International Journal of Clinical Monitoring and Computing **8**: 189–199, 1991.

INFORM: integrated support for decisions and activities in intensive care

James Hunter[1], Marie-Christine Chambrin[2], Paul Collinson[3], Torgny Groth[4], Anders Hedlund[4], Seppo Kalli[5], Aarno Kari[6], George Lenoudias[7], Pierre Ravaux[2], Donnie Ross[1], Jean-Marc Salle[1], Tommi Sukuvaara[5], Ron Summers[7] & Bertil Zaar[4]
[1] *University of Aberdeen, UK;* [2] *Inserm U279, Lille, France;* [3] *West Middlesex Hospital, London, UK;*
[4] *Uppsala University, Sweden;* [5] *Technical Research Centre of Finland, Tampere, Finland;*
[6] *Kuopio University Central Hospital, Finland;* [7] *City University, London, UK*

Accepted 10 July 1991

Key words: intensive care unit, decision support, user requirements

Summary

Many medical decision support systems that have been developed in the past have failed to enter routine clinical practice. Often this is because the developers have failed to analyse in sufficient detail the precise user requirements, because they have produced a system which takes too narrow a view of the patient, or because the decision support facilities have not been sufficiently well integrated into the routine clinical data handling activities. In this paper we discuss how the AIM-INFORM project is setting out to deal with these issues, in the context of the provision of decision support in the intensive care unit.

Introduction

Within Intensive Care, medical staff have to process information in many different ways and at many different levels. At the bedside, the immediate patient data have to be interpreted and abstracted, artifacts must be rejected, alarms recognised, analysed and filtered, all in the context of the known (or hypothesised) clinical state of the patient. In the longer term, this clinical state has to be established (or at least hypothesised), expectations generated for the future, therapeutic regimes established in order to return the patient to an acceptable state, and progress towards this desired state monitored. For the foreseeable future, decisions will continue to be taken by the medical staff concerned. However we now believe that it is possible to provide relevant decision support at all these different levels.

A considerable amount of work has been carried out in relation to decision support in different application areas e.g. Respiratory Management, Fluid and Electrolyte Therapy, Cardiovascular Management, etc. [1, 3–6, 10–14, 16]. However, little attention has been given as to how these individual applications can be integrated into a system to satisfy all of the decision support requirements in a comprehensive manner. In this paper, we describe how the AIM-INFORM project intends to approach these problems. It must be emphasised that at the time of writing, INFORM has been able to undertake only preliminary studies leading to the specification of requirements, and to initial design studies. It is appreciated that many difficult practical problems [6, 15] will need to be resolved before such a system can become an implemented reality.

Firstly, we discuss how we have set out to establish the user requirements for 'decision support'. It is difficult to define precisely what is meant by this

term; we have taken a fairly wide view which includes the display of information, and certain nursing activities and working procedures (which some might not call 'decisions'). Secondly, we set out a framework in which different decision support modules (which deal with different specialist medical areas) could communicate with one another, thus enabling the system to provide support based on an integrated view of the patient.

Finally, it is our firm belief that decision support must be available as an integral part of the complete information handling system, and we will discuss the relevant software and knowledge engineering issues.

Establising user requirements

Before embarking on the design and implementation of any large scale data processing system, it is essential to establish as far as possible the requirements of the users. In a technologically advanced (and advancing) area such as decision support, it is difficult for the users to specify what is required in detail since they will often be unaware of the possibilities. For this reason we tried to establish these requirements in a number of different ways, and then to combine them into a coherent whole. The approaches used were:

(1) to ask the INFORM Clinical Panel (consisting of seven senior clinicians from six European countries), and other clinicians working closely with the project, to list areas in which they felt that medical and nursing staff require decision support; here we were consulting senior clinical experts with some reasonable knowledge of the facilities which decision support can and might offer;

(2) to attempt to understand the areas of uncertainty in decision making in the ICU as perceived by junior doctors and nurses; the first phase of this study involved interviews with 12 members of the nursing staff and 13 junior doctors at one ICU site (Aberdeen); here we were attempting to establish where nurses and junior doctors felt *subjectively* that they would like help in decision making – we did not constrain them by presuppossing any technical solutions;

(3) to widen the results of (2) by administering questionnaires of a similar form to seventeen nurses at three ICU sites (Aberdeen, West Middlesex – London, and Lille); it must be emphasised that resources were only available to conduct a pilot study which would need to be extended if the results are to have any statistical significance.

It should be noted that we are concerned here only with *clinical* decision support. It is arguable that INFORM ought to provide support for *unit management* decisions (e.g. cost evaluation, management of staff, materials and pharmacy) and for *research activities.*

General considerations

We first set out some general requirements which emerged very clearly from our investigations:

- Any decision support system must satisfy one or more of the following requirements:
  - it must improve the use of available resources;
  - it must increase patient safety by preventing mistakes of omission or commission (mainly supporting inexperienced junior doctors and nurses);
  - it must improve the quality of patient care through improved decision making and/or by allowing more time to be spent directly with the patient.
- The system must *assist* people in making decisions, not *replace* them as decision makers. It should support those functions which people do poorly, e.g. keeping track of large volumes of evolving patient data, retaining and applying a constantly increasing volume of medical knowledge, particularly of drugs and other techniques. As a consequence, the system must interact with the user at all times in such a way as to make it clear that the user has the responsibility for making decisions. There are many rea-

sons for taking this view, among them the problem of legal liability, and the need to avoid de-skilling through under-use of critical facilities.

- In designing ICU decision support systems, care must be taken to understand the requirements of all relevant types of staff; the system must support working procedures for different categories in a way which is appropriate to their abilities and responsibilities. Although many units will want to make all of the system's facilities available to all users, it may be that some units will wish to limit the use of specific facilities to certain categories of user, and some means of incorporating such restrictions should be available.
- It is appreciated that a large part of the activities of nurses and juniors consists in following protocols. The system can make it easier to do this by making sure that these protocols are maintained centrally and are easily accessible. It must also be recognised that protocols have to be interpreted for specific patients, and that decision support has a role to play in this interpretation.
- The system must be capable of being adapted to local conditions; it must be flexible enough to cater for differing national practices, and for the individual requirements of different supervising clinicians (even within the same ICU).
- In its initial implementation it must support a minimum level of decision support, and it must be capable of incremental upgrading as new and better decision support techniques are developed.
- It must be flexible in its mode of operation; users must be able to chose those components which they wish to consult; thus a doctor may want to consult the system about what drugs to employ and be able to impose his/her assessment of patient state, even if that assessment is different from that inferred by the system.
- It must provide a time response which is consistent with the decision being supported (e.g. there is no point in taking an hour to generate advice for a situation which requires action in minutes).
- It must be sophisticated enough (i.e. contain enough knowledge) to be of real use; however the necessary level of sophistication will depend on the type of user.
- It must be unobtrusive and well integrated into the overall information handling system. It must be easy to use, with an attractive interface. It must not demand large amounts of training on the behalf of the prospective user; doctors and nurses are there to take care of the patient, not the system.
- The system should not be designed on the assumption that it can provide all of the answers. It is acceptable to say to the user: 'It is not possible to provide advice for this type of patient; it is suggested that you consult Dr. X'.

Organisation of the requirements

We do not have space here to set out our current statement of user requirements in full. However, we can describe how those requirements were organised around the cycle of decision-making activities and actions undertaken by the principal actors in the ICU: nurses, junior doctors, and senior doctors. Since we are setting out the *user* requirements, it seems appropriate to consider user activities as the organising principle (rather than the activities of any computer system). This cycle consists of decision followed by actions:

Decisions
  Assessment of Patient State
  Planning and Ordering
    Monitoring and Investigations
    Therapy

Actions
  Treatment Administration
  Bedside Monitoring
  Checking on Investigations

Assessment of patient state

We use the word 'assessment' rather than 'diagnosis' to indicate that we are concerned with establish-

ing the relevant aspects of the patient's state; this *may* include some form of 'diagnosis' but is *certainly* concerned with identifying those problems which require action.

Planning and ordering
Having made an assessment of the current state of the patient, decisions have to be made as to (i) what further data need to be collected, and (ii) what treatment needs to be given. In both cases a plan needs to be made and instructions given for that plan to be carried out.

Monitoring and investigations
These will include defining those parameters which are to be continuously monitored, and any laboratory investigations to be undertaken.

Therapy
We are concerned here with treatment decisions about actions to be taken e.g. drugs, ventilation, fluid therapy, physiotherapy, etc. Once a treatment has been decided on, a set of expectations will be generated (i.e. what we expect to happen to the patient in the given state as a result of the treatment which is to be applied). These expectations will be expressed in different ways: e.g. the heart rate will remain lower than normal, the blood pressure should return to normal within three hours.

Treatment administration
By this we mean not only the giving of fluids, drugs, etc, but also the carrying out of routine nursing care.

Bedside monitoring
ICU staff have to monitor the state of the patient on a continuous basis. This means checking that the expectations that the clinicians have for the patient are being achieved; expectations may be positive (e.g. the heart rate will return to normal within six hours), or more implicit (e.g. there will be no catastrophic fall in blood pressure). Various actions will be triggered if there is any significant departure from those expectations.

Checking on investigations
If investigations have been ordered then the staff have to check that the results are received at the anticipated time.

It is clear that forecasting is a key aspect of decision making, and is an essential part of all these stages. Once the patient state is identified, the evolution of that state is projected forward (prognosis) to decide if the consequences are acceptable (evaluation of outcome). If they are not, then treatments are proposed, and part of the process of treatment selection is the forecasting of the expected effect of that treatment over time. This cycle has to be initiated and terminated, and special consideration needs to be given to these activities:

Admission
By this we mean the process of deciding whether to admit the patient to the ICU at all.

Post-admission procedures
Once the patient has been admitted, an initial review is undertaken. In fact this consists of the activities of *assessment* and *planning and ordering* discussed above. However these decisions tend to be qualitatively different from those made during the remainder of the stay in the ICU, and we therefore considered them separately e.g. what parameters are to be monitored and whether to ventilate or not. Nevertheless, such decisions might be altered later on, and this possibility should not be neglected.

On-going patient management
Once the initial management decisions have be taken, the basic cycle of actions (bedside monitoring, investigations, and treatment administration) and decisions (assessment and planning) continues during the remainder of the patient's stay in the ICU.

Discharge
Finally the decision to discharge the patient from the ICU has particular aspects which need to be treated separately.

As discussed earlier, different categories of user

will have different requirements; we have identified specific requirements for *nurses, junior* clinicians, and *senior* clinicians. Within a given category there will be different levels of expertise. The precise definitions of 'junior' and 'senior' will vary from country to country and unit to unit. Generally speaking, junior doctors will be those who spend large amounts of time on the ward; seniors will have less frequent contact with the patient. When we consider that seniors would use a particular decision support function, it is understood that they would probably do so less frequently than a junior. In addition, the division of decision-making responsibilities between juniors and seniors will vary. However, even in situations where most decisions have to be ratified by seniors, it is clear that it is essential for the continuing training of juniors that they come to their own independent conclusions, and that decision support systems can be of use in this context.

There are two basic ways in which a decision support system may interact with the user: *user-intiated* and *system-initiated.*

User-initiated

The user initiates the interaction with the system; this may be scheduled as part of the normal routine (e.g. ward round, shift change, etc.) or it may be at some other time; the important distinguishing factor is that it is the user who takes the first step. There are many possible *modes* of interaction under these circumstances: *knowledge base querying,* where the user has free hand to explore the contents of the knowledge base as it applies to the patient in question; *critiquing,* where the user asks for a possible solution to be commented on; *consultation,* where the user asks for the system's opinion. The user may also decide to use whatever general *training* facilities are available. This may be training in the use of the system, or clinical training based on real cases; however the development of such training materials is non-trivial.

System-initiated

Here it is the system that takes the initiative when, while working in a background mode of operation, it 'notices' that something has happened which needs to be brought to the attention of the user. Such a functionality is often called 'watch-dog'. This may be an urgent *alarm* based on a rapid deterioration in patient state, and attention demanding mechanisms (sound, flashing, colours, etc.) will be appropriate. On the other hand, it may be a less urgent *alert.* If unattended to, an alert may become an alarm. Alerts may be posted on the screen automatically, or the system may wait until an appropriate user next has an interaction. One can imagine a situation where the user has entered a drug prescription without actively consulting the decision support system, but the system raises an alert if there is likely to be some undesirable drug interaction, or other consequence.

In general, both types of interaction are possible; however there are circumstances where one type is more appropriate than another.

Priority areas

Our specification is still under revision; however several areas have emerged as being of high priority:

- smart monitoring and alarming by correlating several simultaneous data channels – most importantly heart rate and rhythm, blood pressure, urine output, O_2 saturation, and temperature;
- management (identifation and therapy) of problems in the respiratory and cardiovascular/fluid systems;
- administration of drugs/fluids;
- administration of routine tasks (task lists, task recording, activity-sensitive help screens, etc.).

Integration of different specialisms

In developing decision support systems for a complex domain, there is a natural, and understandable tendency to divide the domain up into a number of smaller, and more tractable sub-specialisms. However, the nurse or doctor on duty is faced with the problem of managing the patient as a whole, and one of our principal objectives is to

establish an architecture in which this patient-wide decision support can take place.

In considering such an architecture, it is valuable to restate the most relevant constraints imposed by the user requirements:

- the decision support facilities must be capable of being used in a variety of different ways by a number of different users; in particular, a user who wishes to access specific features, but not others, must be allowed to override any conclusions that the system has arrived at during earlier stages of processing;
- the system must take an integrated and coherent view of the whole patient; if it has information on one physiological subsystem which is highly relevant to reasoning about another subsystem, then that information must be considered;
- knowledge must be available for use in different ways; the same knowledge can be used both for system-initiated 'watch-dog' functions and for user-initiated consultations:

Further contrainsts arise from considerations concerning the implementation:

- because the development of decision support has to be an on-going task, it must be possible to introduce the decision support facilities on a progressive and incremental basis; our aim is to introduce relatively simple aids, for which we already have the technology, at once, and, as their development matures, provide more advanced facilities.
- because of the scale of the project, different facilities will be developed by groups at different geographical locations.
- it is likely that different components of the system will run on distinct computers; the more local and time critical functions will be performed on a per-patient basis at, or near, the bedside; the longer term functions will be located centrally, probably on shared hardware; the precise partition has still to be decided, but the architecture must allow for it.

All of the above factors point to a highly modular architecture with very clearly defined interfaces between modules. The best example of this type of architecture in Artificial Intelligence is the 'blackboard' model [7], and in fact other groups are using such a model to implement ICU decision support systems [8, 13]. The basic blackboard model consists of two components: knowledge sources (KSs) and the blackboard database. The KSs contain the 'knowledge' needed to solve the problem and are separate and independent. The blackboard data structure stores the problem-solving state. KSs interact only through the blackboard and create data structures on the blackboard that correspond to (partial) solutions. It is very common for the blackboard to be partitioned into a number of levels with different KSs reading from and writing to a limited number of levels. The knowledge in a KS need not be heuristic nor expressed as rules; it could equally well be algorithmic, statistical or consist of some form of qualitative or quantitative model.

In the implementation of a blackboard system, considerable attention has to be given to problems of 'control'. Often a set of control KSs exist that monitor the changes to the blackboard and decide what is to be the next focus of attention. Control in a blackboard system is often opportunistic, i.e. the way in which the computation evolves is highly sensitive to the context provided by the data on the blackboard. At this stage in our design, we are not yet concerned with control issues, and in fact are not committed to the blackboard architecture for implementation. It is quite likely that control will be much more scheduled and user-driven than opportunistic. However the blackboard as a *conceptual* model, which contains a number of modular processes communicating only through a shared data store which is partitioned into a number of identifiable levels, has provided us with a very powerful metaphor which has helped to guide our overall design.

In defining the information on the blackboard, it will be necessary to provide a number of annotations:

- where the information comes from: information may be entered at a particular level by the user; alternatively it may be the result of reasoning by the system; it is clearly necessary to distinguish the two;
- temporal aspects: as a minimum, data must be annotated with the time it was entered, and with the time which it refers to; for example, the

results of a laboratory investigation will be entered at a later time than that at which the sample was taken; a clinician may make a recommendation for a treatment at a specific time, and predict the effects of that treatment for some future time; see [2].

To indicate our overall direction, we will outline the general principles governing the definition of the different levels on the blackboard and the kinds of KSs that will manipulate that data; several of our levels are similar to those defined by others [13].

History

These are mainly symbolised data concerning the history of the patient before entry to the ICU, for example 'heavy smoker', 'history of angina', etc. There is an essential distinction between chronic diseases, and acute diseases that have led to the need to admit the patient into the ICU; this will be made more explicit in the future.

Equipment

Data here are concerned with the characteristics (e.g. type, available settings, etc.) and status (on/off, current settings, fault status, etc.) of the equipment e.g. monitor, ventilator, etc. It can be derived:

(i) from the data-base;

(ii) by inference from other levels (e.g. the status of the ventilator as derived from various respiratory parameters).

Our criterion for distinguishing this level from the raw data level is that the data here are not directly determined by the state of the patient (although they may be eventually, as settings are altered in response to changes in the patient state).

Raw data

These are patient data as derived directly from the data-base. They may come from instrumentation, laboratory results, observation of the patient, etc. Data types can be numerical, boolean, or symbolic – the latter corresponding to any clinical data which may be available. The guiding principle is that these should be data which are collected directly from the patient; terms such as findings, manifestations, signs, symptoms, etc. are often used. Information about the patient which is the result of some reasoning should not appear here; for example, if an initial diagnosis is available, this should appear at one of the interpretation levels. Note that clinical data will normally have to be entered manually, and that it is a design goal that a minimal system should run with very little or no manually entered data.

There is some problem as to exactly where we draw the boundaries of our system. Strictly speaking the true 'raw' data consist of the analogue signals coming directly from the patient. However for our present purposes we will assume that a certain amount of pre-processing takes place. This will depend on the particular type of equipment concerned; for ECG monitors for example, we will assume that we do not have access to the raw signals, but use the normal monitor outputs such as heart rate, etc.

Systematically preprocessed data

These are data which come from the systematic processing of certain elements of the raw data (usually relatively high frequency monitor data); 'systematic' processing means digital processing which is always applied. Different filters, curve fitting, power spectrum estimation, knowledge of signal statistics, and other methods will be used. Different methods will be applied to different signals. Normally each signal will be processed independently of the others, but there may be exceptions.

Cleaned data

These are data which the system has validated. If data is judged to be invalid some form of interpolation may be used to recover the missing values, or the values may be set to 'unknown'. An item of

clean data may be derived solely from the corresponding raw or preprocessed data or there may be a much more complex process involving inferences made on abstractions such as the current patient state. Artifacts and other misleading events will be recognised by KSs which have specific knowledge of the ways in which these manifest themselves.

Initially it is likely that much of the data will be considered valid by default; this will be true of data which are normally considered to be error free (e.g. symbolic data entered by the user, or the results of laboratory investigations), but as the system becomes more sophisticated, it should be capable of recognising occasional errors in data entry (e.g. the misattribution of results from one patient to another).

An important operation at this level is the generation of derived data; it only makes sense to perform calculations on data which are believed to be error free.

Events

This is the level at which we record those external events which have been detected by the system or entered directly by the nurse, which may affect the way the data are interpreted, but which does not intrinsically reflect the internal state of the patient (i.e. they are not interpretations); examples are events which generate artifacts of various kinds (e.g. the taking of a sample from a blood pressure line, turning the patient).

Symbolised data

Symbolisation means the conversion of a numerical value into a symbolic one, normally the comparison with predefined ranges. It is important to be clear that there are different kinds of ranges involved:

(1) the 'normal' state for that patient. The ideal situation would be to use as a reference the state the patient was in before the onset of the current acute episode. However, usually these parameter values have not been measured, and have to be estimated by taking the parameter values for a normal healthy person and modifying them to take account of age, sex, general state of health and chronic diseases, etc. In the simplest case, symbols as *normal, high* and *low* would be used, but, for *high* and *low*, qualifiers such as *slight, moderate* or *excessive* may be necessary.

(2) the goal for the patient which treatment is designed to achieve (the 'optimal' values – as determined by the physician in charge, possibly in accordance with any protocols currently in force). This is not necessarily the same as the normal values; for a patient in shock, the goal for Cardiac Output may be higher than normal in order to maintain adequate delivery to mitochondria in the cells. Such goals will be maintained at one of the treatment levels. Phrases such as *less than desirable* will be used.

(3) the anticipated state of the patient given the treatment that is to be administered. Clinicians realise that treatments take time to act, and that different treatments may produce conflicting effects. We therefore have to represent the knowledge that the clinician has certain expectations for the patient, which can be represented as ranges, and thus the actual data will be symbolised with respect to these ranges. Symbols such as *as expected* and *lower than expected* will probably be appropriate. Thus we might have the statement "the patient's blood pressure is lower than we would like it to be, but given that he is on drug X, the blood pressure is as expected'. Expectations can be generated both for the present and for the future (and possibly even for the past if we are trying to verify a hypothesised condition).

If the data form a time series, then the first derivative can be symbolised as well – this is usually referred to as the 'trend'. Data are given a symbolic trend i.e. *rising, steady* or *falling;* again modifiers may be appropriate such as *slowly* or *rapidly*.

It is clear that such a level on the blackboard is clinically meaningful. Ideally the symbolisation would be carried out once only, independently of

the different knowledge sources which will use the symbolised data. However it remains to be seen whether those knowledge sources will agree on the appropriate symbolisation; it will be a matter of empirical test to see whether this is a real problem.

Much of the data will already be symbolised (e.g. medical history and observations); for the moment we will assume that this is copied unaltered from the clean data level to this one.

Single interpretations

By this we mean statements concerning the state of an individual body system which can be derived without referring to other body systems (e.g. the acid-base system). Single interpretations will normally be made on the basis of data from lower levels on the blackboard. Interpretations may be derived by the system; they may also be entered directly by the clinician concerned.

Once a considerable amount of reasoning takes place, we have to be concerned with problems of representing uncertainty. To a large extent, the uncertainty depends on what measurements are available. For example, when assessing preload of the left side of the heart, the assessment may not be very reliable if only chest X-ray, central venous pressure and systemic arterial pressure are available; however the reliability increases if pulmonary arterial pressure, pulmonary arterial wedge pressure, or even left atrial pressure are measured.

The number of levels of interpretation is still under consideration; to some extent we have to be driven by what is natural when looking at specific domains. We allow for the possibility of defining 'intermediate levels' if necessary.

Combined interpretations

There are a number of ways in which interpretations can become combined:

(i) the interpretation concerns one body system, but requires the use of another system for the interpretation to be made;

(ii) the interpretation consists of more than one simple interpretation with some relationship existing between them e.g. renal failure and left ventricular failure both caused by hypertension.

This is the level at which, ultimately, overall patient status will be represented.

Single treatments

Here we are concerned with treatments which are derived from the interpretation of a single physiological subsystem (e.g. the administration of a cardio-active drug for heart failure). Interactions between treatments are not taken into account at this level.

The first objective at this level is to establish treatment goals; these will include general management strategies, as well as changes in the desired values for specific parameters. The current state of the patient is assessed in the light of these goals and any new therapeutic intervention determined. Finally, any new requirements for monitoring and investigation are established.

Combined treatments

Single treatments have to be combined into an overall treatment plan for the patient; drug interactions have to be checked, and total fluid inputs checked against overall fluid balance requirements. There may be contradicting demands which have to be resolved.

Integration within a single system

Previous experience with decision support systems has shown that users are very reluctant to use them if they involve any significant departure from their routine activities. In the ICU, we have the possibility to embed decision support facilities in a comprehensive and unified computer-based system which will combine the functions of data acquisition, data

storage, data display and access to external data. The decision support components will interface with the patient data management system, and contribute to a coherent user interface; using them will therefore impose no extra overhead on the user.

Software engineering

In a companion paper [9] our colleagues describe the use of CASE tools as the basis for a systematic approach to software engineering in our domain. In order to ensure that the decision support facilities are fully integrated with the rest of the system, we are also using the same tools (notably data flow diagrams) as part of the design process. Furthermore, our aim is to integrate the various types of information specified in the previous section with that required for the conventional data processing. Our objective is to achieve one set of specifications covering all aspects of the system.

Knowledge engineering

Although the use of CASE tools is necessary to achieve overall integration, we feel that they will not be powerful enough when it comes to engineering some of the complex knowledge sources. Knowledge engineering consists of knowledge acquisition and system design and implementation; by knowledge acquisition we mean that activity which encompasses knowledge elicitation (from domain experts) and knowledge modelling (implementation independent description of the domain knowledge). On reviewing the state of the art on knowledge acquisition we decided that the KADS methodology was most relevant and appropriate, and spent some time understanding it in more detail and working with the tool that has been developed to support it (Shelley) [17].

We have tested the adequacy of the KADS conceptual language to model expertise in our domain by using it to describe six existing systems using a common analytical framework in a form of 'reverse engineering'. This consisted of abstracting the conceptual model of an implemented system, in contrast to the normal cycle of definition and then implementation. Our conclusion is that KADS does constitute a suitable approach for knowledge acquisition, and we are waiting to see what tools will become available to assist with system design and implemtation.

Conclusions

It was always our intention that the exploratory phase of INFORM should provide the necessary groundwork to enable a longer term project to start from a firm foundation. We consider that we are close to achieving that objective.

Acknowledgements

We are grateful to the CEC for the provision of funding under AIM project A1029 which has enabled the work described in this paper to be carried out.

References

1. Arturson G, Groth T, Hedlund A, Zaar B. Computer simulation of fluid resuscitation in trauma; first pragmatic validation in thermal injury. Burn Care Rehab, 1989; 10(4): 292–9.
2. Ash D, Hayes-Roth B. Temporal representations in blackboard architecture, Internal Rep, Knowledge Systems Lab, Dep Comput Sci, Stanford Univ, 1990.
3. Autio K, Makivirta A, Sukuvaara T, Kari A, Koski E, Saranummi N, Hunter JRW. Information management and decision support for the critically ill. In: Fox J, Talmon JL, editors Lecture notes in medical informatics, workshop: system engineering in medicine, Springer Verlag, 1989.
4. Chambrin MC, Chopin C, Ravaux P, Mangalaboyi J, Lesteval P, Fourrier F. Computer-assisted evaluation of respiratory data in ventilated critically ill patients. Int J Clin Monit Comput 1989; 6(4): 211–15.
5. Collinson PO, Jones RG, Howes M, Nicholls J, Sheehy N, Boran GR, Cramp DG. Of mice and men – data capture in the clinical environment. Int J Clin Monit Comput 1989; 6(4): 216–22.
6. East TD, Henderson S, Morris AH, Gardner RM. Implementation issues and challenges for computerized clinical protocols for management of mechanical ventilation in ARDS patients. SCAMC-89 1989: 583–7.

7. Engelmore R, Morgan T. Blackboard systems. New York: Addison Wesley, 1988.
8. Hayes-Roth B, Washington R, Hewett R, Hewett M, Seiver A. Intelligent real-time monitoring and control, IJCAI 89, 1989.
9. Kalli S, Leaning M, Ambroso C, Tuomisto T, Heikelä A, Ilomäki A, Yates C, Gregory R, Mereu M, Marraro G. INFORM: Conceptual modelling. In: Information technology in anesthesia, critical care and cardio-pulmonary medicine, 1990.
10. Makivirta A, Sukuvaara T, Koski E, Kalli S. Towards symbolization of intensive monitoring data for knowledge based inference systems. In: IEEE Engineering in Medicine & Biology Society 10th Annual International Conference (New Orleans) 1988: 1282–4.
11. Ross D, Ramayya P. Aberdeen Intensive Care Unit System (ABICUS): a medical information organizer, adviser and auditor. J Clin Monit 1988; 4(2): 158.
12. Shabot MM, LoBue M, Leyerle BJ, Dubin S. Decision support alerts for clinical laboratory and blood gas data. Int J Clin Monit Comput 1990; 7: 27–32.
13. Sittig DF, Pace NL, Gardner RM, Beck E, Morris AH. Implementation of a computerised patient advice system using the HELP clinical information system. Comput Biomed Res 1989; 22: 474–87.
14. Sittig DF, Gardner RM, Pace NL, Morris AH, Beck E. Computerized management of patient care in a complex, controlled clinical trial in the intensive care unit. Comput Meth Progr Biomed 1989; 30: 77–84.
15. Sivak ED, Gochberg JS, Frank DM. Lessons to be learned from the continuing epic of computerizing the intensive care unit, SCAMC-89 1989: 605–8.
16. Summers R, Carson ER, Cramp DG. An intelligent knowledge-based management system for patients undergoing artificial ventillation: AIRS. Medical Informatics 88: International Conference on Computers in Clinical Medicine, Nottingham, 1988.
17. Wielinga BJ, Schreiber G, de Greef P. KADS: Synthesis Report, Deliverable Y3: ESPRIT project P1098, University of Amsterdam, 1989.

Address for offprints:
Dr J. Hunter,
Department of Computing Science,
University of Aberdeen,
King's College,
Old Aberdeen,
AB9 2UB, UK.