International Journal of Clinical Monitoring and Computing **11**: 241–253, 1994.

Intelligent systems in patient monitoring and therapy management

A survey of research projects

Serdar Uckun
Knowledge Systems Laboratory, Stanford University, 701 Welch Road Bldg. C, Palo Alto, CA 94304, USA

Accepted 28 October 1994

Key words: artificial intelligence, intelligent monitoring and control, critical care

Abstract

Although today's advanced biomedical technology provides unsurpassed power in diagnosis, monitoring, and treatment, interpretation of vast streams of information generated by this technology often poses excessive demands on the cognitive skills of health-care personnel. In addition, storage, reduction, retrieval, processing, and presentation of information are significant challenges. These problems are most severe in critical care environments such as intensive care units (ICUs) and operating room (ORs) where many events are life-threatening and thus require immediate attention and the execution of definitive corrective actions. This article focuses on intelligent monitoring and control (IMC), or the use of artificial intelligence (AI) techniques to alleviate some of the common information management problems encountered in health-care environments. This article presents the findings of a survey of over 30 IMC projects. A major finding of the survey is that although significant advances have been made in introducing AI technology in critical care, successful examples of fielded systems are still few and far between. Widespread acceptance of these systems in critical care environments depends on a number of factors, including fruitful collaborations between clinicians and computer scientists, emphasis on evaluation studies, and easy access to clinical information.

Abbreviations: AI – artificial intelligence, ICU – intensive care unit, IMC – intelligent monitoring and control, OR – operating room

1. Introduction

The increasing sophistication of medical devices is often a mixed blessing in health-care environments. Although today's advanced biomedical technology provides unsurpassed power in diagnosis, monitoring, and treatment, interpretation of vast streams of information generated by this new technology often poses excessive demands on the cognitive skills of health-care personnel. Storage, reduction, retrieval, processing, and presentation of information also pose significant challenges. These problems are most severe in critical care environments such as intensive care units (ICUs) and operating room (ORs) where many events are life-threatening and thus require immediate attention and the execution of corrective actions [1, 2].

Several engineering methods and techniques are proposed and put to use in order to alleviate the problems caused by the massive flux of information into critical care environments. These techniques include integrated monitoring devices, central monitoring consoles, more reliable sensors, advanced signal processing techniques, and methods to standardize data exchange between biomedical devices [3]. Despite all these advances, the modern critical care environment is still not immune to problems such as spurious alarms and misplaced or misinterpreted information. Tasks such as information storage, retrieval, and interpretation are entirely handled by overworked medical professionals, thus increasing the possibility and potential severity of information management problems.

During the past fifteen years, the field of artificial intelligence (AI) started tackling problems in real-time information management. To that end, several projects were initiated to help clinicians with their information processing needs in critical care environments. Among these early projects were VM [4] and Compas [5, 6]. VM was developed in the late 1970s as an experiment in extending the MYCIN formalism in order to manage time-varying information. The application area for VM was the management of patients undergoing mechanical ventilation. The Compas system was another advice-giving system which was designed to assist the respiratory therapy of patients with adult respiratory distress syndrome (ARDS). Compas was subsequently modified to serve as a ventilator management advisor for a number of assisted ventilation modes, and its current incarnation (CORE) is in clinical use at the LDS Hospital in Utah [7, 8].

The focus of this article is on intelligent monitoring and control (IMC), or the use of AI techniques to alleviate some of the common information management problems encountered in health-care environments. Over the past fifteen years, we have come to a better understanding of the task domain of critical care monitoring, and of the capabilities of the tools and techniques we have devised. In the next section, I describe the task domain of IMC. In section 3, I present the findings of a survey of over 30 projects which I conducted in 1992. The purpose of the survey was to identify the state-of-the-art in this field, especially with respect to the medical foci of the projects, the reasoning tasks, and the extent of fielding and evaluation. The survey provides a representative sample of current research projects rather than covering the field exhaustively (the responses to the survey are collated in a technical report and are available upon request [9]). Where appropriate, I mention additional influential projects which are not represented in the survey. In the final section, I discuss the current state of the field and hint at future directions.

2. Intelligent monitoring and control: tasks and concepts

Monitoring is the process of observing a physical system, and control is concerned with guiding the behavior of the observed system toward some predetermined objective. Since monitoring is inherently related to control, it has to be performed in real time in order to constrain the utility of particular actions. IMC can thus be defined as the use of AI methodologies in order to perform knowledge-intensive monitoring and control tasks in real time. As this definition implies, the domain of IMC applications is not limited to medicine. In fact, as engineered systems have become more and more complex, sophisticated techniques for keeping these systems operational have become an integral aspect of engineering practice. In engineering domains, IMC focuses on assisting human operators of complex engineered systems by interpreting and explaining system behavior and by providing expert advice on possible actions and their consequences [10]. In this article, I will limit the scope to critical care applications and use the term 'intelligent monitoring and control' throughout.

The related notions of monitoring and control often bring into mind an image of a rapidly-changing system under observation, manipulated by time-critical decisions and actions. However, the real-time constraints in monitoring and control do not necessarily result in significant time pressures on the reasoning process. Many monitored systems do not evolve rapidly in time (e.g., glucose metabolism). The real-time requirement is met in such systems as long as actions can be taken while they still have non-negligible utility, that is, before it is too late. The major challenge in monitoring and control of rapidly-changing systems is timely response whereas the crux of the problem for systems with slow dynamics is filtering, reduction, and interpretation of large amounts of time-varying information.

2.1. Tasks

The tasks that constitute the scope of IMC include interpretation of observed behavior, explanation of causal mechanisms, reasoned response to observed events, and planning courses of action [10]. These tasks can be broadly classified in three categories: diagnosis, prediction, and control.

2.1.1. Diagnosis

A variety of inference methods may be classified under the category of diagnostic inferences. All of these methods are used in existing IMC systems, and a given system often uses more than one of these methods in its reasoning processes. I will discuss these methods along two orthogonal dimensions: level of interpretation (data dimension) and temporal abstraction (time dimension). The level of interpretation dimension relates to the data types that inference methods are concerned with, whereas the temporal abstraction

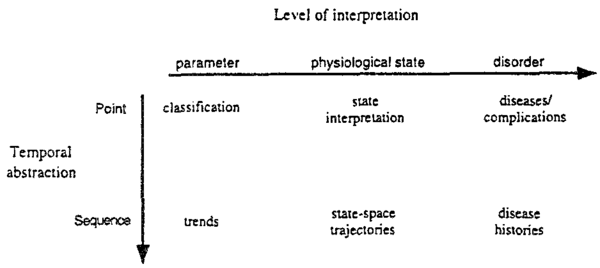

Fig. 1. Diagnostic inferences in IMC.

dimension indicates whether a given inference method interprets single data points or data sequences ordered in time. Figure 1 illustrates these dimensions and the corresponding inference methods. Another orthogonal dimension of analysis (which will not be discussed here) applies to types of abstractions (interpretations): definitional abstractions, qualitative abstractions, or generalizations [11].

At one end of the interpretation spectrum, inference methods interpret values of simple parameters. If the system only considers single data points, the inference operation is called *classification*, typically a number-to-symbol transformation. If the system interprets sequences of data, the inference method may be referred to as *trend detection*. In recent years, new inference methods were introduced to deal with sequences of data which traditionally cannot be dealt with efficiently because of problems such as missing, incomplete, and uncertain information. Such techniques [12, 13] complement established trend detection techniques such as Kalman filtering [14]. A discussion whether such inference methods should be implemented using prescriptive or normative techniques is inconsequential; in this classification we are only concerned with the data requirements, i.e., the input/output behavior.

The next level of abstraction in the interpretation spectrum is state-based abstraction. Inference methods operating at this level synthesize interpretations of parameter values into interpretations of physiological and pathophysiological states. If the interpretation is extended into sequences of states in time, a state-space trajectory is obtained. The identification of state-space trajectories is a useful concept in diagnosis because some disorders can only be distinguished from each other by the sequences of pathophysiological states that they follow in time [15, 16]. In addition to such categorical methods, a recent project for state-space monitoring uses the possibilistic concept of fuzzy automata to capture the uncertainty inherent in state interpretation tasks [17].

The final level of abstraction in the interpretation dimension is that of disorders where a disorder may be defined as a collection of pathophysiological states, either at a certain point in time or as a trajectory of this state vector over time. This is the level at which traditional diagnostic consultation systems operate [18]. A monitoring system does not have to reason at the levels of parameters and states in order to achieve an interpretation at the level of disorders. However, the fragility and limitations of traditional diagnostic expert systems (which typically deal with pattern recognition at the disorder level) suggest that a system which reasons at all three levels will be more robust since it will be able to synthesize its interpretations of parameter values and physiological states into subsequent interpretations of diseases and complications.

Another inference task which may be classified under the diagnosis category is *explanation*. Explanatory inferences typically describe plausible paths from causal agents to diseases to signs and symptoms in a parsimonious and relevant fashion. As such, explanation methods typically (but not necessarily) use the data types and relations which are required by other diagnostic inferences.

2.1.2. Prediction

True predictive inferences require models of the problem domain which is under observation.[1] Such models typically belong to one or more of three categories: structural, behavioral, and functional. Examples of modeling languages used for prediction purposes in medicine are qualitative/quantitative process models [19], qualitative mathematics [20], and numerical models [21]. A comprehensive review of model-based reasoning techniques in medicine may be found in [22].

A major use for predictive inferences in IMC is in treatment planning and management. Possible uses of prediction in the treatment context are:

- evaluation of treatment plans [23]
- optimization of treatment plans [24]

[1] It is possible to achieve some success in prediction by caching certain expectations and outcomes in a database of actions and plans. However such methods are not robust since, for example, competing expectations cannot be resolved unless all possible combinations are enumerated a priori.

- predictive alarming (i.e., anticipation of nondesirable parameter values and pathophysiological states before they occur) [25]
- management of computational resources based on demand estimates [26, 27].

2.1.3. Control

In criticalcare, control inferences are mainly used for therapy planning and management. Depending on the sophistication of domain models the system is equipped with and the competence of its diagnostic skills, the control capability of a system ranges from simple reflex actions (i.e., reactivity) to long-term planning for most possible scenarios including contingencies.

Most of the IMC systems in development today are open-loop systems with respect to planning and control. These systems do not take direct therapeutic actions; rather, they convey their recommendations to clinicians who are ultimately responsible for the well-being of patients. It is up to the physicians, nurses, or technicians to execute or disregard these recommendations. A different twist on the same scenario is that of *critiquing*. In this scheme, the computer system does not offer treatment advice; instead, it criticizes treatment decisions made by clinicians based on its knowledge of the problem domain [28].

There are two fundamental roles for IMC systems in the context of therapy management. One of these is *execution monitoring* of treatment plans devised by either a clinician or by the device itself. In this case, the system observes whether the patient state is evolving according to predictions based on a particular treatment plan. The other is *closed-loop control*: if a system is able to execute treatment plans directly (by manipulating ventilator settings or modifying drip infusion rates, etc.), it is considered a closed-loop system. Not surprisingly, there are very few examples of closed-loop IMC systems in clinical use or trial due to issues such as legal liability and social and cultural factors. Some existing examples are discussed in the survey section of this article.

2.2. Concepts

A number of important concepts differentiate IMC systems from conventional approaches to patient monitoring. Not all of these concepts are exploited in every project. However, many researchers emphasize the importance of these issues in developing the next generation of monitoring and control devices. These concepts are discussed in the following subsections. Other issues such as user interfaces and details of reasoning methods are beyond the scope of this article.

2.2.1. Real-time performance and resource management

Earlier, I stated that real-time performance is crucial for IMC systems. However, the real-time constraint is a loosely-defined constraint. A critical care intelligent agent may face situations where the demands on its cognitive and reasoning skills may exceed its computational capabilities. In order to handle such situations gracefully, an IMC system should be able to manage the use of its internal resources based on interpretations of the external environment [26]. One aspect of resource management is *selective perception*, a method by which more attention is paid to critical data channels in order to distribute resources in favor of situations which may require immediate attention [29].

2.2.2. Handling noisy data

Information available in the real world is often noisy, incomplete, and erroneous. Many ICU monitors today have methods to deal with known failure modes such as disconnected electrodes and with known sources of noise such as AC interference. Fielded IMC systems should be able to provide higher levels of resilience against noise and measurement errors compared to the best of today's commercially-available monitors. Those systems could use techniques such as cross-channel correlation, trend- and model-based expectations, and model-based diagnosis, in order to verify or refute questionable information [27]. Unfortunately, many of the IMC systems in development today are proof-of-principle models and thus do not deal adequately with the problem of noisy data.

2.2.3. Context sensitivity

Interpretation of information in proper context is an important issue which cannot be addressed by monitors devoid of domain models. In contrast, an intelligent monitor may base its interpretations of the patient on a sophisticated knowledge base and all facts relevant to the current situation. Thus, certain features and patterns may be interpreted differently on different patients based on demographics and pathophysiological disorders already known to exist. A trivial example of context sensitivity is the capability to interpret and classify arterial blood pressure measurements based on patient age and existing cardiovascular disease. A more advanced example would be the capability to interpret signs of myocardial ischemia in varying situations, such as during or after cardiac surgery, or on sedated or fully awake patients.

2.2.4. Intelligent alarms

The intelligent alarming idea is somewhat related to the issue of context sensitivity. The purpose of intelligent alarms is to increase the specificity of alarms (to reduce false-positive rates) while maintaining or increasing their sensitivity (reducing false-negative rates). An illustrative albeit simplistic example is an intelligent alarm which can disregard a momentary increase in the heart rate of an ICU patient if it can relate the abnormality to an extraneous cause (patient turning in bed, endotracheal suctioning, etc.) sensed by other means.

3. Survey results

In 1992, I conducted a survey of IMC projects in medicine via the AI in medicine mailing list on internet,[2] via personal contacts, and via questionnaires distributed at conferences. I received information about 32 projects, which is not exhaustive but nevertheless a representative sample of projects from the United States, South America, Europe, and Australia. Appendix 1 lists and briefly describes the projects covered in the survey, including monitoring and control tools which are reported as parts of larger systems. Sections 3.1, 3.2, and 3.3 respectively analyze the application domains, operating modes, and reasoning tasks addressed by these projects. Section 3.4 discusses the issues of evaluation and fielding.

3.1. Application domains

Critical care environments face the greatest challenges in the timely management of information and thus stand to gain the most from IMC applications. Not surprisingly, the majority of projects reported in this survey are applied to critical care problems. Close to two-thirds of reported projects are ICU applications, either in cardiovascular monitoring or in ventilator management. The range of parameters monitored by these systems varies extensively. However, the expertise in most systems is limited to a narrow range of problems typically encountered in the ICU (the Guardian knowledge base is limited to post-operative care of cardiac surgery patients [26]; the SIMON system is developed for a proof-of-concept in respiratory monitoring of infants with respiratory distress syndrome [30], etc.).

Only two of the reported projects focus on problems encountered in the OR [31, 32]. Both of these projects deal with the diagnosis of faults in anesthesia circuits, and none of the surveyed projects is concerned with the important problems of monitoring anesthesized patients and anesthesia alarms. A few earlier projects have addressed anesthesia monitoring [33], and there is a demonstrated need in this field for assistance with information management tasks [32].

With the advent of integrated hospital information management systems, monitoring of clinical events is rapidly gaining importance. Five of the reported projects focus on clinical event monitoring, in other words, monitoring and interpretation of slowly-changing parameters, states, and disorders. These efforts include trend templates [12], temporal abstractions [13], specialized event monitors [34], the Protege-II architecture for therapy management [35], and the GRECC projects from Utah which aim to improve health care delivery for the elderly (see Appendix 1 for details).

Table 1 presents the breakdown of surveyed projects by application area.

Table 1. Application areas of surveyed projects.

Application area	Number of projects
ICU/Cardiovascular monitoring	11
Ventilator management	9
Clinical event monitoring	5
Anesthesia alarms	2
Antenatal monitoring	2
Other biosignal monitoring (ECG, EEG, etc)	2
Insulin management	1

3.2. Operating modes

The prototypical mode of operation for an IMC application is one of a passive observer and controller: the system acquires data from the environment, processes information, and expresses its findings and inference results via a user interface or by means of closed-loop control. As long as data are available on-line, the pas-

[2] To subscribe, send e-mail to Wanda Pratt at <ai-medicine-REQUEST@MED.Stanford.EDU>

sive observer mode is the ideal operating mode for critical care IMC systems since it does not further tax the scarce human resources. Unfortunately, some of the systems mentioned in the survey have to operate with manually-entered data because of technical limitations. Although a serious impediment to fielding, the manual data entry model is reasonable for proof-of-principle studies.

As a positive note, none of the covered systems adhere to the 'consultation system' (or, 'the Greek Oracle') paradigm which has not proven popular or successful among intended users despite extensive efforts in the past couple of decades [36]. The consultation paradigm might have limited merit in clinical event monitoring, but even in that case the trend is toward systems that operate in the background with minimal intervention. IMC systems will have greater appeal as long as they successfully perform their assigned tasks without interfering with or impeding clinicians' busy schedules.

3.3. Reasoning tasks

The majority of surveyed projects perform some form (or several forms) of diagnostic inference. The inference methods utilized by these programs span the whole spectrum of diagnostic inferences as discussed in Section 2.1.1. A smaller number of systems attempt prediction of parameter values, physiological states, or disorders. Over one-third of surveyed systems execute control algorithms. However, most of the control decisions are conveyed to the external environment in the form of open-loop control, that is, advice giving or planning. Only three projects are concerned with closed-loop control. Guardian [26] is a prototype intelligent agent which is capable of closed-loop control in simulation. However, the system has not been tested in closed-loop mode in human or animal experiments. Bedside Pancreas is a closed-loop insulin monitoring and pump control system currently under development (see Appendix 1 for details). Only two closed-loop control systems are actually tested on ventilator management of human subjects [37, 38]. One of these systems, Ganésh, is currently in limited clinical use in Henri Mondor Hospital in Creteil, France, and a successor (NeoGanésh) is under development. Both of these systems focus on a relatively narrow ventilator management problem (weaning) and control two ventilator settings based on the evaluation of a very small number of monitored parameters. Nevertheless, they deserve recognition since they mark the early successful demonstrations of closed-loop ventilator management on humans. Another group which is currently looking into closed-loop ventilator management is the LDS group from Utah, who are using one input parameter (arterial oxygen partial pressure), a PID controller, and the CORE ventilator management protocol in order to control two ventilator settings [39]. To date, this method has not been tested on humans.

Table 2. Categories of inference tasks handled by surveyed projects (the total is greater than 32 since several projects address multiple categories of inference tasks).

Inference category	Number of projects
Diagnosis (incl. explanation)	28
Prediction	8
Closed-loop control	3
Open-loop control (incl. planning)	10

Table 2 outlines the inference categories handled by surveyed projects.

3.4. Evaluation and fielding

There is little consensus or collective wisdom on how to evaluate the performance of IMC systems. In contrast with diagnostic consultation systems the performance of which could be assessed in terms of the accuracy of their diagnoses, real-time monitoring and control systems prove exceedingly difficult to evaluate. A typical critical care monitoring and control system operates continuously in an observe/interpret/(act) cycle. Therefore, the most natural method for evaluating a critical care IMC system is continuous evaluation. The only example of such an evaluation in the survey was performed by Dojat on the Ganésh closed-loop ventilator control system. In this evaluation, a 'comfort range' was described as a multi-dimensional quality surface representing comfort ranges for several respiratory parameters. The percentage of time the patient had spent within the comfort range was then used as a measure of success for the ventilator control system [37].

The more common method for evaluating the performance of a continuous monitoring and control system involves a series of discrete evaluations. A natural interpretation point is the interpretation step of an observe/interpret/(act) cycle. However, depending on how the cycle is triggered, some of the runs through the cycle may be totally redundant (i.e., no 'new' information). So, it may be misleading to evaluate the system at each cycle. The other alternative is to evaluate whenever something 'qualitatively significant' happens, based on some arbitrary or principled criterion. One example is the VentPlan evaluation where the system's therapy suggestions regarding FiO2 (volume fraction of inspired O2) changes were compared to actual decisions made by physicians on recorded patient cases [40]. Another example was the evaluation of the diagnostic performance of YAQ, a quantitative/quantitative reasoner for model-based prediction and diagnosis. In this evaluation, the diagnostic hypotheses proposed by the system were reviewed and criticized by a domain expert on a case-by-case basis [41]. Finally, the clinical trials of the CORE ventilator management protocols at Utah focused on clinicians' compliance with the system's suggestions [7]. In these studies, physician compliance reached as high as 92% of all CORE recommendations.

One method of evaluation noticeably lacking in the literature of the field is the randomized controlled clinical trial, which, incidentally, is the preferred method of evaluation in clinical medicine. Although the researchers from the LDS group claim a 40% survival rate in the most severe category of ARDS patients using the CORE protocol (based on their experience on over 200 patients) as opposed to around 10% without computer assistance, the latter figure reflects the aggregate experience of several centers with widely ranging methods of patient care [1]. The use, useability, and usefulness of IMC systems in the clinical environment is yet to be shown using evaluation methods that are widely accepted in clinical practice.

The lack of fielded IMC systems in routine clinical use is another indicator that the field has yet to reach maturity. Apart from the CORE protocol which is in routine use at LDS and currently installed at a second site in Los Angeles [7] and a few systems tested in clinical environments for limited periods [32, 37, 42], such systems are still largely confined to computer science laboratories. This is largely due to our failure to demonstrate the potential benefits of such systems in clinically significant studies. Arguably, the best measure of success for these systems is the optimization of various quality-of-care criteria in large randomized controlled clinical trials (e.g., length of ICU stay, cost of care, long-term morbidity and mortality, certain measures of patient comfort). The design and execution of such trials should be a critical item on our collective research agenda if we expect IMC systems to gain widespread clinical acceptance.

4. Conclusions

Patient monitoring and therapy management is a relatively new and promising focus of AI in medicine. As evidenced by this survey, a wealth of current projects addresses important issues in critical care and clinical information management. Most of the projects discussed here are either currently under development or are developed as proof-of-principle systems with no intentions of fielding. Thus, almost none of the systems presented in the survey are in actual clinical use. Furthermore, evaluation of these systems, if any, is based on empirical methods and protocols. None of the systems reported here were subject to stringent evaluation protocols such as randomized clinical trials, and no studies were conducted in order to assess the clinical impact of such systems. As a result, identification of methods and protocols for the evaluation of such systems remains a critical issue.

IMC is an interdisciplinary field which lies at the intersection of clinical medicine, computer science, AI, and biomedical engineering. Therefore, comprehensive projects require diverse teams composed of physicians, nurses, medical technicians, computer scientists, and engineers. As many AI in medicine scientists have witnessed in the past, managing such diverse teams is a difficult task because of the vast cultural differences between involved disciplines. The ideal IMC research and development team should have a balanced distribution of clinicians and computer scientists, should have clearly defined clinical goals and evaluation plans, and the clinical goals should never be sacrificed in order to give precedence to exploration of technology.

Finally, IMC systems require easy access to information in order to function adequately. In critical care, this requirement implies robust interfaces to monitoring, treatment, and laboratory instruments, and to heterogeneous clinical databases. These connectivity requirements indicate that IMC research is tightly related to research in medical informatics and computer science, especially in the areas of medical device connectivity standards (MIB, HL7), databases, information retrieval, networking, and shared vocabularies. The success of the LDS group in establishing the CORE protocols as a local standard in ventilator management largely stems from their long-term experience with a sophisticated hospital information management system (HELP). The current trend toward the development of such integrated systems in other centers will pave the way for cost-efficient IMC applications to challenging clinical problems.

Earlier, I stated that IMC systems are developed to alleviate information management problems in critical care. Unfortunately, the field is yet to reach maturity in terms of fruitful applications. The pioneering applications mentioned in this article (and several others which are not represented in the survey) serve to demonstrate the potential applicability of AI technology to critical care information management. With careful attention to evaluating the impact of such systems on health care delivery and cost containment, we should expect to see IMC discussions migrate from esoteric scientific meetings to mainstream medical literature.

Acknowledgements

This work is supported, in parts, by NIH grant HL-38354 and by NASA grant NAG2-581 under ARPA order 8607. I thank Drs. Barbara Hayes-Roth and Edward Feigenbaum for supporting the research at the Knowledge Systems Laboratory. I also thank Benoit Dawant, Alejandro Pazos, Dean Sittig, Mark Leaning, Jerome Soller, Enrico Coiera, and an anonymous reviewer who have provided valuable comments on earlier drafts of this article. Finally, I appreciate the assistance of all who have contributed to the intelligent monitoring survey.

References

1. Morris AH. Paradigms in management. In: Pinsky MR, Dhainaut J-FA, editors. Pathophysiologic Foundations of Critical Care. Baltimore, MD, 1993 193–206.
2. Wright D, Mackenzie SJ, Buchan I, Cairns CS, et al. Critical incidents in the intensive therapy unit. Lancet, September 1991; 338: 676–8.
3. Gardner RM, Hawley WL, East TD, Oniki TA, et al. Real time data acquisition: recommendations for the Medical Information Bus (MIB). International Journal of Clinical Monitoring and Computing 1992; 8: 251–8.
4. Fagan LM, Kunz JC, Feigenbaum EA, Osborn JJ. Extensions to the rule-based formalism for a monitoring task. In: Buchanan B, Shortliffe EH, editors. Rule-Based Expert Systems. Addison-Wesley Reading, MA, 1984: 397–423.
5. Sittig DF, Gardner RM, Morris AH, Wallace CJ. Clinical evaluation of computer-based respiratory care algorithms. International Journal of Clinical Monitoring and Computing 1990; 7 (3): 177–85.
6. Sittig DF, Pace NL, Gardner RM, Beck E, et al. Implementation of a computerized patient advice system using the HELP Clinical Information System. Computers and Biomedical Research 1989; 22: 474–87.
7. Henderson S, Crapo RO, Wallace J, East TD, et al. Performance of computerized protocols for the management of arterial oxygenation in an intensive care unit. International Journal of Clinical Monitoring and Computing 1992; 8: 271–80.
8. Henderson S, East TD, Morris AH, Gardner RM. Performance evaluation of computerized clinical protocols for management of arterial hypoxemia in ARDS patients. In: Proc. 13th Annual Symp. Comput. Appl. Med. Care. Washington, DC, 1989: 588–92.
9. Uckun S. Intelligent systems in patient monitoring and therapy management. Technical Report KSL-TR-93-32. Stanford University, 1993.
10. Hayes-Roth B, Washington R, Hewett R, Hewett M, et al. Intelligent monitoring and control. In: Proc. 11th International Joint Conf. on Artificial Intelligence. Detroit, MI, 1989: 243–9.
11. Clancey WJ. Heuristic classification. Artificial Intelligence 1985; 27: 289–350.
12. Haimowitz IJ, Kohane IS. Automated trend detection with multiple temporal hypotheses. In: Proc. 13th International Joint Conf. on Artificial Intelligence. Chambery, France, 1993.
13. Shahar Y, Musen MA. A temporal-abstraction mechanism for patient monitoring. In: Proc. 16th Annual Symp. Comput. Appl. Med. Care. Baltimore, MD, 1992.
14. Sittig DF, Factor M. Physiologic trend detection and artifact rejection: a parallel implementation of a multi-state Kalman filtering algorithm. Computer Methods and Programs in Biomedicine 1990; 31: 1–10.
15. Uckun S. An ontology for model-based reasoning in physiological domains. Ph.D. Dissertation, Vanderbilt University, Nashville, TN, 1992.
16. Cohn A, Rosenbaum S, Factor M, Miller PL. DYNASCENE: an approach to computer-based intelligent cardiovascular monitoring using sequential clinial scenes. Methods of Information in Medicine 1990; 29: 122–31.
17. Steimann F, Adlassnig K-P. Clinical monitoring with fuzzy automata. Technical Report MES-2/1993, University of Vienna, 1993.
18. Rule-Based Expert Systems. In: Buchanan B, Shortliffe EH, editors. Addison-Wesley Reading, MA, 1984.
19. Uckun S, Dawant BM. Qualitative modeling as a paradigm for diagnosis and prediction in critical care environments. Artificial Intelligence in Medicine 1992; 4 (2): 127–44.
20. Coiera EW. Monitoring diseases with empirical and model-generated histories. Artificial Intelligence in Medicine 1990; 2: 135–47.
21. Rutledge G. A method for the dynamic selection of models under time constraints. In: Proceedings of the AI and Statistics Conference. Ft. Lauderdale, FL, 1993.
22. Uckun S. Model-based reasoning in biomedicine. CRC Critical Reviews in Biomedical Engineering 1992; 19 (4): 261–92.
23. Farr BR, Shachter RD. Representation of preferences in decision support systems. In: Proc. 15th Annual Symp. Comput. Appl. Med. Care. Washington, DC, 1991: 1018–24.
24. Langlotz CP, Fagan LM, Tu SW, Sikic BI, et al. A therapy planning architecture that combines decision theory and artificial intelligence techniques. Computers and Biomedical Research 1987; 20: 279–303.
25. Leaning MS, Gallivan S, Newlands ES, Dent J, et al. Computer system for assisting with clinical interpretation of tumour marker data. British Medical Journal 1992; 305: 804–7.
26. Hayes-Roth B, Washington R, Ash D, Hewett R, et al. Guardian: a prototype intelligent agent for intensive-care monitoring. Artificial Intelligence in Medicine 1992; 4 (2): 165–85.

27. Dawant BM, Uckun S, Manders EJ, Lindstrom DP. Model-based signal acquisition, analysis, and interpretation for intelligent patient monitoring. IEEE EMBS Magazine, 1993.
28. Miller PL. Expert Critiquing Systems: Practice-Based Medical Consultation by Computer. Springer-Verlag, New York, NY, 1986.
29. Doyle RJ, Fayyad UM. Sensor selection techniques in device monitoring. In: Proceedings of the 2nd Annual Conference on AI, Simulation, and Planning in High Autonomy Systems. Los Alamitos, CA, 1991: 154–63.
30. Dawant BM, Uckun S, Manders EJ, Lindstrom DP. SIMON: A distributed computer architecture for intelligent patient monitoring. Expert Systems with Applications, 1993.
31. van der Aa JJ. Intelligent alarms in anesthesia: a real time expert system application. Ph.D. Dissertation, Technische Universiteit Eindhoven, The Netherlands, 1990.
32. Westenskow DR, Orr JA, Simon FH, Bender H-J, et al. Intelligent alarms reduce anesthesiologist's response time to critical faults. Anesthesiology 1992; 77: 1074–9.
33. Garfinkel D, Matsiras PV, Lecky JH, Aukburg SJ, et al. PONI: an intelligent alarm system for respiratory and circulation management in the operating rooms. In: Proc. 12th Annual Symp. Comput, Appl. Med. Care. Washington, DC, 1988: 13–17.
34. Kohane IS. Maintaining alternate interpretations of data from multiple sources in a clinical event monitoring system. In: Proceedings of the MEDINFO-92. Geneva, Switzerland, 1992: 483–9.
35. Musen MA. Dimensions of knowledge sharing and reuse. Computers and Biomedical Research 1992; 25: 435–67.
36. Shortliffe EH. The adolescence of AI in medicine: will the field come of age in the '90s? Artificial Intelligence in Medicine 1993; 5 (2): 93–106.
37. Dojat M, Brochard L, Lemaire F, Harf A. A knowledge-based system for assisted ventilation of patients in intensive-care units. International Journal of Clinical Monitoring and Computing 1992; 9: 239–50.
38. Strickland Jr JH, Hasson JH. A computer-controlled ventilator weaning system. Chest 1991; 100: 1096–9.
39. East T, Tolle CR, McJames S, Farrell RM, et al. A non-linear closed-loop controller for oxygenation based on a clinically proven fifth dimensional quality surface. Anesthesiology 1991; 75 (3A): A468.
40. Rutledge G, Thomson G, Farr B, Tovar M, et al. The design and implementation of a ventilator-management advisor. Artificial Intelligence in Medicine 1993; 5 (1): 67–82.
41. Uckun S, Dawant BM, Lindstrom DP. Model-based reasoning in intensive care monitoring: the YAQ approach. Artificial Intelligence in Medicine 1993; 5 (1): 31–48.
42. Tong DA. Weaning patients from mechanical ventilation: a knowledge-based system approach. In: Proc 14th Annual Symp. Comput. Appl. Med. Care. Washington, DC, 1990: 79–85.
43. Xu J, Hyman S, King P. Knowledge-based flash evoked potential recognition system. Artificial Intelligence in Medicine 1992; 4: 93–109.
44. Jiang A, King P, Smith B. Information interpretation in a real-time knowledge-based respiratory monitoring system. In: Ikeda et al., editor, Computing and Monitoring in Anesthesia and Intensive Care. Springer-Verlag Tokyo, 1992: 47–9.
45. Haimowitz IJ, Kohane IS. An epistemology for clinically significant trends. In: Proc. 11th National Conf. on Artificial Intelligence. Washington, DC, 1993.
46. Shahsavar N, Gill H, Wigertz O, Frostell C, et al. KAVE: a tool for knowledge acquisition to support artificial ventilation. Computer Methods and Programs in Biomedicine 1991; 34: 115–23.
47. Arkad K, Gill H, Ludwigs U, Shahsavar N, et al. Medical Logic Module (MLM) representation of knowledge in a ventilator treatment advisory system. International Journal of Clinical Monitoring and Computing 1991; 8: 43–8.
48. Beech M, Todd S, Tombs V. Knowledge-based techniques for alarm rationalisation in patient monitoring. In: Proceedings of the IEEE Colloquium on AI in Medical Decision Making, 1990.
49. Coiera EW. Qualitative superposition. Artificial Intelligence, 1992.
50. Miksch S, Horn W, Popow C, Paky F. VIE-VENT: knowledge-based monitoring and therapy planning of the artificial ventilation of newborn infants. In: Proceedings of the AIME-93. Munich, Germany, 1993.
51. Widman LE. A model-based approach to the diagnosis of the cardiac arrythmias. Artificial Intelligence in Medicine 1992; 4: 1–19.
52. Dawant BM, Uckun S, Lindstrom DP, Manders EJ. Model-based signal analysis and interpretation in the intensive care unit. In: Proc. 14th Annual Conf. IEEE Engineering in Medicine and Biology Society. Paris, France, 1992: 874–5.
53. Factor M, Gelernter DH, Sittig DF. The multi-trellis software architecture and the intelligent cardiovascular monitor. Methods of Information in Medicine 1992; 31 (1): 44–55.
54. Factor M. The process trellis architecture for real-time monitors. In: Proc. 2nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming. Seattle, WA, 1990: 147–55.
55. Hayes S, Ciesielski VB, Kelly W. A comparison of an expert system and a neural network for respiratory system monitoring. Technical Report TR-92-01. Royal Melbourne Institute of Technology, 1992.
56. Hunter J, Chambrin M-C, Collinson P, Groth T, et al. INFORM: integrated support for decisions and activities in intensive care. International Journal of Clinical Monitoring and Computing 1991; 8: 189–99.
57. Rutledge G, Thomsen G, Farr B, Tovar M, et al. VentPlan: a ventilator management advisor. In: Proc. 15th Annual Symp. Comput. Appl. Med. Care. Washington, DC, 1991: 869–71.
58. Ash D, Gold G, Seiver A, Hayes-Roth B. Guaranteeing real-time response with limited resources. Artificial Intelligence in Medicine 1993; 5 (1): 49–66.
59. Alonso-Betanzos A, Moret-Bonillo V, Devoe LD, Searle JR. Obstetrical decision making based on predictive expert analysis. In: Proc. 13th Annual Conf. IEEE Engineering in Medicine and Biology Society. Orlando, FL, 1991: 1300–1.
60. Moret-Bonillo V, Alonso-Betanzos A, Truemper EJ, Garcia-Martin E, et al. PATRICIA: an expert system that incorporates a patient-oriented approach for the management of ICU patients. In: Proc. 14th Annual Conf. IEEE Engineering in Medicine and Biology Society. Paris, France, 1992: 876–7.

61. Alonso-Betanzos A. A connectionist approach to predict antenatal outcome. In: Proc. 14th Annual Conf. IEEE Engineering in Medicine and Biology Society. Paris, France, 1992: 1004–5.
62. Mora FA, Passariello G, Carrault G, Le Pichon J-P. Intelligent patient monitoring and management systems: a review. IEEE EMBS Magazine, 1993.
63. Larizza C, Moglia A, Stefanelli M. M-HTP: a system for monitoring heart transplant patients. Artificial Intelligence in Medicine 1992; 4: 111-26.

Address for offprints: Serdar Uckun, Rockwell International Science Center, 444 High St. Palo Alto, CA 94301, USA

Appendix 1: projects represented in the survey

The information in this section is up to date as of Fall, 1992. Multiple entries are made for projects which comprise multiple components. Where available, the most comprehensive and recent references are indicated. Further details about projects represented in the survey may be obtained from the author as a technical report [9].

Researchers	Institution	Description, status and references
Orr, Westenskow	University of Utah	Neural-network based anesthesia alarm system. Used to detect breathing circuit faults. Tested using simulators, animal studies, and in the OR [32].
King, Hyman, Xu	Vanderbilt University	Computer-based evoked-potential monitoring in the ICU. Used to evaluate patient status and nervous system function. Experimental [43].
King, Kambam, Smith, Jiang	Vanderbilt University	Respiratory system monitoring: respiratory patterns of dogs monitored during controlled ventilation in order to recognize normal and abnormal events. Project not active at time of writing [44].
King, Patterson	Vanderbilt University	Artificial intelligence vs. neural-network based monitoring: aims to compare the success rates of two methods independently and in combination. In development; will use recorded data sets.
Haimowitz	MIT	TrenDx: intelligent trend detection in pediatric growth data. Develops an epistemiology for trend templates; currently in trial using retrospective clinical data sets [12, 45].
Kohane	Children's Hospital/Harvard Medical School	Clinician's Assistant Project: an architecture for clinical event monitors. Currently uses growth monitors developed by Haimowitz (see entry above) [34].
Shahsavar, Wigertz, Gill	Linkoping University Sweden	KUSIVAR: a knowledge-based system for ventilator management of adult respiratory distress patients. An on-line version named VentEx is currently under trial [46, 47].
Coiera, Higgins, Tombs, et al.	Hewlett-Packard Labs, Bristol, UK	Intelligent alarm technology project: investigates the use of AI technology to enhance current generation patient monitoring systems [48, 49].
Duran, Nunez	Univ. of Barcelona	LATIDO: intelligent monitoring system for congenital heart defects. Focuses on all aspects of health care in the pediatric ICU: diagnosis, treatment, data interpretation, alarming. In development.
Miksch, Paky, Popow, Horn	Univ. of Vienna	A knowledge-based system for monitoring and optimizing artificial ventilation of neonates. In development [50].
Widman, Tong	Univ. of Oklahoma	EINTHOVEN: a model-based ECG interpretation system. Combines model- and rule-based methods to interpret ECGs in critical care, cardiologic diagnosis, and Holter monitoring [51].
Tong et al.	Louisiana Tech. Univ.	WeanPro: a knowledge-based system for ventilator management. Implements a weaning protocol; has undergone a successful clinical testing with encouraging results [42].

Wang, Wischnewsky, et al.	Univ. of Bremen	Intensive-Help: an intelligent monitoring system for ICUs which provides assistance with on-line data acquisition, presentation, interpretation, and documentation. On trial.
Eytan, Pinget, Keipes	Univ. Strasbourg I & II	Bedside Pancreas: a closed-loop insulin monitoring and pump control system. Under development.
Dawant, Uckun, Lindstrom, Manders	Vanderbilt University	SIMON: an architecture for patient monitoring in critical care environments. Addresses the issues of data interpretation, diagnosis, and prediction. Applied to a newborn ventilator management problem [30, 52].
Uckun	Vanderbilt University	YAQ: an ontology for model-based diagnosis and prediction in physiological domains. Used in the context of SIMON (see above). Evaluated using retrospective clinical data [15, 19].
Fertig, Factor, Sittig, et al.	Yale University	ICM: an intelligent cardiovascular monitor based on the process trellis parallel computing architecture. ICM was developed as a prototype patient monitor for post-cardiac surgery patients [53, 54].
Leaning, Patterson	University College London	GAMES-II intensive care exemplar: an EC project focusing on general methodologies and tools for medical KBSs. Current area of interest is the physiological monitoring and control of patients with sepsis [25].
Hayes, Cielsieski, Kelly, et al.	Royal Melbourne Inst. Tech.	Intensive Care Respiratory Monitor: focuses on techniques for combining neural and symbolic processing methods in ventilator management. Project currently not active [55].
Gardner, Hunter, Salatian	Univ. of Aberdeen Scotland	EC/TANIT: Part of INFORM project on providing decision support on drug therapy for intensive care. Currently under development. Same group also interested in interpretation and abstraction of cardiovascular monitoring data [56].
Rutledge, Fagan, et al	Stanford University	VentPlan: an architecture for model-based ventilator management advice. Combines patient-specific models based on belief networks, quantitative constraint models for prediction, and a treatment module. Developed for proof-of-principle [57].
Farr	Stanford University	Treatment plan evaluator for VentPlan (discussed above). Used to recommend ventilator settings based on a decision-theoretic model [23].
Rutledge	Stanford University	Model selection for ventilator management assistance: criteria and methods for selecting the model with the most appropriate level of detail given a patient-specific model. Currently in final stages of development [21].
Hayes-Roth et al.	Stanford University	Guardian: an intelligent architecture for ICU monitoring and ventilator management. Developed as a proof-of-concept system on a blackboard architecture; the application area is post-operative monitoring of cardiac surgery patients [26].
Ash, Hayes-Roth	Stanford University	A reactive response subsystem for Guardian. Used to select tests and treatments under time pressure based on a hierarchy of problems and related therapeutic actions [58].

Shahar	Stanford University	Resume: a formalism for temporal abstraction for interpreting clinical data. This is used in the context of Protege and T-Helper projects, and typically applies to slow clinical data with temporal and value uncertainty. Currently in final stages of development [13].
Fehlauer, Soller, et al.	VA GRECC, Utah	A set of projects evolving around hand-held computer systems for data collection and interpretation in a geriatric care setting. Data integration, analysis, and diagnosis are among major concerns.
Alonso-Betanzos et al.	Univ. de A Coruna Spain	NST-Expert: an expert system for antenatal monitoring and prediction of prognosis. Used for obstetrical decision support [59].
Moret-Bonillo et al.	Univ. de A Coruna Spain	Patricia: an expert system with a patient-specific model of ICU monitoring. Applied to ventilator management support including data interpretation and treatment [60].
Moret-Bonillo et al.	Univ. de A Coruna Spain	A case-oriented approach in ICU monitoring. An extension of the project mentioned above. Shares the same application platform [60].
Pazos-Sierra et al.	Univ. de A Coruna Spain	Application of neural networks to antenatal diagnosis and prognosis. Aimed to compare performance of neural networks with the expert system NST-Expert (mentioned above). Development of a hybrid system is also under consideration [61].
Passariello, Mora	Univ. Simon Bolivar, Caracas, Venezuela	Intelligent instrumentation in cardiology. Focuses on the creation of an infrastructure for knowledge-based instrumentation and decision support in a coronary care unit [62].
Larizza, Berzuini	Univ. di Pavia, Italy	Monitoring and prediction of clinical events and patient outcome using Bayesian networks. Used to monitor and to predict future patient responses to antiviral therapy [63].
Dojat, Borchard, Isabey, Harf	INSERM, Creteil, France	NeoGanesh: a knowledge-based system for closed-loop control of assisted ventilation. Successor to Ganesh. Applied to the weaning stage of pressure support ventilation. Successfully tested on patients [37].
van der Aa, Gravenstein	Univ. of Florida	Intelligent Alarm Project: an expert system used to detect faults in the anesthesia breathing circuit. Tested using a simulator, in animal studies, and in the OR. Under development [31].
Ogrodowczyk	German Heart Center, Berlin	Knowledge-based support for diagnosis and therapy of cardiac disease. Currently in design/early development stage.