

Artificial intelligence applications in the intensive care unit

C. William Hanson III, MD, FCCM; Bryan E. Marshall, MD, FRCP, FRCA

Objective: To review the history and current applications of artificial intelligence in the intensive care unit.

Data Sources: The MEDLINE database, bibliographies of selected articles, and current texts on the subject.

Study Selection: The studies that were selected for review used artificial intelligence tools for a variety of intensive care applications, including direct patient care and retrospective database analysis.

Data Extraction: All literature relevant to the topic was reviewed.

Data Synthesis: Although some of the earliest artificial intelligence (AI) applications were medically oriented, AI has not been widely accepted in medicine. Despite this, patient demographic, clinical, and billing data are increasingly available in an electronic format and therefore susceptible to analysis by intelligent software. Individual AI tools are specifically suited to different tasks, such as waveform analysis or device control.

Conclusions: The intensive care environment is particularly suited to the implementation of AI tools because of the wealth of available data and the inherent opportunities for increased efficiency in inpatient care. A variety of new AI tools have become available in recent years that can function as intelligent assistants to clinicians, constantly monitoring electronic data streams for important trends, or adjusting the settings of bedside devices. The integration of these tools into the intensive care unit can be expected to reduce costs and improve patient outcomes. (Crit Care Med 2001; 29:427-435)

KEY WORDS: intensive care unit; artificial intelligence; expert systems; computer-assisted diagnosis; computer-assisted therapy; decision support techniques; neural networks; algorithms; fuzzy logic; data display; computer simulation; clinical decision support systems; management decision support systems

The amount of data acquired electronically from patients undergoing intensive care has grown exponentially during the past decade. Bedside equipment such as pressure and flow transducers, infusion pumps, pulse oximeters, cardiac output monitors, and mechanical ventilators store electronic data and are equipped with computer interfaces. Modern bedside monitors communicate with a host of devices through data busses and interchangeable, plug-in interfaces. Computerized intensive care systems interface with hospital databases including demographic systems, electronic patient records, order-entry, laboratory, pharmacy, and radiology systems.

It is useful conceptually to recognize that bedside data must be extracted and organized to become information, and that an expert must then interpret this information before it becomes knowledge for diagnostic and/or therapeutic purposes. This review is concerned primarily with the second step in this cognitive

sequence, namely the use of computers to extract information from data and enhance analysis by the human clinical expert. We will argue that an as yet unrealized role for the computer at the bedside is the extraction of information from data rather than mere display. A variety of novel, computer-based analytic techniques have been developed recently. Our purpose is to introduce these techniques and to discuss their potential for clinical applications in the intensive care unit (ICU).

A review of relevant literature cited in MEDLINE was performed between the years 1966 and the present using Ovid software. The search terms "critical care" and "intensive care" were exploded and combined (18 and 423 "hits," respectively). The terms "data warehouse" (18), "neural network" (4215), "genetic algorithm" (238), "fuzzy logic" (431), "case-based reasoning" (55), "belief network" (47), and "data visualization" (42) were then used as keywords to create individual data sets comprising all references to the specific term. Each of these data sets was then individually combined (AND function) with the critical care references. A comprehensive review was made of all references cited in both data sets. In

addition, the smaller data sets (e.g., belief network) were thoroughly reviewed to find abstracts or titles suggesting relevance to critical care practice. Finally, articles of historic significance (e.g., references to MYCIN) and appropriate texts were included for completeness.

Background

Nomenclature. It is useful to begin by defining several terms that are commonly used in describing computer systems used for data analysis. A *management information system* (MIS) has been defined as "a formalized computer information system that can integrate data from various sources to provide the information necessary for management decision making" (1).

The appropriate definition of *artificial intelligence* (AI) is controversial. Alan Turing, the English mathematician, devised what has become known as the Turing test of computer intelligence. He suggested that a computer had artificial intelligence if it could successfully mimic a human and thereby fool another human.

An *expert system* is a computer program that simulates the judgment and behavior of a human or an organization

From the Department of Anesthesia, Center for Anesthesia Research, University of Pennsylvania Health System, Philadelphia, PA.

Copyright © 2001 by Lippincott Williams & Wilkins

with expert knowledge and experience in a particular field.

Data mining is the analysis of data for relationships that have not previously been discovered. The techniques used for data mining can discover hidden associations or sequences in data sets, clustering of data points, and permit visualization of relationships among data or forecasting based on hidden patterns. Data mining is also known as knowledge discovery, and derives its roots from statistics, artificial intelligence, and machine learning.

A *data warehouse* is a central repository for all or significant parts of the data that an enterprise's various business systems collect. An alternative term is a *data mart*.

The Data Stream. The establishment of modern data-intensive ICUs can probably be traced to the introduction of blood gas electrodes in the 1960s. The outstanding developments in bioengineering since then have resulted in an impressive amount of data being reliably available at the bedside. For the most part, however, clinical practice has not taken advantage of the rich, varied, and continuous stream of information. Despite the wealth of electronically accessible data, the synthesis and interpretation of information is done manually in many of today's ICUs with minimal preprocessing. Nurses laboriously transcribe information from monitors onto paper records. Data are lost, creatively interpreted, averaged, and incorrectly transcribed in the process. Additionally, although a sicker patient requires more nursing interventions with consequently less time available for data transcription, it is these patients for whom accurate data are most useful to clinicians for retrospective review and analysis.

There are several reasons for the durability of what would seem to be archaic methods of data acquisition and transcription. Among the most often cited are the legal and financial requirements for documentation of human observations and decisions. These requirements are real and are not likely to change. They mandate written chart entries at regular intervals and at times of change, which must be made and signed by the responsible individual. The second basis for reluctance by clinicians is the failed promise of early expert systems, as discussed below.

Management Information Systems

Management information systems can be categorized as depicted in Table 1. One group of decision support systems is the model-driven or rule-based expert systems (RBS). They are successful to the extent that they are able to represent the subject material accurately and interface well with the user.

Rule-based systems can be thought of as "top-down" systems. "Top-down" programming begins with a complex problem and uses a reductionist approach, breaking the problem down into its constituent parts to arrive at the essential components that characterize it. This approach can be used as a method for simulating the thinking processes of the human medical expert. In reality, however, experts make rapid, often intuitive, diagnoses by beginning with a few hypotheses selected by experience, followed by clinical or laboratory observations that further refine the differential diagnosis.

Rule-Based (Expert) Systems. Some of the first expert systems were developed for medical care. Some notable early expert systems were the MYCIN (2, 3), ONCOCIN (4, 5), and Internist/Quick Medical Reference (6, 7) systems. MYCIN and ONCOCIN were designed at Stanford and simulated the performance of consultants in infectious disease and oncology respectively. The Internist system was designed at the University of Pittsburgh to reproduce the diagnostic behavior of an internist. Unfortunately, these pioneering systems have not been widely incorporated into the practice of medicine. One author suggests, "Artificial intelligence in medicine (AIM) has not been successful—if success is judged as making an impact on the practice of medicine. Much recent work in AIM has been focused inward (sic), addressing problems that are at the crossroads of the parent disciplines of medicine and artificial intelligence. Now,

AIM must move forward with the insights that it has gained and focus on finding solutions for problems at the heart of medical practice." (8)

A variety of impediments have slowed the general acceptance of medical expert systems. They include the small margin of acceptable error in medical practice, the ready availability of experts in most settings, and the complexity of regulatory requirements. Many of these obstacles may disappear in the future. Fewer experts are being trained, and expert systems may become cost-effective in certain environments. The performance of expert systems could improve to the point that they rival or exceed human experts. Finally, increased acceptance of telemedicine may lower the barriers to acceptance of expert systems.

Data-Driven (Intelligent Assistant) Systems. A second, newer generation of decision support systems is data driven (9). Data-driven systems (DDS) take advantage of the large quantity of data that can be acquired electronically to "discover" relationships and assume that future behavior can be predicted from past behavior. They represent "bottom-up" systems in which the data generated by a system is used to describe the characteristics of the system. Data analytic tools such as these are typically less ambitious than expert systems in scope and scale, less expensive to develop and maintain, and well suited to act as intelligent assistants to human experts.

It is useful to contrast the model-driven and data-driven systems using a familiar intensive care construct. A RBS would include rules about the relationship between pulmonary artery occlusion pressure (PAOP) and cardiac output. The rules would be based upon well-understood and accepted physiologic principles such as the Frank-Starling mechanism, whereas a DDS might begin without preconceptions and then discover the Frank-Starling relationship for an individual pa-

Table 1. Decision support systems

Model Driven Decision Support	Data Driven Decision Support
Monolithic	Modular
Incorporation of a body of knowledge	Self-learning
Designed to reproduce the expert	Typically act as intelligent assistants to an expert
All solutions are preprogrammed	Capable of arriving at novel, unexpected solutions or observations
Likely to require substantial maintenance as medical knowledge evolves	Inherently autodidactic and therefore largely self-maintained

tient. Specific characteristics distinguish the two approaches in this example. The RBS embodies general rules about all patients: "If the pulmonary artery occlusion pressure is less than or equal to 5 mm Hg and if the cardiac output is less than 2 L/min, then give 10 mL/kg crystalloid." Conversely, the DDS is used to discover the physiologic behavior of one individual from data acquired during continuous monitoring of that patient: "When Mr. Smith's PAOP decreases to less than 8 mm Hg, his cardiac output decreases to less than 3 L/min." The RBS imposes structure on the data, whereas the DDS derives structure from the data.

Data Mining. New methods of data analysis and decision support have become available in the past decade, which have been euphemistically termed "data mining." The techniques of data mining have evolved from and depend on previous generations of data analysis tools (Table 2). Several different techniques are commonly used in data mining, or data-driven decision support. They include data warehouses, neural networks, genetic algorithms, Bayesian or belief networks, rule induction or case-based reasoning, and machine learning. Fuzzy logic is another relatively new approach to programming that permits ambiguity in descriptions of data. Visualization techniques display large amounts of data in a comprehensive, comprehensible fashion.

Data Warehousing. Data-driven decision support benefits from the creation of a data warehouse and an online analytical processing system (OLAP). Decision support systems consist of a well-organized database (the data warehouse) and an accessible front-end (OLAP) that permits flexible exploration and analysis of data by a nonprogrammer.

The medical director of an ICU could use an OLAP to acquire and analyze data from a data warehouse to answer questions such as, "What is the length of stay of patients admitted (to my ICU) with the diagnosis of respiratory failure who are ventilated for greater than 2 days." Traditionally, questions such as these would require a programmer to query a relational database using a structured query language, which in turn presupposed that the required data were available in a relational database. Using an OLAP, an administrator could answer the question at whim using natural language rather than a specifically designed computer program. The administrator might then follow-up (or "drill down" in the vernacular) with a second question suggested by the answer to the first, such as, "What was the mortality of those patients?"

Although data warehouses are widely implemented in industry, and to some extent in hospital administration, they are essentially unavailable in the intensive care setting. Intelligent algorithms have been developed to enhance the analysis, functionality or display of information in the ICU as described below.

Neural Networks. Neural networks are designed to mimic the performance of the human brain. There are input nodes (or neurodes), output nodes, and a variable number of internal (or hidden) layers. The nodes are connected with different architectures, but typically input nodes are connected to hidden layer nodes and they are in turn connected to output nodes. As the neural network learns from (or "trains on") a data set, the connection weights are adjusted. In effect, important connections are reinforced (positively weighted) and unimportant connections are punished (negatively weighted). Data are fed into

the input nodes, processed through the hidden layer(s), and the connection weights to the output nodes are adjusted.

Neural nets are categorized based on their learning paradigm. In supervised networks, the outputs are known but the importance of the relationship of a given input to an output is unknown before training. In an ICU, a neural network can be used to explore the relationships among several physiologic variables. For example, Buchman used a neural network to evaluate the relationship of several demographic, pharmacologic, and physiologic variables to ICU chronicity (Fig. 1) (10).

In unsupervised networks, the outputs are unknown and the system is encouraged to find interesting, often unsuspected, relationships among the data elements in large data sets. For example, a hypothetical unsupervised neural network might make the novel discovery that a hypotensive episode of greater than an hour's duration immediately following cardiac surgery is highly correlated with subsequent development of pancreatitis.

Neural networks have been used in the ICU setting in a variety of fashions, but most extensively for outcome prediction. Neural networks have been shown to predict length of stay in the ICU (10–12). Other neural network-based systems were successful in predicting ICU mortality (11, 13–15).

Another common application of neural networks in the ICU is the real-time analysis of waveforms such as the electrocardiogram and the electroencephalogram. One neural network-based algorithm identified cardiac ischemia with high sensitivity based on analysis of the ST segment (16), whereas another diagnosed myocardial ischemia in the emergency department patient (17–20). Neu-

Table 2. Evolution of data analytic techniques

Technique	Typical Question Format	Necessary Technologies	Products/Vendors	Characteristics
Data collection (1960–1970s)	How many patients were admitted to the ICU last year?	Data stored on tape/disks	IBM	Retrospective, static data interrogation, programmed queries
Data access (1980s)	What was the mortality in our coronary artery bypass (CABG) last patients year?	Relational database and structured query language	Oracle, Sybase	Retrospective, dynamic data interrogation at record level, programmed queries
Data warehousing (1990s)	What was the CABG mortality last year? → Drill down to those patients with atrial fibrillation	Multidimensional databases and online analytical processing systems	Pilot, Microstrategy	Retrospective dynamic data interrogation at multiple levels, natural language queries
Data mining (emerging)	What's causing the high rate of atrial fibrillation in the ICU?	Massive data bases, multiprocessor computers	Lockheed, IBM	Prospective, proactive information delivery

ICU, intensive care unit.




Figure 1. Supervised neural network: Neural network with three inputs, a hidden layer, and a single output (O). The network is similar to one that Buchman et al. (10) described for the prediction of chronicity in the intensive care unit. This network is supervised, i.e., it is trained through feedback so that connections that create desired outputs (better outcome prediction) are reinforced (thicker connections). HR, heart rate; BP, blood pressure.

ral networks have also been used in the analysis of electroencephalographic patterns in children (21) and adults sedated with midazolam (22). Finally, neural networks have been used to analyze hemodynamic patterns in intensive care patients (23, 24).

Neural networks can reveal unexpected and otherwise undetectable patterns in large data sets. The major weakness in neural network solutions is the fact that the methods by which a relationship is discovered are hidden (or opaque) and therefore not readily understood or explained.

Genetic Algorithms. Genetic algorithms were designed to find near optimal solutions to complicated problems using the principles of Darwinian selection. For example, genetic algorithms have been used to find a near optimal route for the salesperson who needs to travel through several cities on a sales trip. The so-called traveling salesperson problem was once considered noncomputable as the number of cities became large, but genetic algorithms provide a best approximation as an answer. The process of optimization involves the following: a) the creation of a number of possible solutions; b) competition among them using selection criteria (i.e., fastest route, shortest route, least expensive route); and c) the elimination of "bad" solutions. Surviving solutions are then

Degree of Membership




Figure 2. Traditional vs. fuzzy classification schemes: Traditional classification of pulmonary artery occlusion pressures into discrete sets compared to fuzzy classification using fuzzy sets. In the standard classification scheme, any single pulmonary artery occlusion pressure (PAOP) (horizontal axis) has full membership (vertical axis) in only one set (low [L], normal [N], high [H]). Conversely, a single PAOP can have simultaneous partial membership in more than one fuzzy set (very low [VL], very high [VH]). For example, when the occlusion pressure is 5.5 (dashed vertical line), it is classified into a single classic set, whereas it is a "member" of two fuzzy sets (low and normal).

permitted to mutate and cross-breed and compete further. Ultimately, a highly desirable solution is selected from the set of all possible solutions. Note that because they fail to explore all solutions, genetic algorithms are quite efficient but cannot ensure that the surviving solution is the best possible choice.

An example of a hypothetical ICU problem that might be susceptible to this approach is the determination of an optimal staffing configuration for a group of patients with different acuity or nursing requirements. Genetic algorithms have been used to determine the neural net configuration that was most accurate in

predicting prognosis in a group of 258 ICU patients (14).

Fuzzy Logic. Strictly speaking, fuzzy logic is not a data-driven analytic approach. Rather, it is a method of handling data that permits ambiguity, and as a result, it is particularly suited to medical applications. One of the interesting ironies of medical practice is that its practitioners strive for objectivity and precision while dealing with data that are inherently imprecise.

Fuzzy logic has proven to be well suited to a variety of industrial applications, and fuzzy control strategies are, in many cases, more efficient than tradi-

Data set

Case #	Independent variables					Dependent variable
	Beta blocker	Mg ⁺⁺	K ⁺	CVP	Chest tube output	Atrial fibrillation
1	y	2	3.3	10	1250	y
2	y	2.3	4.1	12	500	y
3	n	1.9	3.6	7	400	n
4	y	2.1	3.7	5	750	y
5	n	2.5	5	3	400	n
6	n	1.8	4.1	16	2000	y
7	n	1.7	3	12	1500	y
8	y	2.5	4.3	3	300	n
9	n	2.4	4.7	4	450	y
10	y	2.2	6	11	750	n
200	n	1.8	4.2	8	600	n

Machine Learning




Figure 3. Data parsing by machine learning algorithm: A machine learning algorithm was used to partition this hypothetical database of cardiac surgical patients and demonstrate the correlation between two independent variables (low serum potassium and absence of beta blockade) with the dependent variable (atrial fibrillation). Y, yes; N, no.

tional alternatives. Fuzzy control systems are used in applications as diverse as elevator control, intelligent sensors, climate control, and image stabilization in video cameras.

Figure 2 shows the classically accepted categorization of PAOPs as low, normal, and high. These sets are discrete—any single PAOP is a member of only one set. Fuzzy logic permits the use of overlapping sets, and therefore simultaneous membership in more than one set. Using fuzzy descriptors and control logic in the ICU, very efficient infusion controllers have been designed that use fuzzy logical constructs like “if the blood pressure is low, and the cardiac output is low normal, give dopamine slowly.”

Fuzzy control processes have been used for the administration of anesthetics in the operating room (25–29). In the ICU, fuzzy control strategies have been designed for the administration of fluid (30) and titration of oxygen therapy (31). Fuzzy controllers have also been de-

signed for the administration of vasodilators to control blood pressure in the perioperative period (32–36) and during dialysis (37, 38). Fuzzy logic has been used to control mechanical ventilation (39, 40) and artificial hearts (41). Fuzzy logic diagnostic strategies have also been used in the ICU to analyze physiologic data during a simulated cardiac arrest (42), categorize oxygen destruction (43), interpret EEGs (21), and to distinguish real alarms from artifacts in preterm infants (44).

Fuzzy logical systems are easy to configure and tune. Unlike neural networks, the logical constructs used in these systems are easy to describe and closely approximate the thinking processes used in clinical decision making.

Machine Learning. Machine learning algorithms create rules or classification schemes by searching through data for relevant patterns. Unlike neural networks, they generate rules and patterns that can be evaluated and understood.

The success of machine learning algorithms requires the creation of a data set with independent and dependent variables. The machine learning algorithm is then used to explore the data for interesting or unexpected relationships.

Figure 3 shows the use of a machine learning algorithm to determine the relationship of five independent variables (preoperative beta-blocker therapy, serum magnesium level, serum potassium level, central venous pressure, and mediastinal drain output) to the development of atrial fibrillation (the dependent variable) in postoperative coronary artery bypass patients.

Machine learning algorithms have been designed to derive rules for intelligent alarms on respiratory systems (45, 46).

Case-Based Reasoning. Case-based reasoning is a method of arriving at solutions by analogy. A case consists of a series of attributes describing a situation




Figure 4. Case-based reasoning: In case-based reasoning, a probe (test case) is designed to query a database for cases with similar attributes. Here a hypothetical test case has encephalitis, acute respiratory disease (ARDS), and renal failure. Two other cases with similar attributes are retrieved, (1) encephalitis, ARDS, hepatic and renal failure and (2) encephalitis, renal failure, myocardial infarction and ARDS), whereas two other dissimilar cases (3 and 4) are not.

and a related solution. A series of representative cases are accumulated into a case base. The case base can then be examined by the use of probes or test cases (to find all cases in the database like a test case). The probe may use a subset of the available attributes and constraints on the available solutions to find a set of analogous cases.

Case-based reasoning is well suited to the derivation of solutions when there are discontinuities in the case base. Case-based reasoning might be used to find cases similar to a hypothetical, newly presenting, undiagnosed patient with rapidly evolving encephalitis, acute respiratory distress syndrome, and renal failure (Fig. 4). New diseases such as Hanta virus, acquired immune deficiency syndrome, and Legionnaire's disease are discovered through a process analogous to case-based reasoning. Although a case database might not contain a case with a specific constellation of conditions, case-based systems can extrapolate from other patients with similar histories. Case-based reasoning simulates the reasoning processes of an expert who has a large

catalog stored in his brain and can rapidly recall analogous cases.

Case-based reasoning approaches have been used in the ICU to select a group of patients with similar characteristics from a demographic database (11). They have also been used to select an antibiotic regimen (47, 48) and to project the course of renal function (49).

Bayesian (Belief) Networks. The power of Bayes' Rule has been exploited to develop extremely powerful, readily understood learning algorithms known as Bayesian networks or belief networks. They are versatile and have been used for a wide variety of functions including military applications such as the rapid identification of incoming military targets (missiles, aircraft, vessels) and deployment of counterattacks. Belief networks have also been used to assist computer software support staff in diagnosing problems on help lines. In the latter application, probability trees are constructed describing the potential etiologies of a customer problem.

The initial design of a belief network requires the configuration of a tree of




Figure 5. Architecture of a belief network: A belief network describing the probabilities relating several known variables to oxygen delivery in a patient who lacks a pulmonary artery catheter. In this example, the central venous pressure (CVP), left ventricular ejection fraction (EF), heart rate (HR), hemoglobin (HGB), and oxygen saturation (O2sat) are known and fall into predefined ranges. For example, the heart rate is known and falls within the predefined normal range. Therefore the probability that the heart rate is low or high is zero, whereas the probability that the heart rate is normal is one. The probabilities characterizing the unknowns, e.g., cardiac index (CI), depend on the probabilities of the parent nodes (CVP, EF, and HR) at a given point in time. The network can be configured with initial probabilities by an expert and then tuned with real data. The patient described in the current state of this network has normal CVP, HR, and HGB, a high EF and low O₂sat and therefore has an 82% probability of having normal oxygen delivery (O2del).

nodes describing the relationship of variables to one another. Figure 5 shows a Bayesian logic tree in which the branches describe the relationships among the physiologic variables that determine systemic oxygen delivery. An expert constructs a table describing the pretest probabilities that a given combination of heart rate and wedge pressure will result in a given cardiac output. The network is tuned (and probabilities recalculated) by exposure to real data.

Bayesian networks are well suited to the intensive care environment because of their speed and comprehensibility (transparency) in addition to the fact that Bayes' theorem is well understood and accepted by modern physicians. Belief networks have been used in the intensive care unit to evaluate EEGs (50) and to establish prognosis in patients with head injuries (51).

Data Visualization. Data visualization is a term used to describe the intelligent depiction of information using proximity, grouping, shape, color, animation, and other techniques to enhance data com-




Figure 6. Visualization used to summarize data. Visualization techniques are used here to demonstrate that this hypothetical intensive care unit cares for a large number (*big sphere*) of head-injured patients who are rarely in need of continuous dialysis, mechanical ventilation, or invasive hemodynamic monitoring. *ARDS*, acute respiratory distress syndrome.

prehension by the observer. Figure 6 shows a multidimensional array using three axes (x, y, z) and size to depict the likelihood that a disease process will be treated with a therapeutic modality.

Medical applications of visualization techniques have largely been restricted to three-dimensional anatomical reconstruction for surgical applications. However, it is possible to use data visualization for the display of hemodynamic data (52). Data visualization techniques have also been used in an attempt to improve the depiction of medical data in the patient medical record (45, 53, 54).

CONCLUSIONS

There are a large variety of novel data management tools that have become available over the past decade. They vary in their suitability to a given task, the degree to which the achieved solution is understandable (transparent or opaque), and the ease of configurability. Because of the data rich nature of the ICU, the applications to which these techniques can be applied are varied, ranging from waveform analysis to outcome prediction.

There are obstacles to the rapid adoption of the data analysis tools described in this manuscript. Although OLAPs have been widely adopted as business and financial decision support systems, they have essentially been ignored in clinical decision support. There are examples of medically oriented OLAP tools such as Infomine (Infocure, Atlanta, GA), how-

ever this is a data mining tool designed primarily to allow financial analysis of a medical practice. The query tools available with clinically oriented critical care information systems such as Picis CareSuite (Picis, Arlington, VA) and Sunrise Critical Care (Eclipsys, Delray Beach, FL) permit fixed format reports, retrospective, dynamic data interrogation, or programmed queries (Table 2). As patient information databases mature and become commonly available online, clinically oriented OLAPs will almost certainly become available. Greater physician leadership in the design and implementation of these systems will undoubtedly enhance the speed with which they are adopted.

Many of the studies cited in this review are small and can be characterized as proofs of concept or feasibility rather than true side-by-side comparisons of a computer technique with the current standard of care. In addition, medicine is an appropriately conservative discipline, and regulatory agencies such as the Food and Drug Administration have stringent requirements for the approval of new drugs and technologies to protect the individual patient. Paradoxically, fuzzy control systems and neural networks are being used in automated trains and elevators, helicopters, and even spacecraft.

It is unlikely that intelligent software will replace the clinician in the way that the original medical expert systems were conceived. Future systems are more likely to act as intelligent agents for spe-

Data-driven decision support tools will permit the busy clinician (physician, nurse, respiratory therapist) to function more efficiently, caring for more patients more safely in much the same way that these same tools have been used to enhance the efficiency of business applications.

cialized, complicated problems, and are generally intended to enhance the performance of a human expert. The applications described above vary in their suitability to specific tasks, and it is likely that they will be combined in the smart intensive care unit of the future.

Neural networks and fuzzy systems are particularly useful for waveform analysis. They will be integrated into bedside monitors and continuously analyze waveforms for known patterns (cardiac ischemia, hypovolemia).

Fuzzy controllers will be integrated into bedside devices such as fluid and medication infusion devices, mechanical ventilators, and dialysis machines.

Belief networks and neural networks will be used in the development of smart alarms that integrate multiple data streams (hemodynamics, laboratory data, other monitors), which display event probability (e.g., developing sepsis or acute respiratory distress syndrome).

Data visualization tools will permit the clinician to interrogate and analyze laboratory and hemodynamic trends in an individual patient at a glance.

Case-based reasoning, machine learning algorithms, and visualization tools will be used to analyze information from data warehouses describing the characteristic of an individual ICU. For example, a mini-epidemic of infections attributable to a resistant bacterial organism might be identified and followed using tools such as these.

Data-driven decision support tools will permit the busy clinician (physician, nurse, respiratory therapist) to function more efficiently, caring for more patients more safely in much the same way that these same tools have been used to enhance the efficiency of business applications. Modern assembly lines are more productive and make fewer errors than ever before, partly through the application of data-driven decision support (8). Similar improvements can be expected in and will be demanded of the medical industry in the immediate future.

REFERENCES

- Telem M: MIS implementation in schools: A systems socio-technical framework. *Comput Educ* 1996; 27:85-93
- Shortliffe EH, Davis R, Axline SG, et al: Computer-based consultations in clinical therapeutics: Explanation and rule acquisition capabilities of the MYCIN system. *Comput Biomed Res* 1975; 8:303-320
- Yu VL, Buchanan BG, Shortliffe EH, et al: Evaluating the performance of a computer-based consultant. *Comput Programs Biomed* 1979; 9:95-102
- Hickam DH, Shortliffe EH, Bischoff MB, et al: The treatment advice of a computer-based cancer chemotherapy protocol advisor. *Ann Intern Med* 1985; 103:928-936
- Shortliffe EH: Medical expert systems—Knowledge tools for physicians. *West J Med* 1986; 145:830-839
- Masarie FEJ, Miller RA, Myers JD: INTERNIST-I properties: Representing common sense and good medical practice in a computerized medical knowledge base. *Comput Biomed Res* 1985; 18:458-479
- Miller RA, Pople HEJ, Myers JD: Internist-I, an experimental computer-based diagnostic consultant for general internal medicine. *N Engl J Med* 1982; 307:468-476
- Dhar V, Stein R: *Seven Methods for Transforming Corporate Data into Business Intelligence*. Upper Saddle River, NJ, Prentice Hall, 1997
- Coiera EW: Artificial intelligence in medicine: The challenges ahead. *J Am Med Inf Assoc* 1996; 3:363-366
- Buchman TG, Kubos KL, Seidler AJ, et al: A comparison of statistical and connectionist models for the prediction of chronicity in a surgical intensive care unit. *Crit Care Med* 1994; 22:750-762
- Frize M, Solven FG, Stevenson M, et al: Computer-assisted decision support systems for patient management in an intensive care unit. *Medinfo* 1995; 8:1009-1012
- Tu JV, Guerriere MR: Use of a neural network as a predictive instrument for length of stay in the intensive care unit following cardiac surgery. *Comput Biomed Res* 1993; 26:220-229
- Doig GS, Inman KJ, Sibbald WJ, et al: Modeling mortality in the intensive care unit: comparing the performance of a backpropagation, associative-learning neural network with multivariate logistic regression. *Proc Annu Symp Comput Applications Med Care* 1993; 361-365
- Dybowski R, Weller P, Chang R, et al: Prediction of outcome in critically ill patients using artificial neural network synthesised by genetic algorithm. *Lancet* 1996; 347:1146-1150
- Izenberg SD, Williams MD, Luterman A: Prediction of trauma mortality using a neural network. *Am Surg* 1997; 63:275-281
- Maglaveras N, Stamkopoulos T, Pappas C, et al: An adaptive backpropagation neural network for real-time ischemia episodes detection: development and performance analysis using the European ST-T database. *IEEE Trans Biomed Eng* 1998; 45:805-813
- Baxt WG: A neural network trained to identify the presence of myocardial infarction bases some decisions on clinical associations that differ from accepted clinical teaching. *Med Decis Making* 1994; 14:217-222
- Baxt WG: Application of artificial neural networks to clinical medicine [see comments] [Review]. *Lancet* 1995; 346:1135-1138
- Baxt WG: Use of an artificial neural network for the diagnosis of myocardial infarction [published erratum appears in *Ann Intern Med* 1992; 116:94] [see comments]. *Ann Intern Med* 1991; 115:843-848
- Baxt WG, Skora J: Prospective validation of artificial neural network trained to identify acute myocardial infarction [see comments]. *Lancet* 1996; 347:12-15
- Si Y, Gotman J, Pasupathy A, et al: An expert system for EEG monitoring in the pediatric intensive care unit. *Electroencephalogr Clin Neurophysiol* 1998; 106:488-500
- Veselis RA, Reinsel R, Sommer S, et al: Use of neural network analysis to classify electroencephalographic patterns against depth of midazolam sedation in intensive care unit patients. *J Clin Monitoring* 1991; 7:259-267
- Spencer RG, Lessard CS, Davila F, et al: Self-organising discovery, recognition and prediction of haemodynamic patterns in the intensive care unit. *Med Biol Eng Comput* 1997; 35:117-123
- Hanson CW, Weiss Y, Frasch F, et al: Neuro-fuzzy analysis of hemodynamic data. *Abstr. Anesthesiology* 1998; 89(suppl):474
- Mason DG, Ross JJ, Edwards ND, et al: Self-learning fuzzy control of atracurium-induced neuromuscular block during surgery. *Med Biol Eng Comput* 1997; 35:498-503
- Mason DG, Ross JJ, Edwards ND, et al: Self-learning fuzzy control with temporal knowledge for atracurium-induced neuromuscular block during surgery. *Comput Biomed Res* 1999; 32:187-197
- Martin JF: Fuzzy control in anesthesia [Editorial; comment]. *J Clin Monitoring* 1994; 10:77-80
- Tsutsui T, Arita S: Fuzzy-logic control of blood pressure through enflurane anesthesia [see comments]. *J Clin Monitoring* 1994; 10:110-117
- Zbinden AM, Feigenwinter P, Petersen-Felix S, et al: Arterial pressure control with isoflurane using fuzzy logic. *Br J Anaesth* 1995; 74:66-72
- Hanson CW, Weiss Y, Frasch F, et al: A fuzzy control strategy for postoperative volume resuscitation. *Abstr. Anesthesiology* 1998; 89(suppl):475
- Sun Y, Kohane I, Stark AR: Fuzzy logic assisted control of inspired oxygen in ventilated newborn infants. *Proc Ann Symp Comput Applications Med Care* 1994:756-761
- Ying H, McEachern M, Eddleman DW, et al: Fuzzy control of mean arterial pressure in postsurgical patients with sodium nitroprusside infusion. *IEEE Trans Biomed Eng* 1992; 39:1060-1070
- Ying H, Sheppard L, Tucker D: Expert-system-based fuzzy control of arterial pressure by drug infusion. *Med Prog Technol* 1988; 13:203-215
- Ying H, Sheppard LC: Real-time expert-system-based fuzzy control of mean arterial pressure in pigs with sodium nitroprusside infusion. *Med Prog Technol* 1990; 16:69-76
- Oshita S, Nakakimura K, Kaieda R, et al: [Application of the concept of fuzzy logistic controller for treatment of hypertension during anesthesia] [Japanese]. *Masui* 1993; 42:185-189
- Fukui Y, Masuzawa T. [Development of fuzzy blood pressure control system] [Japanese]. *Iyo Denshi to Seitai Kogaku* 1989; 27:79-85
- Nordio M, Giove S, Lorenzi S, et al: Projection and simulation results of an adaptive fuzzy control module for blood pressure and blood volume during hemodialysis. *ASAIO J* 1994; 40:M686-M690
- Nordio M, Giove S, Lorenzi S, et al: A new approach to blood pressure and blood volume modulation during hemodialysis: an adaptive fuzzy control module. *Int J Artif Organs* 1995; 18:513-517
- Nemoto T, Hatzakis GE, Thorpe CW, et al: Automatic control of pressure support mechanical ventilation using fuzzy logic. *Am J Respir Crit Care Med* 1999; 160:550-556
- Schaublin J, Derighetti M, Feigenwinter P, et al: Fuzzy logic control of mechanical ventilation during anaesthesia. *Br J Anaesth* 1996; 77:636-641
- Kaufmann R, Becker K, Nix C, et al: Fuzzy control concept for a total artificial heart. *Artif Organs* 1995; 19:355-361
- Goldman JM, Cordova MJ: Advanced clinical monitoring: considerations for real-time hemodynamic diagnostics. *Proc Annu Symp Comput Applications Med Care* 1994; 752-755
- Sailors RM, East TD, Wallace CJ, et al: A successful protocol for the use of pulse oximetry to classify arterial oxygenation into four fuzzy categories. *Proc Annu Symp Comput Applications Med Care* 1995; 248-252

44. Wolf M, Keel M, von Siebenthal K, et al: Improved monitoring of preterm infants by Fuzzy Logic. *Technol Health Care* 1996; 4:193-201
45. Muller B, Hasman A, Blom JA: Building intelligent alarm systems by combining mathematical models and inductive machine learning techniques. Part 2—Sensitivity analysis. *Int J Biomed Comput* 1996; 42: 165-179
46. Muller B, Hasman A, Blom JA: Evaluation of automatically learned intelligent alarm systems. *Comput Methods Programs Biomed* 1997; 54:209-226
47. Heindl B, Schmidt R, Schmid G, et al: A case-based consiliarius for therapy recommendation (ICONS): computer-based advice for calculated antibiotic therapy in intensive care medicine. *Computr Methods Programs Biomed* 1997; 52:117-127
48. Schmidt R, Pollwein B, Filipovici L, et al: Adaptation and abstraction as steps towards case-based reasoning in the real medical world: case-based selection strategies for antibiotics therapy. *Medinfo* 1995; 8:947-951
49. Schmidt R, Heindl B, Pollwein B, et al: Multiparametric time course prognoses by means of case-based reasoning and abstractions of data and time. *Med Inf* 1997; 22: 237-250
50. Gade J, Rosenfalck A, van Gils M, et al: Modelling techniques and their application for monitoring in high dependency environments—learning models. *Comput Methods Programs Biomed* 1996; 51:75-84
51. Nikiforidis GC, Sakellaropoulos GC: Expert system support using Bayesian belief networks in the prognosis of head-injured patients of the ICU. *Med Inf* 1998; 23:1-18
52. Hanson CW, Marshall C, Medsker C, et al: A graphical display of clinical cardiopulmonary data. Abstr. *FASEB J* 1996; 10(1698):295
53. Powsner SM, Tuftte ER: Graphical summary of patient status. *Lancet* 1994; 344: 386-389
54. Powsner SM, Tuftte ER: Summarizing clinical psychiatric data. *Psychiatr Serv* 1997; 48: 1458-1461

ATTENTION: ADVERTISERS

If you want your information to reach 15,000+ readers, please contact: Cathy Chapman, Advertising, Lippincott Williams & Wilkins, 351 W. Camden Street, Baltimore, MD 21201-2436; phone: (410) 528-8536; fax: (410) 528-4452.