

An overview of computer- integrated surgery at the IBM Thomas J. Watson Research Center

by R. H. Taylor
J. Funda
L. Joskowicz
A. D. Kalvin
S. H. Gomory
A. P. Guezic
L. M. G. Brown

This paper describes some past and current research activities at the IBM Thomas J. Watson Research Center. We begin with a brief overview of the emerging field of computer-integrated surgery, followed by a research strategy that enables a computer-oriented research laboratory such as ours to participate in this emerging field. We then present highlights of our past and current research in four key areas—orthopaedics, craniofacial surgery, minimally invasive surgery, and medical modeling—and elaborate on the relationship of this work to emerging topics in computer-integrated surgery.

Computer-integrated surgery: an emerging paradigm

Computer-integrated surgery (CIS) procedures will become commonplace in the coming decade, as present trends toward geometrically precise and minimally invasive surgery accelerate. These trends are driven by the desire for better clinical results, lowering overall costs through shorter hospital stays, shorter recovery times, and reduced need for repeated surgery. CIS procedures span the full spectrum of treatment, from diagnosis through preoperative planning and execution to postoperative assessment and followup (see **Figure 1**). Most of the key, enabling technologies are computer-based: 3D anatomy imaging, visualization, modeling, real-time sensing,

©Copyright 1996 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/96/\$5.00 © 1996 IBM

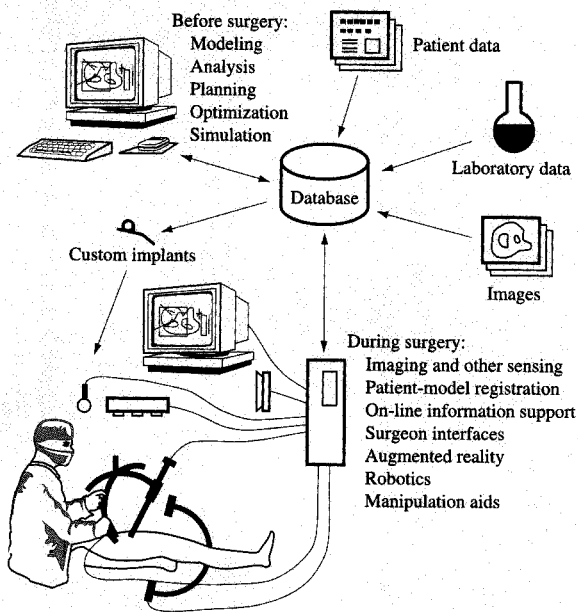


Figure 1

Computer-integrated surgery.

telerobotics, and systems integration. The emergence of very powerful, affordable computer workstations coupled with advances in image processing, graphics, simulation, and robotics has made surgical applications based on these technologies practical. Consequently, the pace of research and clinical activity is increasing sharply worldwide.

Advances in medical imaging technology [computed tomography (CT), magnetic resonance imaging (MRI), positron emission tomography (PET), etc.], together with advances in computer-based image processing and modeling capabilities, have given physicians an unprecedented ability to visualize anatomic structures in live patients and to use this information to improve diagnosis and treatment planning. Systems that combine this information with intraoperative¹ sensing, manipulation devices, graphics, and a variety of other human-machine interfaces are expected to have a comparable effect on surgical execution. A number of systems have been developed for various forms of neurosurgical procedures [1-6], orthopaedics [7-11], eye surgery [12-15], craniofacial surgery [16, 17], and otolaryngology [5, 18-20], among others. One common characteristic of these systems is that they rely on position sensing during the surgical procedure to *augment* the surgeon's ability to manipulate surgical instruments very precisely and to

¹ During surgery.

accurately execute a plan based on 3D medical images. By combining human judgment with machine precision, such systems permit a surgeon to perform critical surgical tasks *better* than an unaided surgeon, and enable the surgeon to do other tasks that could not be done at all.

In conjunction with these trends, the number of minimally invasive surgical procedures, such as laparoscopic and endoscopic surgeries, has grown explosively over the past few years. Two salient characteristics of these procedures are that the surgeon cannot directly manipulate the patient's anatomy with his or her fingers and cannot observe the results of the manipulation. Instead, the surgeon must rely on instruments that can be inserted through a small incision or through the working channel of an endoscope while observing the manipulation on a TV monitor. The surgeon must often rely on an assistant to point the camera while the surgeon performs the surgery. The awkwardness of this arrangement has led a number of researchers to develop robotic devices for assisting in endoscopic surgery. Typical efforts include improved mechanisms with flexible endoscopes (e.g., [21, 22]), specialized devices (e.g., [23]), voice control for existing mechanisms (e.g., [24]), simple camera-pointing systems [25-28], more sophisticated surgical-instrument-tracking systems (e.g., [29]), sensory augmentation systems (e.g., [30]), and full-blown "telepresence" systems [31, 32]. While the details of such systems vary widely, they share the common characteristic of combining computation with novel mechanisms or sensors to significantly extend human capabilities.

As the pace of innovation and clinical application has picked up, there has been a synergistic development of a new research community, reflecting the multidisciplinary nature of this work. This is reflected in books (e.g., [6, 33]), dedicated journals and special journal issues (e.g., [34, 35]), and conferences and workshops (e.g., [36-38]). As often happens with emerging fields, a number of different names for the discipline have been used by different members of the community, with each name emphasizing different overlapping aspects of the underlying technology and the set of applications. Those in common use include "image-guided surgery," "computer-assisted surgery," "medical robotics," "medical virtual reality," "information-intensive surgery," and "computer-integrated surgery." Lately, we have tended to prefer "computer-integrated surgery," since it emphasizes the coupling of presurgical planning with intraoperative execution.

Computer-integrated surgery at IBM Research

For a computer-oriented industrial research laboratory such as ours, active and close collaboration with clinicians is essential if we are to have a significant impact on this emerging field. Computer-integrated surgery (CIS)

applications inherently involve the integration of many different computer-related technologies and their use in an often novel (to computer scientists, at least) environment. Continual experimentation and feedback from clinicians is needed to identify and understand the true problems in a particular application and to determine what will really work in a clinical environment. Similarly, both the understanding of problems and progress toward solutions is often best accomplished in the context of working, integrated surgical systems and systems that provide the infrastructure to facilitate progress on individual research components.

Our strategy is summarized as follows:

- *Conduct applied research in partnership with clinicians* at leading medical centers and university medical schools. Select problems that are important clinically, appear to be solvable, require innovation, provide measurable benchmarks on the way toward solution, and permit us to develop capabilities that can be exploited in further applications. Where a particular capability such as image-based modeling of anatomy is fundamental to a large variety of applications, conduct appropriate research to address it, but use real applications to measure progress and generate test cases.
- *Develop integrated solutions* combining appropriate technologies such as image processing, modeling and analysis, real-time sensing, and manipulation aids to solve the problem. In doing so, try to understand *how* these technologies interact with each other and with the clinical environment in which they will be applied.
- *Aim initially at preclinical² demonstrations* to promote rapid technology development. Once a problem is understood, develop an initial *in vitro* demonstration of a possible solution on plastic models or similar simulators to gain immediate end-user feedback. Then, rapidly iterate generation of prototypes with end-user evaluations to develop more realistic solutions on (possibly) more realistic models and simulators.
- *Transfer solutions* into clinical use in partnership with other groups and enterprises with the requisite complementary skills to manufacture, market, and support products in this area.
- *Participate in the broader CIS research community* through scientific publications, conferences, research collaborations, and similar means.

In implementing our strategy, we pursue a variety of research activities both within our research group and in collaboration with others in IBM and elsewhere. We pursue a mix of “applications” research and “disciplinary” research motivated by identified clinical problems. Work

² Before use on human patients.




Figure 2

Total hip replacement (THR) surgery is an orthopaedic procedure that replaces the ball-and-socket joint between the patient's pelvis and femur. In cementless implants, a very close fit between the femoral implant and the prepared hole is required.

has been concentrated on three surgical disciplines: orthopaedics, craniofacial surgery, and minimal-access surgery. After close consultation with surgeons, we selected these disciplines because of their potential for automation and for the benefits they can provide. In each case, we believe that the use of quantitative information can make a significant difference in treatment outcomes and cost-effectiveness. The following sections describe each in detail.

Orthopaedic surgery

- *Primary total hip replacement surgery by robot*
Our initial involvement in computer-integrated surgery was a joint study between IBM and the University of California at Davis to develop a robotic system for preparation of the femoral canal in cementless total hip replacement surgery.

Total hip replacement (THR) surgery is an orthopaedic procedure that replaces the ball-and-socket joint between the patient's pelvis and femur.³ In THR, the socket is replaced with a cup-like implant inserted into the patient's pelvis. The femoral head is cut off, and an implant containing the ball is inserted into a prepared hole in the femoral canal, as shown in **Figure 2**. In many cases, the implanted components are secured to the bone with biocompatible cement; however, about half of the 300,000 THR procedures performed each year in the U.S. use

³ So-called “primary” total hip replacement (PTHR) procedures are performed on hips that have not previously had THR surgery, and “revision” total hip replacement (RTHR) procedures are performed to replace a failed hip replacement.




Figure 3

Cross-sectional views of holes for implants in human cadaver femurs: (a) manually broached; (b) robotically machined. The average shape error for robotic machining is less than 0.1 mm, compared to 1–4 mm for broaching. (Reprinted with permission from [7].)

cementless implants, which rely on a press fit or growth of bone into the implant to provide “fixation” of the implant relative to the femur. Cementless implants require a very close fit between the implant and bone, both to achieve proper fixation and to provide proper stress transfer from implant to bone. Also, for the implant to function optimally, it should be positioned properly. Unfortunately, the standard manual broaching process to prepare the hole for the femoral component of the implant is rather crude. One study [39] found that only about 20% of the surface area of implants inserted into manually broached femurs actually touched bone, with the average gaps ranging from 1 to 4 mm. Further, the positioning of the hole relative to the femur is rather uncontrolled, and driving the broach too far can split the femur.

The limitations of manual hole preparation led our surgeon collaborators, Dr. William Bargar and Dr. Howard Paul (then of the University of California

at Davis), to ask whether a robot might be able to significantly improve the precision of femoral canal preparation while reducing the chance of surgical complications associated with manual broaching. The resulting joint study involved a number of people at IBM Research, the IBM Palo Alto Science Center, and U.C. Davis.

The system that was developed, called ROBODOC™, demonstrated an order-of-magnitude improvement in the accuracy of femoral canal preparation for cementless PTHR surgery, and has been described extensively elsewhere (e.g., [7, 8]). Briefly, the system had to address the following requirements:

- The surgeon needed the ability to plan the surgery from preoperative data about the patient—i.e., to determine which implant design and size to use and the desired placement of the implant relative to the patient’s femur.

Our planning system, called ORTHODOC™ [8, 40],⁴ provides an interactive environment, permitting the surgeon to do this planning from CT data. The surgeon interactively selects three orthogonal “slices” through the volume described by the CT data, which are then displayed as three 2D images. The surgeon then selects an implant design from a database of implants and interactively manipulates its position and orientation relative to the femur by selecting one of the 2D slice images and specifying a translation or rotation of the implant in the plane of the coordinate system of that slice. The 2D outline (projection) of the implant corresponding to each slice is displayed on the three 2D slice images.

Although this approach is quite simple, it has been very successful and popular with surgeons. First, it presents data in a format (cross-sectional slices) that is familiar to surgeons. Second, manipulation of the implant’s position is intuitively natural, even if occasionally cumbersome. Finally, the surgeon’s judgment is relied on in the interpretation of CT density values, thus avoiding many problems associated with automated image interpretation.

- The system needed a means to register the presurgical plan with surgical reality; i.e., actual placement of the hole with respect to the femur has to correspond with the planned placement with respect to the CT images of the femur.

For simplicity and accuracy, we chose to use implanted markers. Three titanium pins (one in the greater trochanter and two in the femoral condyles) are implanted into the patient’s femur before the CT scan is made. These markers are automatically located relative

⁴ We use “ROBODOC” to refer to either the complete system or the operating-room system, and “ORTHODOC” to refer to the planning system.

to CT coordinates by algorithms in ORTHODOC developed specifically for this purpose [7, 40]. In the operating room, the pins are exposed manually, and the patient's femur is rigidly attached to the base of the robot (described below) by a specialized fixation device. A combination of manual guidance and tactile search with a probe inserted into the robot's cutter collet is used to locate the pins, and the transformation between the coordinates of the CT images and the robot cutter is computed.

This method has proved to be accurate and robust; however, the necessity for a second (though minor) surgical procedure to implant the pins and some associated degree of patient discomfort mean that there is clearly room for additional improvement in this part of the procedure.

- The system had to be able to cut the implant hole accurately and safely under overall supervision of the surgeon.

Our surgical robot is a modified IBM 7576 robot with an added pitch axis, six-degree-of-freedom force sensor, and high-speed cutting tool. During surgery, the force sensor is used for redundant safety checking, for tactile search for the locator pins, and for force-compliant motion-guiding by the surgeon. The robot's incremental precision, i.e., its ability to move short distances (e.g., 10–30 cm) from a known starting point or between two defined landmarks, is very good (typically, 0.05–0.1 mm). However, its uncalibrated ability to position the cutter at a precomputed point relative to its base is limited. To achieve the required cutting accuracy with an incrementally precise but not particularly “accurate” robot, we rely on a combination of calibration to estimate key kinematic parameters and a “constant orientation” cutting strategy to minimize the effect of unmodeled parameters [7]. This strategy has been quite successful. *In vitro* studies demonstrated shape-cutting accuracies better than 0.1 mm and placement accuracies of the order of 0.5 mm. **Figure 3** compares results of manual and robotic machining of the hole for the implant. The system also includes specialized subsystems for redundant safety and internal consistency checking [41] and for keeping the surgeon informed of the progress of the surgery, as well as a simple, sterilizable, hand-held terminal for interacting with the system in the operating room. The overall system architecture is shown in **Figure 4**.

The prototype ROBODOC system (**Figure 5**) was transferred to UC Davis, where Dr. Paul used it to conduct a successful clinical trial on 26 dogs needing hip surgery. Dr. Paul subsequently founded Integrated Surgical Systems (ISS), which reengineered the system to create a version suitable for use on human patients [8],

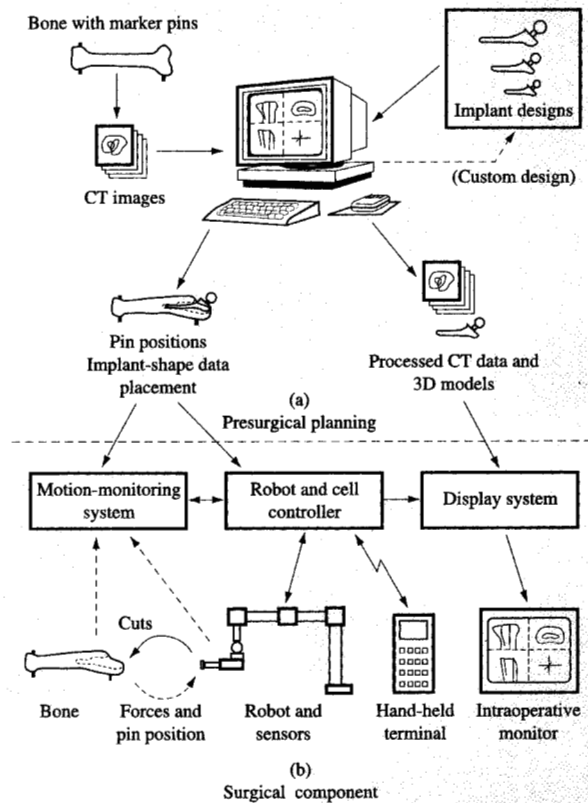


Figure 4

Hip replacement surgery system consisting of (a) a presurgical planning component and (b) a surgical component. (Adapted with permission from [7].)

and conducted a ten-patient clinical trial in late 1992. ISS is currently conducting a clinical trial on 300 patients at multiple sites in the U.S. It is also conducting clinical trials in Europe.

This work illustrates several aspects of our overall research strategy. First, the application concept and requirements came from the surgeons who would use the system. The particular problem was chosen because it was important and challenging, but still feasible. In picking a technical approach, we tried to keep things as simple as possible while still meeting the design objectives. Getting a robot to work accurately and robustly in the operating room was challenging enough without our trying simultaneously to invent new image-segmentation methods, automated surgical-plan-optimization techniques, or other refinements. At the same time, in designing the system, we paid considerable attention to where such refinements might eventually go, and tried to make the




Figure 5

Surgical robot for hip replacement surgery. This photograph shows the ROBODOC system for primary total hip replacement surgery, in a human clinical trial. All of the robot (top) is covered with sterile plastic draping. The robot holds the cutting tool (to which the green hose is attached). (Reprinted with permission from [7].)

system architecture as flexible as possible, yet effectively solving the particular problem.

We found that preclinical demonstrations and rapid iteration were crucial to success. Hands-on experience in the laboratory was essential for the surgeons to become comfortable enough with the system to use it clinically. Perhaps more significantly, such experience was indispensable in promoting mutual understanding between the surgeons and engineers in the research team about what the stated system requirements really meant and also about what the system could and could not be made to do. Similarly, our relationship with ISS, which includes many members of the team that originally developed ROBODOC and which has very substantially refined the original system to make it suitable for human use, has been very useful for our research group at the IBM Thomas J. Watson Research Center, especially in

promoting better understanding of how real-world constraints help define and sharpen research problems. Finally, by demonstrating the practicality of using robotic systems to perform precise orthopaedic surgery, ROBODOC has attracted considerable interest in both the medical and engineering communities and has helped promote further research activities in a number of related areas.

● *Revision total hip replacement surgery*

In cooperation with ISS, we are extending ROBODOC for use in revision total hip replacement surgery (RTHR). In RTHR, a failing orthopaedic hip implant is replaced with a new one. The old implant, usually cemented into the femur, is replaced with a cementless implant. The surgical procedure thus comprises removing the old implant, removing the cement, and fitting a new implant into an enlarged hole broached in the femur. It is both difficult and expensive. RTHR is becoming increasingly common as the population of old implants ages. For example, in 1992, over 28000 RTHR procedures were done in the United States. The number of cases was increasing by approximately 10% a year, and the average cost for each case was approximately \$24000 [42, 43].

RTHR is much more challenging than PTHR. It takes much longer, complications are commonplace, and hospital stays are longer. Intraoperative femoral fracture occurs in 18% of the cases. Removal of cement in the distal femur, especially the area near the remote tip of the implant, is particularly difficult and necessitates the surgeon drilling through the side of the femur in about 10% of the cases. Available instrumentation for distal cement removal is often expensive, works well in only a fraction of cases, and can cause severe difficulties when it does not work. Since the patient has less good bone left, imprecision in preparing the femoral canal for the new implant is more likely to cause serious problems than with primary cases. Our goal is to provide the surgeon with the same advantages in accurate implant placement and "fit" as are provided by ROBODOC PTHR (potentially shorter hospital stays and faster healing), while providing very significant cost savings and patient benefits from substantially reducing the incidence of surgical complications and from reducing surgical times and blood loss.

This work is in its preliminary stages, and we are still planning our detailed technical approach; however, certain aspects of RTHR make it particularly interesting from a research perspective:

- Medical CT images often have significant reconstruction errors, or "artifacts," in the vicinity of metal implants. These artifacts degrade the quality of the image data available for presurgical planning. Consequently, we are

investigating techniques for mitigating the effect of such artifacts, for combining information from CT images with information from standard X-ray images, and for planning in the presence of degraded data.

- RTHR procedures cannot always be preplanned to the same extent as PTHR procedures, both because of the limitations of preoperative CT and because the surgeon cannot always predict exactly what will be required to remove the old implant from the bone. Consequently, we are required to develop robust methods to verify that the assumptions in the presurgical plan remain valid intraoperatively and to give the surgeon the ability to modify the plan, if need be, or to switch to another predefined plan if it is more appropriate. We are currently investigating a number of approaches based on intraoperative X-ray imaging, both for registering, or aligning, preoperative plans and models with the actual robot and patient and as a possible basis for intraoperative planning. One additional bonus of this work is that it will provide a "pinless" registration method that may also be suitable for PTHR and other orthopaedic applications.

This work illustrates once more the problem-driven aspects of our strategy. As with PTHR, our goal is to solve a particular problem in a way that produces solution components that subsequently can be applied to other problems, and to do this through rapid iteration in partnership with the surgeons who will use the system and with the company (ISS) that will develop and deploy the system.

• *Implant insertability analysis*

The ability of the robot to precisely machine complex shapes to very tight tolerances can remove important constraints in the design of customized orthopaedic implants—constraints associated with the limitations of the standard instrumentation used for manual preparation of the femoral canal. This ability potentially permits more optimal placement of off-the-shelf implants while sacrificing minimal good bone. However, a tightly fitting implant must still be insertable to its final working position inside the canal without interference. Determining the insertability of an implant into a canal is thus an essential step in preoperative planning.

Insertability analysis of an implant into a canal is an instance of the so-called "peg-in-hole" problem, a classical and important problem in robotics. The goal is to compute an interference-free insertion trajectory for the peg into the hole from an initial to a final position. Solutions to the peg-in-hole problem have many applications in manufacturing, computer-aided design, and many other areas. In addition to medical implants, some specific examples are the design and analysis of molds so that




Figure 6

Locality principle for implant insertability analysis. Motions of the body point \mathbf{b}_j are constrained by the closest hole surface h_i in the neighborhood Δ_{ij} . The shaded area corresponds to the local configuration constraint $h_i[\mathbf{F}(\mathbf{p}, \theta) \cdot \mathbf{b}_j] \leq 0$. (Adapted with permission from [46].)

their contents can be removed and the design and analysis for assembly of tightly fitting parts in machines.

We have developed a computer program, called Extract, for computing and visualizing the interference-free insertion path of an implant into a hole from computer-aided design models of their shapes. The program formulates the problem as a peg-in-hole insertion problem for complex, tightly fitting, three-dimensional bodies requiring small, coupled six-degree-of-freedom motions in a user-specified direction. The program either finds a successful insertion path or reports the "stuck" configuration, in which case it identifies the surfaces causing the interference, facilitating the redesign of the implant and the hole shapes. Extract allows the user to view the insertion of the implant into the canal and the stuck configurations from different perspectives. The program is reasonably efficient. In about 30 minutes on a RISC System/6000[®] Model 530 workstation, Extract can compute an interference-free insertion trajectory for a tightly fitting implant and a hole shape described with 10000 facets to an accuracy of 0.01 inch. The program has been successfully tested on 30 real cases provided by Biomet Inc., a manufacturer of orthopaedic implants and other medical devices [44–46].

Extract computes an insertion path by formulating local, linearized configuration-space constraints (described below) for small motions and solving a series of linear programming problems. The calculated path is specified by a sequence of configurations (positions and orientations) of the implant during its motion. The implant configurations are uniquely defined by three rotations θ and three translations \mathbf{p} of the implant's coordinate frame with respect to a fixed coordinate frame. The small




Figure 7

Insertability analysis. Example of an insertion sequence of an implant stem (dark body) tightly fit into a canal (translucid body), from the initial inserted configuration (left) to the final working configuration (right). (Reprinted with permission from [46].)

motions between successive configurations in the path are interference-free (to within a prespecified resolution) and define the insertion motion.

To guarantee that the insertion path is interference-free, the program formulates configuration-space constraints derived from the implant and hole shapes. These constraints specify the implant configurations for which the implant surface does not penetrate the hole walls. All configurations in the insertion path must satisfy the constraints. Configuration-space constraints are reduced to simpler local configuration-space constraints by observing that, because of the tight fit between the implant and the hole, the motion of any point on the implant surface is constrained by only a small portion of the hole surface in its immediate neighborhood (see **Figure 6**). Local configuration-space constraints are formulated for each pair of implant-surface point and hole-surface element (i.e., small area of the surface).

To compute a single motion step, we formulate an optimization problem with the local configuration-space constraints in the neighborhood of the current implant configuration. The objective function $\tau(\epsilon, \alpha)$ represents the distance that an incremental translation ϵ and rotation α of the implant will advance it in the insertion direction: along the axis of the hole. The solution of the optimization problem yields the largest interference-free insertion step in a small volume surrounding the implant configuration. The problem for the k th step is thus to

$$\text{maximize } \tau(\epsilon_k, \alpha_k),$$

where (ϵ_k, α_k) represents the incremental motion taken in that step, subject to

$$\mathbf{h}_i[\mathbf{F}(\mathbf{p}_k + \epsilon_k, \theta_k + \alpha_k) \cdot \mathbf{b}_j] \leq 0$$

and

$$\mathbf{F}(\mathbf{p}_k + \epsilon_k, \theta_k + \alpha_k) \cdot \mathbf{b}_j \in \Delta_{ij},$$

for each point \mathbf{b}_j on the implant surface and each element i on the hole wall, where $\mathbf{F}(\mathbf{p}_k, \theta_k)$ represents the implant coordinate system (position and orientation) at the start of the k th step in the insertion path, and each Δ_{ij} is a small volume. There is one constraint of the form $\mathbf{h}_i[\mathbf{F}(\mathbf{p}_k + \epsilon_k, \theta_k + \alpha_k) \cdot \mathbf{b}_j] \leq 0$ for each point \mathbf{b}_j on the implant surface, where $\mathbf{h}_i(\cdot \cdot \cdot) \leq 0$ corresponds to an element of the hole wall that the implant point must not penetrate.

This constraint is valid only as long as the implant point is within a small volume Δ_{ij} ; i.e., $\mathbf{F}(\mathbf{p}_k + \epsilon_k, \theta_k + \alpha_k) \cdot \mathbf{b}_j \in \Delta_{ij}$. In general, \mathbf{h}_i , $\tau(\epsilon_k, \alpha_k)$, and $\mathbf{F}(\mathbf{p}, \theta)$ are nonlinear functions; however, the fact that Δ_{ij} and the incremental displacements ϵ_k and α_k are small permits us to solve instead a linear approximation to this problem:

$$\text{maximize } \tau(\epsilon_k, \alpha_k) = \mathbf{v} \cdot (\epsilon_k, \alpha_k)$$

subject to

$$\mathbf{h}_j \cdot (\alpha_j \times \mathbf{v}_{jk} + \epsilon_k) - \delta_i \leq 0$$

and

$$\mathbf{C}_i \cdot (\alpha_j \times \mathbf{v}_{jk} + \epsilon_k) \leq \mathbf{d}_i,$$

where $\mathbf{v}_{jk} = \mathbf{F}(\mathbf{p}_k, \theta_k) \cdot \mathbf{b}_j$ represents the position of implant point j at the start of iteration k , \mathbf{v} represents the insertion direction, (\mathbf{h}_i, δ_i) corresponds to the constraint $\mathbf{h}_i(\cdot \cdot \cdot) \leq 0$, and $(\mathbf{C}_i, \mathbf{d}_i)$ define a convex polyhedron corresponding to Δ_{ij} .

Extract actually determines a feasible insertion trajectory by first solving the reverse problem; i.e., it starts by placing the implant in the final (inserted) configuration and then computes repeated short-distance steps required to extract it from the hole. At each step, the pairing of implant-surface sample points to hole neighborhoods is updated, and a new (linearized) optimization problem is generated. Once the implant is free, the path is reversed. **Figure 7** shows a typical insertion sequence. Although the trajectory seems simple, feasible insertion trajectories generally require rather complex six-degree-of-freedom motions to permit the implant to conform to the shape of the hole. Details of the algorithm, including a number of important refinements, and a more complete description of the problem formulation can be found elsewhere [44–46].

Craniofacial surgery

Craniofacial osteotomies are often used to correct severe facial abnormalities associated with developmental defects, trauma, cancer treatment, or similar causes. In these procedures, the surgeon cuts the patient's facial bones

into fragments, rearranges them to give a more ordinary appearance, and reattaches them with screws and plates. Such procedures can produce striking results and have the advantage that the bone can grow with the patient; however, these long and tedious operations require extensive preoperative planning and high intraoperative accuracy. Although the number of major craniofacial operations is relatively small (between 3000 and 30000 per year, depending on what is counted), these procedures make an immense difference to the lives of patients. Further, the inherently three-dimensional nature of craniofacial surgery planning has made it an important application in three-dimensional medical imaging, modeling, and visualization [47].

In collaboration with Dr. Court Cutting from the New York University (NYU) Medical Center, we have developed a system for planning and assisting in craniofacial osteotomies. The system, whose architecture is shown in **Figure 8**, consists of a CT-based modeling, analysis, and planning component and an intraoperative tracking system to assist the surgeon in carrying out the procedure accurately. As with ROBODOC PTHR, this system has been extensively described elsewhere (e.g., [16, 47-50]). A brief summary follows.

The planning system, developed by Dr. Cutting's group at NYU, builds 3D polyhedral representations of individual skulls from CT data. These models are analyzed to locate ridge curves, point landmarks, and surface patches corresponding to standard skull anatomy [47]. This process is performed for many normal skulls, and the results are combined to produce a statistical database of normal anatomy. To plan an individual surgical case, the system processes CT images of the patient to produce a 3D model and (with minimal human interaction) locates ridge curves, point landmarks, and surface patches corresponding to an anatomical atlas. Given a possible set of osteotomy fragments, the system computes the rigid-body motion of each fragment that minimizes the difference between corresponding features on the patient and the normative model. For point landmarks, the system seeks to minimize the sum-of-squares distance between corresponding points; for ridge curves between landmarks, it seeks to minimize the sum-of-squares areas; and for surface patches, it seeks to minimize the sum-of-squares volumes. The NYU group has investigated various ways of combining these otherwise incompatible error measures to produce a single optimized result [47].

The planning system has several methods of defining where the osteotomy cuts should be made to produce a proposed fragment set. The most straightforward—but also the most tedious—is simple interactive definition by the surgeon, by means of standard computer-graphics tools. A more convenient alternative relies on the fact that there are fairly standard places for making the cuts. The

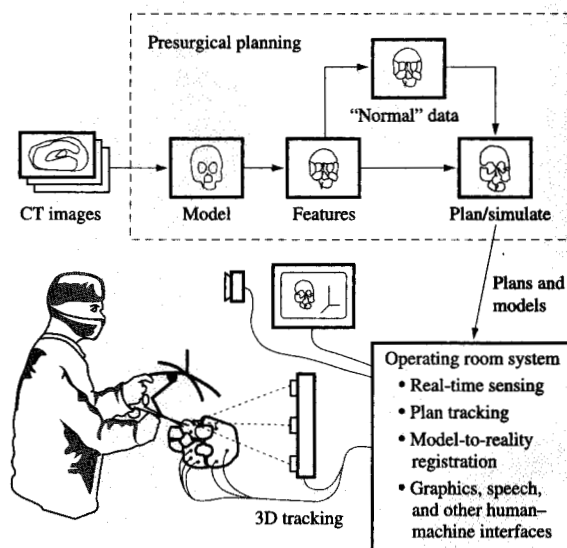


Figure 8

Craniofacial surgery system. This system is joint work between the NYU Medical Center and IBM. (Adapted with permission from [47].)

system is able to use this information to generate a number of different proposed fragment sets, based on standard osteotomy strategies, produce the corresponding optimized relocation plans, and present the results to the surgeon, who can select the strategy that offers the best trade-off between optimized result and surgical complexity. The surgeon can then use this result as the starting point for more detailed interactive planning to adjust where the cuts should be placed or to override the recommended bone-fragment motions. A typical osteotomy plan for a patient with Apert's Syndrome is shown in **Figure 9**.

The operating room system, developed primarily by IBM Research, assists the surgeon in accurately realigning the bone fragments in order to carry out the surgical plan. In present practice, most craniofacial surgery is performed freehand. The surgeon performs the osteotomies to free the bone fragments, realigns the fragments, and reattaches them with screws and plates, based on a visual appreciation of their desired relationship. Although making the cuts is not a problem for a good surgeon, achieving the desired alignment between fragments is very difficult to do without assistance. Even skilled surgeons find it extremely difficult to achieve an alignment better than about 5 mm with such an approach [47]. An alternative method, developed at NYU [51], uses multiple

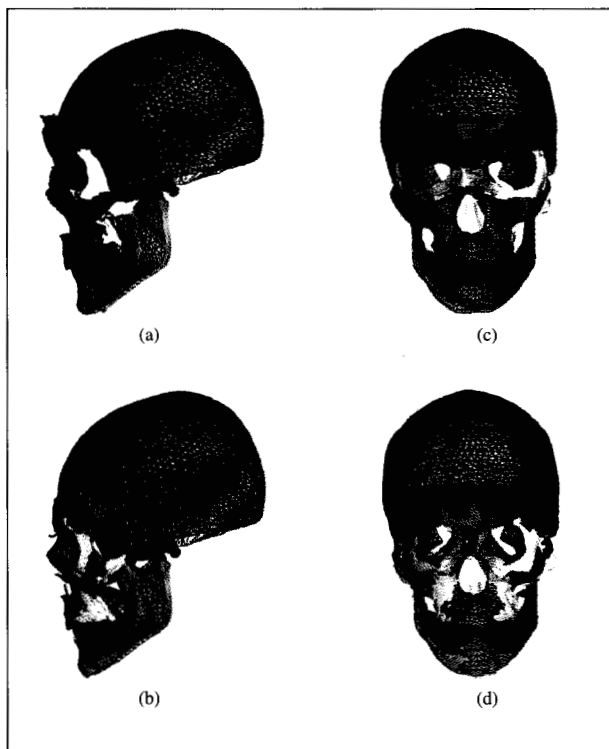


Figure 9

Craniofacial surgery presurgical planning. This component, developed by Dr. Cutting's group at NYU, assists the surgeon in planning optimized osteotomy strategies from patient CT scans. In this display, the wire mesh represents a normal skull, and the solid pieces are computer-optimized bone fragments for a patient with Apert's Syndrome. Diagrams (a) and (c) show the bone fragments in their normal positions, while (b) and (d) show the bone fragments in their optimized positions. (Adapted with permission from [47].)

custom fixtures wired to the patient's teeth to achieve somewhat better accuracy (estimated by Dr. Cutting at about 2 mm); however, this technique is cumbersome and time-consuming, and it sometimes requires compromises in the surgical plan.

The key goals for the IBM intraoperative system were as follows:

- The system should accurately track and report to the surgeon the actual positions of the bone fragments relative to one another and to the skull base.

Our approach relies on optical tracking of light-emitting-diode (LED) markers affixed to the patient's skull, by means of a commercial 3D digitizer such as the Optotrak™ manufactured by Northern Digital, Inc. or the Pixsys™ system manufactured by Pixsys, Inc. Essentially, we affix three or more markers to each

potential bone fragment before any osteotomies are performed, and observe the fragment positions while simultaneously locating known anatomical landmarks on the skull with a pointing device upon which a number of LEDs have been mounted. The relative positions of the bone fragments after the osteotomies can then be determined from the observed marker positions.

- The system should assist the surgeon in manipulating the fragments into the desired configuration and then hold them in place while the surgeon attaches them to one another.

Although, in principle, a robot holding one of the fragments could perform such alignments very quickly and accurately, and could even continuously adjust its position to maintain a desired alignment between fragments, we have chosen a simpler solution based on a passive, remote-center-of-motion manipulation aid, shown in Figure 10. Each natural motion of the mechanism affects only one rotational or translational degree of freedom of a bone fragment or other object rigidly held at the remote motion center (sometimes referred to as the "fulcrum point"). This permits the surgeon to make an adjustment affecting only one or two degrees of freedom of a bone fragment at a time without disturbing the fragments that have already been properly aligned.

The operating room system (see Figure 10), using a Northern Digital Optotrak digitizer with an accuracy of 0.1 mm, has been demonstrated on plastic skull models. We have experimented with a variety of methods, including graphical displays and auditory pitch cues, for informing the surgeon of the relative fragment alignment. Six-degree-of-freedom fragment alignments to an accuracy of the order of 1 mm are achieved with this system in only one to two minutes, seldom requiring more than one or two readjustments of each motion axis.

A version of the planning system is currently in clinical use at NYU Medical Center. Clinical versions of the operating room system are being developed by NYU. We expect that the technologies developed for this procedure will have broad application in orthopaedics, neurosurgery, biopsies, and many other domains, where significant research opportunities exist both in planning (shape analysis, optimization, model-to-patient registration, etc.) and in operating room technology.

Our experience with this application illustrates the importance of working directly with a user in defining system requirements and of picking appropriate technology in addressing the problems. In this case, we were fortunate that our surgeon collaborator (Dr. Cutting) is himself a very good computer scientist who has long been a leader in computer-based surgical planning research. This certainly shortened the time required for

mutual education in the earlier stages of the project and in subsequent work. As more and more clinicians gain computer expertise, we expect that other groups will also find this to be the case. However, the crucial lesson remains. It is at least as important to have early and frequent input from a person who understands the problem to be solved as it is from the researchers who have ideas about the techniques that might be available to solve it. Building this partnership is an important key to success.

Minimal-access surgery

Minimal-access surgical procedures are carried out by inserting instruments through small incisions in the patient's body and using either endoscopic cameras or extracorporeal imaging devices such as ultrasound or fluoroscopy for guidance and feedback. Such "keyhole" surgery has seen remarkable growth in recent years. For example, it is generally estimated that from 60 to 80% of all abdominal surgeries will be performed laparoscopically by the year 2000. This rapid explosion has, to a large extent, been driven by patient demand. Surgery done through small incisions is often much less traumatic than the same surgery done through large incisions. Additionally, there is less postoperative pain, and the patient is able to resume normal activities much sooner than would otherwise be possible after open surgery. Unfortunately, these procedures also are more difficult for the surgeon to master, since direct eye-hand coordination is lost.

We have an active joint study with the Johns Hopkins University Medical School on geometrically precise, image-guided, minimally invasive surgery. The initial focus of this activity has been the development of a robotic system, called LARS, to assist in laparoscopic surgery [29, 52-57]. In laparoscopic abdominal surgery, a camera and a surgical instrument introduced through incisions made in the patient's abdomen are used to locate and manipulate the patient's anatomy. Key tasks in such surgery include camera pointing, tissue retraction, measuring, and positioning a surgical instrument, all under the guidance of images. Even the relatively straightforward task of camera pointing is not always easy for human operators, especially when angled or flexible endoscopes are used. A computer-controlled device can potentially point the camera *better* than the average assistant, provide the surgeon with direct control over the viewing process, and reduce the operating room staff. More importantly, such devices can provide *new capabilities*, such as the ability to locate lesions from real-time images (video, ultrasound, etc.) and then deliver optimized therapy patterns planned from preoperative CT or MRI.

Key attributes of the LARS system (Figure 11) include a novel surgical robot, a highly modular and potentially




Figure 10

Craniofacial surgery intraoperative system. The passive manipulation aid has a remote center of motion, or fulcrum point, that permits the surgeon to align one bone fragment relative to another, one degree of freedom at a time. A clinical version is being developed at NYU. (Reprinted with permission from [47].)

low-cost control architecture, real-time processing of video images for measurement and guidance, and a variety of novel human-machine interfaces. To simplify control and to improve safety and accuracy, the robot has a natural "fulcrum point," or remote motion center, similar to that of the craniofacial system, at which rotational and translational motions are decoupled. For laparoscopic or other percutaneous applications, this fulcrum point is positioned at the point of entry to the patient's body. In other cases (such as radiotherapy or open surgery), it is positioned at the targeted anatomy. The controller is PC-based and incorporates extensive hardware and software safety checking. The present embodiment is designed for rapid prototyping. We have shown that we can readily add additional actuators, reconfigure the manipulator, or even substitute an entirely different mechanical structure, with only *minimal* work. A relatively inexpensive controller



Figure 11

Laparoscopic surgery assistance system. This system incorporates simple image processing, human-machine interfaces, and a manipulator. (Reprinted with permission from [29].)

optimized for a particular mechanism could be developed readily, with only minor software changes.

One unique aspect of the LARS system is its user interface (**Figure 12**). In its present embodiment, the surgeon uses a small joystick device (functionally equivalent to a three-button mouse) clipped onto a surgical instrument held by the surgeon as the primary means of commanding the system. This device is used to designate anatomical features in the field of view of the camera, select operating modes, reposition instruments or cameras held by the robot, remember the positions of and return to key anatomical features, perform measurements, and perform other functions. All motion commands are defined relative to a coordinate system defined by the field of view of the endoscopic camera, and the system is

readily able to accommodate angled-view endoscopes.⁵ Another unique aspect is the system's ability to process video images from the laparoscopic camera and use the results to assist in the control of the robot. For example, the system can use stereo triangulation to accurately locate a designated anatomical feature and then automatically move an instrument held by the robot to the designated feature.

By design, the LARS robot is slow. Since inertia effects and similar dynamic properties of the manipulator are negligible, low-level servocontrol can be extremely simple. Further, the decoupled structure makes standard kinematics calculations relatively straightforward. Nevertheless, design of a motion-control subsystem suitable for supporting the user interface posed some unusual problems, and our solution to them may be viewed as an example of our research strategy of using particular application problems to pose more general research problems and then using working systems as benchmarks to test the solution.

In performing a particular motion task, the controller must consider a number of requirements:

- The LARS robot is kinematically redundant—i.e., it actually has seven actuated (driven by motors) degrees of freedom: translation of the base in three dimensions, instrument yaw, pitch, and roll, and insertion translation (along the axis of instrument insertion). End-of-arm tooling can add more degrees of freedom (e.g., rotation of the eyepiece about the camera). The system must know how to exploit this redundancy to best accomplish each task.
- The requirement that an instrument must pass through a narrow entry portal somewhat constrains the robot's available motion. Thus, when possible motions of the end of the instrument inside the patient's body are considered, the system may be kinematically deficient, even though the robot is redundant. Depending on the particular task being performed, the controller must know how to trade off the available manipulability in order to best achieve a desired result.
- Generally, there are a number of limitations on allowed motion of different parts of the manipulator, including intrinsic constraints such as joint motion limits, and extrinsic constraints such as forbidden regions for an instrument held by the robot (e.g., "Do not poke the scope into the liver").

In principle, for each situation, we could develop specialized code that would also be dependent on specialized knowledge of the manipulator structure; however, such a solution would be complicated and

⁵ Endoscopes whose axes of view are at angles to their mechanical axes.

difficult to maintain. Instead, we require a systematic approach that permits us to combine a number of possibly competing requirements and that also insulates higher-level code from detailed information about the manipulator. Our approach is described more fully in [52]. Essentially, we formulate the problem of determining the manipulator-joint positions $\mathbf{q}(t)$ that will achieve a desired manipulator position and orientation as a quadratic optimization problem:

$$\min \|\mathbf{A}(t) \cdot \mathbf{q}(t) - \mathbf{b}(t)\|$$

subject to constraints

$$\mathbf{C}(t) \cdot \mathbf{q}(t) \leq \mathbf{d}(t),$$

where $\mathbf{A}(t)$ and $\mathbf{b}(t)$ are derived from the relative weights of different goals to be achieved, propagated through the kinematic equations of the manipulator. $\mathbf{C}(t)$ and $\mathbf{d}(t)$ are used to express constraints that must be obeyed, also propagated through the kinematic equations of the manipulator. We solve this problem in a discretized form,

$$\min \|\mathbf{A}(t_i) \cdot \mathbf{q}(t_i) - \mathbf{b}(t_i)\|,$$

subject to

$$\mathbf{C}(t_i) \cdot \mathbf{q}(t_i) \leq \mathbf{d}(t_i)$$

for multiple time steps t_i . This scheme has proved to be both flexible and efficient. We achieve typical solution times of 50–65 ms on a slow (33-MHz '486) PS/2[®]. It has paid off substantially in simplicity of debugging and has been adapted to several very different manipulator designs (e.g., [55, 57]).

The prototype LARS system has been used *in vivo* by our clinical collaborators at Johns Hopkins for both cholecystectomies (gall bladder removals) and nephrectomies (kidney removals). The reaction of the surgeons has been enthusiastic, and we are considering appropriate further steps to exploit the system's capabilities.

Our present main focus is on applications that exploit the ability of the LARS system to accurately and quickly align an instrument on the basis of information obtained from images, as well as its flexibility and programmability. The sequence in **Figure 13** illustrates the ability of the system, using information obtained by processing video images from the laparoscopic camera, to place a surgical instrument on a designated target. Figure 13(a) shows the experimental setup, consisting of the surgical robot holding a Storz therapeutic laparoscope, a rubber simulation of patient anatomy, and a small target to be grasped by a surgical instrument inserted into the working channel of the laparoscope. In this picture, the robot is draped as it would be in surgery. The figure illustrates force-compliant manual guidance of the robot. The robot enters this mode whenever the surgeon depresses two




Figure 12

In vivo video display with superimposed control menus. Shows typical video display (in this case, of a plastic simulated stomach) seen by the surgeon when using the system. The menus on the left-hand side of the screen correspond to control modes or robot functions. The "snapshot" images on the right-hand side correspond to previously saved robot views. Typically, the surgeon selects desired functions or robot positions by using the instrument-mounted joystick to position a cursor over the desired menu item and then "clicking" a button. (Reprinted with permission from [29].)

buttons on opposite sides of the carrier of the surgical instrument. Figure 13(b) shows the display monitor after the surgeon has designated the target, using the instrument-mounted joystick to place cursor crosshairs on the image of the target. Figure 13(c) shows the scene just after the computer has located the target by multiresolution correlation. This view shows the correlation window tree [58]. Normally, this display is used for debugging and would be suppressed in production use. Figure 13(d) shows insertion of the instrument into the working channel. Figure 13(e) shows the scene during the pickup operation. The target appears to be off-center, but it is lined up with the working channel of the scope.

We are extending the image-based navigation capabilities to include the ability to acquire fluoroscopic and other intraoperative images, register them with preoperative plans, and accurately place an instrument or therapy-delivery device (such as an injector for a radioactive pellet). Applications that we are exploring include percutaneous therapy of soft-tissue lesions, minimally invasive biopsy and treatment of bone tumors, and precise percutaneous spinal surgery.

Another possible area of work is the integration of the existing software and system capabilities with either a less expensive LARS manipulator or one of several low-cost

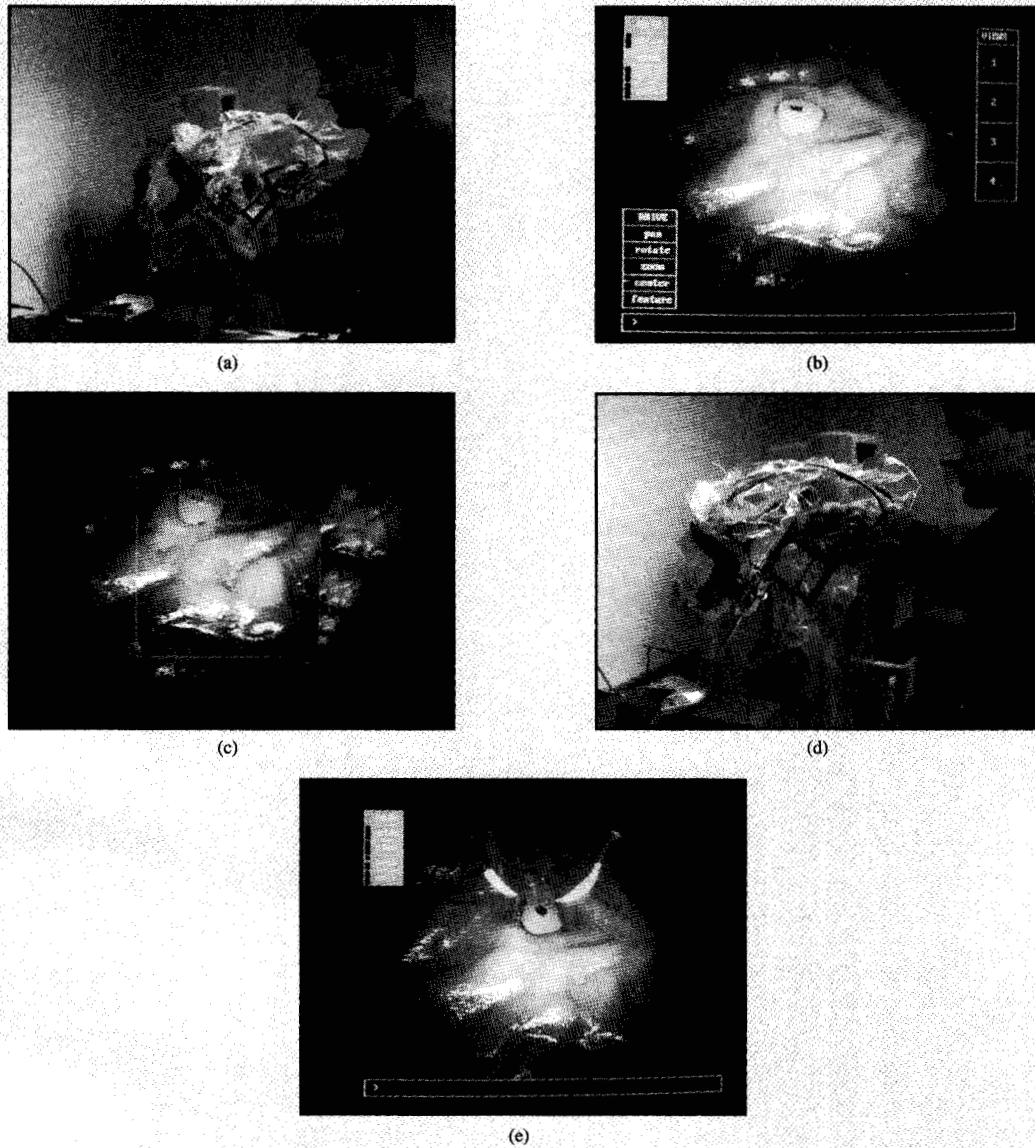


Figure 13

Simulation of an image-guided grasping operation using the LARS robot. (Reprinted with permission from [29].)

176 vendor manipulators specifically designed for camera positioning. Advantages of this approach include significant functional improvement of existing camera-pointing systems at low additional cost, and support for more sophisticated laparoscopic surgical applications requiring coordinated control of several manipulators. In one feasibility test [55], we constructed a simple robot with a passive “wrist” similar to that used in many camera-pointing systems (e.g., [28, 59]). We are

also considering possible application of LARS in a number of “remote surgery” applications.

Image processing and modeling

The extraction of anatomical models from 3D images (e.g., CT and MRI scans), their quantitative analysis, and their registration to the anatomy or to other images, are essential components of many CIS applications. Such

models are necessary for preoperative diagnosis and planning, for intraoperative guidance and execution, and for postoperative follow-up. Our goal is to develop novel, general-purpose algorithms and techniques with broad applicability, both within the context of specific applications, such as orthopaedics and craniofacial surgery, and in nonmedical applications with similar needs, such as anthropology [60, 61].

One typical project of this nature is the development of automatic methods for simplifying complex 3D anatomical models. We have developed an automatic, adaptive, hierarchical simplification algorithm for polyhedral models, called "Superfaces" [62, 63]. The key attributes of this simplification algorithm are that it preserves the topology of the original model, guarantees a provable approximation error bound, does not require any *a priori* knowledge of anatomy, and permits local substitution of higher-resolution models in areas of particular interest. Although a detailed description is beyond the scope of this survey, the basic steps are as follows:

- *Step 1 Superface creation*

A "greedy" face-merging procedure is used to partition the original model into quasi-planar regions called superfaces. **Figure 14(a)** shows a typical output for a skull model. Each colored patch corresponds to a superface. **Figure 14(b)** is a closeup of the area in the box in **Figure 14(a)**. It shows the individual polygonal faces of the original model. The face merging is controlled so that a number of key properties are guaranteed. The most important is that every vertex of every original face subsumed into a superface is guaranteed to lie within a specified distance of a plane associated with the superface.


- *Step 2 Border straightening*

The boundaries between the superfaces are simplified by selecting a subset of vertices shared by adjacent superfaces to be endpoints for "superedges."


- *Step 3 Superface triangulation*

Triangulation points are determined for each superface (i.e., points about which the polygon can be broken into triangles). If necessary, a large superface may be resplit into two or more smaller superfaces, each with its own boundary and triangulation points. The algorithm usually stops here, without actually generating triangles; however, an explicit triangulated model data structure can be produced from the superface boundaries and triangulation points if this is required.

The method is computationally fast and produces reasonably good simplifications. For example, the algorithm simplified a polyhedral model consisting of 196200 polygonal faces, obtained from a CT scan of a plastic skull, to a simplified model with only 6320



(a)




(b)

Figure 14


Phase 1 (creation of superfaces) of superfaces algorithm. Each colored patch corresponds to a quasi-planar "superface": (a) the superfaces obtained by simplifying a model of a plastic skull; (b) a close-up of the output for the area in the box of (a), showing the individual polygonal faces of the original model.

polygonal superfaces, while guaranteeing that every vertex of the original model was within one voxel-diameter⁶ of some facet of the simplified model. Since many graphical and geometric algorithms perform calculations on one triangle at a time, one way of estimating the computational savings that may be achieved from the use

⁶ A "voxel" is essentially the 3D equivalent of a 2D pixel.



(a)



(b)

Figure 15

Superfaces model simplification. The original model (a) has 196 200 polygonal faces, yielding 349 800 triangles. The simplified model (b) has the same topology, but only 14 686 superfaces, yielding 128 040 triangles.

of simplified models is to compare the number of triangles associated with an original polyhedron and the number associated with the corresponding superfaces model. In the example just cited, the triangulated form of the original polyhedral model had 349 800 triangles, while the

simplified model had 78 002 triangles. If the allowable error bound is reduced from one to one-half voxel, the algorithm produces a simplified model with 14 686 superfaces and 128 040 triangles. **Figure 15(a)** shows the original model, while **15(b)** shows the simplified model with 14 686 superfaces. Similar experiments have shown triangle reductions ranging from 4:1 to 12:1 on skull models, and up to 20:1 for skin, while maintaining a one-voxel error bound, with average errors about 10% of the error bound. We are considering a number of ways to extend and exploit this approach.

A second project is the registration of images from different imaging modalities, possibly taken at different times, such as preoperative CT scans and intraoperative X-ray images. Image registration consists of "aligning" the images so as to have a common reference frame. The aligned images can be used to monitor the progress of a disease, to monitor a procedure, to execute a preoperative plan, or to calibrate a robotic device. A number of researchers (e.g., [64–69]) have addressed this problem in various contexts, and it is a current topic of considerable research activity in many groups. We expect to draw upon and extend much of this work, while also focusing on characterization of uncertainties associated with imaging and sensor-alignment errors, computational efficiency, and robustness of registration algorithms in the presence of misalignments. One approach that we are pursuing (similar in some respects to that of [65]) compares observed X-ray images to predicted X-ray images calculated from CT data. In the visualization shown in **Figure 16**, edges from simulated projection X-ray images computed from preoperative CT data for a cadaver femur are shown in yellow. Edges from an actual X-ray are shown in blue. Edge elements appearing at the same place in both images are shown in red. The overlaid image appears red where the two images are both bright.

Another approach we have taken (similar in some respects to that of [66]) generates a model of the surface of the anatomy from 3D CT data and compares 2D silhouettes of this 3D preoperative surface model with 2D contours detected from the intraoperative X-ray data (multiple X-rays may be used). In this approach, tentative 2D correspondences between points from the projected surface model and the detected X-ray contours are automatically generated by means of an algorithm that finds the shortest distance between a point and a polygonal curve. The coordinates of the corresponding pairs of points are given in the coordinate systems associated with the X-ray views.

The object of the registration process is to find the pose of the X-ray source with respect to the 3D anatomy from the planar X-ray information and the 3D surface model of the anatomy derived from CT data. At each iteration of the registration algorithm, an incremental combined


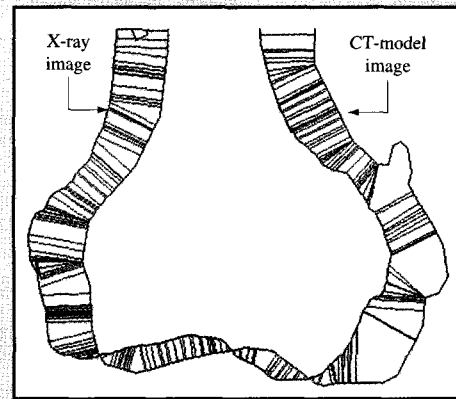



Figure 16

Registered CT and X-ray images. In this visualization, edges from simulated X-ray images based on a predicted X-ray camera model and preoperative CT data are shown in green. Edges from an actual X-ray are shown in blue. Edge elements that appear at the same place in both images are shown in red. The overlaid image appears red where the two images are both bright. The dots are small steel spheres imbedded in the test fixture that holds the bone.

rotation and translation of the CT model with respect to the assumed X-ray source position is computed so as to minimize the objective function, which is the sum of 3D distances between the 3D model surface points and 3D lines joining the X-ray contours and the optical center of the X-ray system. Several computational methods (linear, nonlinear, and statistically robust) have been used and compared for computing the optimal registration. The process of finding correspondences and computing an optimal pose of the X-ray source is repeated until a prespecified registration accuracy (value of the objective function) is obtained, or until the registration is no longer improved. Experiments on simulated data have shown that, with random initial displacements, a registration of the perspective projection of the 3D surface to the actual X-ray data can be obtained with maximum errors of 1–2 mm. The computational expense for the registration was typically 20 seconds of CPU time on an IBM RISC System/6000 58K workstation.



(a)



(b)

Figure 17

Model-based 2D–3D registration using contours: (a) the predicted 2D projected contour of a CT-derived polygonal model of a knee (red) superimposed on contours extracted from an actual X-ray image of the knee (green). The gray lines represent tentative pairings between points on the contours. (b) is a comparable image obtained after the X-ray camera pose estimate has been refined.

Figure 17(a) shows a 2D silhouette of the CT model superimposed on the contours extracted from an actual X-ray image of the knee part of a femur. The silhouette of the CT model is shown in red, whereas the extracted X-ray contours are shown in green. The gray lines between the red and green curves show tentative assignments of corresponding points, which are automatically computed by means of a closest-point finder. **Figure 17(b)** shows improved alignment of the projected CT model to the X-ray image after registration.

Conclusion

Computers and computer-controlled devices have enormous potential to augment the ability of human

clinicians to plan and carry out surgical procedures. For this potential to be realized, close cooperation among engineers, computer scientists, and users is essential. This summary has provided a brief overview of the strategy of the Computed-Integrated Surgery group at the IBM Thomas J. Watson Research Center, and has provided a few examples of our activities in implementing our strategy. This is a relatively young field, and much is yet to be done if this technology is to have its full impact on health care; however, we are excited by the prospects and look forward to exploring the possibilities.

Acknowledgments

In writing a summary of this sort, it is often very difficult for one to properly acknowledge the many people who have contributed significantly to the work discussed. In some ways, the easiest thing to do is to invite everyone who contributed significantly to be a co-author. The problem in this case is that the list would be too long for this to be practical. We have, therefore, limited the author list to the currently active members of the IBM CIS group. At the risk of some injustice, we list here a few of the many other people, both inside and outside IBM, whose efforts have been instrumental in whatever progress we have made to date: James Anderson, William Bargar*, Nils Bruun, Court Cutting*, John DeSouza, Benjamin Eldridge, Dave Epstein, Edward Glassman, Dieter Grimm, David Grossman, Kreg Gruben, Betsy Haddad, William Hanson, Jay Hammershoy, John Karidis, Louis Kavoussi*, Peter Kazanzides, Yong-yil Kim, Deljou Khoramabadi, Robert Lipori, Joseph McCarthy*, Brent Mittelstadt, Bela Musits, Gerald Vicker, Marilyn Noz, Robert Olyha, Howard Paul*, David Rattner*, Josh Rothenberg, Nitish Swarup, Mark Talamini*, Michael Treat*, Roderick Turner*, Bill Williamson, and Joel Zuhars. (Names followed by an asterisk are those of surgeons who have provided much of the crucial clinical feedback.) If we have inadvertently left someone out, please accept both our profound apologies and our thanks.

ROBODOC and ORTHODOC are trademarks of Integrated Surgical Systems, Inc.

RISC System/6000 and PS/2 are registered trademarks of International Business Machines Corporation.

Optotrak is a trademark of Northern Digital, Inc.

Pixsys is a trademark of Pixsys, Inc.

References

1. Y. Kosugi, E. Watanabe, J. Goto, T. Watanabe, S. Yoshimoto, K. Takakura, and J. Ikebe, "An Articulated Neurosurgical Navigation System Using MRI and CT Images," *IEEE Trans. Biomed. Eng.* **35**, No. 2, 147-152 (February 1988).
2. R. L. Galloway, R. J. Maciunas, and C. A. Edward II, "Interactive Image-Guided Neurosurgery," *IEEE Trans. Biomed. Eng.* **39**, 1226-1231 (1992).
3. S. Lavallee and P. Cinquin, "IGOR: Image Guided Operating Robot," *Proceedings of the Fifth International Conference on Advanced Robotics*, Pisa, June 1991, pp. 878-881.
4. Patrick J. Kelly, Bruce A. Kall, Stephen Goerss, and Franklin Earnest, "Computer-Assisted Stereotaxic Laser Resection of Intra-Axial Brain Neoplasms," *J. Neurosurg.* **64**, No. 3, 427-439 (March 1986).
5. S. J. Zinreich, S. A. Tebo, D. M. Long, H. Brem, D. E. Mattox, M. E. Loury, C. A. Vander Kolk, W. M. Koch, D. W. Kennedy, and R. N. Bryan, "Frameless Stereotaxic Integration of CT Imaging Data: Accuracy and Initial Applications," *Radiol.* **188**, No. 3, 735-742 (September 1993).
6. Robert J. Maciunas, *Interactive Image-Guided Neurosurgery*, American Association of Neurological Surgeons, Park Ridge, IL, 1993.
7. Russell H. Taylor, Howard A. Paul, Peter Kazanzides, Brent D. Mittelstadt, William Hanson, Joel F. Zuhars, Bill Williamson, Bela L. Musits, Edward Glassman, and William L. Bargar, "An Image-Directed Robotic System for Precise Orthopaedic Surgery," *IEEE Trans. Robotics & Automation* **10**, No. 3, 261-275 (June 1994).
8. B. D. Mittelstadt, P. Kazanzides, J. Zuhars, B. Williamson, P. Kain, F. Smith, and W. Bargar, "The Evolution of a Surgical Robot from Prototype to Human Clinical Trial," *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, Pittsburgh, September 22-24, 1994.
9. Thomas C. Kienzle, S. David Stulberg, Michael Peshkin, Arthur Quaid, and Chi-Haur Wu, "An Integrated CAD-Robotics System for Total Knee Replacement Surgery," *Proceedings of the 1993 IEEE Conference on Robotics and Automation*, Atlanta, May 1993, pp. 889-894.
10. S. Lavallee, P. Sautot, J. Troccaz, P. Cinquin, and P. Merloz, "Computer Assisted Spine Surgery: A Technique for Transpedicular Screw Fixation Using CT Data and a 3D Optical Localizer," *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, Pittsburgh, September 22-24, 1994.
11. L. P. Nolte, L. J. Zamorano, Z. Jiang, G. Want, F. Langlotz, E. Arm, and H. Visarius, "A Novel Approach to Computer Assisted Spine Surgery," *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, Pittsburgh, September 22-24, 1994.
12. N. Tejima, H. Funakubo, T. Dohi, I. Sakuma, T. Tanashima, and Y. Namura, "A New Microsurgical Robot for Corneal Transplantation," *Precision Machinery*, pp. 1-9 (1988).
13. S. Charles, "Dexterity Enhancement in Microsurgery Using Telemicro-Robotics," *Proceedings, Medicine Meets Virtual Reality II*, San Diego, January 1994, pp. 19-20.
14. P. S. Jensen, M. R. Glucksberg, J. E. Colgate, K. W. Grace, and R. Attariwala, "Micropuncture of Retinal Vessels Using a Six Degree of Freedom Ophthalmic Robotic Micromanipulator," *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, Pittsburgh, September 22-24, 1994.
15. Ian Hunter, Mark Sagar, Lynette Jones, Tilemachos Doukoglou, Serge Lafontaine, and Peter Hunter, "Teleoperated Microsurgical Robot and Associated Virtual Environment," *Proceedings, Medicine Meets Virtual Reality II*, San Diego, January 1994, pp. 85-89.
16. Court Cutting, Russell H. Taylor, Fred Bookstein, Alan Kalvin, Betsy Haddad, Yong-yil Kim, Marilyn Noz, and

- Joseph McCarthy, "Comprehensive Three-Dimensional Cephalometric System for the Planning and Execution of Craniofacial Surgical Procedures," *Proceedings of the Fourth Biannual Meeting of the International Society of Cranio-Maxillofacial Surgery*, Santiago de Compostela, Spain, June 13-16, 1991.
17. D. E. Altobelli, R. Kikinis, J. B. Mulliken, H. Cline, W. Lorensen, and F. Jolesz, "Three-Dimensional Imaging in Medicine: Surgical Planning," *Proceedings of the International Conference of the IEEE Engineering and Biology Society*, 1991, pp. 289-290.
 18. Ludwig Adams, Joachim M. Gilsbach, Werner Krybus, Dietrich Meyer-Ebrecht, Ralph Mosges, and Georg Schlondorff, "CAS—A Navigation Support for Surgery," *3d Imaging in Medicine*, Springer-Verlag, Berlin, 1990, pp. 411-423.
 19. N. Nitsche, M. Hilbert, G. Strasser, H. P. Tuemmler, and W. Arnold, "Application of a Noncontact Computerized Localising System in Paranasal Surgery," *Otorhinolaryngologica Nova* 3, 57-64 (1993).
 20. K. T. Kavanagh, "Applications of Image-Directed Robotics in Otolaryngologic Surgery," *Laryngoscope* 104, No. 3, Pt. 1, 283-293 (1994).
 21. K. Ikuta, M. Tsukamoto, and S. Hirose, "Shape Memory Alloy Servo Actuator System with Electric Resistance Feedback and Application for Active Endoscopes," *Proceedings of the 1988 IEEE Conference on Robotics and Automation*, April 1988, pp. 427-430.
 22. R. H. Sturges and S. Laowattana, "A Flexible Tendon-Controlled Device for Endoscopy," *Proceedings of the 1991 IEEE Conference on Robotics and Automation*, Sacramento, CA, 1991, pp. 2582-2591.
 23. B. L. Davies, R. D. Hibberd, A. Timoney, and J. Wickham, "A Surgeon Robot for Prostatectomies," *Proceedings of the Fifth International Conference on Advanced Robotics*, Pisa, June 1991, pp. 871-875.
 24. R. Sturges, "Voice Controlled Flexible Endoscope," videotape, Carnegie Mellon University, Pittsburgh, 1989.
 25. Yulan Wang, "Robotically Enhanced Surgery," *Proceedings, Medicine Meets Virtual Reality II*, San Diego, January 27-30, 1994.
 26. J. A. McEwen, "Solo Surgery with Automated Positioning Platforms," *Proceedings of the NSF Workshop on Computer Assisted Surgery*, Washington, DC, February 1993.
 27. Joseph B. Petelin, "Computer Assisted Surgical Instrument Control," *Proceedings, Medicine Meets Virtual Reality II*, San Diego, January 1994, pp. 170-173.
 28. R. Hurteau, S. DeSantis, E. Begin, and M. Gagnier, "Laparoscopic Surgery Assisted by a Robotic Cameraman: Concept and Experimental Results," *Proceedings of the 1994 IEEE Conference on Robotics and Automation*, San Diego, May 8-13, 1994, pp. 2286-2289.
 29. Russell H. Taylor, Janez Funda, Ben Eldridge, Kreg Gruben, David LaRose, Steve Gomory, Mark Talamini, M.D., Louis Kavoussi, M.D., and James Anderson, "A Telerobotic Assistant for Laparoscopic Surgery," *IEEE EMBS Magazine Special Issue on Robotics in Surgery*, B. Davies, Ed., April-May 1995.
 30. W. J. Peine, J. S. Son, and R. D. Howe, "A Palpation System for Artery Localization in Laparoscopic Surgery," *Proceedings of the 1st International Symposium on Medical Robotics and Computer Assisted Surgery*, Pittsburgh, September 1994, pp. 250-257.
 31. Philip Green, "Telepresence Surgery," *NSF Workshop on Computer Assisted Surgery*, Washington, DC, February 1993.
 32. B. Neisius, P. Dautzenberg, and R. Trapp, "Robotic Manipulator for Endoscopic Handling of Surgical Effectors and Cameras," *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, Pittsburgh, September 22-24, 1994.
 33. R. Taylor, S. Lavallee, G. Burdea, and R. Moesges, *Computer-Integrated Surgery*, MIT Press, Cambridge, MA, 1995.
 34. J. L. Coatrieux and J. M. Scarabin, Special Issue on Computer Graphics in Medicine, *ITMB Journal (Innovation and Technology in Biology and Medicine)*, 1987.
 35. S. Lavallee, Special Issue on Computer Graphics in Medicine, *ITMB Journal (Innovation and Technology in Biology and Medicine)*, 1992.
 36. R. H. Taylor and G. A. Bekey, organizers, *Proceedings of the NSF Workshop on Computer Assisted Surgery*, Washington, DC, 1993.
 37. *Proceedings, Medicine Meets Virtual Reality II: Interactive Technology & Healthcare*, San Diego: Office of Continuing Medical Education, University of California, San Diego, 1994.
 38. A. DiGioia, M.D., R. H. Taylor, Program Chairman, and T. Kanade, General Chairman, *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, 1994.
 39. H. A. Paul, D. E. Hayes, W. L. Bargar, and B. D. Mittelstadt, "Accuracy of Canal Preparation in Total Hip Replacement Surgery Using Custom Broaches," *Proceedings of the First International Symposium on Custom Made Prostheses*, Dusseldorf, Germany, October 1988, pp. 153-161.
 40. W. H. Hanson, H. A. Paul, Bill Williamson, and Brent Mittelstadt, "Orthodoc: A Computer System for Presurgical Planning," *Proceedings of the 12th IEEE Conference on Medicine & Biology*, Philadelphia, 1990, pp. 1931-1932.
 41. Russell H. Taylor, Howard A. Paul, Peter Kazanzides, Brent D. Mittelstadt, William Hanson, Joel F. Zuhars, Bill Williamson, Bela L. Musits, Edward Glassman, and William L. Bargar, "Taming the Bull: Safety in a Precise Surgical Robot," *Proceedings of the 1991 International Conference on Advanced Robotics I*, Pisa, Italy, June 1991, pp. 865-870.
 42. *Orthopedic Network News*, Mendenhall Associates, Inc., Ann Arbor, MI, 1992.
 43. *Orthopedic Network News*, Mendenhall Associates, Inc., Ann Arbor, MI, 1993.
 44. Leo Joskowicz and Russell H. Taylor, "Hip Implant Insertability Analysis: A Medical Instance of the Peg-in-Hole Problem," *Proceedings of the 1993 IEEE Conference on Robotics and Automation*, Atlanta, May 1993, pp. 901-908.
 45. Leo Joskowicz and Russell H. Taylor, "Preoperative Insertability Analysis and Visualization of Custom Hip Implants," *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, Pittsburgh, September 1994.
 46. Leo Joskowicz and Russell H. Taylor, "Interference-Free Insertion of a Solid Body into a Cavity: Algorithm and a Medical Application," *The International Journal of Robotics Research* 15 (1996).
 47. Court B. Cutting, Fred L. Bookstein, and Russell H. Taylor, "Applications of Simulation, Morphometrics, and Robotics in Craniofacial Surgery," *Computer-Integrated Surgery*, R. Taylor, S. Lavallee, G. Burdea, and R. Moesges, Eds., MIT Press, Cambridge, MA, 1995, pp. 641-662.
 48. Russell H. Taylor, Court B. Cutting, Yong-yil Kim, Alan D. Kalvin, David Larose, Betsy Haddad, Deljou Khoramabadi, Marilyn Noz, Robert Olyha, Nils Bruun, and Dieter Grimm, "A Model-Based Optimal Planning and Execution System with Active Sensing and Passive Manipulation for Augmentation of Human Precision in Computer-Integrated Surgery," *Proceedings of the 1991 International Symposium on Experimental Robotics*,

- Toulouse, France, June 25–27, 1991.
49. C. Cutting, R. Taylor, F. Bookstein, D. Khoramabadi, B. Haddad, A. Kalvin, H. Kim, and M. Noz, "Computer Aided Planning and Execution of Craniofacial Surgical Procedures," *Proceedings of the IEEE Conference on Engineering in Medicine and Biology*, Paris, October 1992.
 50. Russell H. Taylor, Howard A. Paul, Court B. Cutting, Brent Mittelstadt, William Hanson, Peter Kazanzides, Bela Musits, Yong-Yil Kim, Alan Kalvin, Betsy Haddad, Deljou Khoramabadi, and David LaRose, "Augmentation of Human Precision in Computer-Integrated Surgery," *Innovation et Technologie en Biologie et Medicine* **13**, 450–468 (1992).
 51. Court Cutting, M.D., Barry Grayson, D.D.S., and Hae Chun Kim, "Precision Multi-Segment Bone Positioning Using Computer Aided Methods in Craniofacial Surgical Applications," *Proceedings of the 12th IEEE Conference on Medicine and Biology*, Philadelphia, November 1990.
 52. Janez Funda, Russell Taylor, Kreg Gruben, and David LaRose, "Optimal Motion Control for Teleoperated Surgical Robots," *Proceedings of the 1993 SPIE International Symposium on Optical Tools for Manufacturing and Advanced Automation* **2057**, Boston, September 1993, pp. 211–222.
 53. Janez Funda, Russell Taylor, Ben Eldridge, Kreg Gruben, David LaRose, and Steve Gomory, "Image Guided Command and Control of a Surgical Robot," *Proceedings, Medicine Meets Virtual Reality II*, San Diego, January 1994, pp. 52–57.
 54. Janez Funda, Russell Taylor, Steve Gomory, Ben Eldridge, Kreg Gruben, and Mark Talamini, M.D., "An Experimental User Interface for an Interactive Surgical Robot," *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, Pittsburgh, September 22–24, 1994.
 55. J. Funda, B. Eldridge, K. Gruben, S. Gomory, and R. Taylor, "Comparison of Two Manipulator Designs for Laparoscopic Surgery," *Proceedings of the 1994 SPIE International Symposium on Optical Tools for Manufacturing and Advanced Automation* **2351**, Boston, October 1994, pp. 172–183.
 56. B. Eldridge, K. Gruben, D. LaRose, J. Funda, S. Gomory, J. Karadis, G. McVicker, R. Taylor, and J. Anderson, "A Remote Center of Motion Robotic Arm for Computer Assisted Surgery," *Robotica* **14**, No. 1, 103–109 (Jan.–Feb. 1996).
 57. J. Funda, K. Gruben, B. Eldridge, S. Gomory, and R. Taylor, "Control and Evaluation of a 7-axis Surgical Robot for Laparoscopy," *Proceedings of the IEEE International Conference on Robotics and Automation* **2**, Nagoya, May 1995, pp. 1477–1484.
 58. Hans P. Moravec, "Obstacle Avoidance and Navigation in the Real World by a Seeing Robot Rover," Ph.D. Thesis, Stanford University, 1980.
 59. Yulan Wang, "Automated Endoscopic System for Optimal Positioning," company advertising brochure, Computer Motion, Inc., Goleta, CA, 1993.
 60. A. Kalvin, D. Dean, J. Hublin, and M. Braun, "Visualization in Anthropology: Reconstruction of Human Fossils from Multiple Pieces," *Proceedings of the IEEE Visualization '92 Conference*, Boston, October 1992, pp. 404–410.
 61. D. Dean, A. Kalvin, J. Hublin, and M. Braun, "Three-Dimensional CT Composite Reconstruction of Sale and Thomas Quarry Cranial Remains," *North American Journal of Physical Anthropology*, Supplement 16:79 (1993).
 62. Alan Kalvin and Russell Taylor, "SuperFaces: Hierarchical Polyhedral Approximation with Bounded Error," *Proceedings of the SPIE Medical Imaging Conference 1994* **2164**, Newport Beach, CA, February 1994, pp. 2–13 (also IBM Research Report RC-19135, April 1993).
 63. A. D. Kalvin and R. H. Taylor, "Superfaces: Polygonal Mesh Simplification with Bounded Error," *IEEE Computer Graph & Appl.* **16**, No. 3, 64–77 (May 1996).
 64. S. Lavallee, R. Szeliski, and L. Brunie, "Matching 3D Smooth Surfaces with their 2D Projections Using 3D Distance Maps," *Proceedings of the SPIE Conference on Geometric Methods in Computer Vision*, San Diego, July 25–26, 1991, pp. 322–336.
 65. John R. Adler, "Image-Based Frameless Stereotaxy," *Interactive Image-Guided Radiosurgery*, Robert J. Maciunas, Ed., American Association of Neurological Surgeons, 1993, pp. 81–89.
 66. André P. Gueziec and Nicholas Ayache, "Smoothing and Matching of 3-D Space Curves," *Int. J. Computer Vision* **12**, No. 1, 79–104 (February 1994).
 67. S. Lavallee, "Registration for Computer Integrated Surgery: Methodology, State of the Art," *Computer Integrated Surgery*, R. Taylor, S. Lavallee, G. Burdea, and R. Moesges, Eds., MIT Press, Cambridge, MA, 1995, pp. 77–97.
 68. Fabienne Betting, Jacques Feldmar, Nicholas Ayache, and Frederic Devernay, "A New Framework for Fusing Stereo Images with Volumetric Medical Images," *Proceedings of CVRMed 1995*, Nicholas Ayache, Ed., Springer, Berlin, April 1995, pp. 30–39.
 69. Ali Hamadeh, Stephane Lavallee, Richard Szeliski, Philippe Cinqun, and Olivier Peria, "Anatomy-Based Registration for Computer-Integrated Surgery," *Proceedings of CVRMed 1995*, Nicholas Ayache, Ed., Springer, Berlin, April 1995, pp. 212–218.

Received March 10, 1995; accepted for publication May 18, 1995

Russell H. Taylor *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (rht@cs.jhu.edu).* Dr. Taylor received a B.E.S. degree from Johns Hopkins University in 1970 and a Ph.D. in computer science from Stanford in 1976. From 1976 to 1995, he was a research staff member and research manager at the IBM Thomas J. Watson Research Center; since September 1995 he has been a professor of computer science at Johns Hopkins University. His research interests include robot systems, programming languages, model-based planning, and (most recently) the use of imaging, model-based planning, and robotic systems to augment human performance in surgical procedures. Dr. Taylor co-developed the ROBODOC system for hip replacement surgery, and he has been actively involved in a number of other medical robotics and computer-assisted surgery activities both at IBM and at Johns Hopkins. He is editor emeritus of the *IEEE Transactions on Robotics and Automation*, a fellow of the IEEE, and a member of various honorary societies, panels, program committees, and advisory boards. Dr. Taylor has co-chaired a number of conferences and workshops on computer-integrated surgery, including a 1993 NSF Workshop on Computer-Assisted Surgery and the First and Second International Symposia on Medical Robotics and Computer-Assisted Surgery.

Janez Funda *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598.* Dr. Funda received a B.A. degree in computer science and mathematics from Macalester College in 1986, and a Ph.D. in computer science from the University of Pennsylvania in 1991. He joined the Computer-Integrated Surgery group at IBM Research in 1991. Dr. Funda's research interests include robot systems, telemanipulation, man-machine interfaces, and virtual reality. His current research focuses on the use of robotic, sensing, and image processing technology to assist in performing surgical procedures. He holds two U.S. and international patents.

Leo Joskowicz *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598.* Dr. Joskowicz is currently project leader of the Computer-Integrated Surgery group at the IBM Thomas J. Watson Research Center. He joined IBM Research in 1988, after receiving an M.Sc. and a Ph.D. in computer science from the Courant Institute of Mathematical Sciences, New York University. From 1988 to 1993 he was a research staff member in the Artificial Intelligence Department, where he conducted research on qualitative and geometric reasoning, constraint reasoning, and intelligent CAD. Dr. Joskowicz is on the editorial board of the journal *Artificial Intelligence in Engineering*; he is a senior member of the IEEE. His research interests include computer-integrated surgery, medical imaging, image registration, geometric reasoning, and motion planning.

Alan D. Kalvin *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (KALVIN at WATSON).* Dr. Kalvin received the B.Sc. degree from the University of the Witwatersrand, South Africa, the B.Sc. Honours degree from the University of Cape Town, South Africa, and the M.S. and Ph.D. degrees from the Courant Institute of Mathematical Sciences, New York University, all in computer science, in 1975, 1976, 1985, and 1991, respectively. From 1987 to 1990 he worked as a research

scientist at the Institute of Reconstructive Plastic Surgery at the New York University Medical Center. He is currently a research staff member in the Computer-Integrated Surgery group at the IBM Thomas J. Watson Research Center. Dr. Kalvin's research interests include computer vision, medical imaging, geometric modeling, computer-assisted anthropology, and computer graphics.

Stephen H. Gomory *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (STEVEG at YKTVMV).* Mr. Gomory received a B.A. in architecture from Columbia College, Columbia University, in 1988, and an M.Sc. in computer science from New York University in 1995. He is currently a staff programmer at the IBM Thomas J. Watson Research Center, where he works in the Computer-Integrated Surgery group. His computer programming work there has centered around image processing, robotics, and system integration.

André P. Gueziec *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (GUEZIEC at YKTVMV, gueziec@watson.ibm.com).* Dr. Gueziec is a postdoctoral researcher in the Computer-Integrated Surgery group at the IBM Thomas J. Watson Research Center. His research efforts are in object modeling and recognition, which he has applied to medical imaging and computer vision. His research interests also include computational geometry and computer graphics. Dr. Gueziec received his Ph.D. in computer science from the University of Paris 11 in Orsay in 1993. In 1989, he graduated from the Ecole Centrale de Paris. From 1990 to 1993 he was a Ph.D. student at INRIA in Paris. In 1993 and 1994, Dr. Gueziec was with the Courant Institute, New York University. He is a member of IEEE EMBS, IEEE, and SPIE.

Lisa M. G. Brown *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (LISAB at WATSON).* Ms. Brown is currently finishing her Ph.D. in computer science at Columbia University in New York. Her thesis addresses the problem of multimodal medical image registration and deals directly with exploiting the underlying relationships between sensors. She has been working on her thesis while working with the Computer-Integrated Surgery group at the IBM Thomas J. Watson Research Center. Ms. Brown received her B.A. from Johns Hopkins University in 1980, and her M.A. in computer science from Columbia University in 1987. She has worked for several years in industry, applying numerical techniques to model physical systems, and, more recently, in image processing and computer vision. She has made several contributions to the problems of computer vision and medical imaging.