



Initial clinical experience with a partly autonomous robotic surgical instrument server

M. R. Treat,¹ S. E. Amory,¹ P. E. Downey,² D. A. Taliaferro³

¹ Department of Surgery, Columbia University, New York, NY 10034

² Robotic Surgical Tech, Inc., Clinical Coordination, Bronx, NY 10454

³ New York-Presbyterian Hospital, Nursing, New York, NY 10032

Received: 14 April 2005/Accepted: 23 February 2006/Online publication: 17 May 2006

Abstract

Background: The authors believe it would be useful to have surgical robots capable of some degree of autonomous action in cooperation with the human members of a surgical team. They believe that a starting point for such development would be a system for delivering and retrieving instruments during a surgical procedure.

Methods: The described robot delivers instruments to the surgeon and retrieves the instruments when they are no longer being used. Voice recognition software takes in requests from the surgeon. A mechanical arm with a gripper is used to handle the instruments. Machine-vision cameras locate the instruments after the surgeon puts them down. Artificial intelligence software makes decisions about the best response to the surgeon's requests.

Results: A robot was successfully used in surgery for the first time June 16, 2005. The operation involved excision of a benign lipoma. The procedure lasted 31 min, during which time the robot performed 16 instrument deliveries and 13 instrument returns with no significant errors. The average time between verbal request and delivery of an instrument was 12.4 s.

Conclusions: The described robot is capable of delivering instruments to a surgeon at command and can retrieve them independently using machine vision. This robot, termed a "surgical instrument server," represents a new class of information-processing machines that will relieve the operating room team of repetitive tasks and allow the members to focus more attention on the patient.

Key words: Artificial intelligence — Autonomous — Information system — Machine vision — Robot — Surgical

Computer-assisted surgical robots are being used currently in both open and minimal access surgery. The computer-assisted surgical devices available currently are able to mimic the surgeon's motions dexterously (daVinci System; Intuitive Surgical Inc., Sunnyvale, CA, USA) or follow voice commands (AESOP; formerly Computer Motion, Inc., Goleta, CA, USA). However, they are not capable of autonomously planning and executing motions. We believe it would be useful to have surgical robots capable of some degree of autonomous action. The robot's limited autonomy would be in the context of cooperative action with a human surgical team.

We also believe that a good starting point for developing such a machine would be a system that could perform the tasks of delivering, retrieving, sorting, and counting surgical instruments. Difficulties associated with the counting of instruments are a matter of clinical concern [3]. We term this sort of machine a "surgical instrument server." The robot's autonomous motions would be based on its independent appraisal of sensory data relating to the positions of instruments and the requests of the surgeon.

We hypothesized that readily available technologies such as voice recognition, speech synthesis, machine vision, and standard artificial intelligence techniques could be mapped to the clinical requirements of a partly autonomous surgical instrument server. An artificially intelligent predictive ability based on the recognition of surgical instrument usage patterns would allow the machine to anticipate the needs of the operating team and make sure that the right instruments are quickly available for immediate delivery.

The hardware and software architecture of such an instrument handling system could be used in settings other than the conventional open surgery setting, in which we first tested it. For example, this sort of autonomous surgical instrument server would be able to assist a human-driven surgical teleoperator so that surgery could be performed in remote or hazardous

settings without the physical presence of any humans. It would be useful in military and aerospace environments. It also would be a significant step toward the development of an autonomous robotic surgical first assistant. Practitioners of complex laparoscopic procedures have recognized the significance of such a robotic assistant that could intelligently facilitate the performance of procedures requiring the cooperative and highly coordinated interaction of two or more operators.

Materials and methods

The robot, named the Penelope Surgical Instrument Server, is based on a straightforward application of the technologies of voice recognition, speech synthesis, machine vision, a pick-and-place robotic arm, and standard artificial intelligence techniques. As described later, some of the software and hardware implementations of these technologies are commercially available off-the-shelf products. Others were made specifically for this application by Robotic Surgical Tech (New York, NY, USA), the company responsible for the overall design and production of the robot.

The speech recognition system is the standard voice recognition system provided with the off-the-shelf computer operating system we are using (Mac OS X; Apple Computer, Inc., Cupertino CA, USA). This voice recognizer is speaker independent. It recognizes a prescribed limited set of names for the instruments including several common synonyms (e.g., “mosquito clamp” vs “snap”).

The machine-vision system is based on known techniques [5] tailored specifically for this application. It uses two off-the-shelf IEEE 1394 (FireWire) digital cameras housed in the camera post. These cameras look down at two instrument-carrying surfaces, known as the instrument tray and the transfer zone. The cameras locate and identify all the surgical instruments on the instrument tray and in the transfer zone.

These two instrument-carrying surfaces are a key physical feature of the robot’s design. The instrument tray, like the Mayo stand of the scrub nurse, contains an assortment of instruments needed for the case. The transfer zone functions as the part of the surgical field near the actual incision, where instruments are laid down by the surgeon for the scrub nurse to retrieve. The robot can retrieve instruments left in the transfer zone and return them to the instrument tray. The transfer zone also can be used as a ready cache of one or two “favorite” instruments that can be grabbed quickly by the surgeon without the need to go through the process of requesting them from the scrub person.

The robotic arm was custom-made for this application. The arm can rotate and bend at its “shoulder,” “elbow,” and “wrist.” An electromagnet on the end of the arm is used to lift and carry the instruments. The electromagnet has been found to work with most of the standard stainless steel instruments in the operating rooms of the NewYork-Presbyterian Hospital. However, some instruments are made from a type of stainless steel that is not very magnetizable. Other instruments are plastic, and therefore not magnetizable at all. These instruments are modified by attaching a magnetizable steel band to which the robot’s electromagnet can readily attach.

The artificial intelligence software that controls the robot also was developed by Robotic Surgical Tech. The artificial intelligence system is patterned after well-known types of “cognitive architecture.” A cognitive architecture is a description of a mind. The design of a cognitive architecture specifies how the sensory and reasoning sub-systems of the mind work together to provide for intelligent and purposeful information processing. Many cognitive architectures of widely varying design [1] have been created during the past 20 years, primarily for the benefit of two fields: cognitive psychology and artificial intelligence.

The cognitive architecture for the described robot is implemented in software as a rule-based inference engine [4]. Basically, the cognitive architecture constantly examines inputs produced by the various sensors, both external (vision and speech sensors) and internal (arm position sensors). The rules upon which the cognitive architecture acts are “IF-THEN” statements. Here is a simple example: “IF the speech sensor asserts that a request for a Hopkins clamp is made, THEN deliver the Hopkins clamp.” When the cognitive architecture finds that




Fig. 1. The major physical components of the robot are shown. The system stand is designed to straddle the operating room table, over the patient, at the foot of the table.

a given rule’s IF statement is satisfied, the cognitive architecture “fires” that rule, thereby activating the THEN statement of the rule. The THEN statement produces output commands to the speech and/or arm actuators. The cognitive architecture of the robot uses many rules (currently 86). Each rule itself is simple, but the interaction of all the rules can be quite complex.

Our claim that the current robot is partly autonomous is based on the fact that control of the robot’s detailed actions results from decisions made by the cognitive architecture and is not the result of direct human input. The robot has certain programmed assigned tasks such as instrument delivery, counting of instruments, and arranging of instruments on the instrument tray. It of course performs these tasks, but the details of how and sometimes when each task is executed are autonomously determined by the cognitive architecture.

For example, after receiving a request for an instrument, the robot decides the fastest way to deliver that instrument. Usually, the instrument will be on the instrument tray. However, at times, it may be quicker for the robot to pick up an instrument that may already be out in the transfer zone. The cognitive architecture makes that decision on its own.

Another example is the performance of background tasks such as tidying up the instruments on the instrument tray. The robot autonomously decides when the instrument tray needs rearranging and also figures out the best way to move the instruments to achieve this goal. Also, the robot must be able to gracefully interrupt a background task, such as rearranging instruments to honor a new instrument request, then resume the background task without confusion or error.

Another example of partly autonomous behavior can be seen when the surgeon requests an instrument, but then ignores it when it is offered by the scrub nurse. This scenario happens not uncommonly in the operating room, usually because the surgeon is not finished performing the previous step. If the robot finds itself in that situation, the cognitive architecture has to decide whether the surgeon really will want the instrument in the next few minutes or perhaps not want it at all. If the robot decides that the instrument will be needed in the immediate future, it may continue to hold it in its gripper, or it may put it down in the transfer zone, where it can be grabbed quickly by the robot or the surgeon when needed. If the robot decides that the surgeon really is not going to want the instrument after all, it then returns that instrument to the instrument tray.

The cognitive architecture uses a small database of previous instrument requests to determine statistically the likeliest next instru-

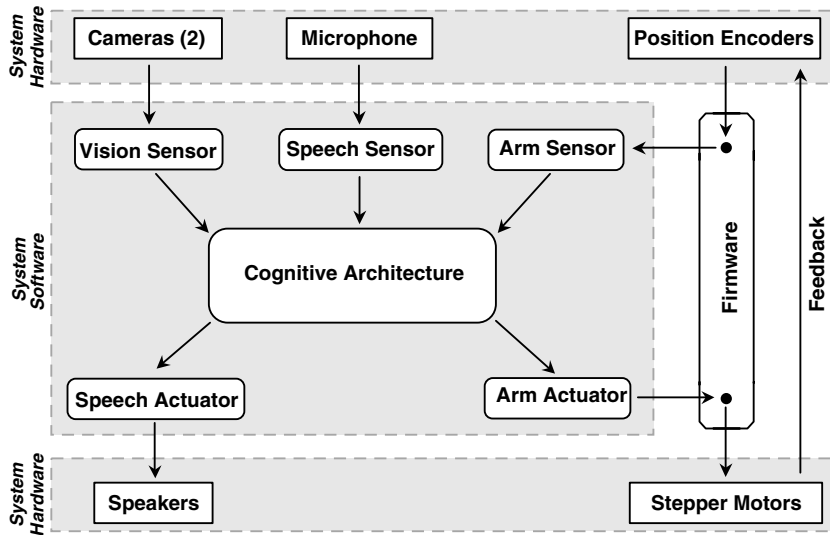


Fig. 2. A diagram of the robot’s software system is shown. The key feature is that all of the robot’s sensor inputs are fed into a central artificial intelligence software construct, known as the cognitive architecture. The cognitive architecture is responsible for producing all outputs to the actuators.




Fig. 3. Layout of the nine surgical instruments used for the robot’s first case. From left to right are the Richardson retractor, needleholder, Mosquito clamp, Hopkins clamp, Metzenbaum scissors, tooth forceps, Allis clamp, and Brown-Adson forceps.

ment requests. If an requested instrument is temporarily ignored but has a high statistical probability of being needed next, the cognitive architecture will decide that sooner or later that instrument will be needed and will continue to keep it readily available.

The classic example of this scenario is seen when the surgeon first requests a needleholder to do some suturing. The surgeon may next request a suture scissors, but then may appear to ignore the proffered scissors because he or she is still tying the last knots of the suture. The cognitive architecture will know it is highly probable that the suture scissors will be needed eventually. On the basis of that “realization,” the cognitive architecture will not return the suture scissors to the instrument tray, but will continue to hold it out for the surgeon to take, unless the cognitive architecture determines that it has other things that need doing, in which case it will leave the suture scissors in the transfer zone for the surgeon to grab when he or she is ready. This is quite analogous to the actions of an experienced scrub nurse or scrub technician.

The physical system is shown in Fig. 1, the software system in Fig. 2, and the layout of instruments as used in the case in Fig. 3.

For this version of the robot, the instruments are laid out by the scrub nurse. They are arranged in prescribed positions, but it is not necessary for the instruments to be set up in precise positions. Once the robot is activated, it takes a picture of the instrument tray, then calculates and remembers the exact positions of the instruments.

Table 1. Chronological sequence of instrument requests

Request no.	Instrument requested
1	Hopkins clamp
2	Tooth forceps
3	Brown-Adson forceps
4	Metzenbaum scissors (“Metz”)
5	Allis clamp
6	Hopkins clamp
7	Mosquito clamp (“snap”)
8	Brown-Adson forceps
9	Needleholder
10	Brown-Adson forceps
11	Metzenbaum scissors
12	Richardson retractor
13	Richardson retractor
14	Richardson retractor
15	Richardson retractor
16	Mosquito clamp (“snap”)

The basic nominal operation of the robot consists of three main steps: (1) the surgeon makes a verbal request for an instrument via a headset microphone; (2) the robot delivers the instrument to the surgeon; (3) when finished with the instrument, the surgeon places it down in the transfer zone. The robot then retrieves the instrument, identifies it, and replaces it in the correct position on the instrument tray.

Approval for the use of the described robot was obtained from the Institutional Review Board of Columbia University Medical Center and the NewYork-Presbyterian Hospital.

Results

On June 16, 2005 at the Allen Pavilion of the NewYork-Presbyterian Hospital, the Penelope Surgical Instrument Server assisted attending surgeon Dr. Spencer E. Amory, scrub nurse Doreen Taliaferro, and circulating nurse Dilcia Burgos-McCollum in excision of a benign lipoma from the right dorsal forearm of a 43-year-old woman. Table 1 shows the chronological sequence of the instruments requested during the case.

Table 2 shows the number of instrument deliveries and returns attempted by the robot, and the number of times these attempts were correctly executed. A correct

Table 2. Instrument delivery and returns

Total number of instrument deliveries to surgeon	16
Number of correct instrument deliveries to surgeon	16
Total number of instrument returns from transfer zone	13
Number of correct instrument returns from transfer zone	13

Table 3. Instrument delivery times

Average time	12.4 s
Range of times	9.5–16.5 s

instrument delivery means the robot correctly identified and physically delivered the instrument without dropping it. The robot achieved 100% accuracy for instrument delivery. However, the 16 instrument deliveries required 25 verbalized requests, meaning that 64% of the time, the instrument was delivered at the first request. The remainder of the deliveries required two or three repetitions of the verbal request until the voice recognition system was able to understand it.

All 13 of the instrument returns from the transfer zone to the instrument tray that the robot attempted were accomplished successfully, for a 100% success rate. However, in 2 (15%) of the 13 instrument retrievals, the instrument, although returned to the tray in the correct spot, was not correctly oriented. There was one other instance (8%) in which an instrument was returned to the wrong spot on the instrument tray. The errors were attributable to software bugs that will be resolved. No instruments were dropped. Most importantly, no instruments were misidentified during the case or missed by the machine-vision software. Three instruments were left out in the transfer zone when the case was deemed over and the robot was turned off.

Instrument delivery time was recorded as the time from a verbal request to the delivery of the instrument at the handoff point (Table 3). The robot's electromagnetic gripper waits at the handoff point for the surgeon to place his hand under the instrument and push slightly upward. This upward push registers with a "bump sensor" in the gripper and informs the software that it is time to turn off the electromagnet, and thereby release the instrument to the surgeon.

Discussion

This case report describes the first time a partly autonomous, machine-vision-guided robot has served as a surgical assistant. The machine has two important and innovative features. The first is the use of machine vision. Machine vision enables the robot to have visual sensory input as to the number, type, and location of the instruments requested and used in the operation. The second feature is the artificial intelligence software that controls the robot. This intelligence means that the robot is not just a "vending machine," but that it can keep track of goals and perform elementary reasoning to ensure that these goals are fulfilled. This goal setting is very important for error recovery. The sense of vision

allows the robot to check whether the goal of delivering an instrument was actually fulfilled. If the vision system sees that the goal was not fulfilled, the software intelligence works backward from that unfulfilled goal and creates subgoals enabling the robot to attempt to figure out what went wrong.

For example, suppose that the robot is trying to deliver an instrument, but its arm is accidentally jostled and the instrument is dropped en route. When the robot sets out to deliver an instrument, it creates an internal goal in its artificial intelligence software. If the delivery goal is not fulfilled, the robot can create a subgoal directing the vision system to examine the entire path the arm traveled during the delivery to see whether the missing instrument can be found along that route. At the very least, if the robot cannot locate that instrument (perhaps if the instrument has fallen off the table), the robot can use its voice to ask the circulating nurse for help or for a replacement instrument.

A key clinical benefit that follows from the combined use of machine vision and artificial intelligence is that the robot can keep track of the instruments requested, used, and returned. This feature may be important for eliminating the difficulties surrounding instrument counts and the potential for leaving instruments behind in the patient. These difficulties have been recognized by members of the Association of periOperative Registered Nurses (AORN):

Some nurses have voiced concerns that the culture in their clinical setting places a higher priority on efficiency and decreasing turnover times than on counting. Some nurses mention that addressing unresolved counts with certain practitioners can be problematic, and others say time pressures and instrument complexity. Reasons for these medical errors remain unclear, but sponge, sharp, and instrument counting remains an error-prone process that often results in pain, disability, and another surgical procedure for patients. These errors also result in costs to the affected patients, clinicians, and health care systems [2].

Currently, the robot is counting only the actual instruments, but the machine-vision system is being extended to keep track of sponges and sutures as well. Although the described robot is physically designed to handle instruments for open surgeries, it will be modified to handle laparoscopic instruments also.

The robot's actual delivery time (average, 12.4 s) for surgical instruments was longer than a human scrub nurse or technician requires under ideal circumstances. We have informally measured the speed of instrument delivery by humans, and have found it to be as fast as 1.5 s. On the other hand, sometimes the speed of human instrument delivery is quite a bit longer, if the scrub person is occupied with other tasks, does not have the requested instrument readily available, or simply is not paying close attention to what the surgeon is doing. The robot's delivery time will improve with refinement of the software and hardware. The robot arm currently moves rather slowly, but as we gain more experience with it, we can increase the physical speed of motion. However, for safety reasons, we do not wish to increase the speed as much as might be mechanically possible. Our planned program of development is aimed at producing a robot

that not only moves (somewhat) faster, but also moves smarter (i.e., has the next instruments very close at hand for speedier handoff to the surgeon). The “prediction engine” software we have written will provide the basis for this smarter movement.

In summary, we believe that the described machine represents a useful and innovative application of technology to perform some of the more mechanical and quantitative tasks found in the operating room. It does not attempt to replace any member of the operating room team, but rather aims to be a helper to them, particularly the scrub nurse. Eventually, the machine may improve patient care by reducing the error rate of certain quantitative tasks such as the counting of instruments. We realize that there are many improvements that can be made, particularly with regard to instrument delivery speed and accuracy of the voice recognition system. We believe that the basic architecture of the machine can be further refined and extended to perform many useful tasks in the operating room.

Acknowledgments. The authors gratefully acknowledge the support of the National Science Foundation Small Business Innovation Research Program, the Telemedicine and Advanced Technologies Re-

search Center of the United States Army Medical Research and Materiel Command, and the NewYork-Presbyterian Hospital. Additional support for work related to this project was received from the Defense Advanced Research Projects Agency, the National Institute of Nursing Research of the National Institutes of Health and the Center for Advanced Information Management at Columbia University, and the New York State Office of Science, Technology & Academic Research. Robotic Surgical Tech, Inc. is a company founded and partly owned by Michael R. Treat, MD. Patrice E. Downey is an employee of Robotic Surgical Tech, Inc. who has a financial interest in the company. Spencer E. Amory, MD, and Doreen A. Taliaferro have no financial arrangements with or financial interest in Robotic Surgical Tech, Inc.

References

1. A Survey of cognitive and agent architectures. Retrieved xx from <http://ai.eecs.umich.edu/cogarch0/index.html>
2. Beyea SC (2003) Counting instruments and sponges. *AORN J* 78: 290–294
3. Gawande AA, Studdert DM, Orav EJ, Brennan TA, Zinner MJ (2003) Risk factors for retained instruments and sponges after surgery. *N Engl J Med* 348: 229–235
4. Laird JE, Newell A, Rosenbloom PS (1987) Soar: an architecture for general intelligence. *Artificial Intelligence* 33: 1–64
5. Whelan PF, Molloy D (2001) *Machine vision algorithms in Java: techniques and implementation*. Springer-Verlag, London, Berlin, Heidelberg