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

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

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

TR96-19 December 1996

Abstract

A system for simulating arthroscopic knee surgery that is based on volumetric object models derived from 3D Magnetic Resonance Imaging is presented. Feedback is provided to the user via real-time volume rendering and force feedback for haptic exploration. The system is the result of a unique collaboration between an industrial research laboratory, two major universities, and a leading research hospital. In this paper, components of the system are detailed and the current state of the integrated system is presented. Issues related to future research and plans for expanding the current system are discussed.

Joint Conference on Computer Vision and Virtual Reality in Medicine and Medical Robotics and Computer Assisted Surgery

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

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

1. MERL -- A Mitsubishi Electric Research Lab, 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

A system for simulating arthroscopic knee surgery that is based on volumetric object models derived from 3D Magnetic Resonance Imaging is presented. Feedback is provided to the user via real-time volume rendering and force feedback for haptic exploration. The system is the result of a unique collaboration between an industrial research laboratory, two major universities, and a leading research hospital. In this paper, components of the system are detailed and the current state of the integrated system is presented. Issues related to future research and plans for expanding the current system are discussed.

Introduction

Computer-based surgical simulation has many applications in education and training, surgical planning, and intra-operative assistance. Given the low availability and high cost of cadaver and animal specimens, surgical simulation can be used in medical education and training to reduce costs, to provide experience with a greater variety of pathologies and complications, and to make it possible to repeat or replay training procedures. In surgical planning, a simulator can enable rehearsal of difficult procedures or planning on patient-specific anatomy. Surgical simulators can enhance communication among medical professionals or between doctor and patient. Intra-operatively, computer modeling can aid in navigation by augmenting the limited endoscope view with a more global view of the patient's anatomy and can provide guidance by preventing the surgical instrument from moving into pre-defined sensitive regions.

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 The ability to perform diagnosis and surgical intervention in joints without open surgery has had a large impact in orthopedics. In the US alone, it is estimated that 1.8 million arthroscopic procedures will be performed in 1996. Of these, 88-90% will be knee procedures [23]. Arthroscopic procedures have been shown to reduce costs, and increase patient recovery rates. However, arthroscopy suffers 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. Because of these technical challenges, it is important that surgeons receive adequate training in arthroscopic techniques. Such training could be enhanced by computer simulation.

In this collaboration, we are developing a computer-based surgical simulator that uses volumetric object models generated from 3D Magnetic Resonance Imaging (MRI) data. Techniques from physics-based graphics are extended for volumetric objects to model physically realistic interactions between both rigid and deformable anatomical models and between anatomical models and surgical instruments. Both visual and haptic feedback are provided to the user via real-time rendering of the volumetric models and a force feedback device.

Figure (1) outlines the components of the system. Patient-specific 3D MRI or CT images are processed to generate volumetric object models. Object models are presented both visually via rendering on the computer monitor and haptically with a force feedback device. Visual parameters such as viewpoint, color and opacity transfer functions, and lighting effects can be interactively controlled and object models can be manipulated with force feedback to change relative object positions, to probe and mold objects, and to simulate surgical procedures such as cutting, tearing, and suturing. This interaction with visual parameters and object models closes a feedback loop between the user and the simulator, enhancing better understanding of both anatomical structure and function in the patient model.

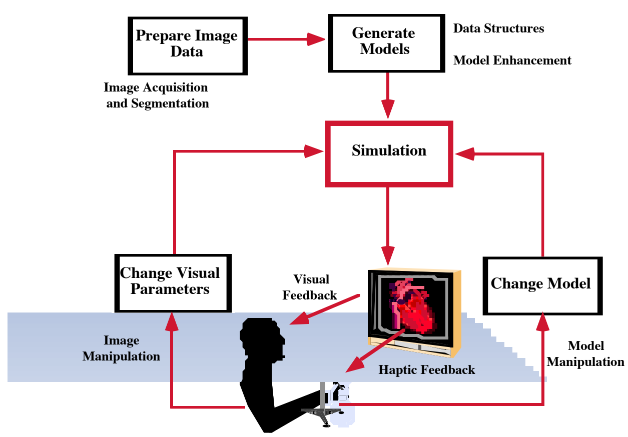

Figure 1: Surgical simulation system components.

The important technologies required by each component in the simulator are listed in Figure (1). These include: acquisition and segmentation of patient-specific 3D medical images; incorporation of the data into efficient and effective data structures; enhancement of the object models with measured material properties, visual and haptic texture maps, and other data; visual feedback through interactive volume rendering; haptic feedback through a surgical instrument controlled by an electromechanical force feedback device; image manipulation -- allowing adjustment of visual parameters; and model manipulation through the haptic interface – allowing interactive manipulation of object models.

Prior Work

There are several related surgical simulation systems under development both commercially and in research laboratories. For example, a number of endoscopy simulators enable navigation through stationary object models [18], [12], [35], [8], [29], [30]. 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. [11], [3], [6]). 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 [32], [28], [4], and in surgical simulation [5], [14], [22].

Manipulation of voxel-based objects [16], has been applied to object modeling [9] and has been combined with a force feedback device for haptic exploration of voxel-based objects [1].

Volumetric Object Representation

A voxel-based volumetric object is a regular or irregular 3D array of data, with each element representing a sampled point (measured or calculated) in the volume. For surgical simulation, this representation has a number of advantages over the use of surface polygons or solid geometric primitives. First, because the data organization is the same as the acquired data, a voxel-based representation is natural for the 3D digital images produced by medical scanning technologies such as MRI or Computed Tomography (CT). Second, since no surface extraction or other data reformatting is required, errors introduced by fitting surfaces or geometric primitives to the scanned images can be avoided. Finally, volumetric objects can incorporate detailed information about the internal anatomical or physiological structure of organs and tissues. This information is particularly important for realistic modeling and visualization of complex tissue volumes.

In addition to accurate representation of anatomy, surgical simulation requires physically realistic modeling of object interactions. This includes detecting and simulating collisions between objects and modeling the response of both rigid and deformable tissues to probing, cutting and tearing. Fast collision detection between voxel-based objects is provided in the prototype system by a straightforward algorithm [9] which will be extended to increase speed and reduce memory requirements using techniques from conventional computer graphics.

Although we intend to incorporate physically realistic dynamics into the surgical simulation system, the current prototype does not yet model reactions to collisions. While there has been significant progress in computer graphics for real-time simulation of collision response for rigid surface models (e.g. [2] and [21]), volumetric object models provide unique challenges for interactive systems. For example, the number of contacts between voxel-based volumetric objects poses a problem for calculating collision responses. With existing physics-based graphics techniques, collision responses can be 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 between two flat surfaces is related to the number of vertices on the contacting surfaces. This can result in tens or possibly hundreds of contact points per collision instance. In voxel-based object models, the number of contacting volume elements is proportional to the area of the contact surfaces – resulting in thousands of contact elements for relatively small contact surfaces. We are investigating ways to integrate the collision effects over a large number of contact points so that collision responses can be calculated at reasonable rates.

Because interior structure is represented, volumetric representations are particularly suitable for modeling deformation, cutting, and tearing of tissues. However, since voxel-based object models can consist of thousands to millions of volume elements, well studied material modeling approaches such as finite element methods can not be used for interactive applications unless the number of volume elements are greatly reduced. An algorithm is described below for deforming voxel-based objects that balances physical realism and mathematical accuracy with speed. The algorithm can be applied to model a range of materials -- including rigid, elastic, and plastic substances -- and can model substances like muscle which have different material properties parallel to and perpendicular to the muscle fiber axis.

Large memory requirements provide an additional challenge of volumetric object representations. An MRI image of size 256x256x256 contains 16M voxels. If visual and material properties, point location, and neighboring connections are incorporated into voxel data structures, as many as 32 bytes of data can be required for each voxel. In addition to obvious data storage issues, these large memory requirements have practical implications for data access that greatly affect speed in both rendering and modeling. We are investigating the use of hierarchical data representations, data compression schemes, and special purpose hardware for reducing data access and storage overhead to help make volumetric methods practical in real-time applications.

Image Acquisition and Segmentation

For the initial prototype system, a T-1 weighted proton density MRI image sequence was acquired of a normal male knee. The image size was 256x256x124 with a voxel size of 0.63x0.63x0.9 mm. These images were 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 (2).

Figure 2. Surface rendered image of segmented knee.

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 (3), 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.

Figure 3. MRI knee image.

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 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 tools specialized for knee anatomy with the goal of requiring less than one hour of segmentation for an MRI image sequence. This time frame would be necessary in a surgical simulation system designed for patient-specific data.

We have recently acquired a high resolution knee image of a normal male of size 512x512x90, with voxel size: 0.25x0.25x1.4 mm. The image acquisition time was 50 minutes. Both T1-weighted proton density images and fat-suppressed images were acquired. Images from these data sets are shown in Figure (3). This data set is currently being hand segmented into bony structures, articular cartilage, menisci, cruciate ligaments, and the quadriceps and gastrocnemius muscles and tendons. Once the segmentation is complete, these models will replace the current models in the prototype system.

Real-time rendering

Volume rendering is a powerful method for visualizing volumetric data [15]. The basic approach in Volume Rendering involves projecting and compositing elements of the volume onto a 2D display. To enable visualization of interior structure, volume elements can be rendered semi-transparently. 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 inherently slow -- with some algorithms requiring minutes of computation per frame. On the other hand, there are many ways to increase the rendering speed. Some examples include: taking advantage of the highly parallel nature of volume rendering algorithms; ignoring or removing elements that will not contribute to the final image; taking advantage of frame-to-frame coherence in animated sequences; and using multiple resolutions when updating moving images.

Real-time visualization of the voxel-based volumetric object models is essential for our surgical simulation system. We are currently using two volume rendering approaches to visualize object models: a modification of Lacroute's shear-warp factorization algorithm [19] and an approach that takes advantage of the 3D texture map memory available on our research platform [7]. In addition, a related project at MERL is investigating the use of special-purpose hardware for real-time volume rendering that combines logic and memory on a single chip and takes advantage of parallel algorithms and high on-chip bandwidths.

In addition to volume rendering, we are using real-time sectional imaging in which the displayed 2D section of the 3D MRI volume tracks the tip of the haptic device [26]. This visualization technique provides the surgeon with precise and detailed information from the MRI image data in addition to visual feedback of the 3D position of the device tip.

The approaches that we are currently using for volume rendering assume that the data is stored in a regular grid of evenly spaced volume elements. However, when object models are deformed by surgical procedures, the relative positions between volume elements can change. Unfortunately, once the volume elements are no longer evenly spaced in a regular grid, many of the techniques used to speed up volume rendering can not be used. In addition, since objects are interactively deformed in the simulator, the amount of preprocessing that can be done for sorting and rendering volumes is limited. In the current system, we use a simple nearest-neighbor splatting technique for rendering deforming tissues. However, this method produces low quality images. We have implemented and are currently investigating several methods for higher quality rendering of irregular volumes including raycasting [25], cell projection [27], [34], and splatting [33]. Some of these techniques may eventually be integrated into hardware approaches.

Haptic Interaction

The surgical simulation system 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). Most systems using force feedback model interactions with surface-based objects. There has been very little research using haptic interaction with volumetric objects. One notable exception is work by Rick Avila and Lisa Sobierajski at GE who use the PHANToM to explore and sculpt synthetically generated voxel-based objects [1].

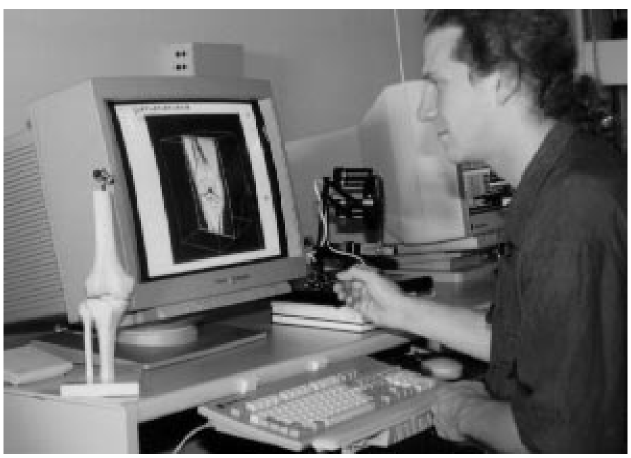

Figure 4. Prototype surgical simulation system.

Unlike the synthetically generated objects of Avila and Sobierajski, where surface normals and edge magnitudes can be calculated analytically, 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 object models resulting from segmentation 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.

We have investigated a number of approaches to data smoothing and surface normal calculation. These include: 1) pre-calculating surface gradients at grid locations and interpolating the gradient at the PHANToM tip from stored neighboring gradients; 2) smoothing the binary data and estimating the gradient at the probe tip using central differences of tri-linearly interpolated smoothed neighbor intensities; and 3) using a smoothing gradient operator filter centered at the location of the instrument tip to estimate the local gradient at the tip location. Each of these methods has tradeoffs in time vs. storage. In general, preprocessing is faster than calculations performed on-the-fly but require more storage. However, pre-processing can not be done if the object is being actively deformed.

Tissue Deformation

As discussed above, one of the challenges of volumetric models is dealing with the large number of elements that make up the volume. In addition to requiring large amounts of memory and special approaches for real-time rendering, the large number of elements make calculations of object deformation a challenging problem. One approach is to perform deformation calculations on a much coarser grid than the resolution of the volume elements. This approach has been taken in [14] for example. Our approach is to use a fast -- though rather simplified -- algorithm to propagate deformation through the volume [10]. In this approach, when the volume is manipulated, the object quickly stretches or contracts according to maximum and minimum distances defined by links between the volume elements. Because the motion constraints are similar to those of a set of linked elements in a chain, this algorithm has been dubbed 3D ChainMail.

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 1D and 2D systems of Figures (5) and (6). 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.

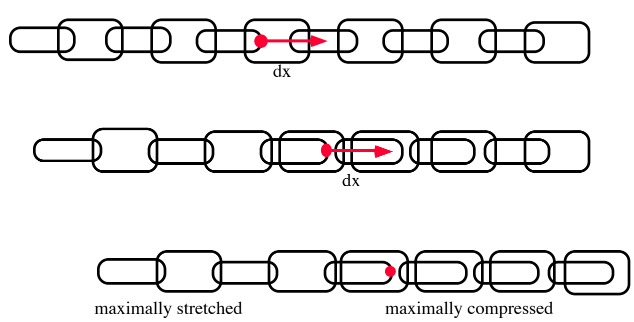

Figure 5. 1D ChainMail.

Tissue elasticity is modeled in a second process that adjusts the distance between local neighboring volume elements to minimize a local energy constraint. This process runs in the background and between applications of the ChainMail algorithm [10].

This approach is particularly fast for tissues with homogeneous (though possibly anisotropic) material properties. In homogeneous tissues, the 3D ChainMail algorithm allows the propagation of a single disturbance through the volume by considering each volume element only once and by comparing each element to only one of its neighbors. In tests on a single processor (R4400) SGI Onyx workstation, volumetric objects with as many as 8,000 elements have been deformed at interactive rates. These tests have been performed without particular effort to optimize or parallelize code. In recent work, we have added a layer of soft cartilage to rigid bone in a 2D model and demonstrated interactive deformation of the cartilage under the constraint that it remain attached to the bone.

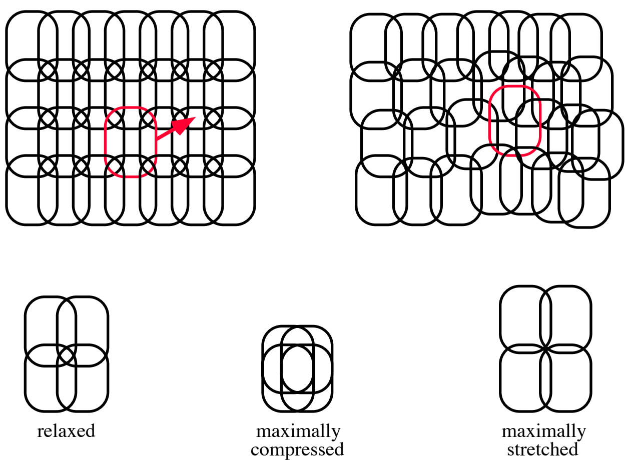

Figure 6. 2D ChainMail.

Status and Results

Progress in each of these technical areas has been presented above. Early work in each of these areas has been integrated into a prototype system for simulating arthroscopic knee surgery. In the prototype, object models consists of the bony structures of the joint, the articular cartilage, the menisci, and the cruciate ligaments. These structures were hand segmented from the low resolution MRI data. The integrated system currently allows probing of the bony structures with real-time visual and haptic feedback. Visual feedback is provided by both volume rendering and sectional imaging. The system will be extended in the short term with recent advances including: a higher resolution MRI image that has been acquired and is currently being segmented; soft-tissue modeling of the cartilage which has been implemented in 2D and is being extended to 3D; improvements in our volume rendering approach; and the addition of models of two surgical tools -- a surgical probe and a cartilage debridement tool -- to the haptic interface.

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. The interdisciplinary nature of the project has enabled team members to become involved in areas outside of their primary training. For example, engineers in the group have observed surgical procedures and surgeons have played a key role in image segmentation and system design. The next prototype system will be installed at the Surgical Planning Lab at Brigham and Women's Hospital to facilitate feedback from other surgeons and medical practitioners.

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 R10,000 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 the Ethernet although this control will be moved to the Challenge once a VME card becomes available from SensAble Technologies.

Discussion and Future Work

Figure (7) 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.

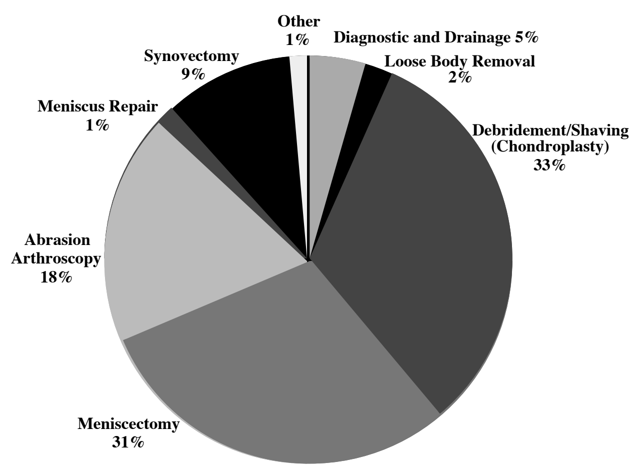

Figure 7. Distribution of 83,294 total Medicare Funded Arthroscopic Knee Procedures, 1992 [24].

We are pursuing many avenues of future work, some of which have been discussed above. In each technical area of this collaborative project we are continuing to work towards both short-term and long-term goals. For example, in volume rendering, we are pursuing hardware approaches to improve rendering speeds and we are investigating ways to render deformed objects at reasonable rates. In object modeling, we are looking at ways to compare results from the 3D ChainMail algorithm with more accurate finite element methods and to verify the efficacy of the algorithm in modeling human tissue. Volumetric object representation should allow us to model tissue cutting and tearing but this still needs to be demonstrated. In haptics, we intend to instrument our PHANToM with surgical tools and to add geometric models of these tools to the system. We will incorporate deformation into the haptic models and are investigating high resolution haptic textures that could be mapped onto the (relatively) low-resolution voxel-based models.

Acknowledgments

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

Bibliography

[1] R. Avila and L. Sobierajski, personal communication, 1996.

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

[3] Boston Dynamics Incorporated: http://www.bdi.com/simulations.html (1996).

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

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

[6] S. Cover, N. Ezquerra, J. O'Brian, R. Rowe, T. Gadacz, E. Palm, "Interactively deformable models for surgery simulation", IEEE Computer Graphics and Applications, Vol. 13, 6, pp. 68-75, 1993.

[7] R. Fraser, "Interactive volume rendering using advanced graphics architectures", whitepaper: http://www/sgi.com.

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

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

[10] S.F. Gibson, "3D ChainMail: a fast algorithm for deformingvolumetric objects", in progress, 1996.

[11] High Techsplanations: http://www.ht.com (1996).

[12] L. Hong, A. Kaufman, Y. Wei, A. Viswambharan, M. Wax, and Z. Liang, "3D virtual colonoscopy", Proceedings of Biomedical Visualization '95, eds. Murray Loew and Nahum Gershon, IEEE Computer Society Press, Las Alimitos, pp. 26-32, 1995.

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

[14] I. Hunter, T. Doukoglou, S. Lafontaine, and P. Charette, "A teleoperated microsurgical robot and associated virtual environment for eye surgery", Presence, Vol. 2, pp. 265-280, 1993.

[15] A. Kaufman, Volume Visualization, IEEE Computer Society Press, Los Alimitos, CA, 1991.

[16] A. Kaufman, "Volume Visualization", CRC Handbook of Computer Science and Engineering, 1996.

[17] U.G. Kuhnapfel, H.G. Krumm, C. Kuhn, M. Hubner, and B. Neisius, "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.

[18] W.E. Lorensen, F.A. Jolesz, and R. Kikinis, "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, 1995.

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

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

[21] B. Mirtich and J. Canny, "Impulse-based simulation of rigid bodies", proc. 1995 Workshp on Interactive 3D Graphics, pp. 181-188, April, 1995.

[22] S.D. Pieper, "CAPS: computer-aided plastic surgery", Ph.D. thesis, Media Lab, MIT, 1992.

[23] Alan Praemer, Department of Research and Scientific Affairs, American Academy of Orthopedic Surgeons, personal communication, 1996.

[24] Alan Praemer, from Health Care Finances Admin., Medicare data, 1992, personal communication, 1996.

[25] S. Ramamoorthy and J. Wilhelms, "An analysis of approaches to ray-tracing curvilinear grids", Tech. report UCSC-CRL-92-07, Computer Research Lab, UC Santa Cruz, 1992.

[26] J. Samosky, "SectionView -- A system for interactively specifying and visualizing sections through three-dimensional medical image data", MS Thesis, Dept. of Electrical Eng. and Computer Science, MIT, 1993.

[27] P. Shirley and A. Tuchman, "A polygonal approximation to direct scalar volume rendering", proc. San Diego Workshop on Volume Visualization, pp. 63-70, 1990.

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

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

[30] D.J. Vining, R.Y. Shifrin, E.K. Grishaw, et al, "Virtual colonoscopy", Radiology 193 (P):446, 1994.

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

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

[33] L. Westover, "Footprint evaluation for volume rendering", ACM SIGGRAPH '90, pp. 144-153, 1990.

[34] P.L. Williams, "Interactive splatting of non-rectilinear volumes", Proc. IEEE Visualization '92, pp. 37-44, 1992.

[35] R. Ziegler, W. Mueller, G. Fischer, and M. Goebel, "A virtual reality medical training system", Proceedings of CVRMed '95, ed. Nicholas Ayache, Springer-Verlag, Berlin, pp. 282-286, 1995.