Semi-automatic Scene Generation Using The Digital Anatomist Foundational Model

Benjamin A. Wong, B.S., B.A., Cornelius Rosse, M.D., D.Sc. and James F. Brinkley, M.D., Ph.D.
Structural Informatics Group
Departments Biological Structure and Computer Science & Engineering
University of Washington, Seattle

ABSTRACT

A recent survey shows that a major impediment to more widespread use of computers in anatomy education is the inability to directly manipulate 3-D models, and to relate these to corresponding textual information. In the University of Washington Digital Anatomist Project we have developed a prototype Web-based scene generation program that combines the symbolic Foundational Model of Anatomy[7] with 3-D models. A Web user can browse the Foundational Model (FM), then click to request that a 3-D scene be created of an object and its parts or branches. The scene is rendered by a graphics server, and a snapshot is sent to the Web client. The user can then manipulate the scene, adding new structures, deleting structures, rotating the scene, zooming, and saving the scene as a VRML file. Applications such as this, when fully realized with fast rendering and more anatomical content, have the potential to significantly change the way computers are used in anatomy education.

INTRODUCTION

The fundamental importance of anatomy is evidenced by the fact that anatomy is one of the first courses taken by students in all the health professions. Traditionally, anatomy is taught by a combination of lectures, hardcopy textbooks, and atlases, with or without cadaver dissection. Increasingly, computer-based approaches have been introduced, primarily through CD-ROM.[8] However, these methods are still seen as little more than adjuncts to the traditional methods, although they do provide valuable supplementary material.

There are several factors that could soon lead to more widespread use of computers in anatomy education. These factors include: 1) the lack of newly trained anatomists, 2) the difficulty in obtaining enough cadaver material, 3) the availability of high quality image datasets such as the Visible Human,[1] 4) high performance computing and networking that can process and deliver high-bandwidth data, and 5) the Web. As these factors converge we can expect to see an increase in the amount of anatomy that is learned online, with a corresponding decrease in cadaver and lecture contact hours. Although cadaver dissection will not and should not be eliminated, it may become a privilege that is only earned on completion of online courses in anatomy.

The Digital Anatomist Project is a long term effort to develop representations and systems that can form the basis for such online courses in anatomy, not only for medical students, but for other types of health professionals as well. As part of that project we have created a series of Web-based interactive atlases that are in widespread use throughout the world.[2] These atlases consist of 2-D images, many of which are snapshots of 3-D anatomical scenes. The images are annotated with labeled regions of interest that allow the web user to click on structures to obtain the name, to outline and label selected structures, and to take self evaluation quizzes.

We recently completed a user satisfaction survey of these atlases in local anatomy courses, and compared the results with a survey of other Web-based anatomy resources.[6] Although users found many features of interest, none were compelling enough to warrant more than a supplemental use of these materials. One of the main deficiencies noted in all the sites, including our own, was the inability to interact directly with the 3-D scenes, and to relate these interactions to corresponding textual information. We believe that this feature, if available for the entire body at real-time speeds, would offer a number of advantages over cadaver dissection in many cases.

Although the Visible Human satisfies the first requirement for such an interactive 3-D system, namely the need for raw data, it is only the first step. The data must be segmented, saved as individual 3-D graphical models, and delivered over the Web. In addition, the models must be integrated with corresponding symbolic information, i.e., the names and semantic relationships of the structures represented by the models. It is this symbolic information, captured by the *Foundational Model of Anatomy*,[7] that provides the "intelligence" that makes sense of the image-based data.

In this paper we describe a prototype Web-based 3-D scene generation and interaction program that could form the basis for new methods of education in anatomy. Building on our earlier work in 3-D reconstruction,[4] Web-based scene interaction,[9] and foundational model development,[7] the system integrates graphical models and symbolic anatomical information in an interface that allows a Web user to construct and navigate scenes in the body, using a standard Web browser. The only other system we have seen that combines spatial and symbolic information in this way is Voxelman.[5] However, Voxelman is CD-ROM rather than Web-based.

METHODS

Architecture

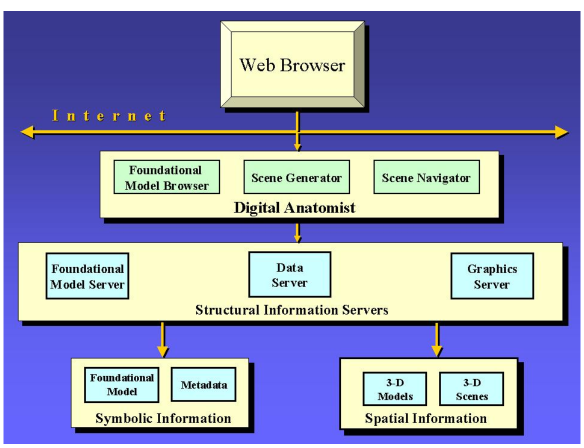

Figure 1: Digital Anatomist components used in the scene generator architecture

Figure 1 shows the scene generator as one module in the Digital Anatomist agent, a program that will eventually know about the available online anatomy resources, the types of users, and the manner of use, and will dynamically generate interfaces customized to each user.[3] The Digital Anatomist is a component of the Digital Anatomist Information Framework, a distributed system in which various client programs access spatial and symbolic anatomy resources by means of several middle-layer structural information servers.[2]

Foundational model

The Foundational Model is a logic-based symbolic description of the structure of the human body.[7] The basis for this model is an ontology that classifies anatomical entities visible to 1 mm resolution, and associates preferred names and synonyms with these concepts. To-date, over 25,000 concepts have been entered for the thorax, abdomen, pelvis and perineum. Additional networks define qualitative spatial and transformational relationships. Currently, over 28,000 links represent *part of*, *branch of* and *tributary of* relationships, and work is progressing on adding additional spatial relationships. Although by no means complete, the model is already rich enough for use in such applications as scene generation.

The model is stored in a relational database, and made available to client programs by means of the Foundational Model server (FM server) that provides a high level application program interface (API) to the model. A Java applet is used by the anatomist authors to add content to the model. A second program, the Foundational Model Browser, provides a rudimentary forms-based interface to

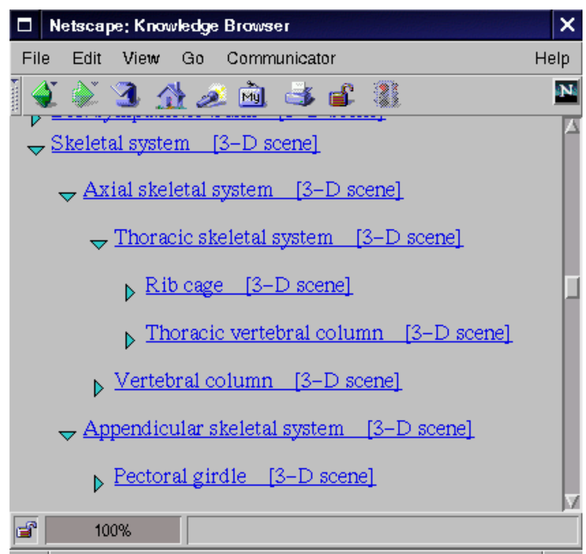

Figure 2: Web forms interface to the foundational model, showing a portion of the *part of* hierarchy

the same model, figure 2. In this case the user can select a given semantic hierarchy (e.g., *part of*), and open up and close subpart hierarchies below a given structure in the FM.

3-D graphical models

Many of the 2-D images in our current online atlases are snapshots of 3-D scenes generated by earlier versions of our in-house graphics toolkit, Skandha, an object-oriented lisp-like language with OpenGL capabilities. The scenes are composed of 3-D surface models created from cadaver sections, like the Visible Human, but predating it.[4] In order to develop the scene generator the 3-D scenes and models were split apart and saved as individual Skandha files, each containing a single 3-D surface mesh corresponding to the smallest structure or structure part that we had previously modeled. A separate data server records the correspondence between these primitive filenames and names from the Foundational Model (*metadata* in figure 1). The primitives were created so as to correspond as much as possible to the most detailed structure parts in the *part of* hierarchy of the FM.

Graphics server

The graphics server (figure 1) is an application written in Skandha, which implements a Web-accessible server with a Lisp-like API. Internally, the graphics server implements an object class (called a "structure") to encapsulate the atomic elements of the scene. Each structure contains a name ("Descending thoracic aorta"), the 3-D model associated with it (if any), transformations to position the model, and filters to alter its appearance ("material"). Structure objects are designed to be arranged in a directed acyclic graph: operations can easily be performed on aggregates of structures recursively. (For example, to render the entire scene the client sends a "display" message to the root structure).

Scene generator

The Scene Generator is another Lisp server that is called by both the FM Browser and the Scene Navigator. A scene can be initiated from the FM Browser, figure 2. For example, if the Web user clicks "3-D scene" for the Thoracic vertebral column, a request to create the scene is sent to the Scene Generator. The Scene Generator first contacts the FM server to find the parts of the Thoracic vertebral column. For each of these parts it contacts the data server and finds out if the parts have associated 3-D model primitives.

If a primitive exists the graphics server is told to create a new "structure" instance, with the pathname (later a URL) to the associated 3-D model. The material properties of the instance are set depending on the type of the structure in the anatomy ontology of the FM. For example, anything that is an artery, a branch of an artery, or a segment of an artery is set to "artery-material', which happens to be red. The structure instance can be apended to a 3-D scene cache for later reuse.

Scene interaction

Once the scene has been generated and loaded into the graphics server the scene is rendered at the server, a snapshot is taken of the scene, converted to a JPEG image file, and included with a dynamically-generated html form that is sent to the Web browser. The Web user can then interact with the scene. Each mouse click is processed by the Scene Navigator program, which in turn contacts the graphics server to manipulate the scene, re-render it, and send a new snapshot. A VRML file of the current scene can be saved at any time.

AN EXAMPLE

To illustrate how the parts of our system work together, here is an example of a user connecting to our servers for a short session. Connecting via any web-browser, the user can access the FM as shown in figure 2. In this example, the user has opened the Skeletal system in the *part of* hierarchy to find the Axial and Appendicular skeletal systems. She has opened the Axial skeletal system and then the Thoracic skeletal system within that. When she clicks on the phrase "3-D scene" next to the term "Thoracic vertebral column" an interactive scene

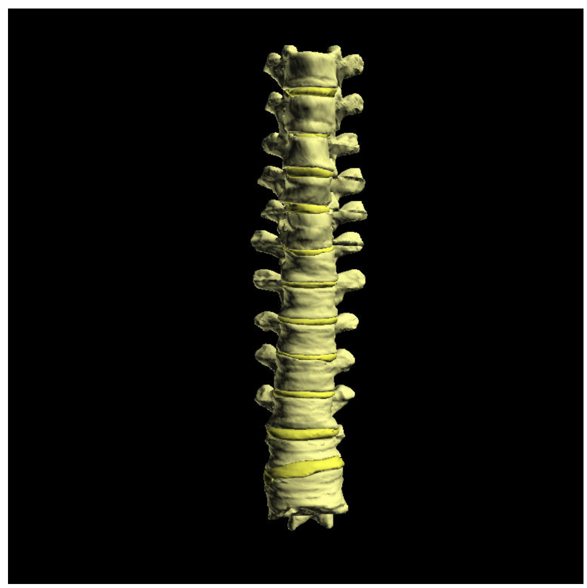

Figure 3: A dynamically generated scene of the Thoracic vertebral column. 3-D models by David Conley.

is created on-the-fly (figure 3) with a panel of controls (figure 4).

This image of the "Thoracic vertebral column" can be manipulated dynamically in 3-D using the same interface described previously.[9] This is a good start, but not a very interesting scene, so the user types in the word "Aorta" in field **A** (figure

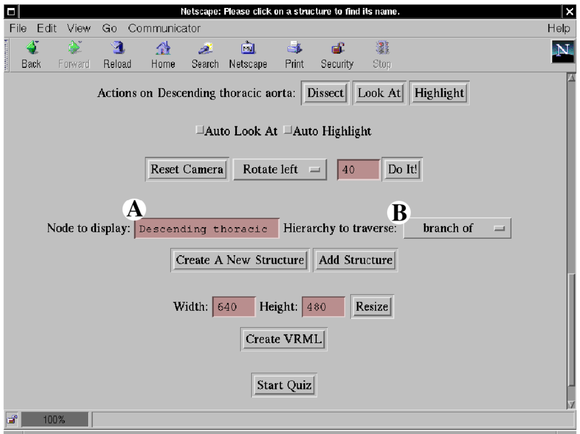

Figure 4: The scene navigator control panel that appears at the bottom of each generated scene.

4), selects *part of* in field **B**, and clicks on the button "Add Structure". This action causes the Scene Navigator to contact the Scene Generator, which in turn contacts the Foundational Model server, to recursively find all the pieces of the Aorta in the *part of* hierarchy, look up the relevant Skandha models in the data server, and add them to the previous scene. (Figure 5).

The user then notices that one of the pieces of the Aorta has odd looking stubs coming out of it. She clicks on the piece and is informed that it is the "Descending Thoracic Aorta". Wondering if the stubs are perhaps branches, she changes selector **B** in figure 4 to read *branch of* instead of *part of* and then clicks on "Add Structure".

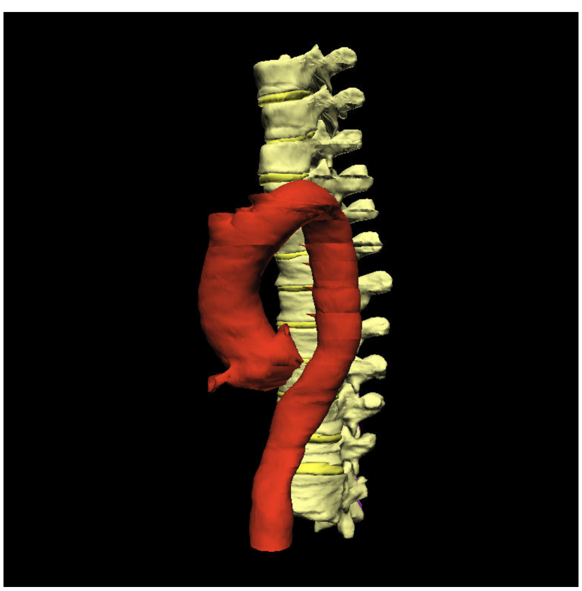

Figure 5: Adding the aorta and rotating

This operation adds a ladder-like structure of horizontal arteries. She clicks on one and is informed that it is the "Right seventh posterior intercostal artery". To fill out the scene she then adds the Trachea (both "part of" and "branches of") in the same way as she added the Aorta. (Figure 6).

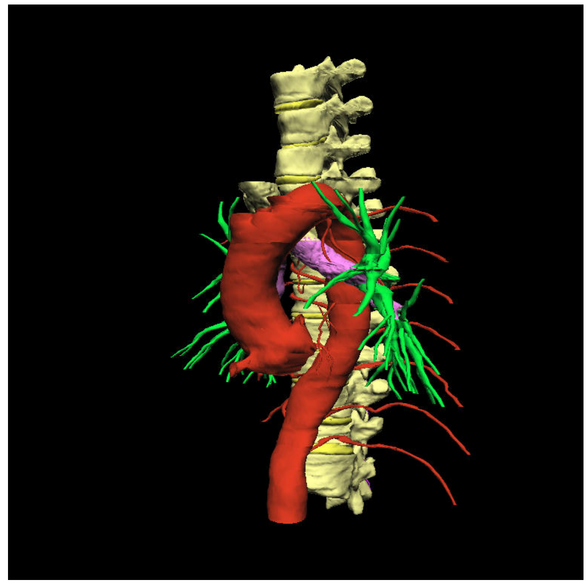

Figure 6: A more complicated scene

However, now there is a piece of the Aorta blocking the view of the Right principal bronchus. She could rotate the view but instead she clicks on the offending piece (the "Ascending aorta" it turns out) and then cuts it away using the dissect button. (Figure 7).

DISCUSSION

In this report we have described a prototype scene generation program that uses the Foundational Model of Anatomy to semi-automatically organize 3-D models into meaningful anatomical scenes. A Web interface allows a user to build up a scene as a series of display trees created from parts or branches in the foundational model, then to interact with the scene, rotating or zooming it, browsing to retrieve structure names, highlighting a selected structure, zooming in on it, and removing it. Once a scene has been manipulated in this way it can be sent to the client as a VRML file.

Although still in the prototype stage, these features demonstrate the potential of combining symbolic and spatial anatomical information in a dynamic and Web-accessible application that could form the basis for new methods of online teaching in anatomy.

In order to take this application beyond the prototype stage several issues must be resolved: 1.

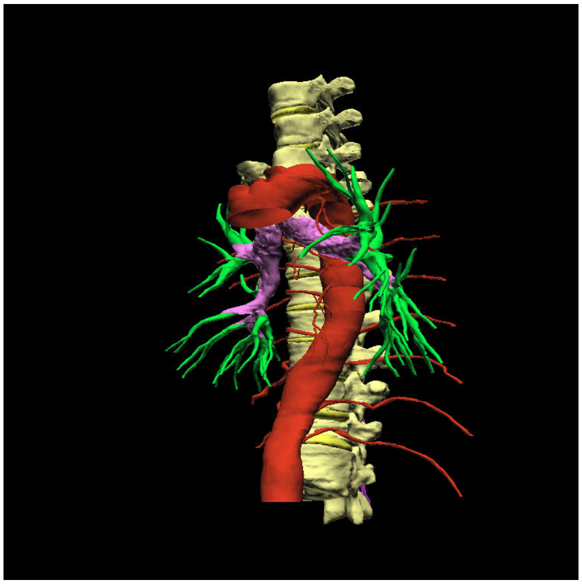

Figure 7: Removing the ascending aorta to reveal the right bronchus

Concurrent use. The reason we have not yet released the URL for the scene generator is that only one user at a time can access the application. Solutions to this problem will require independent threads in the server process so that more than one display tree can be maintained. It will also require that the renderer not take over the screen on the server workstation as it does now. 2. Rendering speed. Our current server hardware requires up to a minute to render each scene before passing the image back to the client. This time can be reduced by faster server hardware and next generation Internet bandwidth, faster VRML viewers at the client workstations, and multi-resolution models. 3. Content. We currently only have models for the thoracic viscera and the brain. Commercial versions of Visible Human 3-D models have been prohibitively expensive up to now. Some mechanism needs to be established to make segmented Visible Human models widely available at little or no cost.

The prototype we have developed shows the potential that could be realized if these issues could be resolved. Once they are resolved, entirely new forms of computer-based tools could gradually become dominant in anatomy education.

ACKNOWLEDGMENTS

This research was supported by National Library of Medicine grant LM06316. 3-D models by David Conley, VRML output routines by Kevin Hinshaw, Skandha4 by Jeff Prothero.

References

1. M.J. Ackerman. The Visible Human Project. *Journal of Biocommunication*, 18(2):14, 1991.
2. J.F. Brinkley, S.W. Bradley, J.W. Sundsten, and C. Rosse. The Digital Anatomist information system and its use in the generation and delivery of web-based anatomy atlases. *Computers and Biomedical Research*, 30:472–503, 1997.
3. J.F. Brinkley and C. Rosse. Requirements for an on-line knowledge-based anatomy information system. In *Proceedings, American Medical Informatics Association Fall Symposium*, pages 892–896, Orlando, Florida, November 7-11 1998.
4. D.M. Conley, K.G. Kastella, J.W. Sundsten, W. Rauschnig, and C. Rosse. Computer-generated three-dimensional reconstruction of the mediastinum correlated with sectional and radiological anatomy. *Clinical Anatomy*, 5:1–17, 1992.
5. K.H. Hohne, B. Pflesser, M. Riemer, T. Schiemann, R. Schubert, and U. Tiede. A new representation of knowledge concerning human anatomy and function. *Nature Medicine*, 1(6):506–510, 1995.
6. S. Kim, J.F. Brinkley, and C. Rosse. Design features of on-line anatomy information resources: a comparison with the Digital Anatomist. In *Proceedings, American Medical Informatics Association Fall Symposium*, Washington, D.C., 1999. Submitted.
7. C. Rosse, L.G. Shapiro, and J.F. Brinkley. The Digital Anatomist Foundational Model: principles for defining and structuring its concept domain. In *Proceedings, American Medical Informatics Association Fall Symposium*, pages 820–824, Orlando, Florida, Nov 7-11 1998.
8. ADAM Software. Adam scholar series. CD-ROM, 1995.
9. B.A. Wong and J.F. Brinkley. Dynamic 3-D scene navigation in web-based anatomy atlases. In *Proceedings, American Medical Informatics Association Fall Symposium*, page 1100, Orlando, Florida, Nov 7-11 1998.