THE UNIVERSITY OF MICHIGAN VISIBLE HUMAN PROJECT (UMVHP)
QUARTERLY PROGRESS REPORT: Y3Q1

Brian D. Athey, Ph.D.
Asst. Professor
Director, UMVHP
April 9th, 2002

UMVHP: THIRD YEAR QUARTER ONE REPORT
TABLE OF CONTENTS

Y3Q1 REPORT
KNOWLEDGE ENGINEERING TEAM

ACCOMPLISHMENTS OF THE QUARTER JUST ENDED

The principal accomplishments pertinent to our Knowledge Engineering team in the quarter just ended were as follows.

1. An authoritative User's Manual for Edgewarp

On March 4, 2002, a draft of a complete User's Manual for the current (1/2002) version of Edgewarp was posted to the internet as a set of 11 "chapters," each a printable postscript file. The directory in which the manual was uploaded also includes the current linux executable of the program and several demonstration image volumes accompanied by annotated Edgewarp scripts (save files, filmstrips, curves), of all the deliverables of this contract, this is the one in greatest demand by colleagues variously local, virtual, or in the extended community. The draft has not yet been critiqued by beta users, and so is posted only to our private web site ftp://brainmap.med.umich.edu/pub/fred/ewsh3.19.man/ and has not been publicized in any form. When it is revised (see Plan 1 below), it will be uploaded to the public site http://vhp.med.umich.edu, in the form of a tar file combining program, test data, and scripts in one easy dataset. This draft of our manual totals 102 pages. It carefully separates the tasks of generating Edgewarp structures, displaying them interactively, and archiving them for future displays or for more intensive editing operations. Under these three headings the manual first reviews three general themes of the Edgewarp world -- operation of the program (mouse and menus), geometry of the high-level windows, and the specific conventions for control of pose (worldview) -- and then systematically surveys all eight of the object classes that currently reside in the Edgewarp world: images, reference trihedra, plane sections, filmstrips, landmark point configurations, polygonal curves, triangulated surfaces, and thin-plate spline grids.

Most of our other accomplishments during the quarter just ended were pursued simultaneously with the preparation of this manual and have already been reviewed in its pages.

2. Further exploration of labels and label systems

The previous progress report announced the creation of a labeling facility for the right-hand (section) Edgewarp window. Labels appear at fixed screen coordinates with a moving leader line connecting them to any point of the section at which a preset curve (for instance, the centerline of a filmstrip) is intersected. Over the quarter, we gained experience in the use of these structures in free combination for insight into sequences of sections as they tumble more freely in the space of Eve. For instance, we now understand how labels corresponding to branching structures should be formulated: as sequences of filmstrips sharing endpoints, together with the representation of those endpoints as landmarks within the larger Edgewarp toolkit. We also explored a variety of algebraic themes for labeling volumes lacking an axial ("filmstrip") structure using this method. Currently the most promising candidate is a label structure taking the topology of a deformable sphere with an equator and a small number of meridional great circles traced upon it. Changes to the Edgewarp program to handle this additional labeling type are in progress.

3. A preliminary filmstrip classification

Filmstrips available at the time of our previous quarterly report included mainly tubular structures such as vessels. We have formalized those structures along with a broad additional range in a list reviewed in one chapter of the User's Manual. The filmstrip types that are likely to prove of value in a variety of pedagogical and informatic applications include, in addition to tubes, the traverse of a brick by parallel planes, the traverse of a cylindrical region by coaxial planes, the traverse of a plane curve by its normal and binormal planes, and the traverse of a ridge curve by its normal and binormal planes. All five of these flavors of filmstrip are explained in detail in the draft Manual.

4. A portable version of Eve

The full power of Edgewarp is clearest in application to the Visible Female, Eve, at full resolution. But demonstrations of this power would often prove frustrating for reasons entirely unconnected with Edgewarp's own operation, to wit, constraints on the "last mile" of the internet connection linking PSC to the demonstration site. To circumvent this frustration, we have prepared a version of Eve that is entirely resident within the local site. The installation requires a just-released version of Linux that handles the 7.1-gigabyte main data file required. The impact of this version of Eve cannot be overstated. (For instance, Cornelius Rosse immediately seized upon it as a candidate for the

visualization engine associated with his Protege-based ontology server at Seattle.) It also makes possible the construction of filmstrips in parallel at sites not particularly internet-capable, such as airplane cabins.

5. Pencils

A completely new Edgewarp display mode has been constructed whereby any list of section planes can be displayed simultaneously in one 3D worldview window. That list can be the entire keyframe collection of a filmstrip or instead a series of disjoint planes at some distance intended to reveal multiple aspects of anatomical organization simultaneously. This facility has been demonstrated in an application to a 3D embryo image that will appear next month in The Anatomical Record – The New Anatomist, the principal journal of imaginative visualizations in contemporary anatomy. The name "pencil" is that of a similar geometric structure in the classic nineteenth-century geometry of space.

Plans For Quarters Y3Q2 and Y3Q3

1. Correcting the User's Manual and uploading it to the project web site

Our highest priority is to disseminate the draft manual to a small list of beta users, some concerned with Eve and others with visual resources of other kinds (e.g., clinical thoracic imaging, embryo MR). Corrections, enrichments, and additional examples for the manual will be assembled as expeditiously as possible and the revised draft uploaded to the public site http://vhp.med.umich.edu as soon as possible. We anticipate a certain load of queries, not always worthwhile, as a result of this posting.

2. Repackaging of the User's Manual in HTML

To ease access to the manual, and in particular to ease directed retrieval, an alternate version of the manual will be prepared that takes the form of html rather than postscript files.

3. Porting EWSH3.19 to Mac and Windows environments

As we write, the current version of Edgewarp has been ported to Macintosh platforms running the MacOS10.1 operating system. The port is not complete, in that performance is enormously slower in that operating system than in the

current preferred Linux environment spite of the superb Mac CPU. Work proceeds to track down the bottlenecks in this implementation -- they seem to center about network issues, not CPU or graphical display. When this port is complete and the issues involved are more effectively understood, a second port will begin, to Windows 2000 and related platforms. These can be expected to be more difficult, in that control of the machine is less extensive under Windows than under Linux.

4. Save file compatibility with VB tracing workbench

To expedite the segmentation computations on which we are about to embark (see below), it will be necessary to combine the navigation capabilities of Edgewarp with the tracing modules of Mr. Pomerance's browser VB at PSC. We determined that the linking mechanism would be Edgewarp's existing capability of archival files for filmstrips. PSC has committed to a VB module that will read the current Edgewarp film file format and convert it to a sequence of scenes in the VB windows from which outline tracings can expeditiously be made. In this way, the contours that seed the segmentations to follow will be optimally informative as a configuration.

5. From sampled surfaces to more authoritative segmentations

The radiation of segmentation styles away from “tubes and blobs” has been recast as a two-phase process. In the first phase, sections through structures are traced as lists of contours (lists of lists of points) in VB, just as at present except that they will be filmstrip-driven for maximum informativeness. In a second phase, intended for prototyping during the remainder of this contract year and a major theme of the years to follow, contours and associated surface normal directions will be treated as structured samples from a distributed non-stationary process whose level sets are the "true" boundary of the structure in question. (This is different from the current implementation of surfaces, in which the same contours and triangles are treated as actually _on_ the surface, an assumption that is demonstrably false in most regions of the volume of Eve.) Using a variety of sophisticated tools, these samples will be extended to a topologically complete representation of the complete boundary wherever possible, and surfaces to be rendered in Edgewarp, VB, and elsewhere, will be renderings of these automatically completed boundaries, not the original hand-tracings that seed them. Algorithms for the extension of traces to complete surfaces include neural nets, quasilinear local signal detection methods, methods based on Morse theory of gradients in RGB space, and several others.

The resulting surfaces will be collected in libraries, labeled with the appropriate Edgewarp auxiliary structures, and released to our pedagogical testbeds for user exploitation and evaluation.

In addition, the following deliverables, originally described in the Y2Q4 report, are still expected for nearterm quarters:

6. Compressed chads

Early in the next quarter a new chad server is expected to come online from our PSC site. The server will supply Eve's contents at a variety of rates, under client (browser) control, corresponding to either lossless or lossy conditions, and Edgewarp will know how to use this facility for optimum user service. Compression will be necessary for supply of voxels from the expected rescanning of the original Eve films at 100u resolution, for instance; the new data set will be too large to reside locally.

7. A filmstrip library

With the porting of EWSH3.19 and the dissemination of an authoritative manual, it is possible to reassign the task of filmstrip production as routine rather than experimental. Working with Dr. Bookstein, a small number of anatomists and anatomy instructors will assemble a library of filmstrips corresponding to one or more of the usual teaching modules at our medical school. For instance, a set of about 40 such filmstrips, all displaying the same branching .cur structures, could implement a tour of the major arteries of Eve's pelvis. Specific curricular content of this task will be determined in consultation with the faculty of our anatomy testbeds.

8. Alternative interfaces

The “pencil” visualization of multiple planes is the first in a series of alternate Edgewarp interfaces that strip out inessential controls from the user's desktop in order to ease interactions with those remaining. Another of these, implemented by Alex Ade, eliminates the landmark features of Edgewarp, leaving pose controls and filmstrip playback; this is suited to classroom use by novices. Over the next quarter, we intend to develop these and additional interfaces further and study their reception by users at different levels of sophistication.

YEAR 3 QUARTER 1
PSC VISIBLE HUMAN SUBCONTRACT STATUS REPORT

1. Description of progress towards completion of quarterly milestones and deliverables:

Primary areas of progress during the quarter include PSC Volume Browser development, surface model construction, volume data services and networking.

Volume Browser:

Volume Browser developments since the December demonstrations and Y2Q4 report fall into three general, overlapping categories; optimizations and extensions of core functionality, interface design, and new features.

A. Optimization and generalization of the core visualization routines has continued. As a result, users experience a faster, more responsive interface. CPU time profiles from user sessions were used to identify the most active code regions. A number of routines were restructured to reduce work by synchronizing to network and display rates. At the same time additional restructuring is taking place for network & data decompression processing in a separate computational thread to minimize the impact on screen update rate. This work has uncovered and repaired several explicit problems and latent bugs.

B. The PSC VB user interface has been revised to improve its efficiency of use. At the December demonstration of the Volume Browser, the PSC-developed cross-platform user interface widgets were not fully integrated into the browser. This integration proceeded by steps and has now been completed. The old flying panel widgets were removed once all of their functions were taken over by their new replacements. A substantially streamlined and flexible user interface is the result and there is fertile ground for further planned improvements.

At the December demonstration, functionality of the 3D, or context, window was limited to a virtual track ball, optionally centering the volume on the plane and zoom. There were new additions of buttons to automatically align the 3D view to the sagittal, coronal, transverse, and any off-axis view plane. A viewing mode which locks the 3D view to the slicing plane window has been added. This affords a 'view of the orientation volume with respect to the slice' in addition to the already implemented 'view of the slice orientation with respect to the

volume'. Toggling between these two views for comparison is also supported. A secondary orientation indicator was added to the 3D window to ease the understanding of the rotation of the view at high magnification.

C. New visualization features were added to the browser to support easier manipulation, navigation, and segmentation of the volumetric and surface data sets. These features feed back into the interface design decisions, which influences core optimization and enhancements. At the request of the Edgewarp team new functionality is now being added to PSC-VB to read Egdewarp's 'sav' and 'film' files. The overall design of VB has been moving toward symmetry of functionality between the 2D-oriented slice window view and the 3D-oriented context window view. It is hoped that the added symmetry will serve to increase ease of use and reduce the learning burden of users since concepts and controls transfer between windows. Symmetry of function should also provide for transparently viewing the data in whatever way is most natural for the user.

Surface model construction:

Surface model construction continues to be the most difficult area. Substantial progress has been made in designing appropriate tracing methodologies and surface reconstruction algorithms for building surfaces of structures present in the VF CCD dataset. The primary objective for the completed quarter was to evaluate as many published surface reconstruction algorithms as possible, with respect to four of the (many) unique characteristics of our own surface reconstruction problem:

A. The initial input to the reconstruction algorithm is manually generated in the form of arbitrarily oriented polygonal and/or curved “contours”.

B. We need to use currently available input data.

C. We have an opportunity to trace new contours which are suitable for input to an appropriate reconstruction algorithm, and

D. The resulting surfaces will be passed to a semi-automated color segmentation algorithm for further refinement and/or returned to the tracer as feedback to indicate areas which require additional tracing information.

In light of just these factors, we have arrived at the following two conclusions:

A. Some of the existing contour data is not amenable to typical point or (parallel) slice-based surface reconstruction algorithms, so we will have to handle the data separately from (but in parallel with) the 'primary' surface reconstruction effort.

B. The current method of tracing contours is sub-optimal with respect to the simplest possible approach for generating a "reasonable" initial surface approximation. In particular, it is counter productive to trace contours whose normal and curvature do not at least approximate the surface normal and curvature along the chosen curve path. (Note that this is one of two primary problems with faithfully reconstructing surfaces from planar slices; the other is the varying density of data within and between contours).

With respect to (A), we are working on ways to "repair'" some of the Photoshop data for which a reasonable surface could not be generated due to anisotropic density of the constituent point set. In several cases, it seems possible to algorithmically reduce the per-slice point set density in an effort to essentially smooth a "density function" defined over the points. This new point set could then be passed to 'cocone' (a successor to 'powercrust'), for example, which uses the point set to approximate surface normals automatically, and consequently extract a surface interpolating the new point set. It may then be possible to re-introduce the missing points into the coarse surface so that 'important' features are respected.

A primary consequence of observation (B) is that turnaround time for new surfaces, even complex surfaces, generated by a trained segmenter could be drastically reduced since -- compared to the current approach – much less information is needed to generate a geometrically and topologically reasonable initial mesh which interpolates the contours. Given contours which locally approximate the surface curvature and (more importantly) the surface normal, it is straightforward to topologically extract regions of the target surface. An initial surface can then be geometrically fit to these regions using a reasonable algorithm for fitting a surface to the face of an embedded planar graph. (The other choice -- generating a dense point cloud on the target surface by tracing contours -- has been shown over many previous discussions to be undesirable for several reasons.) However, this approach has its own difficulties, such as the necessity to specify topological incidences among contours without resorting to toleranced geometric intersection tests. As in (A), some care will therefore be needed to ensure good output.

The fundamental condition that the input contours at least approximate the local surface normal is not demanding on the tracer with the new tools currently in development for the segmentation capability of the Volume Browser. In the following months, we will focus on the following:

A. Training segmenters to generate "good" initial contour data to pass as input to a surface reconstruction algorithm;

B. Implementing the algorithms to reconstruct surfaces from both the Photoshop data and new contours generated by the segmenters; and

C. Augmenting the Volume Browser with tools to simplify segmentation in concert with our surface reconstruction algorithm.

Volume data services:

A great deal of work has gone into improving the volume data service and particularly its compression aspects. This has been done while maintaining the existing service on a 24/7 basis. The volume service is critical because it is the primary network and performance intensive part of this project.

It is interesting to understand the current level of volume traffic by examining the server log files. Over the 2 weeks prior to this report our primary server "vh.psc.edu" delivered ~47Gbytes in the form of ~30Mchads spread over 2007 user sessions. Therefore the average session retrieved 24Mbytes or 15,000 chads. This is somewhat misleading because of the large number of sessions, ~80%, which are short PSC internal development sessions. Aproximately 13% (259/2007), of the remaining sessions originate from the Michigan team. Many of these show a much larger data use and much longer connect time. We presume these are the anatomy segmentation users. They also show bursts of activity followed by relatively long pauses. We infer that this is navigation followed by close inspection or manipulation of a single view. These large sessions often range from 200Mbytes up to 1.2Gbytes in the largest single session. We believe the very longest of these was for filmstrip development.

We currently have two additional backup volume servers in operation. These are the SUN machine "gibson.vhp.med.umich.edu" located at Michigan and the Itanium "it.psc.edu" at PSC. Since these are only reached by special request or by fall over if the primary server is down their usage is very low. The SUN only

had 326 sessions since that server was installed on Oct 3, 2001 and the Itanium only had 425 sessions over a similar period. Nearly all of these can be identified as server or browser development sessions.

More detailed logging in the forthcoming server release will let us probe these user patterns in more detail. The current reason for tracking and analysing usage behavior has been for simulated performance exercises to help with server development. In this role, additional programs query our test servers to establish their ranges of performance under load while trying to mimic reasonable access patterns.

To understand the volume delivery requirements in context, a very dynamic PSC VB session lasting 10 minutes of constant travel to locate and tour the optic nerve followed by bladder to kidney ureter traversal followed by trading the entire length of the spine consumes about 250Mbytes with a reasonably continuous data rate around 3.3Mbits/sec and no significant delays from the uncompressed service which is about 250 chads per second. This is sufficient to present mostly full resolution visualization with a fall back to 1/2 or 1/4 only apparent during fast moves.

Both the uncompressed and the new compressed service are able to deliver almost the same data rates in bytes per second but the compressed method, currently ~3:1 lossless, simply delivers more usable chads in the same number of bits. Therefore the previous tour when performed with compression will reduce to about 1Mbits per second which is unloaded cable modem speed.

Under ideal circumstances we have seen rates from vh.psc.edu over local PSC machine room connections in excess of 800Mbits/sec with large MTUs. (see https://vhserv.psc.edu:8443/docshare/index.cgi/Perf_test/vh_gige.html)
In principle this says one could support ~240 users in uncompressed mode and nearly 800 with the 3:1 compression. In reality however, the sustained network rates with the small MTU size that is supported over the long haul typically produce only 200Mbits/sec and the current external link between the PSC machine room and our PSC offices is limited to 155Mbits/sec. This translates to 45 users uncompressed and 150 with the compressed service.

To show that the delivery rate is really limited by the network and I/O system, local simulated runs on vh.psc.edu itself without passing over an external channel are able to process ~3Gbits/sec per CPU. Since this is a 4 processor machine the actual CPU and gives 12Gbits/sec which could, in theory, drive

~3,500 uncompressed users or 12,000 compressed users. Clearly this is not really feasible with current network connections and our particular installation but it does show the potential of our high performance in-memory approach which avoids all disk bottlenecks and nearly all service time computation. Even the most pessimistic interpretation of these measurements shows that we should have little difficulty with the original target of 40 simultaneous users connecting to the ES-40 server. Of course further improvements to provide higher lossy compression levels in addition to the lossless mode only improves this result. Several other significant conclusions are that individual clients certainly do not need anything beyond a 100Mbit/sec connection to take full advantage of the volume service. The real question for good interactive behavior is network latency rather than bandwidth which is also the reason for avoiding disk seek delays.

In addition to network delivery, compression provides the basis for effectively keeping the entire data set in memory. On of the primary critera for the choice of the Alpha based 64 bit ES-40 is the use of 64 bit addressing to quickly access large data sets in memory. The current Visible Female data set reduces to about 2.7Gbytes of compressed data in memory when one includes the multiple resolution levels and all of the data structure which holds it together. This is coming close the very special 2Gbyte level. At 2Gbytes it would be possible to use a 32 bit server architecture such as a PC. The benefit of course would be much lower unit cost and inexpensive commodity memory. Even though the overall performance would not reach the levels seen with the ES-40 it should be sufficient for ~20 users. Therefore there is incentive to push the compression ratio even though when it does not translate to more users from the current server. Below the multi-user level a disk based single user approach which has been demonstrated with both EWSH and PSC VB avoids the network entirely but does not scale to larger data sets.

Several factors are being examined to push the lossless compression to reach the 2Gbyte level. Although the current method seems to be signal noise limited there are some remaining spatial redundancies which are not used in the current wavelet. These produce patterns of non random codings especially around high contrast edges. By analysing the structural organization of these patterns is appears there is roughly another 30% gain potential which would be enough to squeeze under 2Gbytes. This result is being verified now for possible incorporation with the bundled release described later by using massive compute search of the coding parameters to optimize the coding tables but still operate with a simple Huffman entropy coder rather than an arithmetic coder.

This is significant because of the extreme difference in decompression CPU time. The ~2.5% gain in compression given by arithmetic coding results in almost a 10:1 increase in client side decoding time in a situation which is already CPU limited for screen updates.

As work progresses on very high compression ratio lossy methods another factor which comes into play is the desire to maintain at least a 10:1 ratio between the size of requests sent to the server and the size of the data coming back. This is critical for low bandwith client links such as cable modems or wireless. If this becomes an issue it is a relatively simple matter to compress the request stream.

See the section 4 goals below for further discussion of the compression bundled release.

A small amount of work was done during the quarter towards preparing a reduced resolution raycast projection volume. This is a pseudo volume which trades the depth axis of the memory volume for use as a rotational viewing axis for ray projections. It will be viewed using an adaptation of PSC VB.

Networking:

We have continued to use our involvement with the Web100 project to improve the effectiveness of the existing network interface. Testing of Web100 with our Visible Human application has helped lead to a stable public release of Web100 tools on February 28, 2002. In particular, tests on the volume server machine, vh.psc.edu, helped to improve specific aspects of multiprocessor support.

The Alpha 1.0. software has been divided up into two separate pieces, the linux kernel patch, and a shared library with a set of utilities, called “userland”. Further improvements were incorporated into the Alpha 1.1 release on March 18, 2002. The current software is supported on Linux base kernel 2.4.16, and has been tested on both Intel and Alpha platforms. The software can be downloaded from the web page http://www.web100.org/download.

Other:

Aspects of our work were presented at the First SIAM Conference on Imaging Science held in Boston from March 4-6. Art Wetzel gave the leadoff

presentation “Retrieving Eve, Chad by Chad: A Multiuser Environment for Viewing Visible Human Data” in a session “The Mathematics of Adam and Eve” hosted by Fred Bookstein. This provided an overview of the Visible Human project and some mathematical aspects of representation and compression. It also gave a nice lead in to D.B Karron and Jim Cox’s “The Leveling of Eve: Digital Morse Theory of Isovalued Surfaces in 3D RGB Organ Rendering” and Fred Bookstein’s “Navigating Eve: The Applied Geometry of Fortuitous Shape Descriptors”.

2. Problems encountered during this quarter:

A. The major problem during the quarter continues to be difficulty of surface production from manual anatomy markup and labelings. The anatomy team reports trouble getting good visual feedback and is not entirely satisfied with the quality of some of the current mesh results. There has also been a backlog of manual contour data which had not been sent to PSC for evaluation and surface construction. This is a severe problem since some of that data is not adequate quality or density for decent model production. Manually shifting through this data to sort out the multiple structures they represent is also a time consuming activity which eats into the time for development of planned features that will reduce the need for manual intervention and such dense manual segmentation in the first place.

B. Although we’ve had no problems with the major server hardware we have run out of space for processing and installing additional large data volumes such as the 70mm male using the current data build process. We had planned to simply remove one of the old data sets to bring this up but this was not possible due to the continuous demo requirement. This also caused us to delay installation of the last set of Web100 kernel patches at the server. These were not critical as we've been able to use the time to work on other parts of the project.

C. A grapics compatibility with the Mac OS/X environment was uncovered which prevents window popup and hiding from operating correctly. Therefore, the current Mac release of PSC VB behaves slightly differently from all the other platforms.

D. We performed a “simple” memory upgrade on several of our team’s PC workstations but ran into problems with slightly flakey memory. This took several days to fully identify and caused the temporary loss of some data files.

3. Resolution of problems:

A. Resolution to difficulties in surface construction is an ongoing activity which is discussed as the top goal in the following section 4. We have an active ongoing discussion between groups to work out several solution approaches. We continue to work on the time allocation issue.

B. We've taken several actions to resolve the space problems while simultaneously maintaining service. First, we've replicated the primary visible female service on a SUN machine at Michigan as well as an Itanium here at PSC. This also entailed modifying the PSC Volume Browser to automatically search for an available server when it starts up by cycling through the sequence "vh.psc.edu", "it.psc.edu" and "vhp.med.umich.edu". In this way, we can take the primary server off line and still have service coverage for most uses. This has actually been in place for a number of months. However, we realized that some of the secondary services had not been replicated. We have now reproduced the voxel identity service and mesh service on two other machines at PSC and are fixing a byte ordering problem so they can also be replicated at Michigan on the SUN. Secondly, we are revising the volume data build process to use less space. In principle the process only requires space for one full uncompressed data volume plus the simultaneous compressed representation. For convenience in processing as a series of separate steps, the existing process actually uses three times the full volume space plus the compressed representation. The replacement process is considerably more complex to implement but will eliminate the overhead. Finally, we've also ordered additional disk space for "vh.psc.edu" so that we can continue to run both the old server and the new server and expanded data sets at the same time. Once the new services are fully tested we will replace the backup services on the SUN and Itanium machines with the new versions.

C. A small program which demonstrates the Mac bug, but which runs correctly on all other platforms, was written and sent to the Apple Developer Support group. We were very pleased to receive a response from Apple within a few weeks acknowledging the receipt and reproducibility of our bug report. They also stated that they would direct us to download a test fix when it becomes available and incorporate the fix in Apple's next major OS release. However, there was no firm projected date.

D. Although this intermittent memory problem was difficult to confirm and cost us some lost time the solution was simply to return the memory and replace it from another vendor.

4) Goals for the next quarter & action plan for their completion:

Our top goal is to improve the segmentation and mesh generation process by better coordination with the anatomy team and aggressive implementation of improved tools for manually guided segmentation using PSC VB. Although the Insight Toolkit has some components which could be useful in this area we have found the ramp up time to use its complex coding style is too long for our current needs and its processing mechanism restricts the size of volumes more than we can accept. Therefore, we believe it is best to stay with the current plan to improve the tools we already have and particularly the segmentation parts of PSC VB.

Some of the existing tools in PSC VB intended to assist segmentation are not actually being used by the segmentation team so we need to improve the training of their use and receive better reports of problems so we can implement corrections. The primary existing tool to help with this is the template match contouring we have called the "snap-to" operator. Control parameter settings in the current version are not correctly handling some surfaces which the anatomy team has attempted to produce. A number of additional simple operators which have been deferred in place of other tasks, will be completed during the quarter. These include flood fill and color driven segmentation methods using linear color space partitioning, elliptical color space carving and nearest neighbor color space clustering. Each of these works well in some regions but fails in others so users will have to select those which work best in particular situations rather than having just one standard solution.

The target by the end of the quarter is to have a working surface reconstruction system in place to handle common "simple" surfaces, like tubes and blobs, and some compound surfaces like branches. Non-manifold surfaces will require new, substantially more complex techniques which we will examine in the following months. A complete plan for final data preparation is being written with the help of the anatomy team and the knowledge engineering team.

The second major goal is to complete release of the new server which has been in development and testing for some time. This includes a more elaborate

server data structure, multi-level compression (with the last level providing lossless delivery), priority coding to provide drawing order hints to clients and startup negotiation with clients to transmit data set parameters. Although this includes quite a number of new components, which could in principle be released one at a time, it is most practical to provide the external release as a bundle. This is primarily due to our experience with the requirement for maintaining continuous service and the difficulty of staging incremental steps which are not compatible with current browser releases. For example, the earlier compression server operating since the NASA demo last summer has seen relatively little use, even from PSC, because of the incompatibility despite its reduced network load.

This bundled release will take place right after the planned server OS, memory and disk upgrade with our backup servers handling ongoing usage while new PSC VB versions are distributed and fully tested along with a replacement "hop server" to provide EWSH compatibility. The added disk space will also let us provide the complete 70mm male data service rather than just the portion that was shown in December along with our first attempt to deliver a browsable raycast transparent rendition over a prescribed range of viewing angles.

There are a number of PSC VB changes which have to be released simultaneously with the new data services. Besides the decompression module the browser will accept "meta data" from the server to configure itself to the characteristics of the particular data set. That is, the server will transmit a description of the data set size, number of color channels, physical resolution etc and the browser will adapt its internal data structures accordingly. This has already been demonstrated at PSC by switching between several compiled data descriptions which lets a user dynamically shift between a number of available data sets without leaving the browser. Additional features are being added to PSV VB at the request of the EWSH team to provide compatible "save file" driven filmstrip playback mode. We do not plan to include the EWSH filmstrip production interface.

We would like to perform a full up 40+ simultaneous real user test to confirm the results of multiuser performance simulations. These users should be distributed across the country with likely locations including PSC, UMich, Stanford and the NLM. Ideally this should take place in late June and would use both PSC VB and EWSH users running on the full range of target platforms (ie. Linux, Mac OS/X, Windows). Since the stated target has been 40 simultaneous users from a single server the host would be our high performance ES40 server

at PSC. This trial should be used as an instrumented dress rehearsal to collect statistics and observe problems in preparation for a follow up run(s) near the end of the project year. Network and other statistics will be used to identify problems and bottlenecks. As part of the preparation for this test we intend to release a version of PSC VB to a limited group of friendly users outside of the immediate project who would be willing to participate in the user load test. We also need to plan a realistic distribution of tasks to mimic the load produced by an actual anatomy classroom situation rather than an artificially intense exercise.

Finally, we would also like to perform a small 2 or 3 client test of remote collaboration using both the collaborative feature of PSC VB as well as linkage using the conventional NetMeeting method. The PSC VB collaborative method is designed to avoid the performance delays inherent in the NetMeeting implementation but may have other pros and cons. This test is intended to investigate that issue to provide a data point for the collaboratory project at PSC.

5) Next quarter needs:

We anticipate delivery of disk and memory for vh.psc.edu to handle the full 70mm male data in addition to the currently served volumes. This is on order and should arrive shortly. This is also needed to release the new server and reconstructed data volumes so that we have a fall back mode in case of problems.

PSC has received a response from Apple Computer verifying our bug report related to window popup and hiding behavior. Although we have the email assurance that this will happen soon we do not have a firm date but believe it should arrive during the next quarter. This is needed to produce a fully compatible PSC VB release across all platforms. There is no workaround for this problem other than to inhibit certain window popup and hiding operations.

We need to arrange real users and client machines for the planned 40 seat test. This will require some substantial coordination of both people and machines so we should pick a target date by the end of May for execution in late June. Network routes should be probed to confirm correct operation ahead of time. Similarly, all users should have a least some preliminary experience in driving the browsers so they can be active rather than spending too much time figuring out what to do next.

Y3Q1 REPORT
DATABASING, UIT, ANATOMY, NURSING

Description of progress to date and next quarter goals:

A. Building the VH lesson, on-line
The coupling of databases, VH-3D data, TA labels, multi-media material, and collaboration technology for education is continuing, with the latest addition being the nursing testbed curriculum prototype phase.

B. Building A Lesson:
Development of an on-line module for nursing class that involves all the elements of the VH Knowledge Base is being developed.

Concept Map

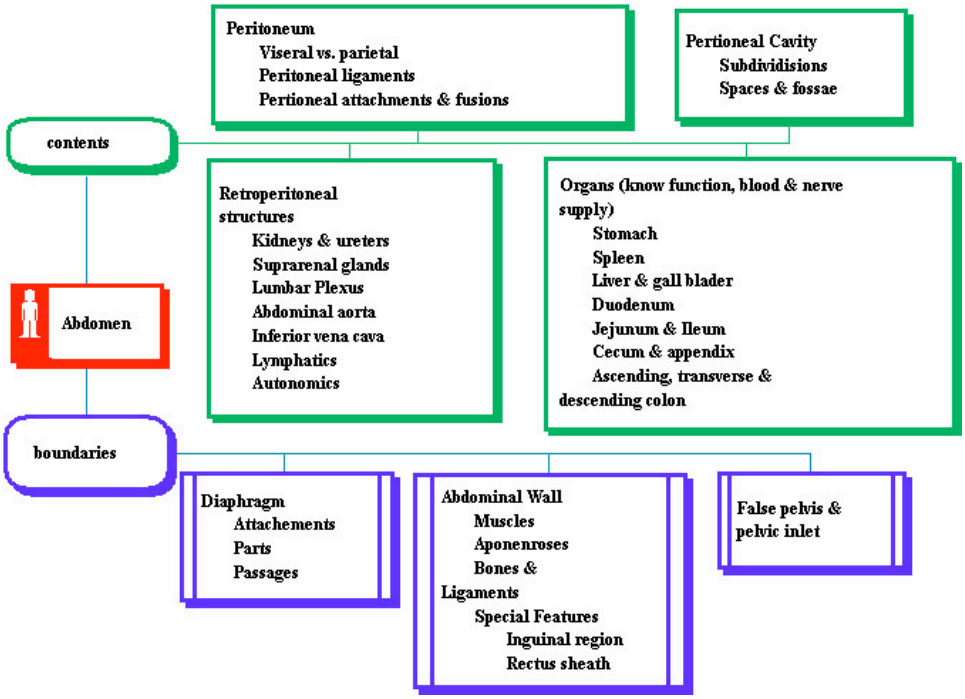

The UIT team has continued to employ the content map scenario for the building of lesson modules. For nursing students, additional features will be added to accommodate the systemic-teaching paradigm of the nursing profession.

C. Prototype Nursing Module: Female Menstruation Through Reproduction:

Placement of two dual boot PC/Linux systems in the School of Nursing, North Ingals Building, was accomplished this quarter, with network upgrades and installation of the Edgewarp and UMPSC browsers. Porting of the iVoxel browser to a windows environment will be. Module development is based on present Nursing coursework. Content available: VH Browser, VH Index Search, digital images, digital multimedia, didactic, class notes, class outline (for sequencing). Content to be generated: nursing terminology, quizzes, Edgwarp flythroughs based on nursing requirements.

Plan of Action:

1. Material gathering
2. Design sequence
    a. A series of pages will be developed following the nursing course outline
    b. Information content placed into the pages as
3: Design prototype
    a. Exploration of content using Browsers and web pages,
4. Evaluation
5. Test deployment
    a. Iterative Development cycle (with implementation for Fall 2002 target).

D. Visible Human Knowledge Base: Contoured Model Development

Preliminary negotiations with the Kendall School of Art and Design, Ferris State University, for additional model development occurred during this quarter. Their output will be directed by our core anatomist group. This test phase is in conjunction with our desire to extend the use of Visible Human data and our testbeds to facilities throughout the State of Michigan.

Tutoring of the core anatomist group in the production of edgewarp filmstrips as guides in the tracing and 3-D model production by nonantomists will begin in Y3Q2.

E. Work on the Anatomy Database:
The structure for the 7679 terms of the TA have been implemented into a Java-based form, and Dr. Heping Zhao continues his translation efforts.
http://vhp.med.umich.edu:808h/kiosk/Tree/index.htm
Chinese input http://vhp.med.umich.edu:8080/kiosk/chinese/index.jsp

F. Collaboration Technology:
Investigation into the use of collaboration tools to enhance teaching and research continues.
The University of Michigan School of Information supports CourseTools/WorkTools (http://worktools.si.umich.edu/) and Work Tools(CHEF): The CompreHensive collaborativE Framework (CHEF) initiative has as its goal, the development of a flexible environment for supporting distance learning and collaborative work, and doing research on distance learning and collaborative work.http://chef.si.umich.edu
The Open Knowledge Initiative (OKI) based at MIT (http://web.mit.edu/oki/) is defining an open architectural specification to be used for the development of educational related software. Over the next quarter, the database team will be determining how the OKI could be implemented into the UMVHP.

Y3Q1 REPORT
IVOXEL BROWSER, ALEXANDER ADE

Progress Report
Efforts this quarter have been directed at enhancing the Web based viewer (iVoxel) and porting Edgewarp to the Apple Macintosh platform.

Ivoxel:

The iVoxel viewer has been divided by function into three unique modules; they are the “Volume View” module (for volume rendered scenes), the “Model View” module (display of and interaction with anatomical models), and the “Slice View” module (arbitrary 2D slices through the dataset). Modules may be used individually or in combination as either stand-alone applications or embedded within Web pages. Modules support both mono and stereo viewing.

1. iVoxel’s Volume View module now supports the addition of arbitrary clip planes to the scene. Combinations of clip planes affords the user a volume

rendered view of regions-of-interest within the dataset. Clip planes are interactive and under user control via the mouse for rotation and translation (figs. 1 and 2). Up to 6 separate clip planes may be added to the scene.

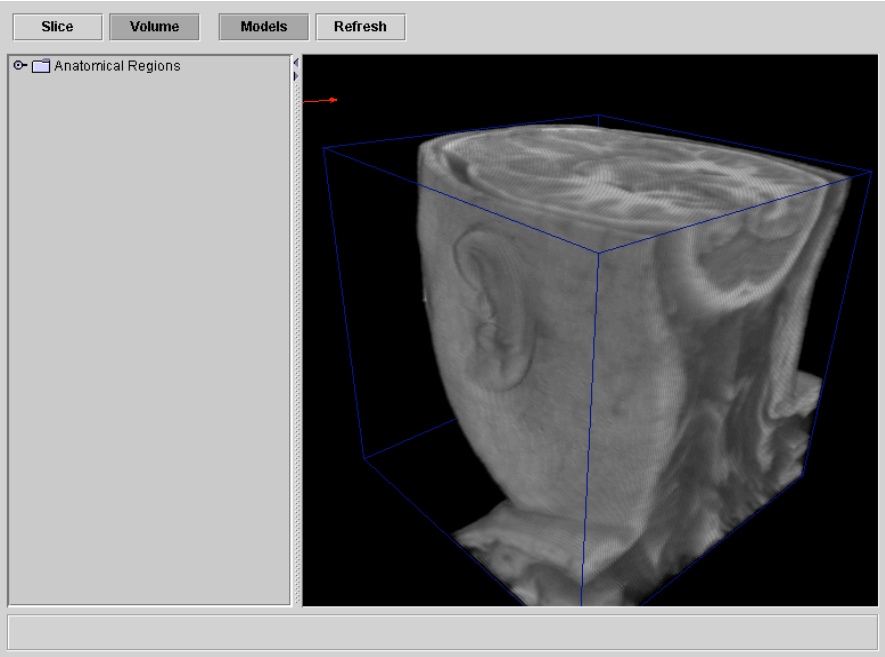

Figure 1

2. iVoxel's Model View module has been updated to support the caching and display of multi-resolution models. Using a set of models at different resolutions, iVoxel supports a "Level of Detail" display system that selects a model at an appropriate resolution for its distance frim the viewer. That is, models that are farther away from the viewer are rendered using their lower resolution versus modles closer to the viewer which are at a higher relolution. With this system, the viewer always sees the highest fidelity data when close to the model. The "Level of Detail" system improves rendering performance by reducing the scene's total triangle count while always displaying the best view possible.

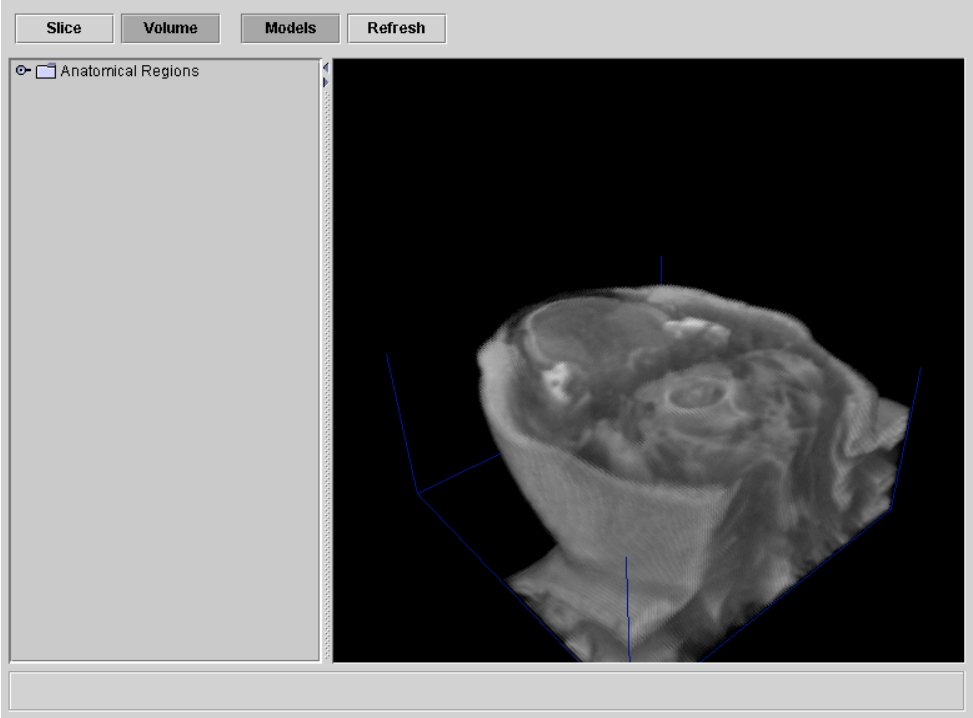

Figure 2

3. The Slice View module requests voxel data from the Pittsburgh Supercomputing Center (PSC) server. Asynchronous downloading and caching of multi-resolution streams is supported. Slice images are constructed on-the-fly from the highest resolution data available. As the slice plane is moved through the volume, image updates use a combination of cached and newly requested data as necessary.

The iVoxel module set has been tested on Sun and Windows platforms. A Macintosh version requires additional system libraries to be provided by Apple. The Apple Graphics Engineering team has assured me that these libraries will be available soon.

Edgewarp:

An Edgewarp port to Apple Macintosh OS 10.1 has been ongoing this quarter. The initial port was successful (fig. 3), but requires hardware acceleration for usable performance levels. Also, network transmission rates are unacceptably slow. Both items will be addressed next quarter.

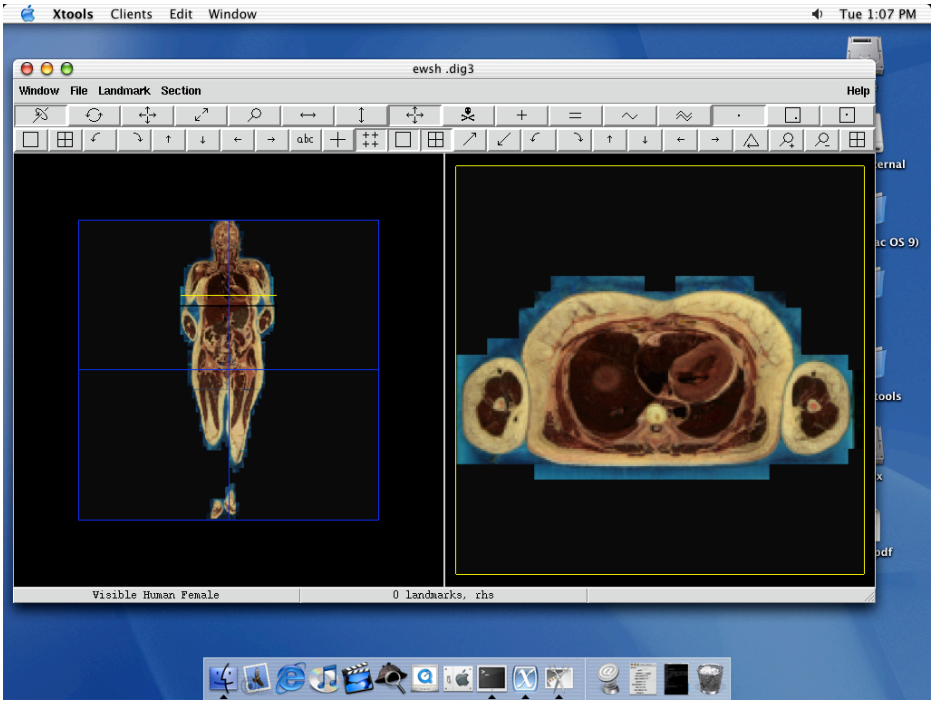

Figure 3

A modified Edgewarp GUI has been created for novice users. The new GUI has a reduced button set that doesn't include the digitizing tools (fig. 4). Buttons may also show text labels that describe the button's function.

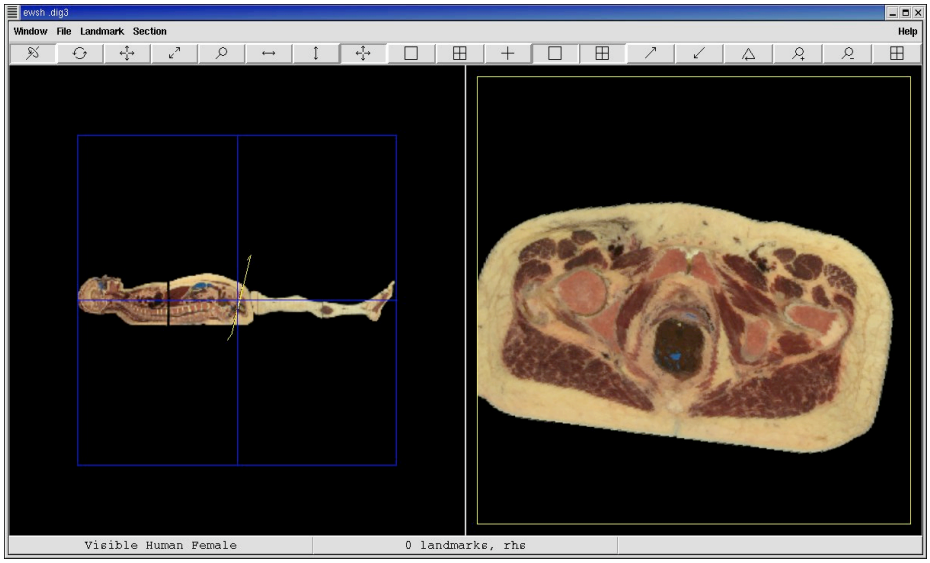

Figure 4

Problems and Resolutions:

We are using Tenon's X-Windows implementation for hardware accelerated Edgewarp on Mac OS X. In the process, we have discovered a bug in their implementation which causes unacceptable rendering artifacts when multiple graphics views are used. A bug report and test case is in preparation.

Network transmission rates on the Mac OS X version of Edgewarp are unusually slow. We are investigating this. Resolution may require tweaking and optimizing TCP/IP settings, or modifying the Edgewarp networking algorithms.

Future Directions:

During year 3, quarter 2, we will continue to enhance iVoxel by supporting spline based labels and animations. Currently, anatomical labels will be stored as curves. To display these, iVoxel will need to interpret these curves and correctly draw the label they represent in the scene.

We will also continue the Edgewarp to Macintosh port. First priority is to address the problems described above. Secondly, we will work to make the Mac version more "Mac-like" in its controls and look-and-feel.