International Journal of Clinical Monitoring and Computing 6: 137–156, 1989.

Decision analysis: A framework for critical care decision assistance

Adam Seiver[1] & Samuel Holtzman
Strategic Decisions Group, 2400 Sand Hill Road, Menlo Park, CA 94025, USA; [1] Surgery Service (112), VAMC, Palo Alto, 3801 Miranda Avenue, Palo Alto, CA 94304, USA

Accepted 21 February 1989

Key words: artificial intelligence, data management, decision analysis, decision-making, ICU, intelligent decision systems

Abstract

The ultimate goal of medical computer systems is to help clinicians make good decisions. Such systems must be based on sound principles. Decision analysis is a 25-year-old discipline that provides the needed rigorous foundation for decision assistance. Decision analysis comprises the philosophy, procedures, and tools that can correct the flaws in existing critical care decision-making practice. Intelligent decision systems – computer-based systems that automate decision analysis – make it practical to apply decision analysis to critical care. *Orchestra* is a pilot intelligent decision system (now under development) that coordinates the efforts of the critical care specialist, the bedside physician, and the bedside nurse in building decision models that can provide recommendations and insight for ventilator management decisions. Decision analysis delivered by intelligent decision systems has great potential for improving critical care decision-making.

Decision analysis: A framework for critical care decision assistance

Over the next five years, computer-based systems will replace the paper chart in many critical care units [1]. Powerful, easy-to-use bedside workstations will acquire, store, and display a comprehensive set of patient data, including the patient's history and physical exam, physiological variables, laboratory results, and radiographic images. Using such a system, the clinician will have vastly better access to facts about his or her patients than he now has, using the existing manual charting methods. However, having ready access to facts leaves the clinician with the difficult task of choosing what to do. It is natural to expect that the computer should provide assistance here, as well. Reed Gardner [2], President of Computers in Critical Care and Pulmonary Medicine has said: 'The ultimate goal of a medical computer system is, after all, to assist physicians in making medical decisions.'

Systems that provide clinical decision assistance will significantly affect clinical practice. However, to be truly useful, these systems will need to be based on sound principles. Decision analysis offers a rigorous framework for designing and implementing computer-based systems that offer decision assistance. In this paper, we begin by examining existing critical care decision-making practices. Next, we present the decision analysis approach to decision-making, including key decision-analytic concepts and techniques. We then introduce intelligent decision systems – computer-based systems that automate decision analysis. Following this, we discuss how decision analysis can be applied in the critical care environment given this environment's special decision-making features. Finally, we briefly describe a pilot intelligent decision system, *Orchestra*, that we are developing for ventilator management.

Current critical care decision practices

Let's start by looking at a typical scenario that illustrates current critical care decision practices.

Mr. A is admitted to the surgical intensive care unit in shock. The physician decides to institute monitoring with radial and pulmonary artery lines. After several attempts, the physician determines that the radial artery line cannot be placed percutaneously. She considers a percutaneous femoral artery line, but assesses the risk of thrombosis as too high. Finally, she succeeds at inserting the catheter through a radial artery cut-down.

One hour after admission, approximately 20 physiological variables are being measured continually. The physician studies the bedside flowsheet and notes that the patient has a mean arterial pressure of 125 mmHg and a very high systemic vascular resistance.

The physician decides to use a nitroprusside infusion to reduce afterload while maintaining preload with a crystalloid infusion. Accordingly, before leaving, the physician writes the following orders for the nurses:

1. Bolus with saline to keep wedge pressure at 10–15 mmHg.
2. Titrate nitroprusside to keep mean arterial pressure at 80–90 mmHg.

When the physician returns four hours later, she notes with satisfaction that the patient's resistance has decreased, that the wedge pressure is 15 mmHg, and that the mean arterial pressure is 85 mmHg.

Twelve hours later, however, the patient has received three liters of saline, is nauseated from the nitroprusside, and is *still in shock*.

At first glance, the patient management illustrated in the scenario appears satisfactory and its ineffectiveness at 12 hours seems surprising. However, when we review the case in greater detail, we can discover at least four different decision-making defects that contribute to this ineffectiveness: 1) sparse alternatives; 2) information overload; 3) superficial objectives; and 4) ineffective delegation of decision-making.

Sparse alternatives

Sparse alternatives lead to missed opportunities. In the scenario, the patient is admitted to the intensive care unit for resuscitation, and the physician is unable to place a radial artery line percutaneously. At that point, she considers her alternatives: percutaneous femoral artery line versus radial artery line placed by cut-down. She decides to perform a cut-down. However, has the physician considered all the pertinent alternatives? Perhaps *no line* is a reasonable alternative, and other monitoring technologies could be substituted. For example, arterial pulse oximetry and end-tidal carbon dioxide monitoring can reduce the need for arterial blood gases. Blood pressure can be measured automatically with self-inflating cuffs. Mixed venous oximetry can provide corroborating evidence about perfusion. These alternative technologies can be substituted singly or in combination. Critical care decision-making suffers when the decision-maker considers only a narrow range of options.

Information overload

Information overload can put the patient in danger. In the scenario, the physician must interpret the large amount of patient data organized in time-oriented fashion on the flowsheet reproduced in Table 1. She correctly notes that the mean arterial pressure, 125 mmHg, and the systemic vascular resistance, 4,232 dyne sec cm^{-5}, are quite high. However, because of the large array of numbers confronting her, she fails to note that the pulmonary artery occlusion pressure (the wedge pressure), which is recorded as 13 mmHg, is 12 mmHg higher than the pulmonary artery diastolic pressure, which is recorded as 1 mmHg. The physician should know that, usually, this situation indicates that the wedge pressure value is wrong, but she fails to take advantage of this knowledge because she is distracted by the large amount of other numbers she needs to consider. Thus, she erroneously concludes that the patient needs afterload reduction. In fact, the patient simply needs volume. Although the physician prescribes saline to maintain preload, this only compensates for the preload-reducing effects of the nitroprusside and the patient remains in hypovolemic shock.

This example illustrates the problem that has been well articulated by Roger Bone 'Possibly the most serious danger . . . is that of drowning a physician in a flood of numbers. The presentation of more data that can be assimilated can contribute to incorrect clinical decisions' [3].

Superficial objectives

Superficial objectives can lead to misguided action. In the scenario, the physician returns after four hours to evaluate the patient's response to therapy and notes that the patient's resistance has decreased, that the wedge pressure is 15 mmHg, and that the mean arterial pressure is 85 mmHg. She feels the therapeutic goals are being accomplished. In fact, the patient remains in shock: cardiac output is low, and there is inadequate perfusion to meet the patient's oxygen needs. From the point of view of survival, the patient has not been adequately resuscitated. This illustrates how clinicians often choose superficial objectives that do not promote ultimate objectives, a problem that William Shoemaker has frequently discussed: 'The traditional approach usually assumes that normal values are the appropriate therapeutic goals ... Therapy should restore physiological defects not just to their normal values but to their optimal values' [4].

Table 1. The flowsheet for the patient in the scenario overwhelms the clinician with a large array of data. Note that the pulmonary artery occlusion pressure, the wedge, is greater than the pulmonary artery diastolic pressure – easily overlooked evidence that the wedge is inaccurate.

Variable	Time							
	7	8	9	10	11	12	13	14
Temperature	34.5	34.5	34.6	34.6	34.6	34.6	34.6	34.6
Respiratory Rate	8	8	8	8	8	8	8	8
Tidal Volume	1200	1200	1200	1200	1200	1200	1200	1200
PaCO2	33	27	30	28	32	32	32	35
PaO2	67	62	65	69	62	70	75	71
pH	7.48	7.57	7.53	7.55	7.50	7.50	7.49	7.45
ETCO2	25	18	20	18	22	22	22	21
SaO2	0.94	0.93	0.93	0.93	0.93	0.93	0.93	0.90
SvO2	0.45	0.36	0.42	0.55	0.56	0.55	0.53	0.52
HR	101	102	101	101	102	104	105	103
Radial Systolic	180	170	150	118	120	117	117	116
Radial Diastolic	100	68	68	69	70	71	71	71
Radial Mean	125	101	94	85	86	85	85	85
Pulmonary Systolic	40	42	43	42	40	42	41	41
Pulmonary Diastolic	1	2	3	3	4	3	4	4
Pulmonary Mean	14	15	16	16	16	16	16	16
Central Venous	17	18	18	17	18	18	18	18
Wedge	13	14	15	15	16	15	16	16
Resistance	4232	3867	3153	2071	2021	2080	2156	2070
Cardiac Output	2.1	1.7	1.9	2.6	2.7	2.6	2.5	2.6

Ineffective delegation

Ineffective delegation of decision-making erodes the quality of patient care. One compelling feature of critical care delivery is that it must be sustained 24 hours a day, 365 days a year. Therefore, those individuals with ultimate responsibility or with valuable expertise simply cannot be continuously present at the patient's bedside. For example, the attending surgeon or the pulmonary specialist have responsibilities in the clinics, in the operating room, in the wards, in the office, and in other critical care units. (And health providers also need time for their personal lives). While a few critical care units now provide in-house intensive coverage around the clock, collaborative decision-making – particularly between physicians and nurses – is essential to high-quality care. In a prospective study of treatment and outcome in 5,030 patients in intensive care units at 13 tertiary care hospitals, death rates (adjusted for severity of illness) significantly diminished when there was improved interaction and coordination between the intensive care unit staff [5].

Currently, two mechanisms exist for accomplishing this coordination: 1) implicit orders; and 2) standing orders. Implicit orders are unwritten guides to decision-making. For example, there is generally no written rule in critical care units that states that if a ventilated patient becomes unstable, the fractional inspired oxygen concentration should be increased to 100 percent while the patient is resuscitated. However, the staff is expected to know enough to increase the oxygen concentration based on training and experience. Likewise, the staff is expected to know that when a surgical patient has a urine output that is less than 0.5 cc kg per hour, it must manage the possibly poor renal perfusion by addressing overall perfusion (usually with volume, occasionally with inotropes) by calling the physician, or by performing both actions.

When all ICU team members are in constant contact and when the team shares a large body of patient experience, unwritten orders are effective, because all members implicitly know what to do and why. However, constant staff turnover and shift-work fragment this background of understanding and severely disrupt coordination based on group intuition.

For example, Mr. Jones, a new nurse, may not know that Mr. A.'s urine output should be maintained at 50 cc per hour. If Ms. Smith, the nurse during the preceding eight hours, does not convey that fact during report at change of shift (a not-too-infrequent occurrence), then Mr. A. may be left dangerously oliguric overnight. In short, implicit orders are susceptible to breaking down under the stress of realistic staffing conditions.

Standing orders are used to fill the gaps that may occur with implicit orders. For example, the physician could obviate unreported oliguria in Mr. A. by writing the following orders.

1. Measure urine output every hour.
2. Call physician for urine output less than 50 cc/hr.

Standing orders aim to decrease errors of *omission*, but they often promote serious errors of *commission*. *Excessive testing* is one consequence of standing orders. For example, to ensure that anticoagulation therapy is closely monitored, the physician might write the order:

> Send blood to laboratory for partial thromboplastin time every four hours.

Such an order will ensure that the laboratory tests are sent off regularly, even if the physician is not physically present to make the request. Unfortunately, the lab ordering system does not always recognize when these orders become obsolete. Joseph Civetta studied testing for coagulation in his unit and found: '[There was] repetitive testing of coagulation parameters in many patients who showed no evidence of a coagulation disorder' [6]. Civetta advocates eliminating standing orders for laboratory tests, but this requires developing a better mechanism for delegating decision-making.

Another consequence of delegating decision-making with standing orders is *tail chasing*. Tail chasing was illustrated in the scenario when the physician left her standing orders.

1. Bolus with saline to keep wedge pressure at 10–15 mmHg.
2. Titrate nitroprusside to keep mean arterial pressure at 80–90 mmHg.

In following these orders, the nurses gave the patient three liters of saline; however, the patient remained in shock, despite the fact that the basic defect was hypovolemia. Administering nitroprusside decreased both the mean arterial pressure and the wedge pressure, which, in turn, caused the nurse to infuse a bolus of saline. This action brought the wedge pressure back up to the desired value, but it also increased the preload and, thus, the mean arterial pressure. As a result, the nurse increased the nitroprusside dosage, which decreased the wedge pressure, which necessitated further saline, and so on. The increased venous capacitance caused by the nitroprusside prevented the massive saline infusion from correcting the hypovolemic defect. These positive feedback situations can be created whenever the management of therapies with counteracting effects (e.g., nitroprusside and dopamine, crystalloid infusion and diuretics) is ineffectively delegated.

The four defects revealed in the scenario – sparse alternatives, information overload, superficial objectives, and ineffective delegation – are symptomatic of a single underlying defect: faulty decision-making. Each of these problems may be viewed as the result of the critical care staff committing to an action that is inconsistent with either what they can do, what they know, or what they really want. Decision analysis provides a framework for avoiding these inconsistencies and for thereby improving the quality of critical care decision-making.

Decision analysis approach to decision-making

Decision analysis comprises the philosophy, methodology, and professional discipline for ensuring high-quality decision-making. While Professor Ronald A. Howard of Stanford University coined the term 'Decision Analysis' in 1964 [7], the roots of decision analysis date back to the work of two great mathematicians, P.S. Laplace [8] and D. Bernoulli [9]. However, it was the 1960s computer revolution that made the professional application of decision analysis possible, and since 1965, decision analysis has been regularly taught at the graduate level at Stanford University and at many other universities. Decision analysis is now a growing professional field, with a two-decade success record.

Science is a *descriptive* discipline: it studies what *is* by describing it. And engineering is a *prescriptive* discipline: it creates what *should be* by designing it. Decision analysis is a *normative* discipline: it is a prescriptive discipline guided by a set of *norms* – that is, principles of right action.

Decision analysis focuses on bringing clarity of action to difficult decisions. By an *action*, we mean the irrevocable allocation of valuable resources. By a *decision*, we mean the commitment to irrevocably allocate valuable resources. Decision analysis can address a wide range of decisions, but it is particularly well suited for decisions involving complexity, dynamics, and uncertainty.

Because of its cost (typically tens or hundreds of thousands of dollars) and the long time necessary to carry out an analysis (around one hundred person-hours), professional decision analysis has been almost exclusively applied within business and industry. However, in an academic setting, decision analysis has also been successfully applied to many medical decisions [10]. Later in this paper, we will discuss how (through the use of intelligent decision systems) professional-level decision analysis can be made much less costly and faster for use in a clinical setting.

Using decision analysis effectively requires understanding its *philosophy*, *procedures*, and *tools*. The decision analysis philosophy fundamentally defines high-quality decision-making in the form of key concepts and distinctions. The procedures of decision analysis constitute an extensive array of techniques to capture and reason about all aspects of a decision. The tools of decision analysis greatly facilitate the decision-analytic process and make it efficient and easy to use.

In decision analysis, *decision theory* provides the general and sound framework for recommending a course of action – *given a decision model.* Decision theory takes as an input a mathematical model of a

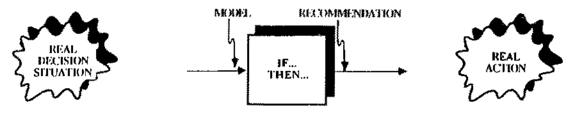

Fig. 1. Decision theory is a conditional statement.

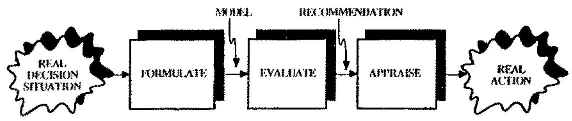

Fig. 2. Formally capturing real decision situations and interpreting formal recommendations are major tasks in decision analysis.

decision problem and provides as an output a prescription for action. Decision theory can therefore be viewed as a conditional statement, as shown in Fig. 1.

Decision analysis is based on decision theory much like medicine is based on biology. Just as physicians must devote considerable time and effort to formulating a medical problem in biological terms and to interpreting any biological conclusions in terms that are meaningful for patient care, so decision analysts must devote the time and effort to formally capture the decision as a decision model and then to interpret the formal recommendations resulting from applying decision theory to the model.

Given this situation, we can thus view the practical use of decision theory as a three-stage process (Fig. 2) whose three stages are *formulation* (i.e., developing the formal decision model), *evaluation* (i.e., computing a recommendation from the model), and *appraisal* (i.e., interpreting the formal recommendation).

However, this strictly sequential approach to using formal decision methods has a major shortcoming-it does not account for the likely disagreement between the decision-maker and the method's recommendation. In fact, such disagreement is almost certain to occur. Given that the decision-maker requires assistance to gain new insight into his problem, we can assume he is having difficulty dealing with his decision. Therefore, formally ana-

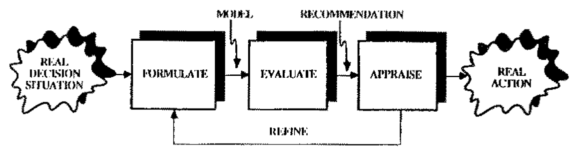

Fig. 3. By producing a sequence of increasingly refined decision models, decision analysis generates the insight necessary for action.

lyzing the decision will probably expose many of the inconsistencies and lack of focus that made the decision difficult in the first place. Moreover, such disagreement is very beneficial, because it exposes important flaws in either the decision-maker's understanding of his decision (i.e., how he preceives and interprets it) or his logic.

A simple way to deal with the possible unacceptability of a formally obtained recommendation – in fact to take advantage of it – is to extend the sequential process by explicitly adding a *feedback* path, as shown in Fig. 3. Such a closed-loop decision process allows the decision-maker to react to any surprising element of the formal prescription by reevaluating and possibly modifying his formulation. Alternatively, if after developing enough insight he agrees with the suggested strategy or if he determines that his disagreement results solely from logical error, he may choose to follow the formal recommendation. Hence, by producing a sequence of increasingly refined decision models, we can help the decision-maker develop the insight necessary for action.

The closed-loop decision process described in Fig. 3 can be viewed as a blueprint for a conversation, which is illustrated in Fig. 4. It involves two key participants; the decision-maker and a decision analyst. During the formulation stage of the process, the decision-maker *teaches* the details of the decision at hand to the decision analyst, who *learns* by building an appropriate decision model. These activities are reserved during the appraisal stage, where the decision analyst *teaches* the decision-maker the implications of the formal recommendation for action obtained during the evaluation stage. Most of the insight developed in the closed-loop decision process shown in Fig. 4 results from

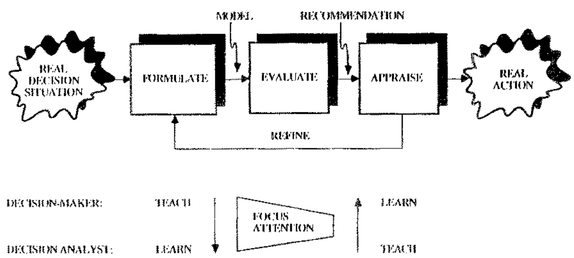

Fig. 4. Decision analysis is a carefully engineered conversation that develops insight by focusing attention on the key aspects of the decision at hand.

this interchange of information and new knowledge between the decision-maker and the decision analyst. Moreover, the formal machinery embodied in the evaluation stage guides and focuses this interchange. This attention-focusing effect assits the decision participants in producing an increasingly simple, yet representative, model of the decision as the process progresses.

Key decision analysis concepts

As the above discussion makes clear, the *decision-maker* plays a central role in decision analysis. The decision-maker either owns the resources to be allocated or is acting in the best interests of their owner. An *expert* is a source of information and alternatives, but he is *not* a source of recommendations. And a *decision analyst* is an expert on process – not content – who guides the conversation toward clarity of action.

We must also be precise about our use of decision terms in medicine. *Diagnosing* consists of thinking about the patient's condition. *Decision-making* consists of thinking about what to do, given a possibly uncertain and incomplete diagnosis. *Treating* consists of action – of doing something to the patient whether that action is diagnostic, therapeutic, or both. Given these definitions, it would thus be incorrect to say: 'I have decided that the patient has appendicitis.' One can *decide* on 'appendectomy,' but one *diagnoses* 'appendicitis.'

Decision-making has three fundamental components: *alternatives* (what you *can do*), *preferences* (what you *want*), and *information* (what you *know*). Key concepts related to each are now considered in turn.

Alternatives

Alternatives should encompass a wide range of possible action, generally both inside and outside strictly medical dimensions. For example, changes in a patient's personal lifestyle and work activities are an important component of radiation therapy and must be part of the definition of that alternative. Alternatives must also account for the patient's specific circumstances. For example, the patient's economic situation, the presence or absence of a living will, and the presence or absence of relatives willing to donate organs are all special constraints that may expand or restrict the set of possible options.

Decisions are often difficult to make because there is no *dominating alternative*. A dominating alternative is an alternative that would be recommended to the vast majority of the patients with a particular set of findings or diagnosis. For example, appendectomy is the dominating alternative for appendicitis. In contrast, there is no dominating alternative for infertility due to blockage of the fallopian tubes. Celiotomy, in-vitro fertilization, and doing nothing are all reasonable options. In such cases, doctors must know more about the problem – principally about the patient's preferences – before they can identify the best alternative [15].

Preferences

Preferences directly represent the desires of the decision-maker, who (as noted earlier) either owns or acts on behalf of the owner of the recources to be allocated. In medicine, preferences almost always concern the patient's length of life (lifetime), personal and work life (lifestyle), overall well-being (comfort), and financial and other economic resources such as health maintenance contracts (wealth). Achieving clarity of action requires explicitly quantifying the trade-offs that the patient wants made among these fundamental attributes.

Measuring preferences is particularly challenging in making medical decisions, because possible outcomes are often unfamiliar to those who will bear them. For example, the patient undergoing radiation therapy for the first time probably has little experience that can allow him or her to translate descriptions of the possible complications into personal terms. Implicit guardianship is often another complicating factor. When the patient cannot make choices for himself or herself, whose preferences should be used? We believe that the preferences used should be those of the patient – the person whose resources are at stake – or, if these are unavailable, those of someone (i.e., a guardian) who understands the patient well enough to act on his or her behalf with his or her best interests in mind. In many circumstances – but not always – this individual would be the bedside physician. Frequently, however, there are multiple stakeholders: the patient, family members, physicians, nurses, and governmental agencies. And even when the source of the preferences is clear, ethical concerns may still make trade-offs difficult. Because of these challenges, decision-analytic methods are the most robust and humane way to deal with difficult preferences, because these methods are explicit, comprehensive, and incisive.

It is important not to confuse direct (primary) preferences with indirect (secondary) preferences. For example, while a patient has an *indirect* preference on his or her mean arterial pressure, he or she has a *direct* preference on survival. In medicine – as in most human endeavors – indirect preferences are a useful means of delegation. However, making decisions requires being aware of primary preferences to avoid pursuing objectives that have become obsolete. For example, keeping the mean arterial pressure normal may be *desirable*, but it is secondary to keeping the patient *alive*.

Information

Information – knowledge about the possible consequences of pertinent actions – is essential to decision-making. In decision analysis, information, both certain and uncertain, is treated explicitly. This information can take one of two forms: Structural and Parametric. *Structural* information specifies relations among decision elements. For example, the relationship between FIO_2 and PIO_2 given by the alveolar gas equation is structural. *Parametric* information specifies decision elements individually. For example, the value of the cardiac output is parametric.

Decision analysis treats uncertainty explicitly. Consequently, decision-analytic recommendations effectively reflect the decision-maker's uncertain situation. In particular, decision analysis can yield optimal recommendations that would be discarded if uncertainty were ignored. These recommendations are often referred to as 'hedging' alternatives. For example, consider a patient who has just suffered a myocardial infarction and is demonstrating second-degree atrio-ventricular block. If we know for sure the patient will develop fixed, complete heart block, then we should put in a permanent pacemaker. And if we know the patient will *not* develop complete heart block, then no pacemaker is indicated. In fact, however, we have only incomplete knowledge about whether the patient will develop fixed, complete heart block, so we hedge by placing a transvenous pacemaker. This alternative is inferior to the permanent pacemaker if heart block is present and inferior to no pacemaker if heart block is absent. However, the hedging alternative is preferable to both if heart block is *uncertain*.

In capturing uncertain information, decision analysis uses the rigorous methods of probability. Specifically, decision analysis views probability from the perspective of its inventors (e.g., P.S. Laplace and D. Bernoulli): probability is a state of *information, not* a state of *nature*. For example, imagine a 35-year-old man who presents to the emergency department with mild chest pain. An EKG is performed and blood work is drawn. You are the physician called to the emergency department to evaluate the patient. What is your probability that the patient has had a myocardial infarction? Based on what you know – mild chest pain, young patient – you would presumably assess this probability to be quite low. However, suppose you discover during your interview with the patient that his father and grandfather both died suddenly at age 35. Now your index of suspicion has increased, and you assess a higher probability. Next, the technician hands you your copy of the EKG, which shows only nonspecific ST-T wave changes. Perhaps now your suspicions are lessened and you decrease your probability of a myocardial infarction. Finally, the laboratory work is returned – the CPK is quite high with a significantly elevated MB fraction. Now your probability has significantly increased and myocardial infarction has become your diagnosis. Throughout all this diagnostic effort, the patient – the state of nature – has, of course, not changed. What *has* changed is your state of information. You revise your probability assessment to accommodate the changes in your state of information. In other words, there is no single 'correct' probability – it depends on what you know.

Because uncertainty is involved in most important decisions, we must distinguish the quality of decisions from the quality of outcomes. An example illustrates the point. Suppose you were offered the opportunity to buy a ticket to a lottery for $1 that offered a 1-in-1,000 chance of winning $100 million (tax free). While this investment would be outstanding for most individuals, it involves a 99,9 percent chance of a bad outcome! Suppose you invested and lost, would purchasing the ticket have been a good or bad decision? Would you invest again in an identical deal? The point is that because the consequences of actions may be uncertain, it is possible to make a good decision and get a bad outcome. (All other combinations of good/bad decision and good/bad outcome are, of course, also possible – including making a bad decision and getting a good outcome.)

The clinical approach to appendicitis provides a medical example that illustrates the difference between the quality of outcomes and the quality of decisions. For example, a 20-year-old man presents with nausea, epigastric pain localizing in the right lower quadrant, and point tenderness at McBurney's point. There is fever and leukocytosis. As a result, the patient undergoes an appendectomy. Most surgeons would consider this a good decision. However, suppose the appendix is found to be normal during the operation. Was it still a good decision to operate? From a decision-analytic perspective, performing the appendectomy would be considered a good decision, because it was consistent with the decision-maker's alternatives, preferences, and knowledge *at the time the decision was made.*

A logical consequence of this distinction is that we should focus on high-quality decision-making, not on high-quality outcomes. Decision analysis shows us that the quality of decisions and the quality of outcomes should be measured separately. However, most individuals' performance is measured in terms of their outcomes, not their decisions. Unfortunately the price of rewarding outcomes and not decisions is *bad decisions.* Defensive medicine illustrates this effect. For example, a thirty-year-old man is admitted for hernia repair. A detailed history and physical examination is unremarkable except for the hernia. Intraoperatively, the patient has a massive myocardial infarction and dies. The surgeon is sued for not obtaining a pre-operative electrocardiogram. Now the surgeon obtains a pre-operative electrocardiogram for all his patients – regardless of indications. Having been sued because of a bad outcome that resulted from an extremely rare event, the surgeon thus makes bad decision-making a routine part of his practice.

Decision models in decision analysis

A *decision model* is a formal representation of a decision problem. A good decision model captures the decision at hand explicitly, succinctly, and unambiguously. Because it is explicit, all important aspects of the decision are available for review. Because it is succinct, only important aspects are represented. Because it is unambiguous, all model elements are clearly defined.

Influence diagrams are the foremost way of representing decision models in modern decision analysis [11]. Influence diagrams are easy to understand, mathematically well defined, very general, and compact. In general, influence diagrams have significant theoretical and practical advantages

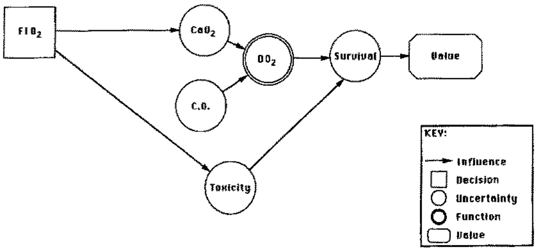

Fig. 5. Influence diagrams can represent FIO_2 adjustment decisions. The choices for FIO_2 (ranging from 21 percent to 100 percent) are represented by the rectangular decision node. The physiological and clinical variables that mediate the effect of FIO_2 on survival are represented by the oval chance nodes (which may be uncertain, e.g., toxicity; or deterministic, e.g., DO_2). The patient's preference for survival is represented by the value node. (Abbreviations: FIO_2, fractional inspired oxygen concentration; CaO_2 oxygen content arterial blood; C.O., cardiac output; DO_2, oxygen delivery.)

over another commonly used decision representation language-decision trees. In addition to representing probabilistic independence effectively, enforcing a clear distinction between informational and probabilistic relationships, and preventing loss of information from asymmetries, influence diagrams grow linearly (as opposed to exponentially) with the size of the problem they represent and, thus, can be used to model much larger decisions than trees can model. These technical advantages are enhanced by the fact that the mathematical concept of influence, conditional as well as informational, is very close to its intuitive counterpart.

As defined by Howard and Matheson [12], an influence diagram is a singly connected, acyclic, directed graph with two types of nodes – decision and chance – two types of arrows or arcs – conditioning and informational – and typically a single sink node of type chance (Fig. 5). The acyclic, singly connected nature of influence diagrams implies that sets such as predecessors and successors of a node are defined in the usual manner.

Decision nodes are usually represented by a rectangle or a square and denote variables under the decision-maker's control. *Chance nodes* – usually represented by an oval or a circle – denote probabilistic variables. A special form of a chance node is a deterministic node, which is usually represented by a double-ringed oval or circle. The value of a deterministic node is known exactly if the value(s) of its predecessor node(s) are specified. *Conditioning arrows* are always directed toward a chance node and denote probabilistic dependence. *Informational arrows* are always directed toward a decision node and denote available information. An influence diagram usually (although not neccessarily) has a single chance node with no successors (i.e., a sink node), which is called the *value node* and which represents the decision-maker's direct preferences. Chance nodes without direct chance predecessors (i.e., chance source nodes) are called *border nodes*.

Figure 5 shows a simple influence diagram that represents the FIO_2 adjustment decision. The possible choices for the FIO_2 (e.g., 21 percent to 100 percent) are represented by the square decision node. The fractional inspired oxygen concentration affects the oxygen content of arterial blood (CaO_2) and the potential for oxygen toxicity. These are treated as probabilistic relationships and represented as chance nodes with conditioning arcs arising from the FIO_2 decision node. Cardiac output (C.O.) is represented as a chance border node. Oxygen delivery (DO_2) is represented as a deterministic node, because it is known for certain given CaO_2 and cardiac output – in other words, it is the product of the values of these predecessors of DO_2. Survival depends probabilistically on oxygen toxicity and oxygen delivery. Finally, we declare survival to be the only significant attribute of the possible outcomes by attaching the value node to survival alone.

Figure 5 shows the structure of the FIO_2 adjustment decision and in particular highlights the trade-offs and the key relationships. An influence diagram, however, can represent a decision completely, not just in terms of its structure. A full description of a decision problem requires that the diagram contain at least one decision node directly or indirectly influencing a value node and that consistent, detailed specifications exist for each node in the diagram. For decision nodes, the set of possible outcomes corresponds to the set of decision alternatives; for chance nodes, this set of outcomes corresponds to the sample space of the variable being represented. Furthermore, for chance nodes, a detailed description should also include a probability measure over the set of possible outcomes. An important, yet subtle, fact about probabilistic specifications of chance nodes is that they must be consistent with the set of direct predecessors of the node and their respective outcomes.

A structurally complete influence diagram whose nodes and relations have not been specified in detail is said to be defined at the level of *structure*. A diagram developed in all the necessary detail is defined in terms of both its *structure* and its *parameters*. An influence diagram is *well-formed* when it has been consistently defined both structurally and parametrically. Algorithms exist for computing the optimal policy from a well-formed influence diagram representing a decision [13, 14] and for deducing other important inferential results (e.g., value-of-information, value-of-control, and other sensitivity measurements).

Intelligent decision systems

As mentioned earlier, decision analysis is not widely used in medicine, because it is too expensive and too slow. Professional decision analysis has been almost exclusively applied in business and industry. By automating the decision analysis process, intelligent decision systems make decision analysis inexpensive and fast. Therefore, intelligent decision systems open the door for the wider medical application of decision analysis.

A *decision system* is a system that makes recommendations for action and is typically implemented on a computer. An *intelligent decision system* is a decision system that delivers expert-level decision analysis assistance [15]. As part of this assistance, the intelligent decision system may provide access to a substantial knowledge base in the domain of the decision.

Intelligent decision systems arise from the use of artificial intelligence technology to automate the formulation and appraisal skills of professional decision analysts in a well-defined decision arena. Intelligent decision systems are made possible by our ability to indentify and analyze in advance the common aspects of the decisions we face.

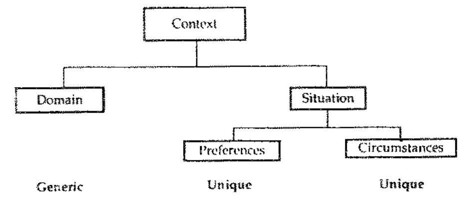

Fig. 6. A decision context is comprised of both generic and unique elements.

We refer to all the elements that are relevant to a decision as the *decision context* As shown in Fig. 6, the decision context can be hierarchically decomposed into its *domain* and its *situation*. The domain consists of the generic subject matter with respect to which the decision is being made. The situation consists of the *preferences* and the *circumstances* pertinent to the decision at hand. Preferences refer to a statement of the satisfaction the decision-maker receives from particular states of the world. Circumstances are the information, constraints, and alternatives in a decision that are unique to a specific decision-maker.

Expert systems technology makes it possible to bring specialized domain knowledge to a decision problem. Decision analysis makes it possible to incorporate circumstances and preferences into decision-making. Intelligent decision systems facilitate the incorporation of all elements – domain, preferences, and circumstances – into the decision. Therefore, intelligent decision systems are ideal for making high-quality decision-making assistance widely available. In an important sense, expert systems increase the *quantity* of decisions (a useful feature when good decision-making relies on the knowledge of a few key individuals whose expertise can thus be made widely available). In contrast, decision analysis increases the *quality* of decisions. By combining both technologies, intelligent decision systems can make *high-quality* decisions available in *quantity*.

Figure 7 shows a plausible architecture for an intelligent decision system [15]. This architecture consists of four interconnected parts: a general-purpose inference engine, a set of data structures, a corresponding set of specialized procedures, and user interface (or front end).

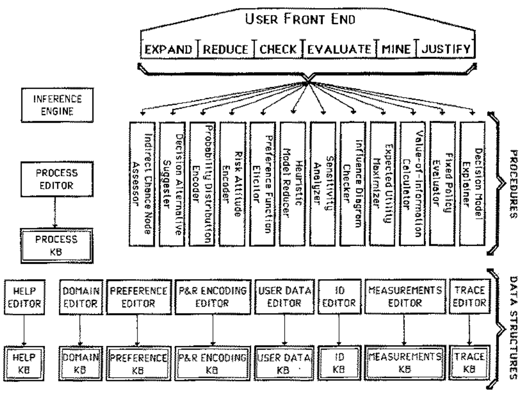

Fig. 7. An architecture for an intelligent decision system. (Abbreviations: KB, knowledge base; ID, influence diagram.)

The general purpose inference engine (illustrated on the top left corner of Fig. 7 – viewed from the side) interprets the knowledge bases throughout the system. In technical terms, its function is to efficiently implement a useful portion of the first-order predicate calculus and associated syllogisms.

The operation of the system revolves around a set of nine data structures. These structures, depicted as double-lined boxes along the left and bottom margins of Fig. 7, can be accessed through structure-specific editors. Editors allow all interactions with the data to occur with consistent syntax and semantics at a high level, which allows the rest of the system to be independent of the physical implementation of the data structures and helps ensure their reliability and integrity.

The data structures in this architecture are manipulated by a set of specialized procedures. Figure 8 depicts a set of twelve such procedures, which are representative of those that should be part of an intelligent decision system. However, a somewhat different set may be better suited in any given implementation. The leftmost five procedures shown deal with decision-model development. In particular, the indirect chance node assessor, the decision alternative suggester, and the preference function elicitor develop the influence diagram model both structurally and parametrically. The probability distribution encoder and the risk attitude encoder further develop the model parametrically.

An important part of the proposed architecture for an intelligent decision system is an interface program to interact directly with the user. This program facilitates the use of the procedures and data structures that constitute the intelligent decision system and adapts the system's interaction to the indentity and expertise of each individual user.

Applying decision analysis to critical care

In applying *decision* analysis to critical care, we must allow for special features of critical care decision-making, which include *delegated responsibility* and *distributed expertise*. Delegated reponsibility governs critical care physician and nurse decision-making. Fundamentally, the patient is the decision-maker, because it is primarily the patient's resources – his or her life and limb – that are at stake. In the critical care setting, the patient delegates decision-making responsibility to the physician, either *explicitly* when the patient is well enough to communicate or *implicitly* when the illness prevents such communication. The physician commits to action by providing a decision strategy. For example, the physician may decide that the patient should be placed on mechanical ventilation and receive hemodynamic life support. The nurse interprets and implements the physician's strategy into specific actions, such as adjusting the settings on the ventilator and the rates of infusion of the cardiac drugs.

Distributed expertise refers to the fact that different members of the critical care team are experts about different things. For example, the critical care physician specialist is most likely to be familiar with life-support technology and the pathophysiology of critical illness. The patient's attending physician is more likely to be familiar with the longitudinal nature of the patient's illness, having had the opportunity to interview the patient pre-operatively and to follow the patient from admission on the ward to the operating room to the intensive care unit. The nurse has the best perspective on the patient's minute-to-minute circumstances. A nurse is physically present at the bedside 24 hours per day and is the first to be aware of changes in the pa-

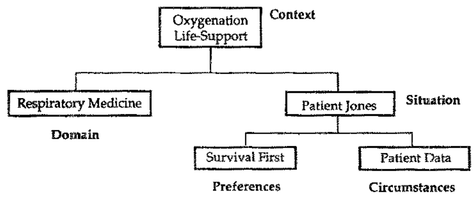

Fig. 8. The Oxygenation Life-Support decision context has a domain-Respiratory Medicine-and a situation defined by the patient's preferences and data.

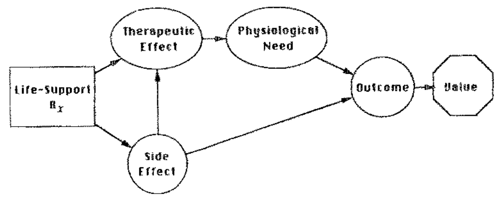

Fig. 9. A decision-framing template expresses the common organizing principles underlying life-support decisions.

tient's condition, as reflected in his or her history, physical exam, monitored observations, and laboratory measurements.

As noted above, we can distinguish within each decision a generic domain and a unique situation. The situation is further divided into preferences and circumstances. Thus, for example, within the context of oxygenation life-support decisions, we identify a body of knowledge applicable for all patients – Respiratory Medicine (Fig. 8). Every oxygenation life-support decision, however, will be made for a specific patient, and it is that patient's preferences and circumstances that define the situation. The patient's preferences may be very easy to describe (e.g., survival at all costs), or they may be difficult to capture, (e.g., subtle trade-offs between a desire to survive and a desire to die with dignity). The patient's circumstances not only include the unique constraints governing what can be done for the patient, but also the structural and parametric information about the patient, which correspond to data and diagnoses now typically recorded in the medical chart.

We can use the taxonomy of the decision context to accommodate the features of delegated responsibility and distributed expertise in critical care decision-making. Delegating decision-making corresponds to assigning responsibility for elements of the decision context. When the patient delegates to the physician, he or she assigns the physician the authority to determine each of the elements of the decision context. The physician may, in turn, wish to assign responsibility for the domain to a consultant critical care specialist and responsibility for the circumstances to the nurse who is continually present at the bedside. And he or she may wish to reserve responsibility for delineating the patient's preferences based on his or her personal contact with the patient. Of course, if the physician believes the nurse has a better rapport with the patient and family based on their continual contact at the bedside, then he or she may delegate responsibility for delineating the patients preferences to the nurse.

Thus, the taxonomy of the decision context allows decision-making to be delegated in a controlled fashion. The fact that different team members have different areas of expertise is gracefully handled at the same time by this 'divide and conquer' approach. Once the various elements of the decision context have been defined, decision analysis provides a methodology for logical synthesis into recommendations and insight.

Decisions about adjusting FIO_2 for a post-operative patient dependent on a ventilator illustrate how the decision analysis approach can be used. We start with the task of encoding respiratory medicine domain knowledge, a task that is assigned to the pulmonary specialist physician. We ask the physician to represent his or her knowledge so the bedside physician and the nurse can easily use it to create influence diagrams for specific patients and decisions. He or she can do this by creating *decision-framing templates* that express common organizing principles of life-support decisions and by creating *knowledge maps* that articulate medical facts and details [16].

For example, Fig. 9 presents a decision-framing template for life-support decisions. We represent life-support therapy generically with a decision node, a rectangle labeled 'Life-Support R_x.' We indicate that life-support therapy has a therapeutic

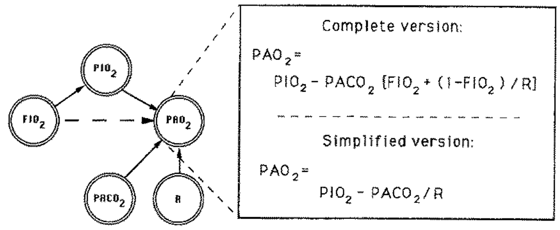

Fig. 10. A knowledge map graphically represents the alveolar gas equation. Note that the simplified version that is commonly used clinically differs from the complete version, because it removes the conditioning arc from FIO_2 to PAO_2. (Abbreviations: see text.)

effect and a side effect with chance nodes – ovals labeled 'Therapeutic Effect' and 'Side Effect' – that are linked to the decision by conditioning arcs from the decision node. The side effect may modulate the therapeutic effect. Consequently, there is an arrow from 'Side Effect' to 'Therapeutic Effect.' The therapeutic effect is directed toward satisfying a fundamental patient physiological need. We represent this with an arrow from 'Therapeutic Effect' to 'Physiological Need.' The patient outcome depends both on how well the patient's physiological need is satisfied and on how great a side effect is generated to accomplish that. This is indicated by the arrows from 'Physiological Need' and 'Side Effect' into 'Outcome.' The patient preferences over the outcomes are represented by the value node and the arrow from 'Outcome' to 'Value.'

The value of this template is that it provides a ready-made structure for building an influence diagram for any life-support therapy – whether it is for oxygen therapy, fluid, inotropes, positive end-expiratory pressure, intra-aortic counterpulsation, or transfusion. Of course, for each specific therapy, the diagram will need to be further developed with the appropriate details for therapeutic effect, side effect, physiological need, outcome and value. The critical care specialist can provide significant guidance for this development and greatly enhance its efficiency by articulating the relevant medical facts and details in the form of knowledge maps.

For example, the immediate therapeutic effect of oxygen therapy is that increasing the inspired fraction of oxygen increases the alveolar partial pressure of oxygen. This effect is described by the

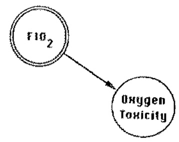

Fig. 11. A knowledge map can represent the key side effect of FIO_2.

alveolar gas equation [17]. Figure 10 is a knowledge map of the alveolar gas equation:

$$PAO_2 = PIO_2 - PACO_2\,[FIO_2 + (1 - FIO_2)/R],$$

where:

$$PIO_2 = 713 \times FIO_2.$$

PAO_2 is partial pressure of alveolar oxygen; PIO_2, the partial pressure of inspired oxygen; $PACO_2$, the partial pressure of alveolar carbon dioxide; FIO_2, the fractional inspired concentration of oxygen; and R, the respiratory quotient.

We graphically indicate that the fractional inspired oxygen concentration directly affects the partial pressure of inspired oxygen. The partial pressure of inspired oxygen, together with the partial pressure of of alveolar carbon dioxide and the respiratory quotient, determine the partial pressure of alveolar oxygen. Note the arrow from 'FIO_2' to 'PAO_2'. This direct influence of fractional inspired oxygen on the partial pressure of alveolar oxygen is deleted in the commonly used approximate form of the alveolar gas equation [3]. This gives the equation:

$$PAO_2 = PIO_2 - PACO_2/R.$$

In a similar fashion, the critical care specialist will be responsible for encoding his or her knowledge about side effects of oxygen therapy and about the underlying physiological needs that oxygen therapy must satisfy. Sample knowledge maps for these elements are shown in Fig. 11 and Fig. 12. We note that the critical care specialist creates the knowledge maps and templates without the time pressure of meeting immediate clinical needs. Such

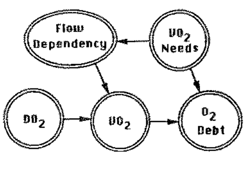

Fig. 12. A knowledge map can represent the physiological relationships related to the need for oxidative metabolism. (Abbreviations: VO_2, oxygen consumption; DO_2, oxygen delivery.)

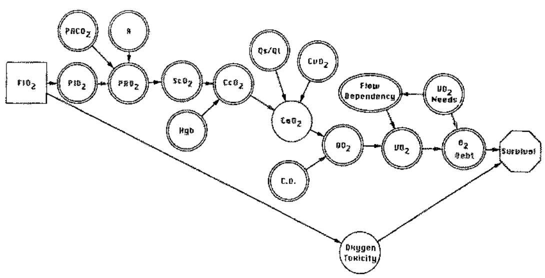

Fig. 13. The bedside physician frames the FIO_2 adjustment decision by using the knowledge maps created by the critical care specialist to expand the decision-framing template. (Abbreviations: PIO_2, partial pressure inspired oxygen; $PACO_2$, partial pressure alveolar carbon dioxide; R, respiratory quotient; PAO_2, partial pressure alveolar oxygen; ScO_2, oxygen saturation pulmonary capillary blood; Hgb, hemoglobin concentration; CcO_2, oxygen content pulmonary capillary blood; Qs/Qt, pulmonary right-to-left shunt; CvO_2, oxygen content venous blood; CaO_2, oxygen content arterial blood; C.O., cardiac output; DO_2, oxygen delivery; VO_2, oxygen consumption.)

decision engineering requires undistracted reflection. As we shall see, the work done by the critical care specialist represents an 'off-line' investment that expedites the 'on-line' influence diagram-building tasks of the bedside physician and nurse.

The bedside physician frames the FIO_2 adjustment decision by using the knowledge maps to expand the decision-framing template. Using oxygen therapy as an example, the physician might build a decision framework as illustrated in Fig. 13. This structurally complete influence diagram characterizes what is relevant to decisions about adjusting the FIO_2 on a specific patient for an interval of time. The diagram explicitly identifies survival as the value node. This may not be appropriate for all cases; here, the bedside physician is simply identifying the appropriate value function for the specific patient situation. Likewise, the remainder of the diagram is not meant to be universal – it simply represents the physician's synthesis of the various knowledge elements into a partial decision model appropriate to the particular patient situation for an interval of time.

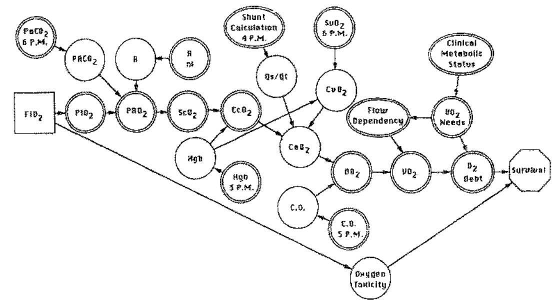

Fig. 14. The nurse details the FIO_2 adjustment at 6 p.m. by linking currently available values for patient data into the physician framework. (Abbreviations, see Figure 14, and $PaCO_2$, partial pressure arterial carbon dioxide; SvO_2, oxygen saturation mixed-venous blood; R nl, normal respiratory quotient.)

The physician provides this decision framework to the nurse as guidance. Then, based on the changing patient circumstances as represented by the continually collected data, the nurse parametrically completes the influence diagram by assessing the values for the 'border' nodes. For example, at 6 p.m., the nurse can enter into the model assessments for the measure partial pressure of arterial carbon dioxide ($PACO_2$), which can serve as an approximation for the partial pressure of alveolar carbon dioxide ($PACO_2$). Likewise, he or she can enter values for the measured hemoglobin. The result is presented in Fig. 14.

Once the nurse has entered the value that describe the patient's circumstances, a complete influence diagram is obtained. It is important to recognize that this diagram not only captures the decision structurally, but because of the mathematical relationships encoded in its constituent knowledge maps and the data provided by the nurse, it also captures the specific details of the decision at hand. Some of these assessments are shown in Fig. 15. Using appropriate algorithms, the fully assessed influence diagram can be evaluated to show how the value node depends on the different possible values of the decision variable. In technical terms.

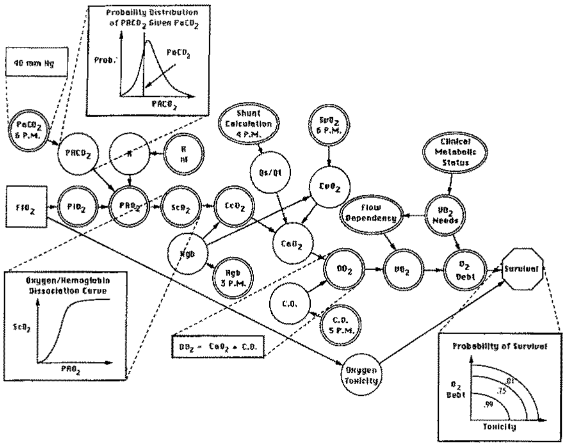

Fig. 15. The completed influence diagram is mathematically well defined, both structurally and parametrically.

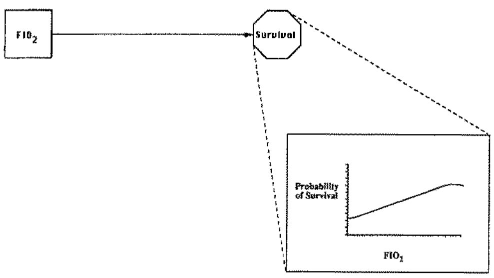

Fig. 16. Using appropriate algorithms, the comprehensive influence diagram can be reduced into a simpler diagram that captures the essence of the FIO_2 decision. The assessment of probability of survival conditioned on FIO_2 is shown in more detail in Figure 18.

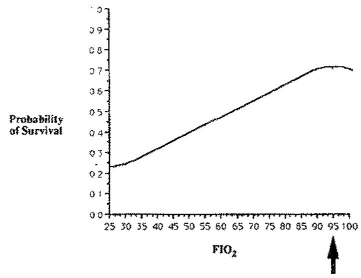

Fig. 17. The graph of FIO_2 versus probability of survival shows that the optimal FIO_2 given the date at 6 p.m. is 95 percent.

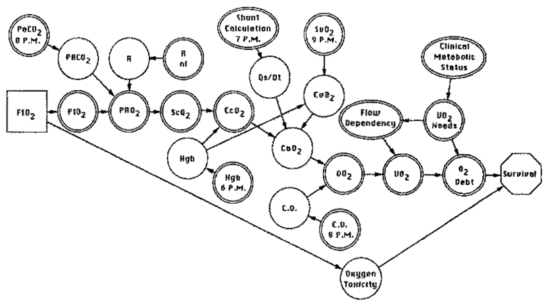

Fig. 18. The nurse updates the influence diagram at 9 p.m. with the recent values for patient data.

the influence diagram of Fig. 15 is transformed to the *minimal influence diagram* in Fig. 16. The multiple assessments underlying the full influence diagram are reduced to the single assessment shown in Fig. 16 that summarizes the essence of the decision – the relationship between the value node and the decision node. The graph in Fig. 17 shows this relationship between FIO_2 and probability of survival in detail for our hypothetical patient given the data available at 6 p.m.

The nurse viewing this graph at 6 p.m. notes that an FIO_2 of 95 percent provides the highest probability of survival – 0.75. The graph also shows that probability of survival is very sensitive to the FIO_2 – the probability of survival markedly diminishes for FIO_2 less than 90 percent.

Three hours later, however, the nurse may have new measurements for hemoglobin and partial pressure of the arterial carbon dioxide. Updating the decision framework provides a new influence diagram appropriate for the patient's circumstances at 9 p.m., as presented in Fig. 18. This influence diagram can be evaluated to generate the graph shown in Fig. 19.

This graph shows that at 9 p.m. the highest probability for survival is obtained with an FIO_2 of 60 percent. However, more important than the recommendation to decrease the FIO_2 from 95 to 60 percent are the *insights* the graph provides. Comparing this graph with the earlier graph shows that the patient has improved – his probability of survival is uniformly higher. Furthermore, survival is no longer sensitive to the value of FIO_2. In addition, other insights can be derived from the decision

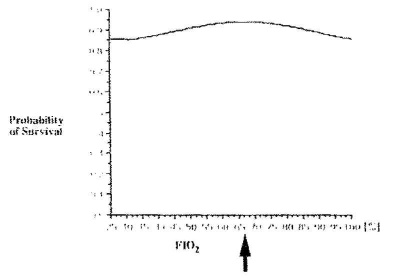

Fig. 19. The new information leads to an updated decision model at 9 p.m., which is evaluated to generate a new graph of FIO_2 versus probability of survival.

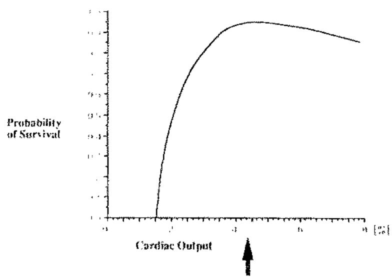

Fig. 20. A graph of the sensitivity of survival to cardiac output suggests that cardiac output is a sensitive variable.

model. For example, the nurse can use the influence diagram to explore the sensitivity of survival to other variables, such as cardiac output. This would generate the graph shown in Fig. 20.

This graph shows that probability of survival is quite sensitive to cardiac output. While the patient presently has a cardiac output (indicated by the arrow) consistent with the highest probability of survival, any change would lead to a significant decrease. Prompted by this discovery, the nurse calls the physician to point out the sensitivity of patient survival to cardiac output. The physician now reassesses the patient and expands the patient's life support to include fluid administration that will ensure that the cardiac output is maintained at the desired level. He creates a new decision framework to guide the nurse, as presented in Fig. 21.

This oxygen therapy example shows how the decision analysis approach can be applied in critical

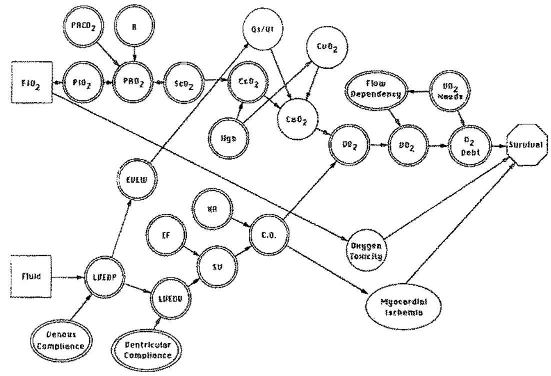

Fig. 21. To buffer the sensitive variable, cardiac output, the physician reframes the problem to include fluid therapy that can maintain cardiac output at the desired level. (Abbreviations, see Figure 14, and LVEDP, left ventricular end-diastolic pressure; LVEDV, left ventricular end-diastolic volume; SV, stroke volume; EF, ejection fraction; HR, heart rate; EVLW, pulmonary extra-vascular lung water.)

care. The critical care specialist, the bedside physician, and the bedside nurse each contribute their special expertise to build an influence diagram that captures what can be done, what is known, and what is desired. This influence diagram can then be evaluated both to generate a recommendation and to provide insight.

This approach goes considerably beyond existing critical care decision practice. For example, rather than calling the physician when the cardiac output is already low and the patient is in trouble, the nurse calls much earlier (while there is still time to act) to inform the physician that cardiac output is an important variable to control and that additional therapy should be included to keep this output in the desired range. The contrast with standing orders is also dramatic. The influence diagram provides the equivalent of on-demand standing orders that are quickly reformulated in response to changing circumstances, that are consistent with the physicians overall strategy for patient care, and that are less brittle or ephemeral. Furthermore, the decision analysis approach provides the tools to appraise the decision recommendation within its appropriate context. Clinicians not only receive recommendations, but also insight. They not only know what to do, but why.

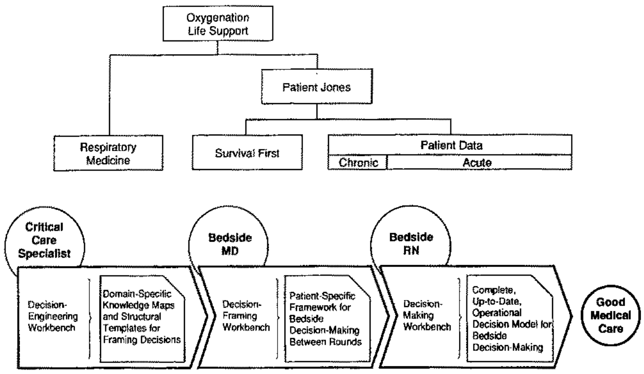

Fig. 22. Using a sequence of computer-based decision workbenches, the critical care team members collaboratively formulate the decision model.

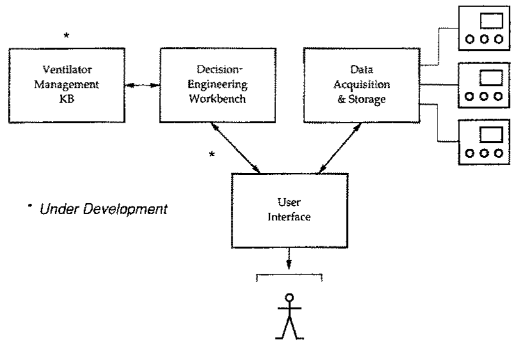

Fig. 23. A pilot system for ventilator management, *Orchestra* has four major components. The major components exist at the prototype level, and efforts are focused on knowledge engineering and integration.

Orchestra: An intelligent decision system for critical care

The critical care team members will need a significant amount of assistance to be able to encode decision templates and knowledge maps, to create decision frameworks, and to complete, evaluate, and interpret influence diagrams. Figure 22 shows a top-level architecture for a computer system that will help each member of the critical care team contribute the appropriate elements to the decision context.

The critical care specialist uses a decision-engineering workbench to encode the decision templates and knowledge maps. (We use the term 'decision engineering' to refer to the task of analyzing a class of decisions – e.g., critical care decisions – in decision analytic terms.) These templates convert the decision-engineering workbench into a decision-framing workbench, which serves as the software tool for the bedside physician. He then uses this program to create the patient-specific decision framework that serves to guide the nurse's decision-making between rounds. And this patient-specific decision framework converts the decision-framing workbench into a decision-making workbench. This is the tool the nurse uses to add to the decision framework the data elements that describe the patient's evolving circumstances. All this results in a completed influence diagram that can be evaluated on the decision-making workbench to generate a recommendation. This diagram can be 'mined' using the decision-making workbench to generate the insights that will lead to good medical care.

Using this approach, we are now building a pilot system, called *Orchestra*, for ventilator management. Its elements are illustrated in Fig. 23. The system runs on the Apple Macintosh II personal computer. The system constituents are a user interface, a data acquisition and storage system, a decision-engineering workbench, and a ventilator management knowledge base. The decision-engineering workbench, called MacAnalyst™, is now complete, and it provides tools for formulating (using a graphical interface), evaluating, and appraising influence diagrams (Fig. 24). It also provides tools for the entry of decision-framing templates and knowledge maps. The data acquisition system, also completed, is called Respirator Workstation. It provides a programming environment, called WISP, that allows for the interactive creation and execution of software drivers for medical instrumentation that supports the RS232 protocol. The Respirator Workstation thus allows flexible data acquisition from a wide variety of medical instrumentation. Drivers have now been written for the Puritan-Bennett 7200a microprocessor ventilator,

™ MacAnalyst is a trademark of IDS Partners.

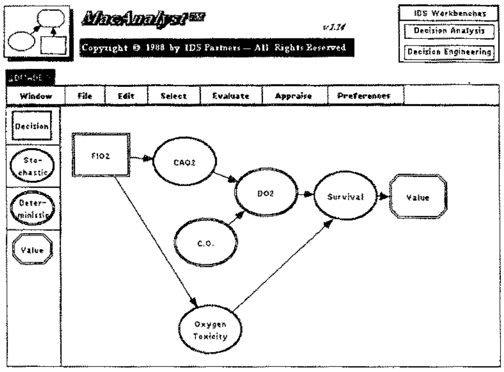

Fig. 24. The decision-engineering workbench in *Orchestra*, called MacAnalyst, provides tools for formulating, evaluating, and appraising influence diagrams.

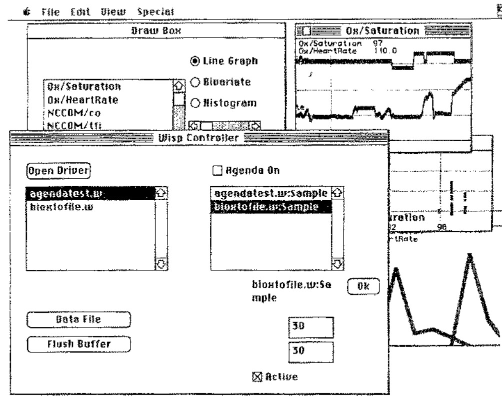

Fig. 25. *Orchestra's* user interface provides easy control of the data acquisition system and allows display of the data in multiple formats.

for the Ohmeda 3700 pulse oximeter, and for the Bard urine-output measuring device. Drivers will be written for the Siemens 1281 physiological monitor and for the Oximetrics mixed-venous oximeter when interface boards for those instruments are released by their manufacturers. The *Orchestra* user interface allows graphic display of the data, entry of noninstrument data by the clinician, and easy control of the data acquisition functions (Fig. 25). A link to the decision-engineering workbench is now under development. This will allow the nurse to easily incorporate the acquired data into the decision framework. Also under development is the ventilator knowledge base, which running on the decision-engineering workbench, forms an expert system that in response to inputs from the critical care team effects the sequential transformation of the decision-engineering workbench first to a decision-framing workbench and then to a decision-making workbench.

Summary

We feel that decision analysis can address all major ICU decision defects through a consistent, comprehensive, and efficient decision-making methodology. In addition, intelligent decision systems can make professional-level decision analysis available at the bedside by greatly reducing its cost and by increasing its speed. *Orchestra* illustrates how intelligent decision systems technology supports the application of the decision analysis approach to critical care. This new clinical decision-making approach has great potential for improved critical care decision-making.

Acknowledgements

Figure credits: Figures 2, 3, 4, and 5 – Copyright© 1988 by Strategic Decisions Group. Reprinted by permission of Strategic Decisions Group. Figure 8 – Copyright© 1985 by Samuel Holtzman. Reprinted by permission of Samuel Holtzman. Figure 25 – Copyright© 1988 by IDS Partners. Reprinted by permission of IDS Partners.

References

1. Brimm JE. Computers in critical care. Crit Care Nurs Q 53–63 March 1987.
2. Gardner RM. Computerized management of intensive care patients. MD Computing 1986, 3: 36–51.
3. Bone RC. Monitoring respiratory function in the patient with adult respiratory distress syndrome. Seminars in Respiratory Medicine 1981; 3: 45–55.
4. Shoemaker WC. Relation of oxygen transport patterns to the pathophysiology and therapy of shock states. Intensive Care Med 1987; 13: 230–43.
5. Knaus WA, Draper EA, Wagner DP. An evaluation of outcome from intensive care in major medical centers. Annals of Internal Med 1986; 104: 410–18.

6. Civetta JM, Hudson-Civetta J. Cost-effective use of the intensive care unit. In Eiseman B, Stahlgren L, eds. Cost-effective surgical, 1987.
7. Howard RA. Decision analysis: applied decision theory. In: Hertz DB, Melese J, eds. Proceeding of the Fourth International Conference on Operations Research. New York: Wiley-Interscience, 1986: 55–71.
8. Laplace PS. A Philosophical essay on probabilities. 6th edition. New York: Dover, 1951.
9. Bernoulli D. Specimen theoriae novea de mensora sortis. Papers of the Imperial Academy of Sciences in Petersburg 1738; 5: 175–92.
10. Pauker SG, Kassirer JP. Decision analysis. New England Journal of Medicine 1987; 316: 250–8.
12. Howard RA, Matheson JE. Influence diagrams. In: Howard RA, Matheson JE eds. Readings on the principles and applications of decision analysis, Vol. II. Menlo Park: Strategic Decisions Group, 1984: 719–62.
13. Olmsted S. On representing and solving decision problems. PH.D. dissertation. Stanford, Calif.: Engineering-Economic Systems Department, Stanford University, 1983.
14. Shachter RD. Evaluating influence diagrams. Operations Research 1986; 34: 871–82.
15. Holtzman S. Intelligent decision systems. Reading, Mass.: Addison-Wesley Publishing Company, Inc. 1989: 90–107.
16. Howard RA. Knowledge maps. To be published in Management Science, 1988.
17. Comroe JH, Forster RE, DuBois AB, Briscoe WA, Carlsen E. The lung. 2nd edition. Chicago: Year Bood Medical Publishers. Inc., 1962: 339–41.