

CASI DMT HARDWARE AND SOFTWARE DEVELOPMENT EFFORT

RESEARCH PROGRESS REPORT
FOURTH QUARTER 4Q

COMPUTER AIDED SURGERY INCORPORATED

300 EAST 33RD STREET, SUITE 4N
NEW YORK, NEW YORK 10016 USA

<http://www.casi.net>

TELEPHONE: 1 (212) 686 8748

FAX: 1 (212) 448 0261

E-MAIL: KARRON@CASI.NET



Date of Publication: 10/31/2002 11:24 PM.

Address: 300 E. 33rd St., Suite 4N, New York, NY 10016

USA.

Telephone/Fax: (212) 686-8748, (212) 448-0261.

Research Technical Progress report 01 for

NIST ATP COOPERATIVE AGREEMENT

70NANB1H3050

Prepared by P. M. Ross, PhD, J. L. Cox, PhD, M. B. Rothman, D. B. Karron, Ph.D. PI with help
from C. DaSalla and R. Peralta


D. B. Karron, Ph.D. President and C.T.O.

ABSTRACT

Our primary goal is the application of the novel mathematical theory (DMT and BMT) to segment and eventually delineate tumors and anatomical landmarks with greater precision, accuracy, and reproducibility, in less time than possible using current medical imaging systems. Because the results are based on mathematical analysis of raw scan data, they are objective. This is a considerable advantage over existing technology, which relies on observers' judgments, with attendant difficulty in reproduction resulting from imprecision and inaccuracy.

In this phase of the project there are 3 separate processes to attain our milestone objectives.

1) **Pre-process an image or image stack.** In this study we are using public domain data from the Visible Human Project. We process CT and MRI scans using Digital Morse Theory (DMT) into a DMT tree (a graph).

2) **Navigate through the DMT segmented nodes using an advanced GRAPH** (not Graphic) interface to identify a target. We want to identify the interesting objects from a possible palette of several million objects identified by our pre-processing. The large number of objects is due to the hierarchical nature of objects. Many objects are assumed (incorporated or included) in parent objects and one parent can have a very large progeny of child objects. Noise features tend to make a very large number of small objects with a shallow graph relationship; strong features are characterized by deep nesting of fewer objects.

3) **Highly Optimized Display of object with 3d movement and color.**

The work has proceeded on or somewhat ahead of schedule despite the need for significant reworking of our software base. Specifics of milestone and task completion are given in the body of the report. Some changes in personnel are described and justified

TABLE OF CONTENTS

Abstract	2
Progress Metrics.....	5
Difficulties	6
Personnel Issues and Changes	8
Deletions	8
Additions	9
Accomplishments.....	9
Glossary	20

PROGRESS METRICS

The appended Microsoft PROJECT database contains details on project progress metrics and tasks. We are using this software system to record tasks completed, estimate completion times, and to help us refine our ability to cost out tasks in terms of man hours and productive code lines.

Appended are also code metrics and a guide to the jargon and terminology in these reports.

Also appended is a sample data set now animated on our web site. This can be played from our web site but is subject to bandwidth limitations. A local copy is appended on the electronic/computer submission of this report. The reader is advised to run the local copy but readers without access to the report source code can run the web site links. This movie loop shows graphically the level of progress we have achieved.

The principal goal of the project is to develop software that can accurately detect, localize, and visualize in 3 dimensions, tumors and other surgically relevant anatomical features present in MRI and CT data. To accomplish this we need geometrically precise, reproducible and semi-automatic soft tissue segmentation from CT and MRI data. Our approach is to analyze and organize the raw image data by means of Digital Morse Theory (DMT). We preprocess the raw image data to produce a DMT tree together with a segmentation of the data into corresponding DMT zones. This organization is referred to as “pre-processing” of the data and forms the basis for our system. The preprocessed data may be displayed visually. Successful development of processing and visualization software for these aims completes the first phase of our semiautomatic anatomical feature and tumor identification paradigm. To date, we have implemented a unique visual navigation system based on DMT, with novel capabilities, for exploring 3 dimensional medical data. The starting point of the DMT organization and our system involves density threshold defined objects bounded by isosurfaces where the threshold grey values represent different tissue densities. The raw 3-dimensional medical data generally represent the density of tissue at various points in space. The major problem that we are

addressing is that in soft tissue, 3 dimensional density gradients are extremely weak, necessitating computer analysis. Existing paradigms are inadequate for this purpose, as described in our initial grant application.

The need for precise segmentation, in areas of a medical image with low gradient, requires that we move beyond both traditional isosurface technology and volumetric gradient based techniques. With these needs in mind we have already implemented our first segmentation surface capability using DMT, with the ability to systematically separate and render 3-dimensional objects bounded by surfaces that are not the isosurfaces that are often used in surface-based feature detection. In year one of our project, we have already succeeded in producing a system with unique, state-of-the-art capabilities.

DIFFICULTIES

The original GANTT changed from the time of this initial implementation, as demonstrated during our first site visit as we elaborated and refined the task steps required to achieve the required milestones. The design changes embodied by these GANTT chart modifications are a necessary evolution in order to successfully complete the major first year milestones, as set forth in the original proposal. The following describes the design changes and subsequent software needed to complete tasks 8 and 9 of milestone 1, and tasks 1 through 8 of milestone 3. Milestone 2, with the exception of tasks 9 and 10, is embodied in the NodePlot module described below.

One challenge in the first year of the grant was the need to rewrite and refine part of the project software base. We discovered, after experimentation, that the original design would be inadequate to achieve our broader goals for a variety of reasons. The original representation of the zones included a voxel list for each zone. This list contained every voxel that intersects a zone. This was extremely redundant resulting in space and processor inefficiency. This redundancy is caused by the fact that a single voxel may

straddle the boundary between two or more zones. The original design also fails to identify precisely those triangles that belong to a surface bounding an object within a zone, as there may be more than one surface portion passing through a single voxel. One of these surface portions may actually belong to an isosurface bounding an object in a neighboring zone. We thus require a zone implementation that is both more succinct and that facilitates the tracking and construction of individual boundary surfaces and that insures no “leakage” onto nearby objects. We also desired to implement SpiderWeb in a manner that would allow consistent construction of non-isovalued surfaces. This is to enable the implementation of our segmentation surfaces. The segmentation surfaces will be, in some cases, non-isovalued surfaces that separate tumors and other anatomical features from surrounding tissues.

One final issue involved changes to the object display and zones necessitated by the redesign of the DMT tree display (see NodePbt module description). We realized that we needed a different DMT tree display. Dr. Wolberg had implemented 3 distinct and impressive DMT tree displays. For our purposes we decided that, for effective visual navigation and segmentation of the data, we needed to see the DMT tree displayed by each node’s isovalue and the volume of the corresponding object. With this new display we would be able to identify objects by their tissue density and volume. Integration of the new object display code with NodePlot and subsequent validation has satisfied milestone 3.

We completely rewrote the object display and redesigned the zone representation with these issues in mind. The representation of the zones now contains no redundancy. Only those points that are inside a zone are listed with each zone. SpiderWeb was rewritten so that it scans the points for a zone and tracks surfaces only from hit points (surface intersections) that occur on voxel edges that emanate from a point within the zone. In this way we insure that there is no leakage into neighboring zones and hence onto nearby objects. We further implemented the surface tracking/construction in a manner that will allow the hit points to be non-isovalued in the future. Specifically, each hit point on a voxel edge is shared by 4 voxels sharing this edge. Once a hit is calculated (interpolated) it will be the same for each of the voxels sharing this edge. As long as these

constraints are maintained, we can vary the value of the hit points as we track a surface. Incorporating ideas from SpiderWeb, this will result in consistent and provably topologically correct non-isovalued segmentation surfaces. We are presently developing the theoretical and practical details of these ideas.

Personnel Issues and Changes

It was found that some initial personnel choices were not optimal for the project. Specifically, use of off-site consultants led to persistent challenges in software compatibility with the project base computer platforms, name midrange SGI super workstations. It was felt that personnel willing to work on site would strengthen the project and also cost less. Consequently it was required that all effort be consolidated to a concentrated on site effort and an additional SGI computer was purchased that could support two programmers simultaneously. Also, It was felt that having a scientist with business experience would be more valuable to the team than a businessperson with scientific experience in the role of office manager. The initial selection of project personnel was therefore modified to reflect better appreciation of the needs of the project.

Deletions

Mr. Elisha Gurfein completed a year contract as business manager at CASI. His contract was not renewed after much soul searching and angst. It was mutually agreed that Mr. Gurfein and this project were not a good fit because CASI did not have the size and overhead funding to pay for the expensive overhead Mr. Gurfein required. Mr. Gurfein's NIST ATP program functions were being handled by students and contract personnel who were retained into the second year. Computer support for Mr. Gurfein became expensive and required too. Further problems were caused by the lack of state co-funding promised by ITAC. Mr. Gurfein worked diligently to attempt to obtain these funds but the Sept 11 statewide budget crunch meant that there would be no new significant new state co funding for that year.

Additions

Dr. Isaac Dimitrovsky is anticipated to improve code optimization. Dr. Dimitrovsky has vast experience with writing highly optimized code and in real-time image processing. Dr. Dimitrovsky worked for both Montage Group Ltd. (producing digital video editing products) and Visible World (a recent startup including members of Montage and Philips). Dr. Dimitrovsky has a PhD. in Computer Science from Courant Institute of NYU and has done significant work toward a second PhD. in BioChemistry at NYU Medical Center.

Mr. Matthew Rothman has supplanted part of Dr. Wolberg's role in graphic programming. His background includes extensive experience in scientific programming, including econometric models on large data sets, which is in many ways similar to DMT analysis in the present project. His sports physiology and personal training with emphasis in functional anatomy further enhance his value to the project.

Dr. Peter Ross has an extensive background in academic research on cancer, including 10 years as a laboratory director at Rockefeller University. He has considerable experience in industry and holds patents based on his research. In addition, he has started a number of medically related businesses.

Accomplishments

Accomplishments to date are based on our building experience writing DMT code and novel optimizations as we learn the nature of reducing this algorithm to code implementation.

Examination of the appended GANTT charts contained in the Microsoft PROJECT database give detailed analysis of the retrospective and prospective project milestones and tasking with tasking assignments.

We now briefly discuss the ideas behind our DMT organizing technology, in order to clarify the scientific, design, and software development issues that we dealt with in year one of the project..

The basic idea behind DMT is as follows. The two main techniques for viewing and apprehending meaningful images derived from volume data have been volume imaging and isosurface extraction. In isosurface extraction surfaces of constant scalar value segment the image into regions of data containing points of density greater than the isovalue (objects), and regions of lower density (called complementary objects). The starting point of DMT is isosurface based, but we use this to move beyond traditional isosurface based methods. DMT studies how the isosurface based images vary as one varies the isovalue or density threshold. DMT studies precisely how the induced objects change as the isovalue. The starting point is a topologically correct isosurface algorithm called SpiderWeb. The values at which the topology of the SpiderWeb scene changes are called DMT criticalities. For example, as we lower the threshold, or isovalue, a new object can be detected at the new, local maximum value. A complementary object is destroyed at the local minimum. At a saddle value, objects (or complementary objects) can merge or split or one of the boundary surfaces of an object can change in topological (homotopy) type. In three dimensions this means the object changes in its connectivity so as to induce a genus change. For example, a sphere can become a torus (donut) or vice versa. Besides providing an important organizing principle, feature topology provides significant information; human hearts may vary in shape and size, but they all have 4 chambers. Topology represents the number of objects in the scene, their gross geometry, and their interconnectivity in a precise mathematical way. Other attempts to apply Morse theory to medical data have been made but they have missed important criticalities and failed to correctly chart the topology changes, or else they have made unrealistic and incorrect assumptions about the data, leading to incorrect or imprecise results, as detailed in our original proposal.

We can represent the topological evolution of the objects (the creation, division and merging of objects), at any particular the isovalue, by means of a DMT tree. Each node in the tree is a DMT criticality and represents a topologically distinct change in an embedded object in the data. The region between criticalities is the range of isovalues for which this object is topologically invariant (all changes can be described in terms of pure geometry and affine transformations). Corresponding to each node in the tree is the region of space contained in this object, which is called a zone. Parentage (and progeny) is defined in the DMT tree by topological change. The first criticality that induces a topological change in an object is made the parent of the object's node. For example, when two objects merge at a saddle point, each node is made the child of the saddle node. A new object is created, and we trace its evolution.

The DMT tree and zones hierarchically organize the entire dataset and form the basis of our system. As described below, we are not limited by the isosurface paradigm. Indeed, our segmentation techniques will not rely solely on isosurfaces. In this way, we will provide a new volume data paradigm that remedies the flaws of volume and pure isosurface based imaging.

Here is a description of the design, development, testing and integration of the three major components of our system and how they realize the milestones set forth in the GANTT chart representing the first year objectives of the project. We describe the development of the three major software components, where design decisions of each component influenced, and were influenced by, the need for the integration of the overall system.

The first component is the DMT preprocessing module. This module preprocesses medical volume data and is being extended to RGB data. A portion of the preprocessing software converts raw medical data, acquired by different imaging modalities, including CT, MRI, and color data, to a standard format, which we shall call "volume data". This data is then input to the main preprocessor module. This completes tasks 1 and 2 of milestone 1 on the Gantt chart.

The next phase of our preprocessing module identifies the DMT criticalities in the volume data. These are the criticalities that cause topological changes in the isosurface defined objects as the isovalue is varied. There are three different classes of criticalities: voxel vertex criticalities that occur at a data reading, interstitial or interpolated criticalities that occur inside a voxel (all other Morse theoretic schemes ignore these), and isoset criticalities, where a group of identically valued data readings function as a higher level structural criticality such as a ridge or a cirque. Each vertex criticality can be maximum, where a new object is created, a minimum, where a complementary object (a hole or bubble) vanishes, or a saddle point. At a saddle point objects can merge, complementary objects can split, or an object can change in topological type (for example change from a sphere to a torus). The interstitial or interpolated criticalities are saddle points where the connectivity of the isosurfaces produced by the SpiderWeb algorithm change. Finally we identify isosets. These are voxel edge connected sets of identically valued data readings. We determine if each such set is a maximum, a minimum, or a saddle set.

After we have identified the DMT criticalities, the final phase of the preprocessor builds the DMT tree from the criticalities, and builds zones associated with each node in the DMT tree. Each criticality is a node in the DMT tree. Each node corresponds to an object in the data containing the criticality. As the isovalue is decreased from this critical value this object grows and deforms until it comes to contain another criticality, which in some manner changes the topology of the object. This new criticality is made the parent of the original and will represent a new object (or more precisely, a topologically invariant family of objects). The zone of the node is the volume swept by the boundary of the node's object, as the isovalue is varied from its critical value to its parent's critical value. The data structure for each node keeps track of this range of isovalues and contains a pointer to the zone. Thus, the DMT tree serves as an index for a zone organization of the data.

We tested the preprocessor on a variety of simulated datasets and real medical image data obtained from the Visible Human Project. This completed task 3 of milestone 1 of the Gantt chart.

We initially implemented an object display based on the following scheme. The preprocessor computed a voxel list for each zone. When a node in the DMT tree is selected, the voxels for that zone are scanned. An active voxel is a voxel containing an isosurface intersection. Each active voxel is fed to the SpiderWeb triangle construction algorithm. Thus our initial implementation of SpiderWeb constructed triangle tilings within each voxel independently. For a specific isovalue, all the nodes in the DMT that correspond to objects at that isovalue are selected and processed individually. Prof. Wolberg implemented the initial object display using the SpiderWeb routine and the DMT tree processing that we had implemented. He also implemented our initial DMT tree display.

As described above, it was necessary to rewrite SpiderWeb. In addition to rewriting SpiderWeb, we rewrote the object display with enhanced features to coordinate with our new DMT tree display (see the description of NodePlot elsewhere in this document). This was done to successfully complete tasks 1 through 7 of milestone 3. We designed several interactive modes for spatially navigating the objects and for simultaneously varying the isovalue, primarily for testing purposes. Then we developed software to integrate the object display with the new NodePlot DMT tree display. In addition to displaying an object corresponding to a node in the DMT tree, we wanted the ability to begin testing our DMT segmentation paradigms. We added the capability to select any collection of nodes in the DMT tree and see the corresponding objects. In this manner we can simulate non-isovalued segmentation.

The DMT tree is displayed on a coordinate grid in which one axis is the isovalue and the other axis is the volume of the corresponding objects. A traditional isosurface consists of selecting all nodes in the DMT tree that lie on the vertical line defined by the isovalue in this coordinate frame. We now have the ability to select the nodes on any curve or series of curves through the coordinate frame of NodePlot. A union of isosurface components of differing values will define the resulting object or objects. In this way we can view objects whose boundary varies in threshold value. We can now vary the thresholds uniformly across the boundary. This is accomplished by varying the isovalue of each selected node simultaneously within the range of isovalues valid for that node.

For example we can tune so that each isovalue is increased by one third of the range for each node, simultaneously, according to the range of each node.

We are able to perform spatial navigation of the object display in real-time. This includes rotation, translation, and fly-through of the objects. We are working on tuning and optimizing the surface construction, so that the construction of selected objects and variation of isovalue can also be done in real time. For large images, the tiling of complex objects can take seconds. We desire an order of magnitude improvement in response time. One of the optimizations will be to implement our Level of Detail management tools based upon DMT in this regard. Once objects are tiled and initially displayed, the affine transformations required for spatial navigation are performed in real-time through the use of triangle lists (a feature of the GLUT library for OpenGL).

We wish to emphasize that at the present time our system has features that no other volume visualization system possesses. Furthermore, the precise object identification that we can presently perform is completely reproducible, as the selected nodes can be recorded. This differs from ad hoc schemes that involve pixel modification and or dithering. We are not modifying the volume data, but rather modifying, in a controlled and reproducible manner, the surfaces that segment the data to reveal distinct objects.

Thus, we have completed milestones 1-3. We have in fact exceeded the goals of milestones 2 and 3 and made some significant progress on year 2 milestones.

The subtasks of milestone 4 on the GANTT chart that correspond to year one deliverables have been completed. We have made progress with several approaches to the DMT decomposition of RGB data for task 5 of milestone 4. Optimization of the preprocessing, task 3 of milestone 4, and the optimization of the object display, for production system needs, now becomes a priority. We have begun rewrite of the backend preprocessor. The critical detection has been highly optimized, and now runs an order of magnitude more quickly than our initial code. We have verified that the results agree with

the original testing We have developed a nice user interface for running the front-end visualization and exploration code to satisfy task 4 of milestone 4.

Description of NodePlot Program

NodePlot is the program that was written to search and display data from the DMT tree (Milestone 2). The development of nodePlot was achieved as follows:

The first step was to determine the correct algorithm, language and hardware to support sizeable DMT trees. In order to do that, the capacity of various algorithms & languages (C++ and Java) were evaluated on multiple hardware platforms with test (random) and real data from Visible Human Project. The C++/SGI Unix platform was selected because Java/PC & Java/SGI implementations become too slow at around 10,000 nodes. The latest version of nodePlot using C++/SGI/OpenGL technology works well up to 400,000 nodes. Typically, 500 unfiltered CT slices will produce 400,000 nodes.

The next step was to design and provide a real time interface for the end user. We therefore implemented sophisticated de-crowding & navigation techniques, described under tasks 3 thru 7 and 9 thru 11 of Milestone 2 on the Gantt chart. In its present version, the NodePlot user can select any size rectangular area and zoom (magnify) that area using a rubber band (see glossary); undo last Zoom; restore graph to original state; select contiguous nodes (same mouse/key combinations that Microsoft uses); select non-contiguous nodes (again same as Microsoft); and double-click to select and view nearest node (1 node). In addition, we have added support for esoteric hardware (Milestone 2, task 9), which permits the end user to move object ONE axis (X, Y or Z) at a time with Dial box or manipulate in 3-d with 3-d space ball (see glossary). To improve image quality, we have applied advanced OpenGL animation techniques so rubber band & zoom function fast and flicker-free (Task 3 thru 7, and 9 thru 11, milestone 2). Finally,

NodePlot feeds the list of nodes to the program objDisplay for representation (Task 4-10 of Milestones 2 and tasks 4 thru 6 of Milestone 3). A Reset command was added and a log file of transactions between systems was implemented for improved utility.

A second version of NodePlot has been mostly completed. The new version contains added de-crowding techniques so user can find the “needle” (tumor) in a “haystack” (medical image scan), integrating ANN search trees; implementing tree-level navigation through nodes; supporting alternate segmentation criteria; displaying dynamic fly-over statistics about closest node and neighboring nodes as the cursor moves around in the graph (Using ANN search trees); and working on more sophisticated 3d mouse SpaceBall.

This page intentionally left blank.

Code metrics of the C++ system written for DMT node navigation and selection

File Name	Lines	Words	Characters	Purpose
annLookup.cc	158	518	4079	ANN tree
annLookup.hh	14	42	326	
boxLookup.cc	27	83	602	Divide & Conquer Nearest Neighbor Lookup (Needed for fly-overs)
boxLookup.hh	28	70	573	
calcDistance.cc	20	60	432	Optimized Euclidean Distance Calculations.
calcDistance.hh	20	60	432	
config.cc	380	777	9197	Default Program Settings stored/read from a file, plus ability to change during program execution
config.hh	82	214	2076	
displayColorLegend.cc	198	411	4807	Displays a side box listing all the colors and their meanings in nodePlot
displayColorLegend.hh	11	30	260	
displayMenu.cc	220	479	5391	Displays side box listing keyboard commands and mouse commands
displayStat.cc	361	822	8507	Display summary statistics of various

File Name	Lines	Words	Characters	Purpose
				sorts
displayWorking.cc	97	188	1997	Display working or busy indicator for when program is thinking and user is waiting.
neighborhood.cc	168	547	3982	Support functions for fly-over and summary statistics.
neighborhood.hh	29	74	658	
neighborhood_Point.cc	69	145	1132	
neighborhood_Point.hh	35	86	753	
nodePlot.cc	1628	3916	40429	Main program which integrates and executes everything else
nodePlot.hh	17	49	481	
utilColor.cc	71	145	1695	Code to set color consistently throughout my program
utilColor.hh	33	70	564	
utilFormat.cc	81	191	1820	Functions to nicely print user information
utilFormat.hh	16	60	444	
utilMouse.cc	405	990	10205	Implementing mouse commands such as zoom, contiguous & non-contiguous select.
utilMouse.hh	45	128	1040	
utilNodeData.cc	204	548	5551	Functions to read & manipulate DMT trees
utilNodeData.hh	72	168	1623	
Total	4489	10871	109056	

GLOSSARY

Full Name & Abbreviation	Meaning
Approximate Nearest Neighbor (ANN)	<p>Tool to increase capacity of system (At least 5 times greater) and greatly improve navigation through complex data by solving the nearest neighbor problem Developed under NSF grant CCR-9712379.</p> <p>http://www.cs.sunysb.edu/~algorithm/files/nearest-neighbor.shtml</p> <p>http://www.cs.sunysb.edu/~algorithm/implement/ANN/implement.shtml</p>
De-crowding	<p>Techniques to search through large amounts of data for relevant objects while maintaining its context and, by this means, so to speak find a “needle in a haystack” (See also Fly-over and Zoom)</p>

Full Name & Abbreviation	Meaning
Dialbox	Set of 8 dials which control movement in one plane (X, Y, Z) at a time.
Digital Morse Theory (DMT)	The mathematical basis for data organization and for image processing in this project.
DMT Tree	DMT Tree is a collection of nodes into zones segmented (organized) by isovalue (see also Tree)
Fly-over	As cursor moves over an area (flying), Some information about that area is also displayed. (See Also De-crowding & Zoom)
Isovalue	A contiguous object within a specified density.
Node	Distinct isovalued object.
NodePlot	Program written under this grant to view & navigate nodes isovalues
ObjDisplay	Program written in this project which displays picture of nodes selected from nodePlot program. ObjDisplay can also work indecently from nodePlot.
OpenGL	OpenGL is Hardware & Software for high end graphics. Developed by SGI both hardware & software are in use for this project.

Full Name & Abbreviation	Meaning
Rubber band	Method to mark area of graph for such purposes as display and zooming. After you select an anchor point with a mouse click, as you move mouse away a rectangle forms on screen. The rectangle stretches hence the term “rubber band”
Silicon Graphics, Inc. (SGI)	Company providing very high-end UNIX workstations for fast computation & fast/accurate graphical display that is used in this project.
Spacemouse	3d pointing device with expanded functionality
Visible Human Project	Public access anatomy data set used in this project.
Voxel	Cubic cells formed from 8 adjacent data readings in 2 contiguous slices from a medical image CT or MRI.
Zone	Set of objects within a particular isovalue range (see Isovalue, DMT Tree)
Zoom	Allows user to select an arbitrary rectangular region of graph and then produce a magnified graph of that region. (See Also De-crowding & Fly-over)