HR05]. This system uses a combination of visual and inertial tracking methods to locate the users viewing position. It simulates a person running into the path of the vehicle being driven by the trainee, rather than throwing an object into the vehicles path. Another training tool is the MR MOUT system described in [CHE05]. This system trains soldiers in urban combat, and uses a Head Mounted Display (HMD) to present the mixture of reality and computer generated imagery to the user. One medical application of AR is to train paramedics in endotracheal intubation [LDR02a] [JRK] [LDR02b]. This approach typically uses a head mounted projective display connected to a PC and employs an optical tracking system or equivalent. The aim is to increase the success rate of this life-saving procedure. Another medical training example involves the simulation of Laparoscopic surgery [MD98]. In this system the user has at his disposal two virtual laparoscopic joysticks, which are designed to feel like the instruments used during the real procedure. The winner of the first Eurographics Medical Prize was also an AR application [AB03]. They provide a tool for liver surgery planning in which the use of AR contributes to a user-friendly design and simplifies complex interaction with 3-D objects. However, no general purpose anatomy teaching tool that uses AR has been reported.

Whilst the facility to create volume renderings of patient anatomy is becoming more accessible within a hospital - the medical scanner manufacturers routinely provide software to do this - the user interface still relies on a keyboard and mouse. Further, the interface tends to be non-intuitive with many parameters to be set, and therefore difficult to use. A different approach was taken by the neurosurgical visualization system developed by Hinckley [KHK94]. They use a doll's head and a plastic cutting-plane to aid interaction with volumetric brain data. Both props are tracked using magnetic tracking technology, which is prone to errors from metal distortion. A similar idea has also been used in an AR system for overlaying virtual 3D representations of molecular structures onto autofabricated models of molecules [AG05]. This approach uses video-based tracking but also displays irrelevant objects such as a user's hands. This will not be a problem in our anatomy teaching tool. In this paper we utilise and improve on some of the best components from the work described above. Our solution does not strictly conform to the definition of AR given earlier, as there is minimal superimposing of real images and computer-rendered images - currently just text labels are superimposed onto the video feed. Instead we focus on using a conventional AR interface to control the orientation of a high definition volume rendering. Thus we use the term Mixed Reality to describe our approach. Further, we investigate the hypothesis that an effective Mixed Reality anatomy teaching tool can be produced using inexpensive off-the-shelf components.

3. Methods

3.1. Object Detection and Tracking

The ARToolkit [Com] is a public domain software library for building AR applications and we have utilised this software for our prototype anatomy teaching tool. The ARToolkit typically makes use of paper markers that can be tracked in the video footage of the real world scene. Figure 1 shows the plastic male torso model that is the main user interface component of our teaching tool [AHP]. It depicts musculature on one side and surface skin on the other, and is dissectible into seven parts: torso, lung (two parts), heart, liver, stomach, and intestine. It costs just $26 and is primarily aimed for use with school children studying human biology. However, similar models with more body parts are also commonly used props for teaching anatomy to medical students and nurses.

Figure 1: *Human Torso Model*

To be compatible with the ARToolkit, small markers were attached to the removable organ models using sticky tape on the flattest surfaces that we could find on each piece of the model - see Figure 2. The ARToolkit derives its tracking information from a standard USB webcam (Creative Labs Ltd.) attached to the PC. The ARToolkit tracking works by first finding black square shapes in the scene, and then by analysing the pattern inside the square and comparing it to a set of patterns to determine the best match.

Our initial tracking effort was an adapted version of one

Figure 2: *Organ models with markers attached*

Figure 3: *Tracking program showing stroke font organ name*

of the demonstration programs provided with the ARToolkit. Instead of overlaying a 3-D computer generated object onto the video footage, we use a 2D stroke font to display a label for the particular anatomical object being manipulated. This label is fixed relative to the position and orientation of the physical representation of the object, as demonstrated in Figure 3. Our tracking system was initially written using the "C" programming language as it made the adaptation of the demonstration program quite straightforward.

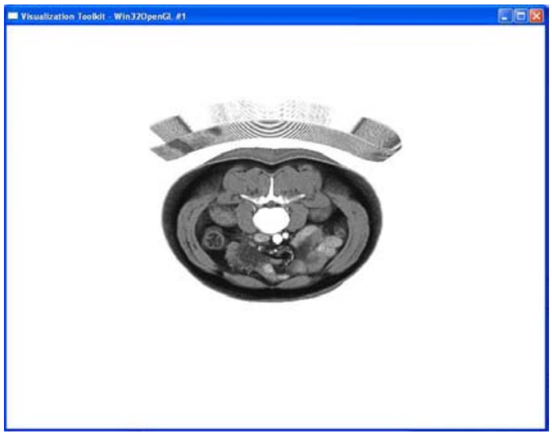

Figure 4: *A volume rendering from CT Data of the Human Torso*

3.2. Volume Rendering

The open source Visualization Toolkit (VTK) has been used to provide volume rendering functionality. VTK is implemented in C++, but it also has bindings for Tcl, Python and Java. We use the C++ version for our software.

We begin with anonymised CT data in DICOM format provided by the Royal Liverpool Hospital. Equally we could use the data available from the Visible Human project. The raw image data is extracted from the DICOM files and then converted into the vtk structured points format using a rawToVTK utility [Vid]. A typical volume rendering created for the teaching tool is illustrated in Figure 4. This has been optimised by first setting all of the black background voxels to be completely transparent as they would otherwise obscure most of the organs within the volume. Secondly, a colour transfer function is used to increase the brightness of the volume, which helps to emphasise the different features of the anatomy.

The next step is to manipulate the position and orientation of the rendered volume according to the position of the tracked anatomical object in the video footage.

3.3. Integration

The original version of the tracking software was now reimplemented using C++ to facilitate integration with the VTK code. In particular, the volume renderer can thus be made a member variable of the tracking software, providing easy access to all of its public methods. The flow of the program was not altered at this stage, and remained functionally identical to the original code. One issue, however, is that VTK stores its transformation matrices (in homogeneous form) transposed when compared to the ARToolkit. We therefore had to alter the volume rendering source code to transpose the volume transformation matrix. In addition,

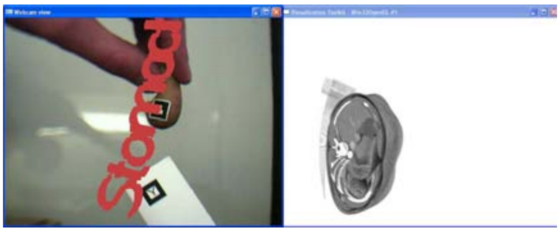

Figure 5: *The program running, showing the volume being clipped*

many of the elements in the matrix must be negated to compensate for the fact that the ARToolkit measures the position of the camera relative to the tracked pattern, whereas VTK measures the position of objects relative to the camera. Without this adjustment the volume will be transformed in the opposite direction relative to that which is intended.

The current version of the education tool uses two side by side windows as illustrated in Figure 5.

3.4. Interactive Volume Clipping

Following the integration of our tracking software and our volume rendering code we require additional functionality to allow the displayed volume to be clipped using an arbitrary plane. The aim is for the student and/or instructor to use a plastic rule to control the clipping plane. The plane can be specified within the software using a homogeneous transformation matrix supplied by our tracking component. However, VTK requires that a plane be specified relative to the volume's co-ordinate system, and not the camera or the plane pattern co-ordinate systems, and again we also have to account for the differences in the representation of the transformation matrices in the tracking and volume rendering components. Using our matrix for the plane and the matrix for the current volume transformation, the origin for the plane can be derived using simple vector geometry as illustrated in Figure 6. To find the normal vector we obtain and reverse the rotations of the current volume transformation by transposing the 3x3 matrix representing the rotations within the homogeneous transformation matrix. This is allowable as there are no scaling or shearing operations contained within the matrix. This inner matrix was then multiplied with the corresponding portion of the matrix representing the plane, which provides us with a rotation matrix that represents the rotation between the object marker and the plane marker.

This rotation data is represented in a homogeneous transformation matrix, and multiplied by a new matrix that contains a z-translation of value 1. The x, y, and z translations from this matrix are then used as the values for the normalised normal of the plane. As all multiplications were by one and we only require three values, a full matrix multiplication is not required. Instead we just calculate the values

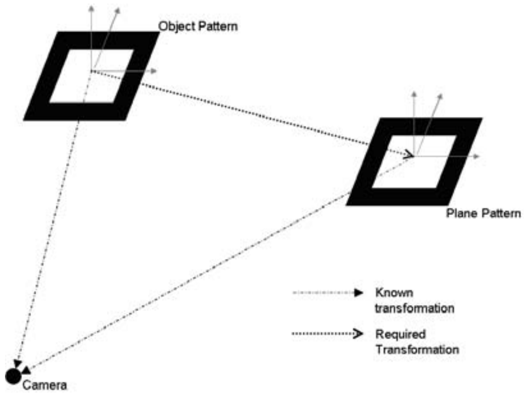

Figure 6: *Deriving the plane origin*

needed directly, which are the values within our rotation matrix, thus saving some processor time.

$$\begin{bmatrix} a & e & i & m \\ b & f & j & n \\ c & g & k & o \\ d & h & l & p \end{bmatrix} * \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} a & e & i & i+m \\ b & f & j & j+n \\ c & g & k & k+o \\ d & h & l & l+p \end{bmatrix}$$

As d, h, l, m, n and o are all zero, and p = 1

$$\begin{bmatrix} a & e & i & m \\ b & f & j & n \\ c & g & k & o \\ d & h & l & p \end{bmatrix} = \begin{bmatrix} a & e & i & i \\ b & f & j & j \\ c & g & k & k \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

We only need the values i, j, and k to specify our normal. The plane defined in this way forms the input for the volume clipping method. To accommodate this the flow of the tracking software is changed so that it only calls the volume rendering method once it has the matrices for both the object to be rendered and the plane, rather than transforming the object and then creating the plane in a separate step. This ensures that if a plane is not needed, this will be known at the time of rendering, and so the object would only have to be rendered once.

4. Results

The hardware platform used to develop the teaching tool is a standard Windows XP PC with a 3.2 GHz Pentium 4 processor, 1GB of RAM, an ATI Radeon X600 graphics card with 256MB of memory, and an 80GB hard drive. The software is always run with an optimum display resolution for this application of 640x480 pixels. The initial organ tracking software runs at a frame rate of 12.5fps when a single organ model is in view, with a minimal deterioration when the number of organs is increased. The software is able to maintain an acceptable interactive performance for the tracking and registration tasks of the AR toolkit.

When the volume rendering code is integrated into the application, the frame rate achieved for different volume sizes is summarised in Table 1. Even when using a relatively small volume, the performance obtained on our test platform is disappointing and is not sufficient for interactive rendering rates. Frame rates improve when the clipping plane is introduced as the size of volume being manipulated is reduced. Ongoing work is currently optimising the code so that real time manipulation can be achieved. Use of high performance computing resources is another alternative that is being investigated within our research group to achieve real time for medical AR applications [CH06].

The current version of the tool has been demonstrated to an anatomy lecturer and to practicing surgeons. Despite the slow frame rates, their comments on the value of mixed reality for anatomy teaching have been positive. Potentially, they see that this approach could bring back some of the value of dissection as a means for exploring and learning anatomy. Separating out the organs allow relationships to be understood, which is key to the learning process. The relatively low cost is also particularly attractive to anatomy departments.

Volume Size	Application Frame Rate	Using Clipping Plane
512x512x246	0.545fps	0.820fps
256x256x123	1.302fps	1.868fps
128x128x61	3.040fps	6.586fps

5. Conclusions and Future Work

The results obtained augur well for the utility of this mixed reality anatomy teaching tool. The final configuration of the system needs to be optimised for interactive performance, but will remain cost effective and affordable to medical schools and home users alike. Although many anatomy departments provide access to computer based anatomy learning tools, they are often unused due to lack of student motivation. Contrarily, the system described in this paper has been greeted with much enthusiasm.

The project is still work in progress. In the future we intend to add several new features to our system; the first being to allow the user to iterate through several segmented volumes, thus allowing them to view internal portions of the organs being viewed. Level of detail is also important and must use high resolution to pick up all relevant anatomical structures. Another possible addition is to have multiple volumes so that several organs could be rendered at the same time and manipulated independently. However, this will impact on the speed of the application, and therefore is not a priority at present. Other requests from the end users are to provide additional learning points such as histology, and common pathologies for more advanced students.

Support for optical see-through glasses is also being explored. This will allow the user to interface more naturally with the model organs and clipping plane prop, and will avoid the conflicting viewpoints problem that results from using a static webcam and monitor set-up. This will add significant cost to the system, however, and so will not be a compulsory component of the tool.

6. Acknowledgements

We would like to thank Dr John Delieu (Lecturer in Anatomy and Licensed Teacher of Anatomy, Oswestry Orthopaedic Hospital), and Dr Derek Gould (Consultant Radiologist, Royal Liverpool University Hospitals) for their invaluable comments in developing and assessing this project.

References

[AB03] A. BORNIK R. BEICHEL B. R. G. G. E. S. F. L. M. S.: Computer-aided liver surgery planning: an augmented reality approach. *Medical Imaging 2003: Visualization, Image-Guided Procedures, and Display 5029* (2003), 395–406.

[AG05] A. GILLET M. SANNER D. S. A.: Tangible augmented interfaces for structural molecular biology. *IEEE Computer Graphics and Applications 25*, 2 (2005), 13–17.

[AHP] ARTOOLKIT HOME PAGE L. V. S. M. .:. http:http://www.hitl.washington.edu/artoolkit/.

[Azu97] AZUMA R.: A survey of augmented reality. *Presence: Teleoperators and Virtual Environments 6, 4* (1997), 355–385.

[BT02] B. TEMKIN E. ACOSTA P. H. E. O. A. T.: Web-based three-dimensional virtual body structures: W3d-vbs. *J Am Med Inform Assoc. 9(5)* (2002), 425–436.

[CH06] C.J. HUGHES N. J.: A flexible infrastructure for delivering augmented reality enabled transcranial magnetic stimulation. 219–224.

[CHE05] C.E. HUGHES C.B. STAPLETON D. H., E.M.SMITH: Mixed reality in education, entertainment and training. *IEEE Computer Graphics and Applications 25*, 6 (2005), 24–30.

[Com] COMPANY A. C.:.

[Gra18] GRAY H.: *Anatomy of the Human Body*. 1918.

[HR05] H. REGENBRECHT G. B.: Augmented reality projects in the automotive and aerospace industries. *IEEE Computer Graphics and Applications 25*, 6 (2005), 48–56.

[JRK] J. ROLLAND L. DAVIS F. H.-L. J. N. C. I., KERNER K.: Merging augmented reality and anatomically correct 3d models in the development of a training tool for endotracheal intubation. ((Electronic) Proceedings ISBI2002).

[KHHT] K. H. HOHNE B. PFLESSER A. P. M. R. T. S. R. S., TIEDE U.:.

[KHK94] K. HINCKLEY R. PAUSCH J. C. G., KASSELL. N. F.: A three-dimensional user interface for neurosurgical visualization. *In Proc. SPIE Conf. on Medical Imaging 2164* (1994), 126–136.

[LDR02a] L. DAVIS F.G. HAMZA-LUP J. D. Y. H. S. F. C. M. G. M. J. N. K. L., ROLLAND J.: Application of augmented reality to visualizing anatomical airways. *SPIE AeroSense 4771-Helmet and Head Mounted Displays VII*, 3 (2002), 400–405.

[LDR02b] L. DAVIS Y. HA S. F. G. M. C. M. B. P. J. N. K. L., ROLLAND J.: Augmented reality and training for airway management procedures. *Medicine Meets Virtual Reality 02/10* (2002), 121–126.

[MD98] M. DOWNES M.C. CAVUSOGLU W. G. L. W. F. T.: Virtual environments for training critical skills in laparoscopic surgery. *Medicine Meets Virtual Reality VI* (1998), 316–322.

[oMTvhpw] OF MEDICINE: THE VISIBLE HUMAN PROJECT WEBSITE N. L.: Robust rendering of general ellipses and elliptical arcs. `http:http://www.nlm.nih.gov/research/visible`.

[Vid] VIDHOLM E.:.

[vox] (Voxel-man) `http:http://voxel-man.de/`.

[VSS] V. SPITZER G. SPITZER C. L. K. R. L. G. K. G., SMYTH P.: Vh dissector: a platform for curriculum development and presentation for the anatomical arts and sciences.