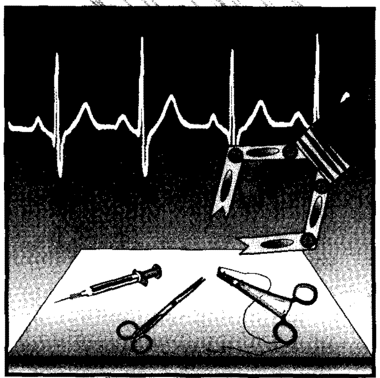

A Telerobotic Assistant for Laparoscopic Surgery

Russell H. Taylor[1], Janez Funda[1], Ben Eldridge[1], Steve Gomory[1], Kreg Gruben[2], David LaRose[1], Mark Talamini[2], Louis Kavoussi[2], James Anderson[2]

[1] IBM T. J. Watson Research Center
[2] Johns Hopkins University School of Medicine

The goal of this work is to develop a new generation of "intelligent" surgical systems that can work cooperatively with a human surgeon to off-load routine tasks, reduce the number of people needed in the operating room, and provide new capabilities that complement the surgeon's own skills. An underlying premise of this work is that machine capabilities coupled with human judgement can accomplish many tasks better than either could do alone. A further premise is that such a partnership is synergistic with present trends toward geometrically precise, image guided, and minimally invasive therapies. The net result will be better clinical results, lower net costs through shorter hospital stays and recovery times, and reducing the chances for repeated surgery.

Most of the key enabling technologies such as 3D imaging, modelling, visualization, realtime sensing, telerobotics, and system integration, is computer based. The emergence of very powerful, affordable computer workstations together with scientific advances in imaging, modelling, and telerobotics, mean that critical cost/capability thresholds have been crossed, and the pace of research and clinical activity is increasing sharply. Much of this activity takes advantage of the increased precision with which computer-controlled mechanical devices can position and maneuver surgical instruments. This aspect of machine capability has been exploited in a number of orthopaedic and neurosurgical applications, e.g., [1-6]. Some of the other work in this area has concentrated on exploiting computer and robotic technology either to reduce fatigue, restore hand-eye coordination, and improve dexterity of human surgeons, or to reduce the number of personnel required in the operating room, e.g., [7-14]. This dichotomy is by no means absolute. Some of these systems, e.g., [1], clearly incorporate aspects of both types of functionality. The system described in this article has aspects of both types of functionality. Although the initial application domain is laparoscopic surgery, using relatively simple tasks such as camera pointing and instrument positioning, the system is capable of operating both under the surgeon's direct control and more autonomously under the surgeon's supervision, while extracting targeting information from realtime images. We anticipate eventually applying this robotic system to a very broad range of surgical tasks.

Laparoscopic surgery has seen remarkable growth over the last five years. In 1992, 70 percent of all gall bladder surgery in the U.S., Europe and Japan was done laparoscopically. By the year 2000, it is estimated that from 60 to 80 percent of abdominal surgeries will be performed laparoscopically [15]. Flexible endoscopy is similarly becoming more and more prevalent. Two salient characteristics of these procedures are that the surgeon cannot directly manipulate the patient's anatomy with his (or her) fingers and that he cannot directly observe what he is doing. Instead, he must rely on instruments that can be inserted through a canula or through the working channel of an endoscope. Often, he must rely on an assistant to point the camera while he performs the surgery. The awkwardness of this arrangement has led a number of researchers to develop robotic augmentation devices for endoscopic surgery. Typical efforts include improved mechanisms for flexible endoscopes (e.g., [7], [16]), specialized devices for particular applications (e.g., [10]), voice-control for existing mechanisms (e.g., [9]), full blown "telepresence" systems [11, 17], and simple camera pointing systems [14, 18-20].

Of these efforts, the most ambitious in some ways is the telepresence surgery system of Green, et al. at SRI International (Menlo Park, CA) [11, 12], whose aim is to use a force reflecting manipulator, stereo visualization, and other "virtual reality" technology to give the surgeon the sensation of doing open surgery. Although the system reported in this article has some of the same capabilities as the SRI system and, indeed, its image guidance functions may make it in some ways be better suited to remote telesurgery, where time delays are large, our primary goal is somewhat different. We view surgical robotic devices as being most valuable in their ability to aid and augment the surgical team, allowing more efficient use of of available surgical talent and enhancing the ability of

0739-5175/95/$4.00©1995

surgeons to work quickly and accurately. Our goal is not so much telepresence surgery, as the provision of an intelligent "third hand," operating under the surgeon's supervision that can off-load routine tasks, reduce the number of people needed in the OR, and provide new capabilities (such as accurate targeting) that complement the surgeon's own abilities.

At the other extreme are systems [14, 19, 20] whose goal is to do the very simple task of aiming a laparoscopic camera. This action can possibly reduce the number of people required in the operating room while leaving the responsibility for manipulating the patient's anatomy completely up to the surgeon. These systems typically provide a very simple teleoperation interface, allowing the surgeon to directly steer a robot holding a laparoscopic camera. Camera pointing has some obvious attractions as an entry-level application, since it is relatively simple, participates in the surgery only passively, and does not require a fundamental change in other aspects of the surgical procedure.

Our system includes a specially designed remote-center-of-motion robot that holds a laparoscopic camera or other instrument, a variety of human-machine interfaces, and a controller. The controller provides robot-control, image processing, and display functions. Our system has some aspects in common with the previously discussed laparoscope holding systems. In particular, we provide direct teleoperator control of camera positioning as one mode of operation, although, perhaps, with more flexibility and convenience in controlling the view, and a richer set of human-machine interfaces. For example, our system is able to maintain an "upright" image while panning an angled-view laparoscope. A more crucial difference is that we provide alternatives to direct teleoperation for guiding the system.

In particular, the system is capable of capturing images from the camera and processing them to obtain geometric information about the patient's anatomy, which may then be used to assist in aiming the camera or positioning other instruments held by the robot. Our eventual goal is a suite of functional capabilities including retraction, countertraction, hemostasis, suturing assistance, simple dissection, etc. that a surgeon might reasonably expect from a human assistant. We also expect the system to be able to combine information coming from the camera with information obtained from other imaging modalities (CT, MRI, ultrasound, fluoroscopy, etc.) to perform tasks, such as accurate positioning of therapy delivery devices, which are better suited to machine than to human capabilities.

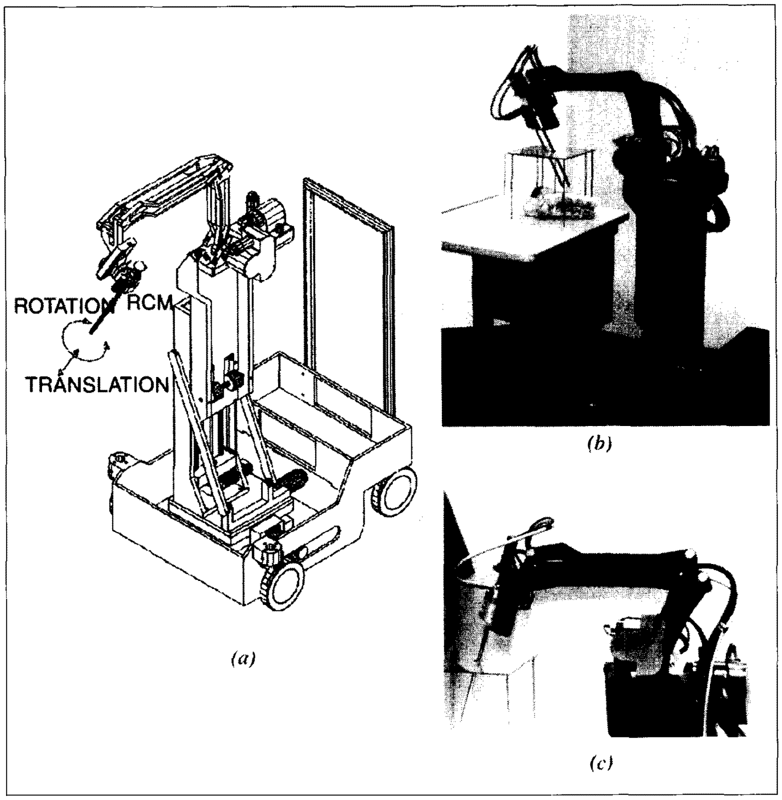

1. Remote center-of-motion robot: (a) design drawing, (b) photograph of whole robot, and (c) photograph of distal four axes. All motions are kinematically decoupled at the point where the laparoscopic instrument would enter the patient's body.

The present system prototype was developed as part of a joint study between IBM and the Johns Hopkins University Medical School. In subsequent sections, we describe the robot, the human machine interfaces, and operational characteristics of the system.

Surgical Robot

Manipulator Design

Safety, control convenience, and flexibility for use in a wide variety of surgical applications were important factors in determining the manipulator design. In laparoscopic applications, rigid instruments are inserted into the patient's body through small canulas inserted into the abdominal wall. This arrangement creates a "fulcrum effect," so that the instrument has only four significant motion degrees-of-freedom (three rotations and depth of penetration) centered at the entry portal. Only very constrained lateral motions are acceptable.

If a robot is holding an instrument, it is very important that its motions obey these constraints. A conventional industrial robot can, of course, be programmed to move an instrument about such a fulcrum. Unfortunately, such motions usually require several manipulator joints to make large, tightly coordinated excursions. Thus, even relatively slow end-effector motions can require rapid joint motions. Any control or coordination failure can thereby represent a potential safety hazard both for the patient and for the surgeon. Simply slowing down the actuators can cause the overall functioning of the robot to be painfully tedious. Consequently, we have a strong preference for manipulator designs that require only low velocity actuation, do not have motion singularities in the normal working volume, and permit simple stable controls. Similarly, the motions required to perform a task should be reasonably intuitive for the surgeon. Even if the control computer is handling all the details, it is desirable not to surprise the surgeon with unanticipated complex motions. Finally, we want a great deal of modularity to allow us to reconfigure the system for different procedures.

Our solution is to construct a kinematically redundant manipulator composed of a proximal translation component, along with a distal remote center-of-motion component that provides angular reorientation about a fixed point and a controlled insertion motion that passes through the remote motion center. Our present embodiment, shown in Fig. 1, consists of a 3-axis linear xyz stage, a 2-axis parallel four bar linkage providing two rotations (R_x and R_y) about the remote motion center, and a 2-axis distal component providing an insertion motion, s, and rotation R_s about the instrument axis, which passes through the remote motion center. Thus, the robot's distal four degrees of freedom are kinematically decoupled about the remote motion center, whose position may be translated in space by the proximal three-axis linear stage. In addition to mechanically enforcing the fulcrum constraints, this design has the important benefit that "natural" motions of the manipulator (i.e., those that can be accomplished by motion of a single actuator) correspond to common primitive task motions, such as insertion of instruments into the patient's body. For use in laparoscopic camera navigation, we have also implemented an additional motorized degree-of-freedom to rotate the camera "head" about the eyepiece of an "angled-view" laparoscope, thus making it possible to keep the image on the screen upright as the laparoscope is rotated about its axis.

For laparoscopic surgery, the remote motion center would be positioned to coincide with the point of entry into the patient's body. Similarly, for a frameless stereotaxy application involving multiple biopsies at a single puncture site, the remote motion center would also be positioned to coincide with the puncture site. The distal parts of the robot might then be used to aim a needle guide along multiple biopsy paths. We have also speculated on possible uses of the robot for more open surgeries. In an orthopaedic bone machining application, for example, the instrument carrier could either be replaced by a specialized cutting tool or could be adapted to hold such a tool so that the tip of the cutter was located at the remote motion center.

The range of motion of the present manipulator is ±100 mm for the base x and y translations, ±200 mm for the base z translation, ±60 degrees for the R_x and R_y rotation axes, ±160 degrees for the instrument rotation, R_z, and ±80 mm for instrument insertion, s. The detachable camera head rotation element allows ±160 degrees of rotation of the camera head about the eyepiece of the laparoscope. The instrument carrier (Fig. 2) can be disconnected easily from the robot to facilitate cleaning and to provide a convenient sterile boundary. The instrument carrier is sterilized before surgery and the remainder of the surgical robot is covered with a sterile drape. Interchangeable collets in the instrument carrier accommodate cylindrical instruments (such as laparoscopes) up to 17 mm in diameter.

2. Detail of instrument carrier, showing force sensor: The carrier in mounted to the instrument translation stage by a keyed dovetail and is readily removable for cleaning and sterilization. The force-torque sensor is mounted just proximal to the point of detachment. In the present embodiment, the instrument rotation motor and bearings are not sealed, and gas sterilization would have to be used. However, these components could be redesigned for other, more convenient, sterilization methods.

The entire robot is on lockable casters and can be wheeled up to the operating table. This approach was chosen to provide maximum flexibility in positioning the robot and in allowing it to be easily introduced into and removed from the surgical field. We have also considered alternative designs in which the robot is simply mounted on the operating table rail.

Modularity has been emphasized in both the kinematic structure and the detailed implementation of the manipulator and controller. This approach should make it fairly simple to customize subassemblies as more experience is gained or new requirements emerge. For example, we are already considering design modifications to the four-bar linkage component to reduce bulk, further increase stiffness, and provide adjustability in the lengths of the links.

The robot is designed to be non-backdrivable. All linear axes are driven by dc motors acting through lead screws. The major revolute axes (R_x and R_y) are driven by dc motors acting through a combined harmonic drive and worm gear transmission. One important safety consequence of kinematic decoupling and high reduction drive trains is that only small, low power motors are required and that no axis drive needs to be capable of any faster motion than required for the corresponding task motions. A second safety consequence is that the mechanism will not move when the motors are de-energized. We can absolutely prevent unwanted motion or stop a "run away" situation simply by turning off the power. Furthermore, since joint motions are relatively slow, there is more time available for safety monitoring and appropriate actions (such as shutting off power) should such intervention become necessary. The very high reduction ratio and non-backdrivable transmission elements cause any motion to stop very quickly when power is removed.

One potential difficulty with non-backdrivability is the problem of what to do after a "safety freeze" that occurs while the robot is holding an instrument inserted into the patient. Since the robot will become rigid, rather than floppy as would be the case if backdrivable actuators were used, it will not be possible for the surgeon simply to grasp the robot to withdraw the instrument. Instead, the surgeon would loosen the collet in the instrument carrier and withdraw the instrument, after which the robot can be wheeled away. Alternatively, the entire instrument carrier can be disconnected from the robot using the quick release mechanism provided. One significant advantage of this approach is that it avoids possible damage to the patient caused by the uncontrolled instrument motions, such as can result if the robot simply becomes floppy or continues to move because of inertia after a "safety freeze" is initiated. If additional passive compliance is needed, the most appropriate place to provide it is either in the laparoscopic instrument itself or in the instrument carrier.

The robot has a six degree of freedom force-torque sensor placed just proximal to the instrument carrier, as shown in Fig. 2. This sensor allows the controller to monitor external forces exerted on the instrument during surgery and then take appropriate action (e.g., freeze the robot and issue a warning message to the surgeon) to prevent the robot from exerting excessive force on the patient. The force information provided by the sensor can also be integrated into the motion control law, giving the robot the ability to comply with (i.e., move away from) external forces. This mode can be used to take hold of the instrument and manually guide the robot (by exerting forces against the instrument) into the initial position for surgery or to move it to a different portal during the procedure. We also anticipate future uses of this capability for tissue retraction and similar surgical tasks, although friction on the instrument as it passes through the cannula seal may limit sensitivity. If this becomes a serious problem, additional distal force sensing (e.g., [21]) could easily be interfaced to the controller for greater sensitivity.

Robot Motion Control Subsytem

Low-level motion control, joint servoing, and basic safety monitoring are performed by a fast rack-mounted personal computer equipped with a combination of off-the-shelf and custom interface electronics. Higher level control is performed by an IBM PS/2 workstation connected to the low-level controller through a shared memory interface.

Safcty is a fundamental design goal for the system, and many interfaces are provided to support this requirement. For example, the controller electronic design monitors power supply and cable integrity and anticipates the provision of redundant position encoders on each actuated joint, although such encoders are included only on the R_x and R_y axes of the present (non-human-rated) robot. Both computers, but especially the low-level controller, perform extensive consistency checks to verify system integrity. Other checks are performed by dedicated electronics within the controller itself. If any inconsistency or out-of-tolerance condition is detected, the controller turns off the robot power and initiates appropriate actions to notify the surgeon and application software. Additionally, the power drive electronics incorporate a safety timeout feature as well as "power enable" interlocks. The controller software includes a realtime process that performs consistency checks every 5 ms. If a check fails, the controller can immediately disable manipulator power. If all checks are passed, the controller then re-enables the safety timeout. If the safety timeout is not re-enabled within 10 ms, manipulator power is automatically turned off and appropriate status indicators are set. Our experience with this approach, both in industrial [22, 23] and surgical [24] robots, has shown that it provides a high degree of confidence in basic hardware and software integrity of the control system.

Although our present manipulator design is very well suited for "keyhole" surgeries, we have tried to insulate higher levels of application software from dependency on any particular kinematic structure, to an extent that goes somewhat beyond what is found in a typical industrial robot. Instead of simply specifying desired position goals for the surgical instruments and solving the corresponding kinematic equations, the control software sets up and solves nonlinear optimization problems to most closely achieve a desired instrument-to-patient relationship, subject to task and manipulator design constraints.

Consider a simple camera pointing task, in which the goal is to achieve a particular view of a body organ using a rigid 30 degree angle-of-view laparoscope. In general, this is a six degree-of-freedom task. Unfortunately, the laparoscope is constrained by the cannula, so that only four degrees-of-freedom (three rotations and insertion depth) are available. A fifth rotational degree of freedom may be added by rotation of the camera about the eyepiece of the laparoscope optics. This camera rotation is redundant with instrument rotation if a 0 degree laparoscope is used. However, for angled-view scopes it can be used to rotate the image to maintain some preferred view orientation.

Clearly, trade-offs are necessary, based on what is most important for a particular task. For example, if one is simply aiming the camera for the purpose of viewing the patient's anatomy, one may wish to minimize apparent rotation about the axis-of-view at the expense of some variation in lateral displacement of the image or distance from the end of the laparoscope. On the other hand, if the intent is to project laser energy along the optical path of the laparoscope, then only very small lateral aiming errors can be tolerated, but image rotation may be less important.

It is often necessary to place bounds on the motion of different parts of the robot or surgical instruments and to guarantee that these bounds are rigorously enforced. For example, it may be very important to tell the robot to keep the end of the laparoscope out of the patient's liver. Similarly, we have so far been assuming that no lateral motion of the canula is permitted. If only the most distal four axes of the robot are being used, and the remote center of motion is placed at the canula, this constraint will be met trivially. However, there is a certain amount of "give" in the patient's abdominal wall, and there are some circumstances, such as stereo ranging or precise subsidiary motions for tissue manipulation, in which it would be desirable to use the proximal xyz stage to displace the canula laterally by a small amount, so long as the patient's anatomy is not stretched too far and unmodelled innstrument deflections caused by lateral forces do not interfere with accuracy.

Finally, additional motion capabilities can be added to the robot or instruments. For example, a steerable prism [25] can be added to the laparoscope to vary its angle-of-view. Or the rigid instrument may be replaced by some sort of steerable snake. In such cases, it is important to be able to take advantage of whatever manipulation capabilities exist, without at the same time requiring that substantial software libraries be rewritten.

Our approach, described more fully in [25], is to express the problem of determining manipulator joint positions $\mathbf{q}(t)$ to achieve a desired motion task as a quadratic optimization problem:

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

Such that:

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

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. Similarly, $\mathbf{C}(t)$ and $\mathbf{d}(t)$ express constraints that must be obeyed, again propagated through the kinematic equations of the manipulator. In our present solution method [25], we do not attempt to minimize the integral error, i.e., the value of $\min \|\mathbf{A}(t)\mathbf{q}(t) - \mathbf{b}(t)\|$ integrated over time, t. Instead, we solve the minimization problem for multiple time steps, using linearized expressions for $\mathbf{A}(t)$, $\mathbf{b}(t)$, $\mathbf{C}(t)$, and $\mathbf{d}(t)$. This formulation permits task-step dependent optimization criteria and constraints, such as "minimize image rotation" and "guarantee that the view axis passes within 0.5 mm of the defined target point," to be combined with standing instructions, such as "minimize joint motion" and "guarantee that the remote motion center stays within 3 mm of the canula center." It is also possible to have compound instructions such as "minimize the displacement of the remote motion center from the center of the canula, but in all cases guarantee that the displacement never exceeds 3 mm." Weighting factors are used to specify the relative importance of different optimization criteria. If the constraints cannot all be satisfied, appropriate software exceptions are generated to be handled by higher levels of the application software.

This formulation does not make assumptions about the kinematic structure of the robot or even the number of controllable degrees of freedom available. In cases where redundant degrees of freedom are available, additional optimization criteria (typically, minimization of total joint motion or maximization of available free motion) are used to control how the extra freedom is to be used, subject to constraint satisfaction. In cases where insufficient degrees of freedom are available to force the optimization criterion to zero, the optimizer does the best it can, again subject to constraint satisfaction.

In practice, this scheme has proved to be quite flexible and acceptably efficient, with typical solution rates of 15-20 Hz using a relatively slow (33 MHz '486) IBM PS/2 mod 90. It has been implemented both for kinematically deficient (four-degrees-of freedom) and kinematically redundant (seven-degrees-of-freedom) manipulalator configurations, including systems with somewhat different physical designs from our current robot [25]. Our experience has been that this formulation is indeed very successful in promoting a high degree of functional portability between manipulator designs. For example, functions developed on the highly redundant (six rotations plus instrument insertion) experimental remote center of motion (RCM) manipulator described in [25] were successfully ported to the four degree-of-freedom distal portion of our present robot in just a few days. Furthermore, the tradeoffs made by the optimization software proved to be quite sensible, so that the apparent performance of system functions remained quite acceptable. Similarly, addition of a camera-rotation motor to keep the view upright when rotating an angled-view laparoscope was quite easy. Extension of the paradigm to accommodate another experimental manipulator which used a passive-linkage universal joint to enforce the fulcrum constraint was also relatively straightforward [26].

Human-Machine Interfaces

During laparoscopic procedures, the surgeon's gaze is most often centered on the television monitor displaying the live video image transmitted by the laparoscopic camera. This image is the surgeon's primary feedback in controlling the surgical instruments in relation to the the patient's anatomy. It also is frequently the basis for his or her communication with people assisting in the procedure. If the robotic system is to function as an effective assistant, rather than as a simple teleoperated slave, it is important that it have access to this important information source and communication channel. Consequently, the controller has the ability to capture and extract information from the laparoscopic images and to superimpose simple graphical overlays on the live video images. Typical overlays include cursors, simple graphical displays, icons, and text indicating distances, other quantitative information, and system status. We are also considering, but have yet to implement, a number of other display functions, including peripheral display of patient status information, computer enhanced presentation of the color video signal, registration and overlay of preoperative models.

Similarly, a primary means for the surgeon to instruct the system is by pointing to objects displayed on the video monitor. Although we have demonstrated the ability of the system to track visual markers on the surgical instruments, thus allowing the surgeon to designate anatomical features of interest simply by pointing at them directly [27], in practice it has proved much more convenient for the surgeon to use a mouse or joystick to position a cursor on the display screen. An obvious difficulty is that it can be quite inconvenient for the surgeon to let go of a laparoscopic instrument in order to grasp a conventional pointing device. Foot pedals are an often-suggested alternative, but have mixed popularity with surgeons. Feet are inherently more clumsy than hands for precise tasks, and there are sometimes a number of other foot switches already in use, so that adding one more can be confusing. Our approach has been to provide a small (gas or soak sterilized) joystick device that can be clipped to a laparoscopic instrument and operated without requiring the surgeon to release the instrument. We have evaluated a number of different designs; our current embodiment, shown in Fig. 3, is functionally equivalent to a three button mouse. It has a single TrackPoint (tm) joystick adapted from an IBM ThinkPad (tm) computer, and three push-buttons in a package about 35 mm across. We have also combined three such joysticks into a single surgeon interface that can be gas sterilized or placed inside a sterile drape and clipped to a convenient position in the surgical field.

3. Instrument mounted joystick: The embodiment shown is the functional equivalent of a three button mouse.

Synthesized speech has proved to be extremely useful as a means of providing information and short instructions to the surgeon. On the input side, speech recognition systems are just beginning to be reliable and fast enough to be useful as a "hands off" command interface. In an earlier embodiment of the system [27, 28], we constructed such an interface using an experimental speech recognition system developed at IBM Research. As expected, we learned that speech recognition is clearly the most convenient modality for many surgeon inputs; but that (a) it cannot substitute for pointing in many situations, and (b) recognition accuracy and response time are critical to surgeon acceptance. We are planning to apply these lessons to the present system in the near future, using recent product-level IBM speech recognition systems as the basis.

As with all aspects of the system, we have emphasized modularity in designing these interfaces and, to the extent possible, have tried to insulate application software from detailed dependencies on any particular hardware embodiments or configuration. One obvious advantage of this approach is the ability to take advantage of the rapid evolution of new technology in this field, such as head mounted displays, haptic interfaces, and other "virtual reality" devices, and we have already begun to explore some of these possibilities. Another advantage is that modularity also tends to improve system robustness, both from a software engineering viewpoint and by making it easy to provide redundant interfaces. For example, if a speech synthesizer fails, the same information can be displayed (albeit more annoyingly) as text superimposed on the video monitor.

Operating Modes

Direct Teleoperation

In direct teleoperation, the surgeon interactively controls the motion of the robot by directly commanding individual motions. Perhaps the most direct form is force compliance. The surgeon grasps the laparoscopic instrument and pulls on it; the controller responds to the force/torque values sensed by the force sensor in the robot's "wrist", and moves the robot in the direction that the surgeon is pulling. Two modes are provided: one uses the proximal xyz stage to translate the remote motion center, and the other uses the distal four axes to control instrument orientation and insertion. Although we have yet to implement such a mode, it would also be quite straightforward to implement a remote force controller, in which the surgeon exerts forces on a detached six degree of freedom "force joystick." In this case, the center of motion compliance could be set to produce teleoperation modes analogous either to the "anatomy centered viewpoint" or "viewpoint displacement" modes described below.

In other modes, the surgeon uses the instrument mounted joystick to specify motions of the laparoscope or other instrument held by the robot. When a single joystick is used, one of the push buttons is used to select pairs of motion directions (e.g., "xy", "zR_z", "R_xR_y") to be controlled by the joystick. When multiple joysticks are active, this multiplexing is not needed. We provide two basic joystick controlled modes. In "anatomy centered viewpoint" mode, a particular anatomical feature remains centered in the camera's field-of-view. The sensation to the surgeon looking at the television monitor is one of flying about an imaginary sphere centered on this feature, zooming in and out (i.e., shrinking or enlarging the sphere's radius) or rolling about the camera's axis of view. Most often, the anatomical feature is located by triangulation from a pair of video images, as discussed in the next section. "Viewpoint displacement" mode is used to move the camera to view different parts of the patient's anatomy. In this mode, the sensation is more nearly one of flying through the patient's anatomy.

Vision Guided Operation

The surgeon has the ability to designate anatomical features of interest by pointing at them. As discussed above, the most common pointing means is to use the instrument mounted joystick to control a cursor superimposed on the video display, although other modes are also possible. Once a feature has been designated, the controller can determine the 3-D position of the anatomical feature by image processing. When a monoscopic video source, such as a standard laparoscopic camera is in use, the controller captures one image, moves the robot to displace the camera a small amount perpendicular to the view axis, and acquires a second image. Multiresolution correlation [29] is used to locate the feature in the second image, and the feature's spatial position is computed by triangulation. If a stereo laparoscope is available, then the subsidiary motions may be dispensed with. We are exploring the acquisition of such a laparoscope, and have already demonstrated the use of the image processing software for a simulated biopsy experiment using two standard TV cameras.

Once the feature's position is determined, the controller can easily solve an aiming problem and move the robot so that the feature is centered in the camera's field of view. If desired, additional correlation steps can be performed to "zero in" on the feature, although this hasn't proved to be important in practice. One useful capability, which we have demonstrated, is the ability to designate a viewpoint and save it for later recall. For example, the surgeon may define two or three views of the anatomy being observed, together with views of the entry portals for hand-held surgical instruments. Subsampled video "snapshots" of these views are aligned along the edge of the TV monitor, and the surgeon can at any time return to a stored view by pointing at it with the on-screen cursor and "clicking" a button to select (see Fig. 4).

Guided Autonomy: Assistive Functions

One of the key attributes of a good assistant is the ability to perform simple tasks autonomously, under the general supervision of the surgeon. An important goal for our surgical robot is that it be able to do much the same thing. The system should be able to perform a simple task without requiring detailed control by the surgeon. In fact, vision guided camera pointing is one example of such a function. The surgeon simply designates the anatomical feature to be viewed, and the robot automatically centers the feature. In fact, in the case of angled-view laparoscopes, the robot can usually do a better job than can an average human assistant, since the controller is not confused by coordinate transformations and the robot is both more accurate and more steady than a human.

4. In-vivo video display with superimposed control menus: shows typical video display 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 viewing positions. Typically, the surgeon would select desired functions or robot positions by using of the instrument mounted joystick to position a cursor over the desired menu item and then "clicking" a button. In some modes (e.g., "pan") pushing on the joystick causes the robot to shift viewpoint seen through the camera.

We have begun to explore applications in which the robot positions a surgical instrument, rather than a simple diagnostic laparoscope. In many of these applications, the robot positions a therapeutic laparoscope so that a surgical instrument inserted into the working channel will be accurately placed on a particular anatomical feature. One such example is shown in Fig. 5. In this example, a small pellet represents a gall stone that has spilled out of a broken gall bladder during a cholecystectomy, and must be retrieved. The surgeon selects "go to" mode from a menu by pushing a button; the controller uses the speech synthesizer to inform the surgeon that it is in "go to" mode and asks the surgeon to designate the feature to be grabbed (in this case, the pellet). The surgeon uses the instrument mounted joystick to point at the pellet and pushes a button. Then, the controller acquires a stereopair of images, locates the anatomical feature, and shows the surgeon where it thinks the feature is. The controller then uses the speech synthesizer to ask the surgeon to confirm that it has located the feature correctly and waits for permission to move the robot. After the surgeon confirms the desired motion by pushing a button, the controller moves the robot so that the laparoscope's working channel is properly aligned with and at the correct "standoff" distance from the pellet. The surgeon then inserts an appropriate tool through the working channel and grasps the pellet. In the future, we anticipate extending this capability to a number of different assistive tasks, such as biopsy sampling, multiple drug injections, retraction, hemostasis, and suturing. Such assistive capabilities tend to follow a common general paradigm. The surgeon will select a specific action to be performed and will designate the appropriate anatomical target. The system will accurately locate the designated target, obtain confirmation if needed, and maneuver the instrumentation into position, often performing additional subsidiary sensing and control. It will then perform the desired task under the surgeon's general supervision, again often performing additional sensing and control steps on its own, within constraints determined for the task. It should be noted that this paradigm has many potential advantages for remote surgery applications, in which delays can make simple teleoperation impractical. An additional extension is the incorporation of anatomical models obtained from preoperative imaging, such as CT or MRI, or from other intraoperative modalities such as fluoroscopy or ultrasound. One of the key advan-

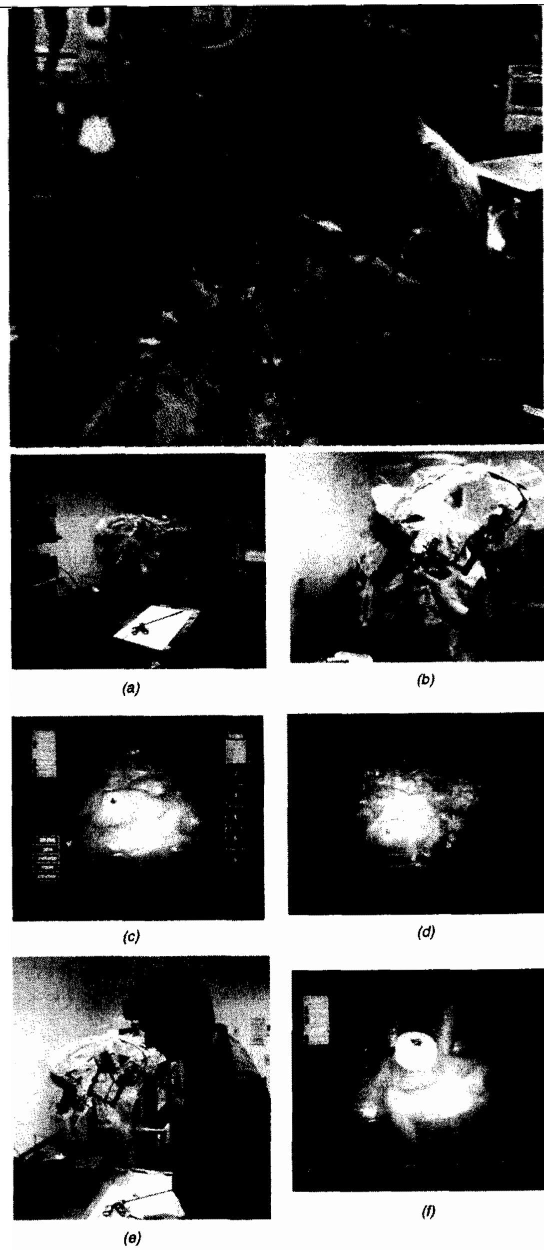

5. In-vitro demonstration of point and grab application: (a) experimental setup, consisting of the surgical robot holding a Storz theraputic laparoscope with a 6 mm working channel, a rubber simulation of patient anatomy, and a small target to be grasped by a surgical instrument inserted into the working channel of the laparoscope. The robot is draped as it would be in surgery. (b) force compliant manual guiding of the robot. The robot enters this mode whenever the surgeon depresses two buttons on opposite sides of the instrument carrier. (c) display monitor after the surgeon has designated the target using the instrument mounted joystick to place cursor crosshairs on the image of the target. (d) scene just after the computer has located the target by multiresolution correlation. This view shows the correllation window tree. Normally, this display is used for debugging and would be suppressed in production use. (e) insertion of the instrument into the working channel. (f) the scene during the pickup operation. The pellet appears to be off-center, but is lined up with the working channel of the scope.

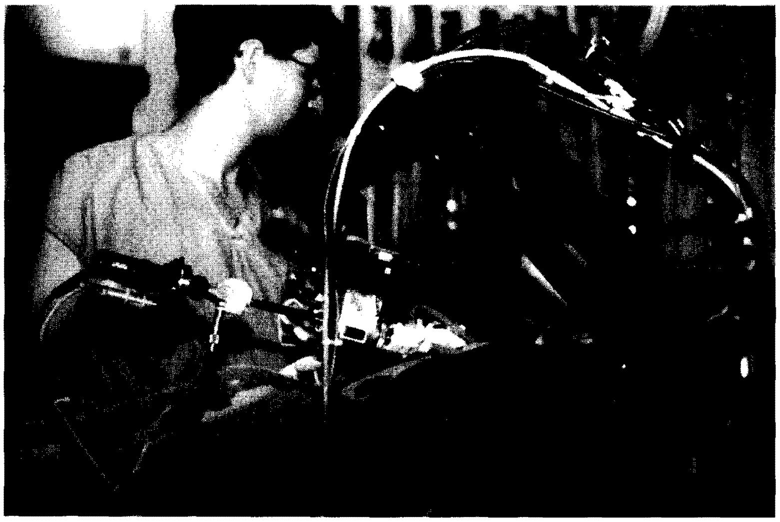

6. Use of LARS system in surgery: Shows early cadaver evaluation of LARS robot to hold a laparoscopic camera. The system has also been used in in-vivo evaluations on pigs, using approved protocols following all applicable University, US Government, and IBM animal care and use guidelines. F●r sterility, the robot would be covered with a sterile drape in normal clinical use.

ages of the robot, relative to a human, is that it is very accurate and stable. This makes it a natural candidate for brachytherapy, biopsies, and other "frameless stereotaxy" applications. Again, we are beginning to explore such applications.

Status

At the present time, the prototype system described here is fully functional and performing well in our laboratory at IBM Research. A second system (Fig. 6) has been installed at Johns Hopkins University Medical School, where in-vivo preclinical testing has begun. Our collaborating surgeons, Dr. Mark Talamini and Dr. Louis Kavoussi, have successfully used the system to perform both laparascopic cholecystectomies and nephrectomies. Initial surgical feedback has been very positive, and we are beginning to consider additional ways to exploit the precise positioning and image guidance capabilities of the system.

Summary and Conclusion

We have described a robotic system designed to function as an intelligent "third hand" in laparoscopic and other general surgical procedures. The system includes a specially designed robot, a variety of human-machine interfaces, image processing capabilities, and a modular controller that supports a number of operating modes. Preliminary experience with the system indicates that it is capable of easy and intuitive navigation of arbitrarily-angled laparoscopic telescopes inside a patient, reliable extraction of 3-D information from intraoperative images, and safe and accurate positioning of surgical instruments relative to patient anatomy. The user interface has proved to be sufficiently powerful to allow convenient access to all system functions and sufficiently intuitive to allow novice users to learn quickly to operate the system effectively. Although considerable work remains to be done, our early experience with the system prototype and the feedback from the surgeons are very encouraging.

Acknowledgments

We wish to thank a number of people both at IBM Research and elsewhere who contributed substantially to this work. Dr. Michael Treat of Columbia Presbyterian Medical Center suggested the initial application of laparoscopic camera pointing to one of the authors (Taylor) in 1989 or 1990 and provided useful input and feedback during early phases of this work (e.g., [30]). David Grossman and John Karidis were key participants in many conceptual design discussions for the remote center of motion manipulator, and Jerry McVicker contributed early designs of key components. Bob Lipori, Jay Hammershoy, Bob Krull, and other members of IBM Research's Central Scientific Services (CSS) built the prototype manipulator, which, notably, worked as soon as it was assembled and wired up. Bob Olyha, of CSS, developed key electronic components and interfaces, and so should share the credit for the unusual ease with which the robot was debugged. Nils Bruun and Dieter Grimm and other CSS members designed and built an earlier remote center of motion robot with a very different mechanical design that nevertheless proved very useful for software and control system development. MIT Coop students Nick Swarup and John DeSouza contributed to various aspects of human-machine interface and software development. Ted Selker and Joe Rutledge of IBM Research contributed early Track-Point(tm) prototypes for use in developing the instrument mounted joystick. Finally, we owe special thanks to Mr. John Tesar of Karl Storz Endoscopy, US, who provided the laparoscopic cameras and ins●ruments used in developing the system.

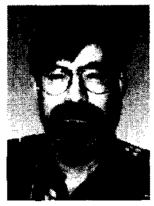

Russell H. Taylor received a B.E.S. degree from Johns Hopkins University in 1970 and a Ph.D. in Computer Science from Stanford in 1976. He joined IBM Research in 1976, where he developed the AML language. Following a two year assignment in Boca Raton, he managed robotics research activities at IBM Research from 1982 until returning to full time technical work in late 1988. Since March 1990, he has been manager of Computer Assisted Surgery. 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. 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 can be reached at IBM, Thomas J. Watson Research Center, PO Box 704,Yorktown Heights, NY 10598. e-mail: rht@watson.ibm.com

Janez Funda received a B.A. degree in Computer Science and Mathematics from Macalester College in 1986, and a Ph.D. degree in Computer Science from the University of Pennsylvania in 1991. He joined the Computer Assisted Surgery group at IBM Research in 1991. His research interests include robot systems, telemanipulation, human-

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 US and international patents.

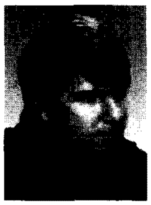

Benjamin Eldridge received an M.S. degree in Physics from Rennsalaer Polytechnic Institute in Troy NY. His specialty is instru mentation development and integration. He has published 19 papers, and holds 1 patent. In 1993 and 1994, he was a member of the Computer Assisted Surgery Group at IBM Research, where he worked on electromechanical and electronic design and implementation of robotic devices for surgery.

Stephen Gomory is a Senior Associate Programmer at IBM T. J. Watson Research Center. He received his M.S. degree in Computer Science from NYU in 1994, and wrote his M.S. thesis on computer vision. His activities include design and implementation of image processing methods and human-machine interfaces for image-guided surgery applications.

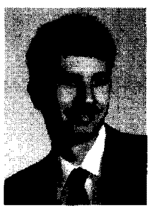

Kreg G. Gruben received a Ph.D. in biomedical engineering from the Johns Hopkins University, School of Medicine in 1993. From 1993-4 he was a postdoctoral fellow in the Johns Hopkins Department of Radiology where, in collaboration with IBM Research, he worked on the development and testing of surgical robots. He is currently an Assistant Professor of Kinesiology at the University of Wisconsin- Madison. His current research interests include the biomechanics of human motion.

David LaRose holds a Bacheolor's degree in Cognitive Science from Brown University, and is currently enrolled as a graduate student at the Carnegie Mellon University Department of Electrical and Computer Engineering, where he is studying robotics and Computer vision. From 1990 to 1993, he was a member of the Intelligent Robotics and Computer Assisted Surgery Groups at IBM Research, where he worked on development of autonomous mobile robots and design and implementation of surgical augmentation systems.

Dr. Mark Talamini is an Assistant Professor of Surgery at The Johns Hopkins University School of Medicine in Baltimore Maryland. He received a Bachelor of Arts in Natural Sciences from the Johns Hopkins University, and his M.D. degree from the Johns Hopkins University School of Medicine. He completed his surgical residency at the Johns Hopkins Hospital and is a fellow of the American College of Surgeons. He is the director of minimally invasive surgery at the Johns Hopkins Hospital, as well as an Attending Surgeon there. His research interests include advanced minimally invasive surgery, robotic surgery, and advanced imaging. He serves on the Editorial Board of Surgical Laparoscopy and Endoscopy.

James H. Anderson is Professor and Director of Diagnostic Radiology Research at the Johns Hopkins University School of Medicine in Baltimore, MD. Dr. Anderson received his Ph.D. degree in Physiology from the Universidty of Illinois at Champaign- Urbana. His primary areas of research intrerest include Interventional Radiology, Medical Imaging, Minimally Invasive and Image Guided Robotic Assisted Therapy.

Dr. Louis Kavoussi is Chief of Urology at the Johns Hopkins Bayview Medical Center and Associate Professor of Urology at the Johns Hopkins Medical Institutions. He attended medical school at the State University of New York in Buffalo and did his resident training at Washington University in St. Louis. Within the field of endourology, and he has helped create many urological laparoscopic procedures, including laparoscopic nephrectomy. His other current interests include development and application of telerobotic surgical systems.

References

1. **PJ Kelly, BA Kall, S Goerss and F Earnest:** "Computer-assisted stereotaxic laser resection of intra-axial brain neoplasms," *J. Neurosurg*, pp. 427-439, March 1986.

2. **AL Benabib, P Cinquin, S Lavallee, JF LeBas, J Demongeot and J de Rougemont:** "Computer-driven robot for stereotactic surgery connected to CT scan and magnetic resonance imaging," *Proc. of American Society for Stereotactic Functional Neurosurgery*, pp. 153-154, 1987.

3. **YS Kwoh, J Hou, E Jonckheere and S Hayati:** "A robot with improved absolute positioning accuracy for CT guided stereotactic surgery," *IEEE Transactions on Biomedical Engineering*, pp. 153-161, February 1988.

4. **MA Lewis and GA Bekey:** "Automation and Robotics in Neurosurgery: Prospects and Problems," in Michael LJ Apuzzo, MD, editor, *Neurosurgery for the Third Millenium*, American Association of Neurological Surgeons, Spring 1992.

5. **H Paul, B Mittelstadt, W Bargar, P Kazanzides, B Williamson, J. Zuhars, et al:** "Accuracy of implant interface preparation: hand-held broach vs. robot machine tool," *Proc. Orthopaedic Research Society*, Washington, D. C., 1992.

6. **TC Kienzle, SD Stulberg, M Peshkin, A Quaid, and C Wu:** "An integrated CAD-robotics system for total knee replacement surgery," *Proc. 1993 IEEE Conference on Robotics and Automation*, pp. 889-894, Atlanta, May 1993.

7. **K Ikuta, M Tsukamoto and S Hirose:** "Shape memory alloy servo actuator system with electric resistance feedback and application for active endoscopes," *IEEE Robotics and Automation Conference*, 1988.

8. **S Charles, RE Williams, and B Hammel:** "Design of a surgeon-machine interface for teleoperated microsurgery," *Proc.IEEE EMBS Conference*, pp. 883-884, 1989.

9. "Voice controlled flexible endoscope," videotape, R. Sturges, Carnegie Mellon University, 1989.

10. **BL Davies, RD Hibberd, A Timoney and J Wickham:** "A surgeon robot for prostatectomies," *Proc. Fifth International Conference on Advanced Robotics*, pp. 871-875, Pisa, June 1991.

11. **P Green:** "Telepresence surgery," *NSF Workshop on Computer Assisted Surgery*, Washington, DC, February 1993.

12. **P Green:** "Advanced teleoperator technology for enhanced minimally invasive surgery," *Proc. Medicine Meets Virtual Reality Conference*, San Diego, June 4-7 1992.

13. **Y Wang:** Automated endoscopic system for optimal positioning, *Computer Motion*, Inc., 1993. Company advertising brochure.

14. **Y Wang:** "Robotically enhanced surgery," *Proc. Medicine Meets Virtual Reality II*, San Diego, Jan 27-30 1994.

15. World Medical Device and Diagnostic News, August 5, 1992.

16. **RH Sturges and S Laowattana:** "A flexible tendon-controlled device for endoscopy," *Proc IEEE Robotics and Automation Conference*, pp. 2582-2591, Sacramento, 1991.

17. **B Neisius, P Dautzenberg and R Trapp:** "Robotic manipulator for endoscopic handling of surgical effectors and cameras," *Proc. Medical Robotics and Computer Assisted Surgery,* Pittsburgh, Sept. 22-24 1994.

18. **JA McEwen:** "Solo surgery with automated positioning platforms," *Proc. of NSF Workshop on Computer Assisted Surgery*, Washington, D.C., Feb. 1993.

19. **JB Petelin:** "Computer assisted surgical instrument control," *Proc. Medicine Meets Virtual Reality II*, pp. 170-173, San Diego, January 1994.

20. **R Hurteau, S DeSantis, E Begin and M Gagnier:** "Laparoscopic surgery assisted by a robotic cameraman: concept and experimental results," *Proc. 1994 IEEE Conference on Robotics and Automation*, pp. 2286-2289, San Diego, May 8-13 1994.

21. **WJ Peine, JS Son and RD Howe:** "A palpation system for artery localization in laparoscopic surgery," *Proc. 1st Int'l Symposium on Medical Robotics and Computer Assisted Surgery*, pp. 250-257, Pittsburgh, Sept. 1994.

22. **RH Taylor and DD Grossman:** "An integrated robot systems architecture," *IEEE Proceedings*, July 1983.

23. IBM Corporation, AML/2 Manufacturing Control System User's Guide, 1986.

24. **RH Taylor, HA Paul, P Kazanzides, BD Mittelstadt, W Hanson, et al:** "Taming the bull: safety in a precise surgical robot," *Proc. 1991 Int. Conference on Advanced Robotics*, Pisa, Italy, June 1991.

25. **J Funda, R Taylor, K Gruben and D LaRose:** "Optimal motion control for teleoperated surgical robots," *Proc. 1993 SPIE Intl. Symp. on Optical Tools for Manuf. & Adv. Autom.*, Boston, September 1993.

26. **J Funda, B Eldridge, K Gruben, S Gomory and R Taylor:** "Comparison of two manipulator designs for laparoscopic surgery," *Proc. 1994 SPIE Int'l Symposium on Optical Tools for Manufacturing and Advanced Automation*, Boston, October 1994.

27. **RH Taylor, J Funda, D LaRose and M Treat:** "An experimental system for computer assisted endoscopic surgery," *Proc. IEEE Satellite Symposium on Neurosciences*, Lyons, November 1992.

28. Computer Assisted Surgery at IBM Research. Videotape showing excerpts of past and current work on computer assisted surgery at IBM T. J. Watson Research Center. Queries should be directed to Russell H. Taylor, Manager of Computer Assisted Surgery Research, IBM T. J. Watson Center.

29. **HP Moravec:** Obstacle Avoidance and Navigation in the Real World by a Seeing Robot Rover, Ph.D. thesis, Stanford University, 1980.

30. **RH Taylor, J Funda, D LaRose and M Treat:** "A telerobotic system for augmentation of endoscopic surgery," *Proc. 14th IEEE Medicine & Biology Conf.*, Paris, November 1992.