MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

Simulating Surgery using Volumetric Object Representations, Real-Time Volume Rendering and Haptic Feedback

Sarah Frisken Gibson, Christina Fyock, Eric Grimson, Takeo Kanade, Ron Kikinis, Hugh Lauer, Neil McKenzie, Andrew Mor, Shin Nakajima, Hide Ohkami, Randy Osborne, Joseph Samosky, Akira Sawada

TR97-02 December 1997

Abstract

Surgical simulation has many applications in education and training, surgical planning, and intra-operative assistance. However, extending current surface-based computer graphics methods to model phenomena such as the deformation, cutting, tearing, or repairing of soft tissues pose significant challenges for real-time interactions. In this paper, the use of volumetric methods for modeling complex anatomy and tissue interactions is introduced. New techniques for modeling soft tissue deformation and tissue cutting at interactive rates are detailed. In addition, an initial prototype for simulating arthroscopic knee surgery that has resulted from an ongoing collaboration is described. Volumetric models for the knee simulator were derived from 3D Magnetic Resonance Imaging. Visual and haptic feedback is provided to the user via real-time volume and polygon rendering and a force feedback device.

Simulating Surgery using Volumetric Object Representations, Real-Time Volume Rendering and Haptic Feedback

Sarah Gibson[1], Christina Fyock[1], Eric Grimson[2], Takeo Kanade[3], Ron Kikinis[4], Hugh Lauer[1], Neil McKenzie[1], Andrew Mor[1,3], Shin Nakajima[4], Hide Ohkami[1], Randy Osborne[1], Joseph Samosky[2,4], Akira Sawada[1,5]

1. MERL, 201 Broadway, Cambridge, MA 02139, USA
2. Massachusetts Institute of Technology, Cambridge, MA, USA
3. Carnegie Mellon University, Pittsburgh, PA, USA
4. Brigham and Women's Hospital, Boston, MA, USA
5. Mitsubishi Electric Corporation, Hyogo, Japan

Abstract

Surgical simulation has many applications in education and training, surgical planning, and intra-operative assistance. However, extending current surface-based computer graphics methods to model phenomena such as the deformation, cutting, tearing, or repairing of soft tissues pose significant challenges for real-time interactions. In this paper, the use of volumetric methods for modeling complex anatomy and tissue interactions is introduced. New techniques for modeling soft tissue deformation and tissue cutting at interactive rates are detailed. In addition, an initial prototype for simulating arthroscopic knee surgery that has resulted from an ongoing collaboration is described. Volumetric models for the knee simulator were derived from 3D Magnetic Resonance Imaging. Visual and haptic feedback is provided to the user via real-time volume and polygon rendering and a force feedback device.

1 Introduction

Computer-based surgical simulation has many applications in education and training, surgical planning, and intra-operative assistance. In medical education and training, surgical simulation can reduce costs associated with cadaver specimens, provide experience with a greater variety of pathologies and complications, and provide the ability to repeat or replay training procedures. In surgical planning, simulators will enable rehearsal of difficult procedures or planning on patient-specific anatomy and will enhance communication among medical professionals or between doctors and their patients. Intra-operatively, computer modeling can aid in navigation by augmenting the limited surgical field with a more global view of a patient’s anatomy and could provide guidance, for example, by preventing the surgical instrument from entering sensitive regions.

In order to provide useful feedback to the user, surgical simulators must provide adequate realism. Tissue models should respond in a realistic way when they are manipulated. Rendered images of the surgical field must be realistic enough to be compelling. Haptic or force feedback must mimic forces experienced in real-life because the sense of touch provides important cues in surgery. These requirements impose significant demands on the surgical simulator. These challenges include a need for physically realistic modeling (such as soft tissue deformation and tissue cutting or tearing) and tradeoffs among physical and visual realism, cost, and the demands of real-time interaction.

In computer graphics, objects are commonly represented by surface-based polygonal models. Because graphics workstations have special-purpose hardware for fast polygon rendering and because algorithms and systems have been developed for modeling physical interactions between (rigid) polygonal objects, there are advantages to using polygonal models in a surgical simulator.

However, surface-based models do not represent object interiors and they are inadequate for modeling objects and tissues with complex interiors, for modeling the deformation of arbitrary or heterogeneous volumes, and for simulating cutting through tissues. In addition, surface-based methods cannot model fluid flow, tissue tearing or bone fracturing. In contrast, volumetric models can incorporate a great deal of information about interior structure and mechanical properties of heterogeneous tissues. By adjusting forces that hold volumetric elements together, it may be possible to model fluids, tearing, and fracturing with volumes. In this paper, we discuss techniques for deforming volumetric object models, for detecting collisions between volumetric objects, and for modeling tissue cutting. In addition, we describe an initial prototype simulator for arthroscopic knee surgery.

The ultimate goal of our collaborative project is the development of a computer-based surgical simulation system that incorporates volumetric models generated from patient specific 3D Magnetic Resonance Imaging (MRI) data. As illustrated in Figure (1), the system will provide both visual and haptic feedback to the user via real-time rendering and a force feedback device. Physically realistic modeling of anatomical models and surgical instruments will provide a closed feedback loop between the user and the simulator.

Some of the important technologies required by the simulator include: patient-specific 3D image acquisition and segmentation; incorporation of the data and measured material properties into anatomical object models; visual and haptic feedback through real-time rendering and an electromechanical force feedback device; image manipulation -- allowing adjustment of visual parameters; and physically-realistic model manipulation through the haptic interface.

2 Prior Work

There are several surgical simulation systems under development both commercially and in research laboratories. For example, a number of endoscopy simulators enable navigation through stationary object models (Lorensen et al, 1995), (Hong et al., 1995), (Ziegler et al., 1995), (Geiger and Kikinis, 1995), (Vining et al., 1993), (Vining et al., 1994). In some of these, path planning and collision avoidance between the virtual endoscope tip and static object surfaces are incorporated. Deformation of surface models is used to model soft tissues by some surgical simulations systems (e.g. (Cover et al, 1993)). These systems use either mass-spring models or control points on spline-like surfaces for interactive deformation of surface-based object models.

A number of groups have used volumetric methods for modeling deformation and cutting of tissue volumes. Finite element analysis, a computationally expensive but well studied technique for modeling volumes under applied forces, has been applied to facial and muscle modeling (Waters, 1987), (Terzopoulos and Waters, 1990), (Chen, 1991), and in surgical simulation (Cotin et al., 1996), (Hunter et al, 1993), (Pieper, 1992). Manipulation of voxel-based objects (Kaufman, 1996) has been used in object modeling (Gibson, 1995) and combined with force feedback for haptic exploration of voxel-based objects (Avila and Sobierajski, 1996) Sculpting of volumetric objects is described in (Wang and Kaufman, 1995) and (Galyean and Hughes, 1991).

3 Volumetric Methods for Surgical Simulation

Modern 3D medical image scanners can provide a great deal of information about internal anatomical structure and function in individual patients. Figure (2) shows a single slice of a 3D volume MRI image with resolution 0.25x0.25x1.4 mm^3. Internal structure, such as the heterogeneous muscle and fatty interstitial tissues indicated by arrows in the image, can be important in modeling tissue behavior. Unlike surface-based methods where objects are modeled by hollow shells, volumetric object representations provide a way to incorporate all of the information available from the MRI data into a functioning anatomical model. This information is useful for modeling soft tissue behavior -- such as tissue deformation or tearing, and when cutting or joining tissues.

3.1 Volumetric Representation

A volumetric graphical object is a regular or irregular 3D array of data, with each data element representing a sampled point (measured or calculated) in the volume. Examples of regularly spaced volumetric objects include 3D images produced by medical scanners -- such as MRI or CT scanners. Examples of irregular volumes include data resulting from finite element calculations or data from geological survey samples. The deformation algorithm described below generates a changing, irregular data volume because the distances between volume elements vary when objects are deformed.

Elements in the volume can encode information about visual and/or mechanical properties in a single value or within a more complex data structure. For example, each element could encode only the element's color or it could encode color, transparency, edge strength, elasticity, tissue type, and connections to neighboring elements. Figure (3) is an example of a data structure used in one 3D object deformation implementation.

The ability to encode many different attributes within each element enables the modeling of complex materials, structures and behaviors. However, since volumetric objects can consist of thousands to millions of elements, (an MRI image of size 256 x 256 x 256 contains 16 million voxels) this representation poses significant challenges for rendering, data storage and retrieval, and tissue modeling. This paper presents approaches for dealing with these large volumes in rendering and tissue modeling. In addition, future research into hierarchical data representations and data compression will help to make volumetric methods more practical for real-time applications.

3.2 Soft Tissue Deformation with Volumetric Models

Volumetric object models can consist of thousands to millions of elements. Unfortunately, the two most common volumetric methods for modeling deformable systems -- mass-spring models and finite element methods -- are too computationally complex to handle systems of this size interactively. The computational complexity can be reduced by a significant reduction in the number of volume elements used in the deformation model (Hunter et al., 1993) and/or pre-calculation of deformation modes (Pentland and Williams, 1989). However, both of these approaches limit the possible range of deformation and they are not suitable for modeling systems with changing topology, as would occur during tissue cutting.

In surgical simulation, realistic interactions with object models are very important. However, realism and speed must be balanced to enable real-time manipulation of object models. We use a two-stage approach to modeling tissue deformation which allows us to simulate a range of materials -- including rigid, elastic, and plastic substances -- as well as substances like muscle that have different material properties parallel to and perpendicular to the muscle fiber axis. The first stage is applied during direct manipulation. In this stage, tissue elements are moved only when they violate a set of simple linear constraints. This first order deformation is quickly propagated through the volume from the point of manipulation. Because the motion constraints are similar to those of a set of linked elements in a chain, this algorithm has been dubbed 3D ChainMail (Gibson, 1997).

In the 3D ChainMail algorithm, volume elements are linked to their 6 nearest neighbors. When one node of the structure is pulled or pushed, neighboring links absorb the movement by taking up slack in the structure. If a link between two nodes is stretched or compressed to its limit, displacements are transferred to neighboring links. In this way, small displacements of a selected point in a relatively slack system result in only local deformations of the system, while displacements in a system that is already stretched or compressed to its limit causes the whole system to move. This concept is illustrated in the 2D system of Figures (4). Much like the links in a chain, neighbors only respond to a given node's movement if the constraints on distances between nodes are violated. Changing the constraints on link lengths allows us to model both rigid and deformable objects.

In the second stage of the deformation process, we define the system energy according to the relative positions of neighboring elements. When distances between neighbors are within some specified range, the system is in a low energy state. When the distances are farther from their equilibrium distances, the system has higher energy. Whenever the system is idle and between applications of the 3D ChainMail algorithm, the positions of all elements in the object are iteratively adjusted to reduce the system energy. The result of these two deformation stages is that when an object is directly manipulated, it quickly deforms to an approximate shape and then the object relaxes to a more natural shape as the simulation proceeds.

This two stage approach is particularly fast for tissues with homogeneous (though possibly anisotropic) material properties. In homogeneous tissues, the 3D ChainMail algorithm propagates a disturbance through the volume by considering each volume element only once and by comparing each element to only one neighbor. This allows small deformation times that increases approximately linearly with the number of object elements. In tests on a single processor (R10K) SGI Onyx workstation, volumetric objects with as many as 125,000 elements have been deformed at interactive rates. These tests have been performed without particular effort to optimize or parallelize code.

3.3 Collision Detection and Response

In addition to modeling the deformation of individual objects, a surgical simulation system must model interactions between objects. This requires detecting and reacting to collisions among (possibly deformable) objects. Collision detection for voxel-based objects is relatively straight forward. We use a method described in (Gibson, 1995) in which the interaction space is represented by an occupancy map -- a regular 3D grid of cells. Each element in an object is mapped into a cell in the occupancy map. As the object is moved, cell contents in the occupancy map are updated. If an object element maps into an occupied cell in the occupancy map, a collision is detected. Figure (5) diagrams this algorithm in 2D.

In the current system, collisions are detected and object interpenetration is prevented but the system does not model collision response. When a collision is detected, the paths of moving objects are retraced until no overlap between the objects occurs. In future systems, we intend to add dynamics so that physically realistic reactions to collisions can be simulated. However, while there has been significant progress in computer graphics for real-time simulation of collision response for rigid surface models (e.g. (Baraff, 1989) and (Mirtich and Canny, 1995)), these methods cannot be directly transferred to volumetric methods. In existing physics-based graphics techniques, collision responses are calculated by solving systems of equations with dimensions proportional to the number of contact points between the objects. In surface-based methods, the number of contact points is related to the number of vertices on the contacting surfaces -- with possibly tens or hundreds of contact points per collision instance. However, for volumetric objects, contacting volume elements can number in the thousands for relatively small contact surfaces. Solving systems of linear equations of this size are prohibitive for interactive applications. We have begun to investigate ways to integrate the collision effects over many contact points so that collision responses can be calculated at reasonable rates.

3.4 Tissue Cutting, Tearing and Suturing

Modeling surgical interventions requires the simulation of tissue cutting, tearing, and repairing or suturing. Using a volumetric model in which elements encode the positions of neighboring elements, cutting and tearing can be performed by breaking connections between neighbors. For example, in the data structure of Figure (3), the connections between 2 neighbors such as element

A and its right neighbor B can be broken by setting the appropriate neighbor pointers, here the right pointer of A and the left neighbor pointer of B, to NULL. Similarly, two elements A and B can be joined together by setting their appropriate neighbor pointers to point to each other.

For cutting, neighbor connections are broken along the path of the knife instrument as it is passed through the virtual object. Intersections between the knife path and the object are detected using a strategy similar to that used for detecting collisions between two objects. A fast, Bresenham-like line drawing algorithm (Foley et al, 1992) is used to draw edge lines into the occupancy map. If the knife path encounters a cell occupied by and element or an edge line, the appropriate neighbor connections are broken.

Tearing occurs when the distance between two elements is stretched beyond the allowable limit. This occurs, for example, if one side of the object is tacked in place and the other side is grasped and pulled in the opposite direction. When such a limit violation is detected and cannot be resolved by moving neighboring elements, the connection between the two elements is broken.

For joining or suturing objects, elements along the path of the joining instrument which have NULL neighbor connections are paired and joined. These edge elements are detected by searching occupancy map cells within the vicinity of the joining instrument. If two elements with complementary missing neighbor connections are found, then the two elements are joined. For example, if element A, with a missing right neighbor and element B, with a missing left neighbor are both found in the vicinity of the glue brush, then the right neighbor of A is set to B and the left neighbor of B is set to A. Figure (6) is the user interface for an implementation of a 2D system that allows object translation, deformation, arbitrary cutting, joining (or gluing), element erasing, and tacking of elements into place. An application of this system for simulating cutting through knee ligaments is described below.

4 Simulation of Arthroscopic Knee Surgery

As part of an ongoing collaboration between MERL, Carnegie Mellon University, Massachusetts Institute of Technology, and Brigham and Women's Hospital, we have built a prototype simulation system to implement and test some of the ideas that have been presented in this paper (Figure 7.). The current system consists of 3D volumetric models of knee anatomy created from MRI image scans, a haptic interface device for interacting with the knee model, and real-time rendering for visualization of the interaction. The current status of this system is described in more detail below and future plans for integrating more of the techniques introduced in this paper are described.

4.1 Motivation

Volumetric techniques have applications in many areas of surgical simulation. For our initial prototype system, we have focused on arthroscopic knee surgery, a relatively new and minimally invasive procedure that is often used to diagnose and treat knee injuries. In arthroscopy, the joint is visualized and accessed through small portals. An optical endoscope equipped with a video camera allows visualization of the procedure through one of the portals, while surgical probes and other instruments are inserted into additional portals (see Figure (8)).

We have focused on arthroscopic knee surgery for a number of reasons. First, in the US alone, it is estimated that of approximately 1.8 million arthroscopic procedures performed in 1996, 88-90% were knee procedures (Praemer, 1996).

Second, while arthroscopic procedures have been shown to reduce costs and increase patient recovery rates, they suffer from specific technical limitations: namely limited visibility through the arthroscope, difficulty in orienting the camera with the surgeon's viewpoint, and restricted motion of the surgical tools. Computer simulation will enhance the education and training of surgeons and help them deal with these technical challenges.

Third, the knee offers an ideal initial platform for this project. Important structures lie in a limited volume so that the size of a volumetric knee model is reasonable. Many of the important structures in the knee are rigid so initial haptic and rendering implementations could be performed on static data sets. In addition, the soft tissues that are of primary importance -- cartilage layers, the menisci, and cruciate ligaments -- are relatively small and hence can be used to test deformation and tissue cutting algorithms. In addition the deformable tissues are small enough that real-time rendering of these irregular volumes is currently feasible. As more efficient algorithms, better data representations and faster hardware are developed, techniques developed for the arthroscopic knee simulator will be extended for other simulator systems.

4.2 Prototype System: Current Status

In the prototype system, object models consist of the bony structures of the joint, the articular cartilage, the menisci, and the cruciate ligaments. These structures were hand-segmented from MRI data. The integrated system currently allows probing of bony structures with real-time visual and haptic feedback. Visual feedback is provided by both volume and polygon rendering and sectional imaging (Samosky, 1993). The system will be extended in the short term with recent advances including: a higher resolution MRI image that has been acquired and segmented; soft-tissue modeling of the cartilage and cruciate ligaments; and the addition of two surgical tool models to the haptic interface: a surgical probe and a cartilage debridement tool.

While the ultimate goal of this research collaboration is to produce a system that could run on affordable and accessible hardware, the research platform for technology development and the current simulator prototype consists of an 8-processor SGI Challenge with MIPS R10K processors, Infinite Reality graphics, 4 RM6 raster manager boards with 3D texture mapping (for interactive volume rendering), and 512 Mbytes of RAM. Haptic feedback is provided by a SensAble Technologies' PHANToM with 3 degrees-of-freedom force reflection and 6 degrees-of-freedom sensing of position and orientation. The PHANToM is currently controlled by an SGI Extreme connected to the SGI Challenge via an Ethernet connection.

4.2.1 Image Acquisition and Segmentation

For the initial prototype system, a T-1 weighted proton density MRI knee image sequence was acquired. The image size was 256x256x124 with a voxel size of 0.63x0.63x0.9 mm^3. This image was hand segmented into bony structures (femur, tibia, fibula, patella), cartilage (femoral, tibial, and patellar), lateral and medial menisci, and anterior and posterior cruciate ligaments. These structures are illustrated in the surface rendered image of Figure (9).

Unlike CT images, where thresholding techniques can be used effectively for segmentation, there are no general automatic segmentation tools for MRI images. As can be seen in Figure (10), in MRI, image intensities within a single tissue can vary significantly while adjacent tissues can have very similar or identical intensities -- making it difficult to define surfaces between the structures.

Although segmentation techniques can be fine-tuned for particular MRI sequences or specialized for specific anatomies to enable semi-automatic segmentation of MRI images within a limited application, we do not currently have such a system customized for the knee. As a result, knee images for the current simulation system were segmented by hand -- requiring an expert to trace and label individual structures in each cross-sectional image of the 3D MRI image. While this procedure is tedious and time consuming, it has provided us with the models required by other components of the surgical simulation system. Future plans include the development of a set of image segmentation tools specialized for knee anatomy that would reduce the segmentation time and facilitate surgical planning based on patient-specific data.

We have recently acquired a high resolution knee image of size 512x512x90, with voxel size: 0.25x0.25x1.4 mm^3. The image acquisition time was 50 minutes. Both T1-weighted proton density images and fat-suppressed images were acquired. A 2D sectional image from the proton density scan is shown in Figure (2). This data set has been hand segmented into bony structures, articular cartilage, menisci, cruciate ligaments, and the quadriceps and gastrocnemius muscles and tendons and will eventually be incorporated into the knee arthroscopy prototype.

4.2.2 Real-time rendering

Volume rendering is a powerful method for visualizing volumetric data (Kaufman, 1991). The basic approach in volume rendering involves projecting and compositing elements of the volume onto a 2D display. Volume elements can be rendered semi-transparently for visualization of interior structures. However, because of this transparency, volume elements must be projected in order (either back-to-front or front-to-back) and every element in the volume may contribute to the final image. The need to access, order, and render millions of voxels (a volume of size 256x256x256 contains 16 million voxels) makes volume rendering computationally expensive -- with some algorithms requiring minutes of computation per frame.

As part of a parallel project, MERL and Mitsubishi Electric are developing special purpose volume rendering hardware for interactive PC-based applications. The near term goal of this project is hardware acceleration for volume rendering. A longer term goal is the integration of volumetric and surface-based rendering in 3D graphics hardware. This project complements our activities in surgical simulation since fast rendering hardware will facilitate real-time, high-quality rendering of volumetric models.

Interactive visualization of the surgical simulation is essential for the prototype system. We are currently using two approaches for visualizing object models: volume rendering using the 3D texture map memory available on our research platform (Fraser, 1995) and hardware-accelerated polygon rendering of surface models that were generated from the segmented images. A cursor indicating the position of the tip of the force feedback device is also rendered in both approaches. Figure (11) shows two sample renderings.

Polygon rendering allows us to improve the visual appearance of the rendered image by adding shading and texture. However, the surface model differs somewhat from the volumetric model that is used for haptic rendering and tissue modeling. On the other hand, the current volume rendering implementation lacks shading and surface definition. We hope that advances in volume rendering hardware will address this problem, allowing us to produce high quality shaded volume rendered images like that shown in Figure 12 in real-time.

Both the high quality volume rendering algorithm and the texture map volume rendering approach assume that the data is stored in a grid of regularly spaced volume elements. This is sufficient for visualization of the rigid bony objects or static data. However, when objects are deformed, changes in the relative positions of volume elements result in an irregular grid. We have implemented and are currently investigating several methods for higher quality rendering of irregular volumes including raycasting (Ramamoorthy and Wilhelms, 1992) cell projection (Shirley and Tuchman, 1990), (Williams, 1992), and splatting (Westover, 1990). Unfortunately, while some of these techniques may eventually be integrated into hardware approaches, they are currently too slow to provide feedback for interactive deformation. In the meantime, we use OpenGL to render element points or lines connecting elements in the deformed object. While the resultant images do not have the visual quality of volume or polygon rendered images, they provide enough visual feedback to guide the interaction.

4.2.3 Haptic Interaction

The surgical simulation prototype uses SensAble Technologies' PHANToM to provide force reflectance and feedback for haptic interaction (Figure (4)). The PHANToM provides 6 degrees of sensing (3D position and orientation) and 3 degrees of force feedback (position only). Use of the force feedback device allows the user to explore object models using the sense of touch. Eventually, the force feedback device will provide valuable sensory feedback to the surgeon during simulation of tissue deformation and cutting.

The object models in the surgical simulation system are derived from segmented MRI images. Because the resolution of individual voxels is much lower than the haptic resolution of the PHANToM, the binary segmented object models must be smoothed before surface normals are calculated. If smoothing is not done, the binary nature of the data can cause direction of the surface normals to vary widely with small changes in position, causing unstable force feedback.

In our prototype implementation, we have used a haptic model that is similar to the density field method of (Avila and Sobierajski, 1996). A discrete 3D density field map was generated by smoothing the binary segmented data with a 5x5x5 Gaussian filter. During haptic interactions, the applied force is calculated to be proportional to the gradient of the density field at the PHANToM tip. The gradient is calculated using a central difference method from trilinearly interpolated densities of 6 positions located one unit distance above, below, behind, in front, to the right and to the left of the tip position:

$$F(x,y,z) \propto \big(d(x+1,y,z)-d(x-1,y,z), d(x,y+1,z)-d(x,y-1,z), d(x,y,z+1)-d(x,y,z-1)\big),$$

where $d(x,y,z)$ is the density field at the PHANToM tip position, (x,y,z).

In our current system, the PHANToM is controlled by an SGI Extreme connected to the SGI Challenge via a Ethernet connection. This is necessary because we cannot interface the PHANToM directly with the Challenge at this time. In order to enable interactions at >1KHz rates we store two models of the knee, a haptic model on the Extreme and a visual model on the Challenge. The Ethernet connection is used to send information about the PHANToM tip location to the graphical interface and for communicating the global position and orientation of the model.

4.2.4 Deformation, cutting, and suturing of a 2D knee model

Tissue deformation, cutting and suturing have not yet been integrated into the prototype surgical simulation system. However, we have implemented the 2D system shown in Figure (5) that allows simulation of these actions on a 2D volumetric model. Figure (13) is a sequence of images captured during interaction with this system. The 2D objects consist of rigid models for the tibia and femur and deformable soft tissue models for the tibial and femoral cartilage and the posterior cruciate ligament. In Figure (14b) and (c), the grasping tool is used to grab and pull on the posterior cruciate ligament. In (d) the posterior cruciate ligament is stretched and tacked in place. In (e) the ligament is cut with the knife tool. Note in (f) that after the cutting, the elastic nature of the ligament causes its shape to relax away from the cut edge to a lower energy state, mimicking what would happen in real life.

Before these techniques can be extended to 3D and integrated into the surgical simulation system, there are several technical issues that must be addressed. The 3D ChainMail algorithm has been used to interactively manipulate an object of size 50x50x50 -- which is significantly bigger than the volumetric ligament and cartilage models. Hence, we believe that deformation calculations will not be prohibitively expensive. However, our current use of separate haptic and visual object representations is not practical for interactive model manipulation since both representations must be updated whenever the object changes. In addition, model changes must be communicated to the

Extreme over the (relatively slow) Ethernet connection. We are investigating ways to control the PHANToM directly from the SGI Challenge to reduce this problem. Secondly, we are testing a number of rendering strategies so that we can provide real-time visual feedback of sub-volumes of deformed objects. Finally, the current density field technique that has been used for the haptic feedback requires preprocessing of the data by smoothing with a Gaussian filter. This preprocessing may be too expensive when the data is being interactively manipulated. Hence we are exploring new haptic modeling methods.

4.3 Future Directions and Focus

We have gathered initial feedback about the prototype system from physicians and surgeons that will be used to guide the next phase of the system. Important goals include the incorporation of deformation and cutting into the simulator, the addition of surgical tools such as a surgical probe and a cartilage debridement tool to the haptic interface, and validation of tissue properties.

Figure (14) shows the distribution of the 83,294 Medicare funded arthroscopic knee procedures performed in 1992. This chart is organized in such a way that, starting at 12:00 and proceeding in a clockwise direction, the difficulty of simulating these arthroscopic procedures increases. Our existing system, which allows visualization, navigation and haptic probing of rigid structures in the knee model, would be a useful training tool in all of the listed procedures. However, we intend to extend the system to simulate more complete procedures. In the near term we will start with the goal of simulating the shaving or debridement of damaged cartilage.

5 Discussion

Volumetric methods are a powerful tool for modeling objects with complex interior structure and tissue properties. In this paper, we have discussed many of the advantages and technical challenges that are faced in using volumetric methods in surgical simulation. We have also introduced several techniques for modeling and interacting with volumetric objects and have described an initial prototype simulator that uses volumetric methods for modeling arthroscopic knee surgery.

While we advocate the use of volume methods, we readily admit that there are many times when surface models are more suitable. For example, if a CAD or polygonal model of a surgical instrument exists, there may be no advantage in converting the model to a (possibly) less accurate volumetric model. Hybrid methods that combine both polygonal and volumetric models for rendering, haptics, and physically-realistic modeling method would be ideal. However, there are many challenges in modeling and rendering that must be overcome before the merging of polygonal and volume graphics becomes feasible. In the meantime, volumetric methods provide a powerful complement to surface graphics for modeling, visualizing, and interacting with computer models.

10 Acknowledgments

We gratefully acknowledge the Wellness Business Development Center, Mitsubishi Electric Corporation, Japan, for financial support in this collaboration.

11 Bibliography

R. Avila and L. Sobierajski, (1996) “A haptic interaction method for Volume Visualization”, Proc. Visualization ‘96, ed. R. Yagel and G. Nielson. p.p. 197-204.

D. Baraff, (1989) “Analytical methods for dynamic simulation of non-penetrating rigid bodies”, (proc. SIGGRAPH), Computer Graphics, Vol. 24, pp. 19-28.

D. Chen, (1991) "Pump it up: computer animation of a biomechanically based model of muscle using the finite element method", Ph.D. thesis, Media Lab, MIT.

S. Cotin, H. Delingette, J.M. Clement, L. Soler, N. Ayache, and J. Marescaux (1996) "Geometrical and physical representations for a simulator of hepatic surgery", Proceedings of Medicine Meets Virtual Reality IV.

S. Cover et al., (1993) "Interactively deformable models for surgery simulation", IEEE Computer Graphics and Applications, Vol. 13, 6, pp. 68-75.

J. Foley, A. vanDam, S. Feiner, and J. Hughes, (1992) “Computer graphics principles and practice”, Addison-Wesley, pp. 72-81, 2nd edition.

R. Fraser, “Interactive volume rendering using advanced graphics architectures”, whitepaper: http://www/sgi.com.

T. Galyean and J. Hughes, (1991) “Sculpting: an interactive volumetric modeling technique”, (proc. SIGGRAPH), Computer Graphics, Vol. 25, pp. 267-274.

B. Geiger and R. Kikinis, (1995) "Simulation of endoscopy", Proceedings CVRMed '95, ed. Nicholas Ayache, Springer Verlag, Berlin, pp. 277-281.

S.F. Gibson, (1995) "Beyond volume rendering: visualization, haptic exploration, and physical modeling of voxel-based objects", in Visualization in Scientific Computing , eds. R. Scateni, J. van Wijk, and P. Zanarini, Springer-Verlag, pp. 10-24.

S.F. Gibson, (1997) “3D ChainMail: a fast algorithm for deforming volumetric objects”, proc. Symposium on Interactive 3D Graphics, ACM Siggraph, (in press), April.

L. Hong, et al., (1995) "3D virtual colonoscopy", Proc. of Biomed. Vis. '95, eds. M. Loew and N. Gershon, IEEE Comp. Soc. Press, Las Alimitos, pp. 26-32.

W. Hsu, J. Hughes, H. Kaufman, (1992) "Direct Manipulation of Free-form deformations", Computer Graphics (proc. SIGGRAPH), Vol 26, 2, pp. 177-184.

I. Hunter et al, (1993) "A teleoperated microsurgical robot and associated virtual environment for eye surgery", Presence, Vol. 2, pp. 265-280.

A. Kaufman, (1991) Volume Visualization, IEEE Comp. Soc. Press, Los Alimitos, CA.

A. Kaufman, (1996) “Volume Visualization”, CRC Handbook of Comp. Sci. and Eng..

U.G. Kuhnapfel et al, (1995) "Endosurgery simulations with KISMET: a flexible tool for surgical instrument design, operation room planning, and VR technology based abdominal surgery training", Proc. VR '95 WORLD Conference.

W.E. Lorensen, F.A. Jolesz, and R. Kikinis, (1995) "The exploration of cross-sectional data with a virtual endoscope", in Interactive Technology and the New Medical Paradigm for Health Care, eds. R. Satava and K. Morgan, pp. 221-230.

P. Lacroute and M. Levoy, (1994) "Fast volume rendering using a shear-warp factorization of the viewing transform", proc. SIGGRAPH, Computer Graphics, pp. 451-457.

Y. Lee, D. Terzopoulos, and K. Waters, (1995) "Realistic modeling for facial animation", Computer Graphics (proc. SIGGRAPH), pp. 55-62.

B. Mirtich and J. Canny, (1995) “Impulse-based simulation of rigid bodies”, proc. 1995 Workshp on Interactive 3D Graphics, pp. 181-188, April.

A. Pentland and J. Williams, (1989) “Good vibrations: modal dynamics for graphics and animation”, (proc. SIGGRAPH), Computer Graphics, Vol. 23, pp. 215-222.

S.D. Pieper, (1992) "CAPS: computer-aided plastic surgery", Ph.D. thesis, MIT.

Alan Praemer, (1996), Department of Research and Scientific Affairs, American Academy of Orthopedic Surgeons, personal communication.

Alan Praemer, (1996) from Health Care Finances Admin., Medicare data, 1992, personal communication.

S. Ramamoorthy and J. Wilhelms, (1992) “An analysis of approaches to ray-tracing curvi-linear grids”, Tech. report UCSC-CRL-92-07, Comp. Res. Lab, UC Santa Cruz.

J. Samosky, (1993) “SectionView -- A system for interactively specifying and visualizing sections through three-dimensional medical image data”, MS Thesis, MIT.

P. Shirley and A. Tuchman, (1990) “A polygonal approximation to direct scalar volume rendering”, proc. San Diego Workshop on Volume Visualization, pp. 63-70.

D. Terzopoulos, K. Waters, (1990) "Physically-based facial modeling, analysis, and animation", J. Visualization and Computer Animation, Vol. 1, pp. 73-80.

D.J. Vining, A.R. Padhani, S. Wood, et al, (1993) “Virtual bronchoscopy: a new perspective for viewing the tracheobranchail tree”, Radiology 189 (P):438.

D.J. Vining et al, (1994) “Virtual colonoscopy”, Radiology 193 (P):446.

S. Wang and A. Kaufman, (1995) "Volume sculpting", ACM Symposium on Interactive 3D Graphics, Monterey, CA, pp. 151-156, April.

K. Waters, (1987) "A Muscle model for animating three-dimensional facial expression", Computer Graphics, Vol. 21, 4, July, pp. 17-24.

L. Westover, (1990) “Footprint evaluation for volume rendering”, ACM SIGGRAPH ‘90, pp. 144-153, 1990.

P.L. Williams, (1992) “Interactive splatting of non-rectilinear volumes”, Proc. IEEE Visualization ‘92, pp. 37-44.

R. Ziegler et al (1995) "A virtual reality medical training system", Proc. CVRMed '95, ed. N. Ayache, Springer-Verlag, Berlin, pp. 282-286.

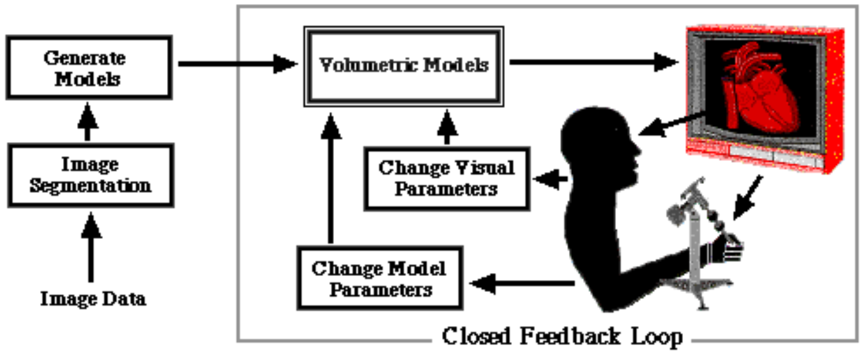

Figure 1: Surgical simulation system components. This system provides a closed feedback loop between the user and the simulator. The user can manipulate object models and observe the results both visually and haptically.

Figure 2. High resolution MRI image with voxel resolution: 0.25x0.25x1.4 mm. Arrows indicate structure in muscle and fatty tissue that are important in modeling tissue behavior but that are hard to incorporate into surface-based models.

```
typedef struct 3DElement {
        char type;                                             /* tissue type */
        char r, g, b;                                  /* element r,g,b color */
        float x, y, z;                                    /* element position */
        struct 3DElement *right, *left, *top, *bottom, *back, *front;
                                               /* pointers to element neighbors */
}
```

Figure 3. Structure definition for a single element in a volumetric element for a tissue deformation application.

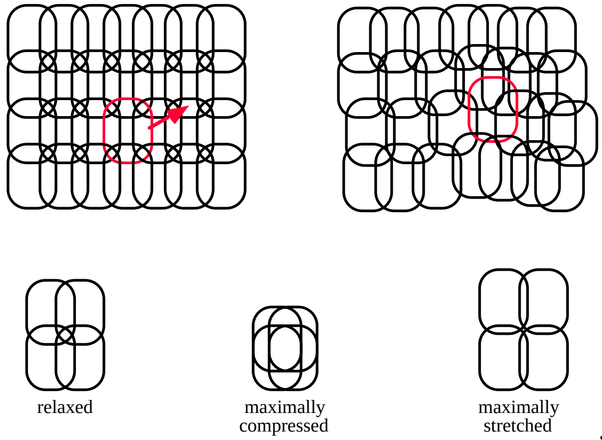

Figure 4. 2D ChainMail. a) and b) When one link of the 2D object is moved along the path of the arrow, its neighboring links move to satisfy maximum and minimum distance requirements between elements. Sets of 4 links at the bottom of the figure show relative link positions when the 4 links are in c) a relaxed state, d) a maximally compressed state, and e) a maximally stretched state.

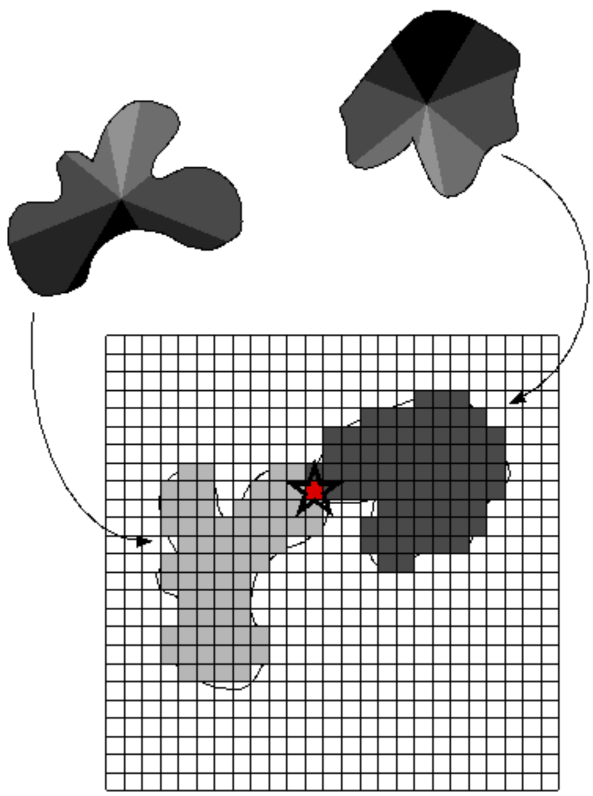

Figure 5. Simple collision detection for volumetric objects. Object elements are mapped into cells of the regular grid of the 3D occupancy map. As objects are moved, cells of the occupancy map are updated. If an element is written into a cell that is already occupied by a different object, a collision is detected.

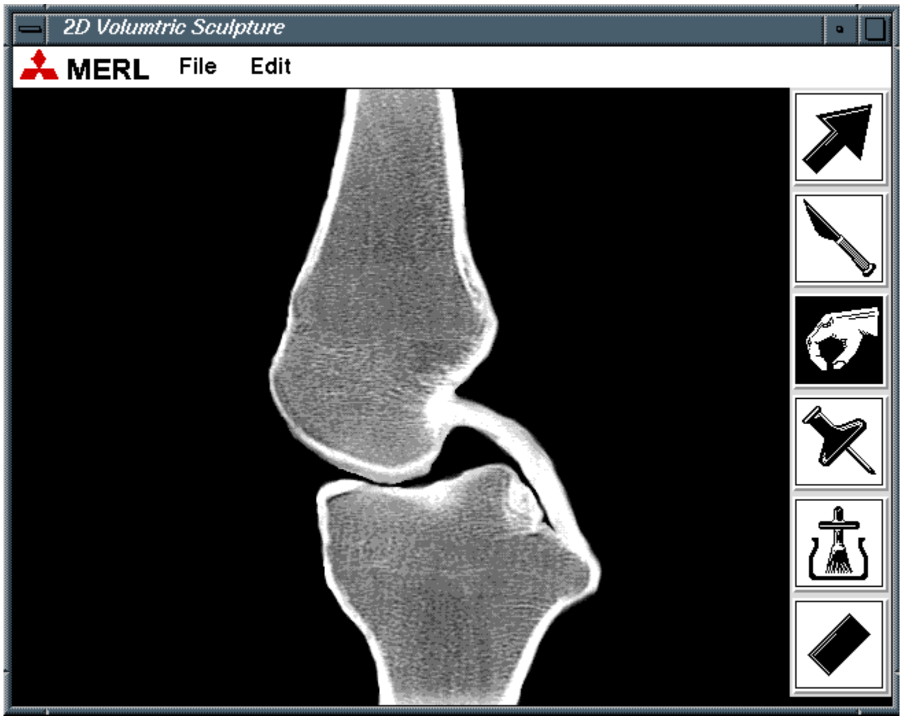

Figure 6. An interactive 2D application for manipulated 2D volumetric objects. Object elements can be moved (resulting in object translation), cut, grasped and moved (resulting in object deformation), tacked into place, glued together, and erased interactively using the computer mouse. Pointing and clicking on the buttons on the right side of the user interface switches between these modes.

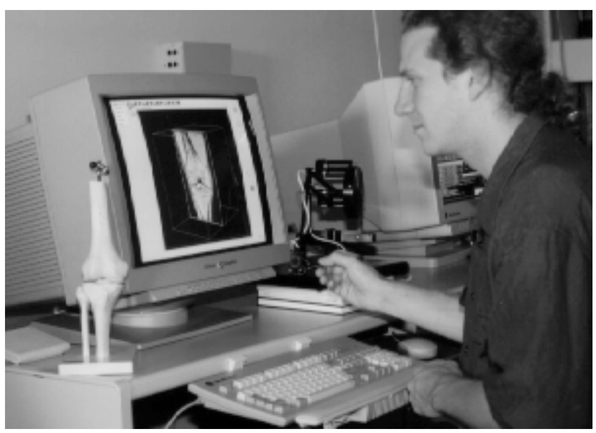

Figure 7. Prototype surgical simulation system. A 3D volumetric knee model has been generated from MRI images. A force feedback device allows the user to haptically explore the boney surfaces of the knee model while real-time polygon and volume-based rendering provide visual feedback to the user.

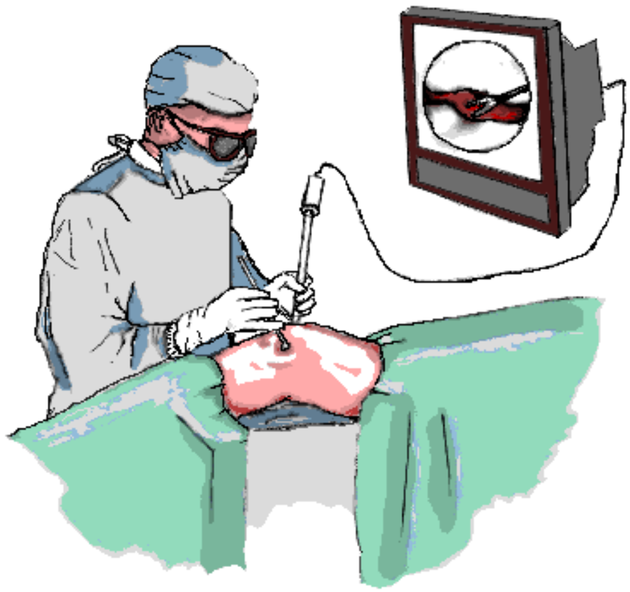

Figure 8. Illustration of arthroscopic knee surgery. The knee is accessed through small portals. A video endoscope inserted through one portal captures a video image of the surgical field and projects it onto a video monitor. The surgeon manipulates surgical instruments through additional portals.

Figure 9. Surface rendered image of segmented knee. The knee was segmented into bony structures (femur, tibia, fibula, patella), cartilage (femoral, tibial, and patellar), lateral and medial menisci, and anterior and posterior cruciate ligaments

Figure 10. MRI knee image. Automatic segmentation of MRI images is difficult because intensities within a single tissue can vary significantly while adjacent tissues can have very similar or identical intensities

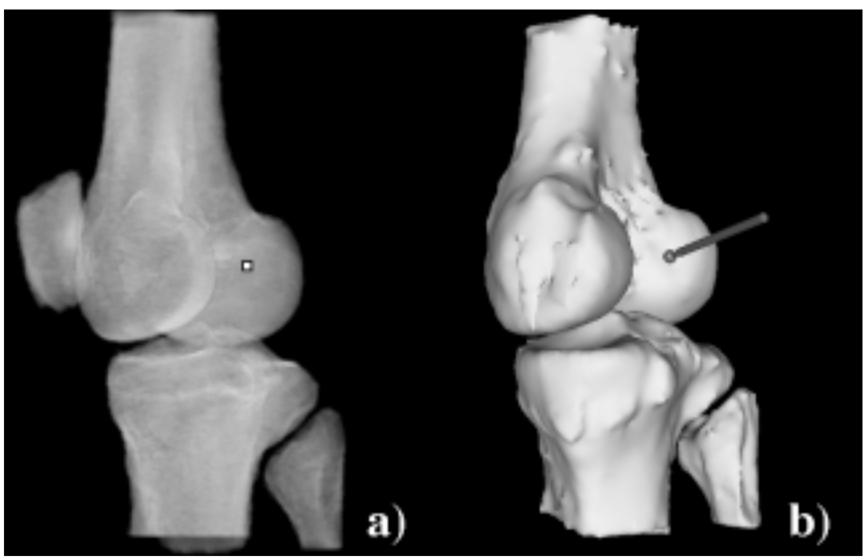

Figure 11. Rendered images of knee bones in the surgical simulation prototype system. a) texture mapped volume rendered image. b) polygon rendered image of surface models.

Figure 12. High quality volume rendered image of the volumetric knee model with diffuse and specular shading. Voxel opacity and shading were calculated from the binary segmented image. Color values were derived from the original image intensity values. This image took several minutes to compute on an SGI workstation with an R10K processor.

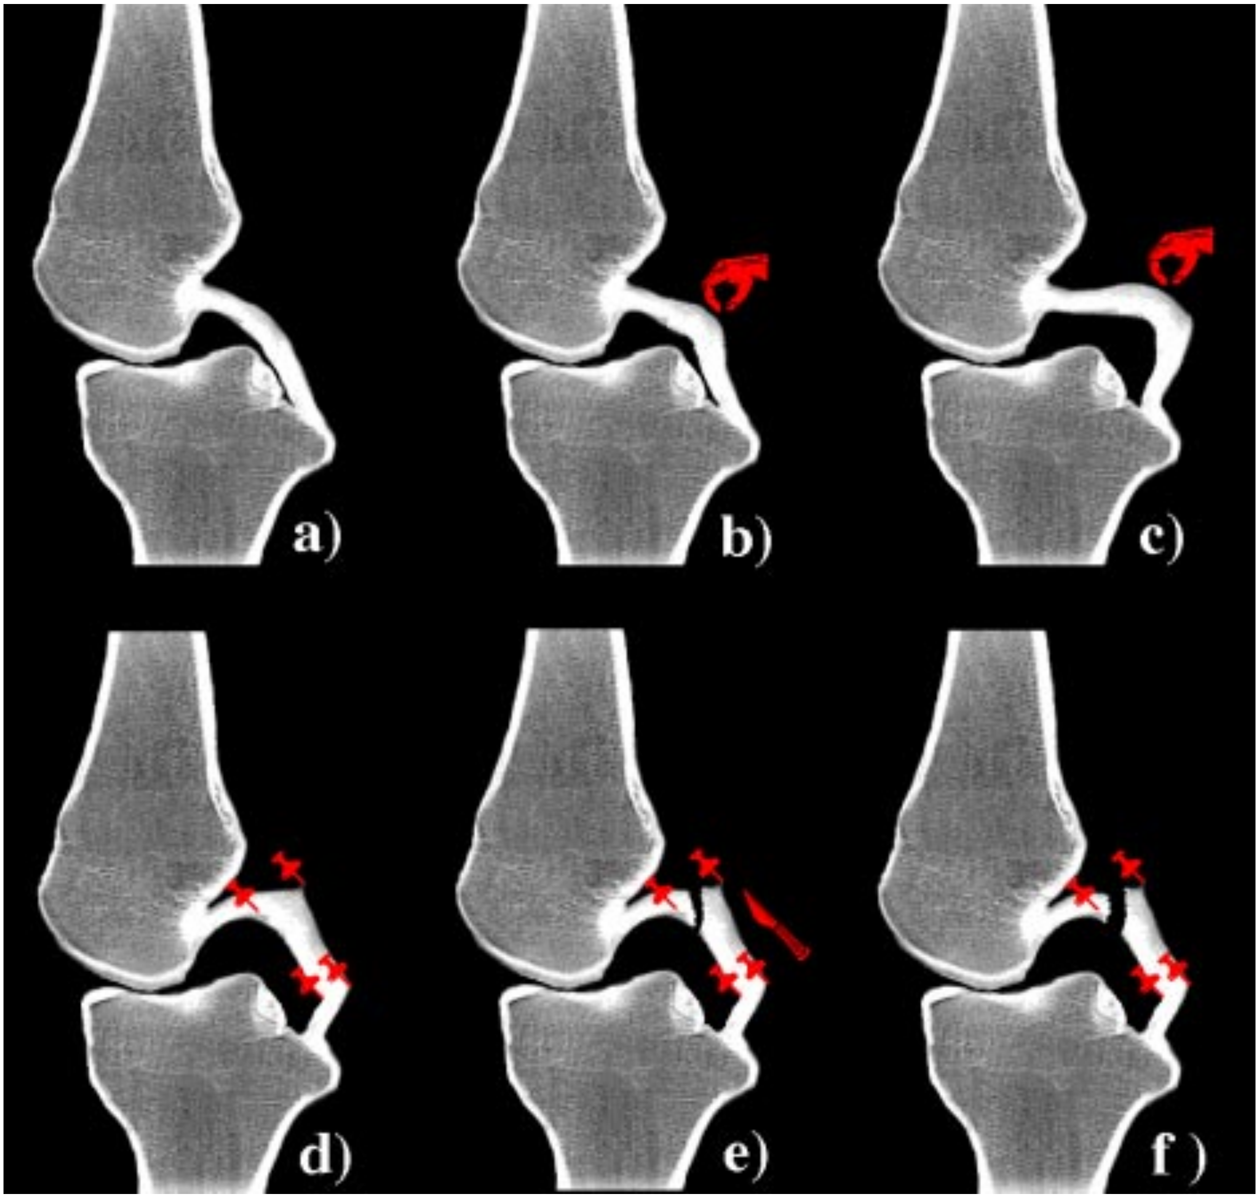

Figure 13. 2D system showing soft tissue deformation and cutting. In b) and c), the posterior cruciate ligament is grasped and pulled. In d), the ligament is stretched and tacked into place. In e) the ligament is cut. f) shows how the elastic ligament tissue pulls back from the cut.

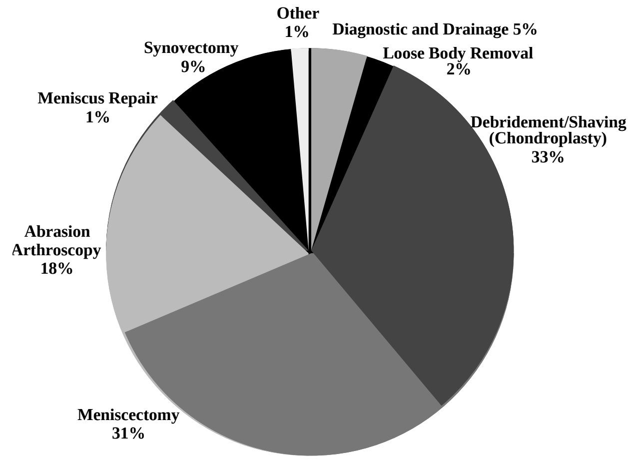

Figure 14. Distribution of 83,294 total Medicare-funded arthroscopic knee procedures, 1992 [24].