SPECIAL ARTICLE

Anesthesiology
58:362–369, 1983

Critiquing Anesthetic Management: The "ATTENDING" Computer System

Perry L. Miller, M.D., Ph.D.*

ATTENDING, a computer system under development using Artificial Intelligence techniques, is designed to critique an anesthetist's preoperative plan for anesthetic management. The system receives as input: 1) a list of a patient's medical problems; 2) a planned surgical procedure; and 3) an anesthetic plan outlining the agents and techniques to be used for premedication, induction, intubation, and maintenance of general or regional anesthesia. ATTENDING critiques this plan, discussing the risks and benefits of the proposed approach and of other reasonable approaches.

To critique a physician's plan, the ATTENDING system must be able to 1) explore alternative approaches to a patient's management in a flexible fashion; 2) assess the relative risks and benefits of these alternatives; and 3) produce a readable, easily understood analysis written in English prose. This article describes how these design issues are addressed in ATTENDING's current implementation. The ATTENDING system is presently being used in a tutorial mode which allows self-evaluation by an anesthetist. (Key words: Anesthesia: management. Anesthetic techniques: selection. Education. Equipment: computers.)

THE APPLICATION OF COMPUTERS in medicine has been increasing steadily. This trend will likely accelerate as computer prices continue to fall and as their processing power continues to rise. Despite this growing use of computers in many areas of medicine, however, they have yet to assume a major active role assisting the physician in planning the medical management of his patients. This article describes ATTENDING,[1] a computer system being developed using Artificial Intelligence (AI) techniques,[2] to assist an anesthetist by critiquing his preoperative plan for anesthetic management.

To use ATTENDING, an anesthetist must first evaluate a patient and formulate a tentative management plan. The anesthetist then inputs to ATTENDING: 1) a list of the patient's medical problems, 2) the planned surgical procedure, 3) the anesthetic plan, outlining the agents and techniques to be used for *premedication, induction, intubation,* and *maintenance* of general or regional anesthesia.

* Assistant Professor of Anesthesiology, Department of Anesthesiology, Yale University School of Medicine.

Received from the Department of Anesthesiology, Yale University School of Medicine, New Haven, Connecticut 06510. Accepted for publication September 3, 1982. Supported in part by NIH Grant 03978 from the National Library of Medicine. Presented in part at the annual meeting of the American Society of Anesthesiologists, Las Vegas, Nevada, October 1982.

Address reprint requests to Dr. Miller.

The ATTENDING system then critiques this plan, discussing the risks and benefits of the proposed approach and of other reasonable approaches. The system thereby serves as a source of feedback to help the anesthetist evaluate and optimize his proposed approach.

In undertaking to *critique* a physician's plan of management, ATTENDING differs from other medical decision-making systems. Most other systems attempt to *simulate* a physician's decision-making process. They gather information as a physician would, and try to reach similar conclusions. As a result, these systems effectively try to tell the physician how best to manage his patient.

The ATTENDING system, in contrast, expects its physician user to evaluate the patient, grapple with the management issues, and only then turn to the computer for assistance. Indeed, this is how doctors customarily ask other doctors for advice. They evaluate a patient, formulate at least some thoughts as to management, and only then do they ask a colleague or a consultant for advice.

Background

Artificial Intelligence (AI) is an emerging discipline which in recent years has received substantial coverage in the popular press. Research in AI focuses on designing computer systems which can perform tasks that appear "intelligent." Often this research involves the analysis of human reasoning processes, and the attempted integration of similar capabilities into the computer.

AI computer systems have been developed 1) to play sophisticated games such as chess; 2) to analyze complex visual scenes; 3) to "understand" (respond appropriately to) English; and 4) to perform medical diagnosis. Each of these is an ambitious undertaking, and there is wide acknowledgment that it will be years before computer systems develop capabilities which even begin to approximate human intelligence in its flexibility, generality, and power.

As a consequence, AI research frequently cannot have the realistic short-term goal of implementing fully developed systems for practical use. Rather, the emphasis is more often on 1) defining the fundamental problems involved and 2) developing increasingly pow-

0003-3022/83/0400/0362 $01.20 © The American Society of Anesthesiologists, Inc.

erful tools for flexible processing of complex real-world knowledge by a machine. Practical systems, while acknowledged as an ultimate goal, often are deferred by necessity to the indefinite future.

During the past decade, a number of systems have been developed applying AI in Medicine (AIM). Several surveys of these projects exist.[3,4] In this AIM research, a major emphasis has been on *diagnosis*, perhaps because diagnosis is perceived by the layman as the prominent intellectual process of medicine. Examples of AIM systems which focus primarily on diagnosis are EXPERT[5] which has been applied to ophthamologic and endocrine problems, PIP[6] dealing with renal disease, MYCIN[7] which is oriented towards infectious disease, and CADUSEUS[8] which embraces the whole of internal medicine.

The ATTENDING system, in contrast, does not deal with *diagnosis*, but rather with *medical management* (taken here to mean treatment in its broadest sense). Several other AIM systems have addressed problems of management. Examples are the Digitalis Advisor,[9] the VM system[10] designed to assist ventilator management, and ONCOCIN[11] which assists in oncologic protocols. In fact, there is a growing perception that medical management may prove a more fertile field for computer assistance than diagnosis.

There are several reasons for this perception. First of all, computer-assisted diagnosis has certain practical limitations. Many medical patients have chronic problems which fluctuate in severity and whose diagnoses therefore are already known. Also, many missed diagnoses result from a physician not pursuing deeply enough the underlying etiologic factors of known problems. Certainly when a patient presents clear-cut findings which a physician cannot integrate, a computer could be very helpful by matching those findings against a wide range of rare syndromes and unusual presentations of common disease. But such patients are relatively rare.

Medical management, on the other hand, may offer interesting opportunities for computer assistance even in diagnostically simple patients. The optimal treatment of a medical problem can vary significantly depending on a patient's other problems. It may happen, however, that a physician responds to a given problem without fully considering the patient as a whole. The methodical computer can play a useful role if only by checking for overall consistency of management approach, thereby helping the physician better tailor his treatment to the patient's individual needs.

Anesthesia offers a number of advantages as a domain in which to explore these issues of computer-assisted *management*: 1) Anesthesia has a central management component. 2) There is a single major decision point (the start of the operation) when a number of major management decisions must be made at once. 3) There are a variety of management alternatives to consider, with enough complexity to be interesting, but sufficiently circumscribed to be practically dealt with. 4) Although there is often considerable leeway in choosing among management options, there are also clear-cut risks and benefits in the presence of particular diseases. 5) Interesting risk tradeoffs frequently occur in complicated patients.

As a result, anesthesia provides a fertile domain for studying computer-assisted management, and for exploring ATTENDING's approach of critiquing a physician's plan. It is anticipated that this approach will ultimately prove extendable to other areas of medical management as well.

Example

The ATTENDING system can be used either in a *consultation* mode, or in a *tutorial* mode that allows anesthetist self-evaluation. In the tutorial mode, as demonstrated below, the system can be used for clinical education while its data base is still small. Here, an anesthetist sits at a computer terminal, and has the option of 1) allowing ATTENDING to select a hypothetical case from its list of teaching cases, or 2) choosing one of these hypothetical cases himself. (In the consultation mode, he defines an arbitrary patient by listing a set of medical problems known to the system.)

For instance, ATTENDING might select the following hypothetical case:

> A 34-year-old male with a history of asthma requiring several past hospitalizations comes to surgery after an auto accident for repair of a likely fractured spleen. His blood pressure is 90/50. He recently ate lunch.

In this tutorial mode, the patient's underlying problems (asthma, possible hypovolemia, and a full stomach) have already been input to ATTENDING. Therefore, the anesthetist need only outline his proposed plan. He might, for instance, propose 1) induction with thiopental, 2) rapid sequence intubation with succinylcholine, and 3) maintenance with enflurane, curare, nitrous oxide, and oxygen. Typing this plan into the computer takes approximately 45 s. ATTENDING currently critiques this plan as follows:

> There are a number of basic principles involved in this patient's management. First, in a patient with asthma, one must consider the problem of bronchospasm. Second, in the presence of hypovolemia, one must attempt to prevent hypotension. Third, a patient with a full stomach must be protected from aspiration.
>
> Intubation of this patient is clearly important because of a full stomach. Rapid sequence intubation is a reasonable approach. This does, however, have the possible risk of bronchospasm. An alternative approach, perhaps, to rapid sequence intubation would be intubation after mask induction with cricoid pressure.

In regard to induction, induction using thiopental could have two possible risks. First, in a patient with asthma, there is the conceivable risk of bronchospasm. Second, there is the possible risk of hypotension in the presence of hypovolemia. An alternative approach to induction using thiopental would be induction with ketamine, This has the advantage of helping avoid hypotension since ketamine is supportive of blood pressure, and of suppressing bronchospasm since ketamine is a bronchodilator.

From the standpoint of maintenance, inhalational technique has the advantage of preventing bronchospasm by keeping the anesthetic level deep. On the other hand, it could have the risk of hypotension in the presence of hypovolemia. The choice of enflurane involves the conceivable risk of initial bronchospasm resulting from airway irritability. An alternative would be halothane. This has the advantage of helping prevent bronchospasm since halothane is a bronchodilator. An alternative approach to inhalational technique would be balanced anesthesia. This involves the risk of possible bronchospasm due to light anesthesia. It does, however, have the advantage of maintaining blood pressure.

Looking at other aspects of the proposed plan, there are two considerations worth mentioning: 1) Curare would have the conceivable risk of bronchospasm triggered by histamine release. 2) Curare might have the risk of hypotension in the presence of hypovolemia, again, since curare can induce histamine release.

Several features of this example critique merit discussion: 1) The critique starts with a brief discussion of the underlying management principles involved. 2) In the analysis of the plan itself, ATTENDING is programmed to mention positive aspects of the plan, as well as any possible criticisms, on the theory that advice is best received when so presented. 3) The system must be able to accept *any* plan (good, bad, or indifferent), and still produce an organized, focused critique.

The following sections discuss how these capabilities are incorporated into the ATTENDING system.

Methods

The ATTENDING System Design

This section provides an overview of the design of the ATTENDING system. To allow ATTENDING to critique a physician's plan appropriately, three basic problems must be confronted: 1) The system must be able to *explore all the alternative approaches* to a patient's management in a flexible fashion. 2) It must be able to *assess the relative risks and benefits*, so it can focus on the most clinically appropriate approaches. 3) It must *communicate* its comments and suggestions in a critique which is readable, easily understood, and diplomatic.

The remainder of this section addresses each of these design problems in turn. ATTENDING is implemented in the LISP programming language.[12]

Exploring Alternative Approaches

ATTENDING uses "Augmented Decision Networks" (ADNs) to represent alternative approaches to anesthetic management. These are based on the Augmented Transition Network (ATN) formalism which has been used widely in natural language research.[13]

Figure 1 shows several ADNs currently used by ATTENDING. Each network consists of a number of *states* (circles) connected by *arcs*. Starting from the *initial state* (ρ) of a network, a *path* can be traced following arcs from one state to another. Such a path ends whenever a *POP* arc is traversed. The process of tracing such a path corresponds to the process of making a set of anesthetic management decisions.

The arcs are labeled with the names of anesthetic techniques and agents. Some arcs are labeled with the name of a *drug* (succinylcholine, halothane), some with the name of an *"elemental" technique* (nasal intubation), and others with the name of a *"non-elemental" technique* (general anesthesia, maintenance, inhalational anesthesia). These are "non-elemental" in the sense that they do not fully specify how the technique is to be implemented. Further subdecisions must be made. These are outlined by subnetworks which define each non-elemental technique.

For instance, starting at the "highest" network (ANES), two arcs leave its initial state, labeled GA (general anesthesia) and REGIONAL. If the GA arc is chosen, then before the path can continue, a subpath first must be traced through the GA network. In the process of tracing this path, further subnetworks must in turn be investigated. This corresponds to the process of making subdecisions within subdecisions within subdecisions, etc. A completed path through the ADNs corresponds to a completely formulated anesthetic plan. Figure 2, for instance, shows the path which corresponds to the example plan.

Using these decision networks, the exploration of alternative approaches becomes straightforward. For instance, if a patient with renal failure (RF) is to be anesthetized, and the anesthetist has proposed using enflurane, this implies that he has chosen the arc labeled ENFLURANE in figure 3.

For a patient with RF, however, enflurane involves a certain risk. A list of all risks associated with enflurane is attached to the ENFLURANE arc. As a result, ATTENDING can recognize the risk, and will therefore explore alternative paths through the networks, attempting to minimize the risk involved. In so doing, it first searches for a different path through the INHAL network (*i.e.*, for an alternative inhalational anesthetic). In figure 3, there is one alternative, halothane. If halothane involves no risk, then ATTENDING need search no further. If all inhalational anesthetics involve risk, however, then the system will explore other approaches to general anesthesia, and only if all these involve risk, will it explore more global alterations of approach, such as a regional technique.

In this way, the ADN formalism allows ATTENDING to tailor its advice to the physician's plan in a natural way, looking first for alternatives that involve the least change from the proposed approach.

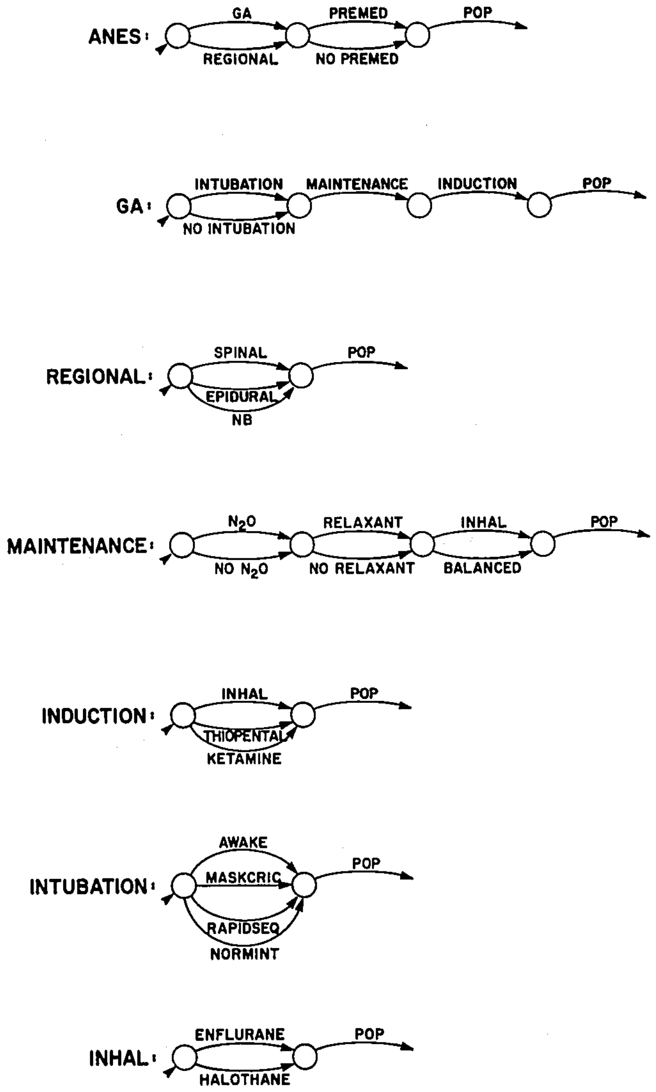

FIG. 1. Sample Augmented Decision Networks used by ATTENDING to structure its knowledge of the alternative approaches to anesthetic management.

Risk Analysis

As ATTENDING explores alternative approaches to anesthetic management, it must be able to compare the relative risks and benefits involved. Without this ability, it will not be able to focus on the most appropriate choices in its critique of the physician's plan.

ATTENDING *might* handle risk statistically. This is the conventional approach to risk analysis. This approach would require the collection of a large amount of data: likelihood and morbidity figures for all the risks to be included in the system. The collection of these data would be a formidable, if not impossible, task.

Such statistical risk analysis attempts to reduce the risk of alternative approaches to *numbers*. For instance, in conventional clinical decision analysis,[14] all possible outcomes are first defined, and each is given a likelihood and a "morbidity" value. The "expected value" of a given choice is then the sum, for all its possible outcomes, of each outcome's value times its likelihood.

The drawbacks inherent in this approach to risk are widely recognized. First, it is difficult to assign numeric values to outcomes. How does one assign a precise value to bronchospasm, hypotension, or the loss of an eye, a limb, or a life? Also, the gathering of likelihood statistics is a major task, and different studies may obtain substantially different results.

Granting these difficulties, and recognizing the massive data gathering that would be required, it is reasonable to ask whether ATTENDING need take a statistical approach to risk. Physicians certainly make management decisions without precise numbers to guide them. Could not a machine do this? Indeed, ATTENDING does *not* handle risk statistically.

Instead of reducing anesthetic risks to numbers to allow precise comparison, ATTENDING takes a *heuristic* approach to risk analysis. (The concept of a "heuristic" is fundamental to AI. In general, any solution whose ad-hoc nature does not lend itself to a precise equation or algorithm is a heuristic. Much real-world knowledge has this character. Indeed, AI might be called the science of heuristic knowledge.)

In ATTENDING, each risk is assigned a *rough estimate* of its magnitude (LOW, MODERATE, HIGH, or EXTREME). This corresponds to a rough estimate of the risk's likelihood times its morbidity. Clinically, these estimates seem to correspond quite naturally to how an anesthetist thinks about the risks, and are therefore easy to assign. Several examples of risks of different magnitude are given: 1) *LOW*: The risk of using enflurane in the presence of renal failure; 2) *MODERATE*: The risk of using morphine in a patient with asthma; 3) *HIGH*: The risk of using succinylcholine in the presence of a penetrating eye wound; and 4) *EXTREME*: The risk of not intubating a patient with a full stomach.

Benefits are estimated similarly, and are treated internally by the system as "negative risks." In using rough estimates of risks and benefits, ATTENDING's approach parallels that of several AI diagnostic systems which use *rough estimates of diagnostic likelihood* (rather than Bayesian statistical techniques) as the basis for their analysis.[15]

In its internal analysis of risk, ATTENDING makes further simplifications. In particular, if an alternative

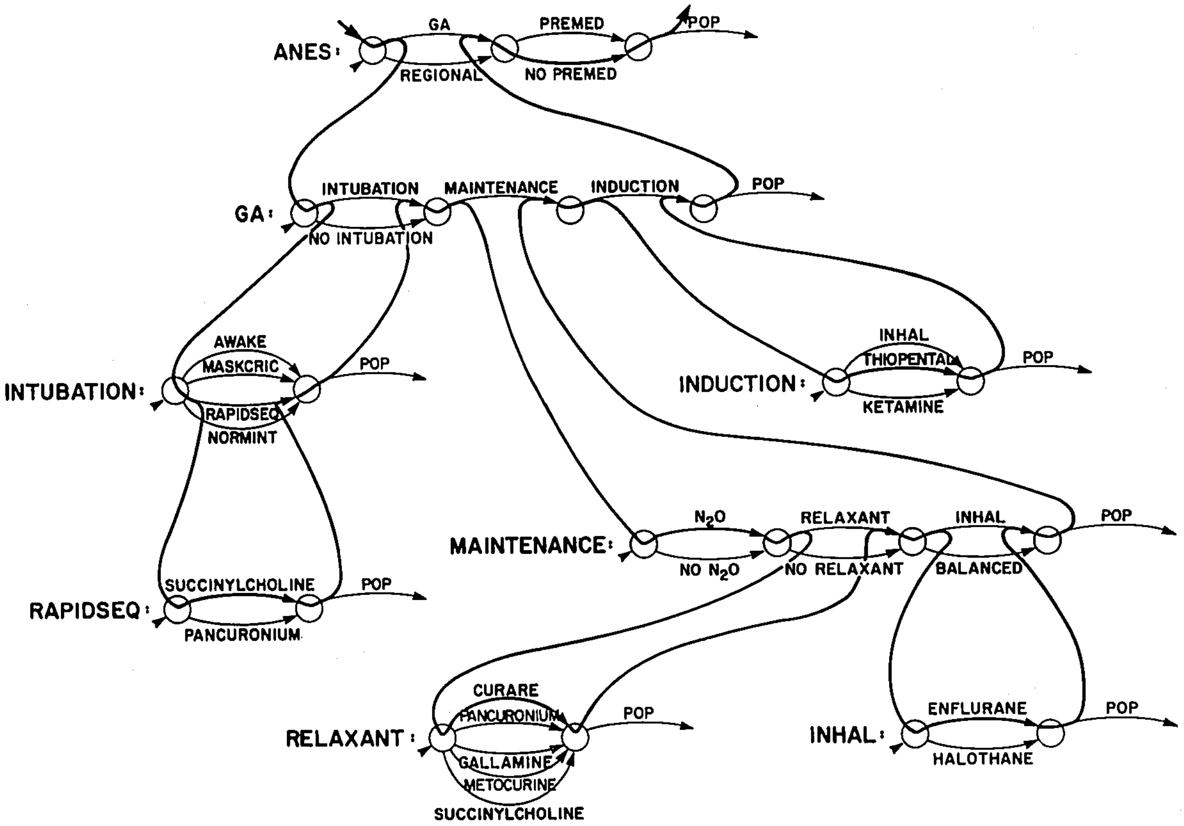

FIG. 2. The complete path through the ADNs, corresponding to the example anesthetic plan. (The order of the various decisions in the ADNs is largely arbitrary.)

involves several subrisks, then for purposes of comparison, that alternative is assigned the *highest* subrisk value. In other words, if a given approach involves three subrisks (of LOW, MODERATE, and MODERATE magnitude), then that approach is considered to have MODERATE risk. This simplification makes it easy to tell whether or not two alternatives have *roughly* equal risk.

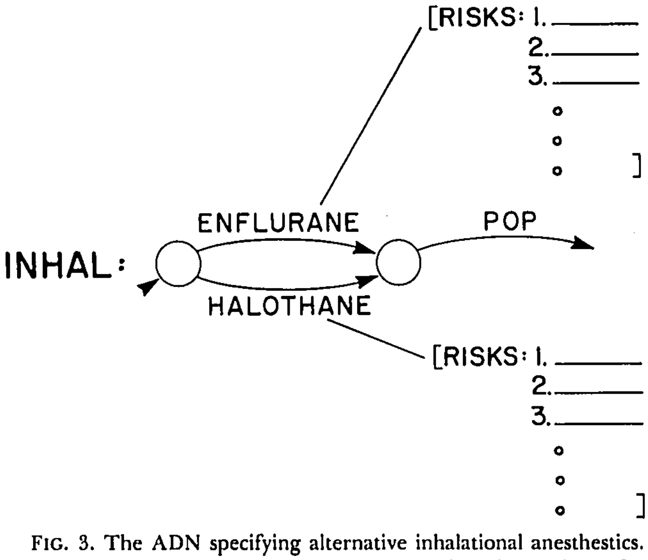

FIG. 3. The ADN specifying alternative inhalational anesthetics. Each arc has an associated list specifying the possible risks and benefits of using that technique in different circumstances.

Using these simplifications, ATTENDING can eliminate management choices which are clearly inferior and can thereby focus on the clinically relevant alternatives. (The system also employs preference rules which allow it to further refine and focus its analysis on particular techniques in certain circumstances.)

By taking this heuristic approach, ATTENDING is able to compare anesthetic risks and benefits without first compiling massive statistical data. In figure 3, for instance, the system can simply inform the anesthetist that enflurane involves the theoretical risk of fluoride ion toxicity in a patient with renal failure, and that he might want to consider halothane instead. It is not necessary to document such advice with statistics. The anesthetist can appreciate the issues involved, and can do as he chooses. The goal of ATTENDING's analysis is to outline the clinically relevant alternatives, and leave the final decision with the physician.

CREATING A PROSE CRITIQUE

The final step in ATTENDING's analysis is the creation of its prose critique of the physician's plan. This

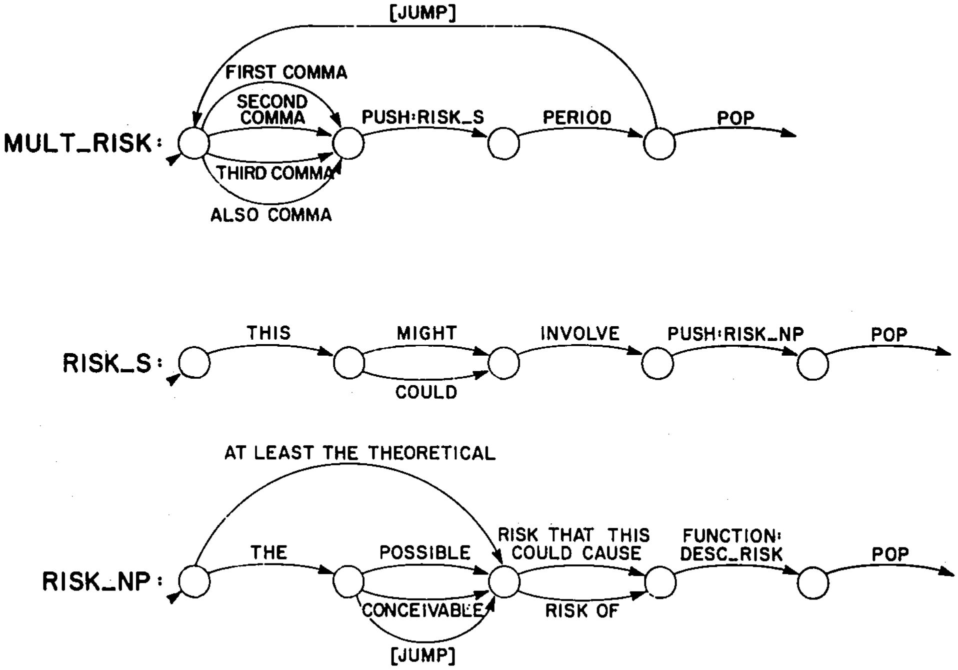

FIG. 4. Several simplified ATNs used by ATTENDING to produce its descriptions of risks.

step is crucial, since no matter how apt the content of the analysis is, if it is not clearly and logically expressed, its effect will be diminished greatly. This section outlines how the ability to produce such a discussion is programmed into the machine.

In producing its prose discussion, ATTENDING again uses the ATN formalism upon which its decision analysis is based. Thus, ATTENDING uses the ATN in two different ways: 1) to structure management decisions, and 2) to facilitate the production of its prose discussion.

Figure 4 shows three example ATNs, *simplified versions* of networks used by ATTENDING in describing risks. Here again, starting at the initial state of a network, one can trace a path from state to state, ending whenever a POP arc is traversed. (Whenever a PUSH arc is encountered, before it can be traversed, a subpath must first be traced through the specified subnetwork.)

As these paths are traced, prose fragments are output producing ATTENDING's discussion. For instance, different paths through the RISK__NP network can produce the following possible prose segments:

THE RISK OF ____

AT LEAST THE THEORETICAL RISK OF ____

THE CONCEIVABLE RISK THAT THIS COULD CAUSE ____

The exact path taken through the network depends on the material being discussed. Certain arcs are activated and inactivated depending on the characteristics of the risk being described. (The blanks in the prose segments are filled in, by a function named DESC__RISK, with a description of the particular risk involved, *i.e.*, "HYPOTENSION," "BRONCHOSPASM," etc.).

Once ATTENDING has determined the *content* of its analysis as described previously, it then activates the appropriate ATN networks. These produce the paragraphs which *express* the material in English prose.

Discussion

A central emphasis in developing ATTENDING has been on exploring the computer science design issues, outlined above, involved in critiquing a physician's plan of management. In particular, the system must be able to accomodate *any* proposed plan and still produce an appropriate critique, using that plan as the basis for its analysis. To allow this, the system's medical knowledge must be *more complete, more explicit*, and *more flexibly organized* than if the system merely attempted to formulate a single approach to a patient's management.

In addition to these interesting technical design issues, there are a number of *clinical* advantages to ATTENDING's approach of critiquing a physician's plan: 1) It casts the computer in the role of an ally rather than a potential competitor. 2) It requires the physician to evaluate a patient and grapple with the management issues before turning to the computer for assistance. 3) It leaves the primary decision-making with the physician. Since the physician must ultimately be responsible for the patient's care, this makes both medical and medicolegal sense.

In addition, there are frequently several possible approaches to a patient's management. No single approach is necessarily "right" or "wrong." Different physicians often have their own particular styles of practice. Many would not tolerate a system which did not allow them to practice in their accustomed way. For all these reasons, therefore, it makes sense to let the physician take the lead, and let a computer advisor tailor its advice to the plan he proposes.

Thus, the development of a system to critique a physician's plan involves interesting computer science design issues, and also may have significant potential clinical advantages.

LIMITATIONS

In implementing a system like ATTENDING, limits must be placed on the scope of the problems addressed. In this regard, ATTENDING has been designed to deal with a *very central* set of anesthetic management considerations (the agents and techniques to be used for premedication, induction, intubation, and maintenance of general or regional anesthesia). These considerations have the advantage that they involve interesting *risk tradeoffs* in complicated patients, and therefore, lend themselves well to ATTENDING's type of analysis.

There are, however, many anesthetic considerations that ATTENDING currently ignores. These include the management of intraoperative problems, fluid and electrolyte considerations, preoperative preparation of a patient for surgery, management of the patient's chronic medication, etc.

There are also other restrictions to the system's current medical knowledge.

1) *Anesthetic Techniques:* ATTENDING's ADNs currently represent 30 elemental techniques, and 20 nonelemental techniques. These cover most of the commonly used agents and techniques, but nevertheless may be modestly expanded as system development continues.

2) *Underlying Medical Problems:* The system's data base currently contains the anesthetic implications of 25 underlying medical problems, including asthma, hypovolemia, coronary artery disease, increased intracranial pressure, carotid artery disease, chronic renal failure, abdominal aneurysm, cerebral aneurysm, full stomach, and liver disease. Thus, a number of interesting problems are covered, but many others are not. This constitutes the major current data base limitation.

The ultimate goal is to extend ATTENDING's database of anesthetic implications of underlying problems to allow reasonably general use of the system. There will, however, always be medical problems ATTENDING is unfamiliar with. The system is *not* envisioned as an "automated reference" to be used to explore the anesthetic implications of rare diseases. A textbook is perfectly adequate for this task. Rather, the system is envisioned as a source of *interactive feedback* to help a physician deal with a large fraction of the patients he sees every day, to help prevent inadvertent oversights, and to help optimize his management approach.

In its present limited form, however, the system can still be tested and evaluated in its tutorial mode. Here the system itself describes hypothetical cases for analysis. These hypothetical cases, of course, involve only medical problems with which ATTENDING is familiar. This mode could be used for resident education, giving experience and feedback in the formulation of different approaches to anesthetic management. It also could be used for continuing education of practicing anesthetists. In addition, the tutorial mode could help introduce residents to anesthetic subspecialties, such as neuro-anesthesia, pediatric anesthesia, or cardiac anesthesia.

Thus, despite the current data base limitations, ATTENDING is nevertheless clinically useful as an educational tool. This use will allow the design philosophy to be fully worked out while the data base is still modest in size. Expansion can come later, after more experience with the system has been gained.

References

1. Miller PL: ATTENDING: a system which critiques an anesthetic management plan. Proceedings of the American Medical Informatics Association CONGRESS 82, New York, Masson, 1982, pp 36–40
2. Winston PH: Artificial Intelligence. Reading, Addison-Wesley, 1977
3. Kulikowski C: Artificial intelligence methods and systems for medical consultation. IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume PAMI-2, 1980, pp 464–476
4. Shortliffe EH, Buchanan BG, Feigenbaum EA: Knowledge engineering for medical decision making. A review of computer-based clinical decision aids. Proceedings of the IEEE 67:1207–1224, 1979
5. Weiss S, Kulikowski C: EXPERT: A system for developing consultation models. Proceedings of the Sixth International Joint Conference on Artificial Intelligence, Tokyo, 1979, pp 942–950
6. Pauker SG, Gorry GA, Kassirer JP, Schwartz WB: Towards the simulation of clinical cognition: Taking a present illness by computer. Am J Med 60:981–996, 1976
7. Shortliffe EH: Computer-based medical consultations: MYCIN. New York, American Elsevier, 1976
8. Miller RA, Pople HE, Myers JD: INTERNIST-I, an experimental computer-based diagnostic consultant for general internal medicine. N Engl J Med 307:468–476, 1982
9. Gorry GA, Silverman H, Pauker SG: Capturing clinical expertise: A computer program that considers clinical responses to digitalis. Am J Med 64:452–460, 1978
10. Fagan LM, Kunz JC, Feigenbaum EA, Osborn JJ: Representation of dynamic clinical knowledge: Measurement interpretation in the intensive care unit. Proceedings of the Sixth International

Joint Conference on Artificial Intelligence, Tokyo, 1979, pp 260–262
11. Shortliffe EH, Scott AC, Bischoff MB, Campbell AB, Van Melle W, Jacobs CD: An expert system for oncology protocol management. Proceedings of the Seventh International Joint Conference on Artificial Intelligence, Vancouver, 1981, pp 876–881
12. Winston PH, Horn BK: Lisp. Reading, Addison-Wesley, 1981
13. Woods WA: Transition network grammars for natural language analysis. Communications of the Association for Computing Machinery 13:591–606, 1970
14. Weinstein MC, Fineberg HV, Elstein AS, et al: Clinical Decision Analysis. Philadelphia, WB Saunders, 1980
15. Szolovits P, Pauker SG: Categorical and probabilistic reasoning in medical diagnosis. Artificial Intelligence 11:115–144, 1978

APPENDIX

Availability of the ATTENDING System

ATTENDING currently is implemented in the LISP programming language on the SUMEX-AIM computer system (Stanford University Medical Experimental Computer-Artificial Intelligence in Medicine).†

LISP is a computer language that facilitates the experimental implementation of large AI systems. For practical implementation and distribution, however, such a system may well have to be rewritten in a conventional, more efficient programming language. Since ATTENDING is a developmental research system, still undergoing design, refinement, and expansion, it is not at present ready for this type of conversion to practical use.

The present developmental system can, however, be used experimentally for teaching: within our department, and possibly on a collaborative basis at other institutions as well. The SUMEX-AIM computer system facilitates such collaboration since it allows nationwide use via a readily accessed computer network. SUMEX-AIM, funded by the NIH Division of Research Resources, has been developed to promote nationwide collaborative research on Artificial Intelligence applications in Medicine.

† Freiherr G: The seeds of artificial intelligence: SUMEX-AIM. NIH Publication No. 80-2071, March 1980.