ELSEVIER

Computer Methods and Programs in Biomedicine 72 (2003) 65–80

Computer Methods and Programs in Biomedicine

www.elsevier.com/locate/cmpb

Towards knowledge-based systems in clinical practice: Development of an integrated clinical information and knowledge management support system

Dimitris A. Kalogeropoulos [a], Ewart R. Carson [a,*], Paul O. Collinson [b,a]

[a] *Centre for Measurement and Information in Medicine, City University, Northampton Square, London EC1V OHB, UK*
[b] *Department of Chemical Pathology, Mayday University Hospital, Surrey CR7 7YE, UK*

Received 1 April 2002; received in revised form 1 July 2002; accepted 1 July 2002

Abstract

Given that clinicians presented with identical clinical information will act in different ways, there is a need to introduce into routine clinical practice methods and tools to support the scientific homogeneity and accountability of healthcare decisions and actions. The benefits expected from such action include an overall reduction in cost, improved quality of care, patient and public opinion satisfaction. Computer-based medical data processing has yielded methods and tools for managing the task away from the hospital management level and closer to the desired disease and patient management level. To this end, advanced applications of information and disease process modelling technologies have already demonstrated an ability to significantly augment clinical decision making as a by-product. The wide-spread acceptance of evidence-based medicine as the basis of cost-conscious and concurrently quality-wise accountable clinical practice suffices as evidence supporting this claim. Electronic libraries are one-step towards an online status of this key health-care delivery quality control environment. Nonetheless, to date, the underlying information and knowledge management technologies have failed to be integrated into any form of pragmatic or marketable online and real-time clinical decision making tool. One of the main obstacles that needs to be overcome is the development of systems that treat both information and knowledge as clinical objects with same modelling requirements. This paper describes the development of such a system in the form of an intelligent clinical information management system: a system which at the most fundamental level of clinical decision support facilitates both the organised acquisition of clinical information and knowledge and provides a test-bed for the development and evaluation of knowledge-based decision support functions.

Keywords: Medical decision support; Knowledge-based systems integration; Integrated clinical information management support

* Corresponding author. Tel.: +44-20-7040-8370; fax: +44-20-7040-8364
E-mail address: e.r.carson@city.ac.uk (E.R. Carson).

0169-2607/02/$ - see front matter
PII: S0169-2607(02)00118-9

1. Background

1.1. On health care delivery

Dealing with the increasing cost of health care delivery and the detrimental consequences on the quality and effectiveness of clinical outcome is the primary concern of health care delivery system (HCDS) administrators and policy makers throughout Europe. Until recently, HCDS control primarily involved economic measures and activities. Today, such an approach is obviously insufficient and often dangerously erroneous. HCDS control may be both effectively as well as efficiently performed qualitatively, to both the indirect as well as the direct or immediate benefit of the patient, via the proper implementation of the comprehensive spectrum of knowledge, techniques, and technologies developed by research in measurement and information in medicine (MIM). These tools allow supportive and often corrective action to be taken much closer to the source of the problem rather than in pursuit of compromising costs and symptoms of an ailing control philosophy.

During the past decade a number of relevant methodologies have emerged, as dictated by key qualitative cost–benefit control protocols. The rapidly disseminated practices of evidence-based medicine [1–3] and outcome-based medicine [4,5] or disease management [6], are concepts which were born and developed within the realm of MIM and associated technologies, have led to the proliferation of quite a number of proprietary approaches to clinical decision making support. Some of these include the use of advanced information technologies, while some have negligently avoided the use of the underlying enabling tools. Evidence-based clinical guidelines [7,8] and care pathways [9] are but a taste of these approaches.

MIM researchers have for long maintained that the effective introduction of IT to the task of supporting and facilitating clinical decisions will help improve the quality of patient care, in relation to the management of disease as well as individual conditions and patients, optimise the cost–benefit equation, and ultimately transform the traditional structure of health care provision. Furthermore, substantial evidence has been produced that strongly supports these claims. Clinical information is the core of this (re)evolution towards effective and efficient HCDS control, and supporting the management of clinical information is the key to evidence-based decisions and disease management. Systems designed to perform this function use a range of type of disease process [knowledge] models, some of which have been adequately assessed in terms of their potential to improve decision making in real time. It is time to accept the fact that MIM produces working results in this respect and to grant the introduction of safe and useful systems into routine clinical practice, thereby enabling the effective and efficient application of evidence-based and disease management methodologies in clinical practice while at the same time providing the necessary means for improving both the quality of an HDCS as well as of the experience these systems embody.

1.2. Review of the underlying decision support technologies

Researchers in the fields of MIM, medical informatics and artificial intelligence in medicine (AIM) have for long maintained that the effective introduction of information technology to the task of supporting and facilitating clinical decisions, will help improve the quality of patient care, optimise the cost–benefit equation, and ultimately transform the traditional structure of health care provision [10–13]. Significant advances have indeed been achieved in these fields in the past three decades [14,15] and, having demonstrated an ability to effectively support clinical decisions [16–19], knowledge-based systems (KBS) are becoming increasingly ubiquitous in various clinical settings.

Nonetheless, few systems have so far been successful in entering routine use. The two main problems underlying this failure originate in that KBS are evidence-based medical decision support systems. This means that they are designed to model diseases based on heuristic and therefore imprecise and incomplete knowledge representations and are built to reason with partial belief and

incomplete information using such representations. The difficulty faced in this context is twofold. On the one hand, their successful operation results from linking together pieces of knowledge as the result of considering patient measurements and observations at various levels of abstraction, in order to reach a logically valid diagnostic or therapeutic conclusion given the data and their knowledge content. In order for this to work, KBS are mainly developed using AI programming languages that are specifically designed for this task. However, these languages are part of an overall systems development environment which does not lend itself to complex information management systems development. As a result, KBS are perceived primarily as experimental entities and are essentially isolated from the clinical environment itself. On the other hand, KBS must be evaluated with regard to their ability to effectively and safely support decisions and this process is both extremely lengthy [16] as well as nearly impossible as long as they are isolated from the clinical environment [20]. With respect to evaluating their correctness and accuracy, there is neither an unambiguous standard against which to compare the accuracy of the generated decision-supporting information, nor a precise and absolute method for establishing the correctness of the process by which it is generated [21,22]. With respect to evaluating their clinical effectiveness, firstly their purpose in the clinical setting must be clearly defined and secondly their performance must be constructively assessed [23] within this context. Since medical decision support problems are broadly defined as determining, how, when, and in what manner to provide information to health care professionals in order to increase the quality of their decisions with respect to individual patients or populations of patients [21], the purpose of developing KBS must be clinical information management rather than direct and pure diagnostic or therapeutic support.

Over the past decade or so it has become clear that in order to provide the means to assess the above problems and to develop effective solutions, KBS and other clinical decision support systems must be integrated within the information processing activity of the clinical user, for the development of systems geared toward supporting the management of clinical information [13–15,24]. A good example of this is the well known QMR system which was originally designed to function as a standalone consultation system in the pioneering INTERNIST-1, II and CADUCEUS AI incarnations, and which became successful in clinical practice as a diagnostic aid following its conversion into an integrated clinical information management tool [25–29].

2. The integrated system

2.1. Aims and objectives

The intelligent clinical information management support (ICIMS) system described below has been developed with the aim of overcoming the aforementioned obstacles encountered in the development and dissemination of clinical decision support KBS. More specifically, ICIMS has been designed and developed with the aim to integrate a prolog prototype blood–gas analysis interpretation KBS called BGAS into the routine clinical information management environment of the critical care unit (CCU) and to provide a test-bed for the development and evaluation of clinically effective intelligent-agent decision support functions. Specific objectives were:) >

1) To design a system which combines the computer-based clinical decision support tasks of the acquisition, organisation, storage, update and review of the information generated in the process of monitoring the ICU patient, as well as of the domain knowledge-base required for the contextual interpretation of the acquired clinical information, within a singular system architecture.
2) To use the clinical information management support system in order to develop and constructively assess the integration of the cognitive, clinical information processing tasks comprising the prototype KBS interpretative problem-solving task-domain into the ICIMS system and consequentially into clinical practice, in order to incorporate the

computational intelligence necessary for the interpretation of the patient data acquired in the process being supported.

3) To provide the means to assess specific problems encountered in the integration process, and to develop effective solutions, by employing an approach which enables the active participation of a clinical advisor who acts as an assessor of the functional, cognitive and ergonomic effectiveness of the KBS integration process, and of the overall decision support provided by the ICIMS system during its development.

2.2. The KBS prototype

An earlier system, BGAS, was developed in prolog to provide computational tools for the acquisition, representation and manipulation of the domain knowledge-base required for interpretative decision-making in the domain of acid–base balance, and to assess the performance of a singly connected hierarchical belief network in providing assistance with the interpretation of blood–gas laboratory analysis data [30,31]. A knowledge editing environment, named Framebuilder, was developed to enable clinicians to construct a strict hierarchy of probabilistic classification knowledge frames, and to specify expected patterns of evidence for the recognition of 16 simple and complex disorders of acid–base metabolism, by choosing clinical parameters from a vocabulary of laboratory data, signs and symptoms, relations between data variables and clinical history.

Fig. 1 depicts the prototype belief network which was constructed using Framebuilder for the probabilistic classification of the evidential information generated by blood–gas analyses. Each clinical parameter specified in the disorder profile frames was accompanied by the conditional probability of the particular piece of evidence being observed, given the disorder represented in the frame. Furthermore, each frame was assigned an a priori value of the probability of the occurrence of the represented disorder given no evidence had been observed. Table 1 lists the 16 disorder profiles represented in the prototype knowledge base, along with their basic definitional features. The hierarchical belief network constructed using Framebuilder was processed in the manner suggested by Ref. [32], in order to assess and propagate the effect of each piece of evidence given in a case, using a blackboard controlled, task-specific reasoning module for the construction of patient-specific models (PSM) of probabilistic classification, from the general hierarchical model [33]. Fig. 2 depicts the resultant dual-panelled blackboard architecture with its associated interpretative task-domain knowledge sources.

PSM blackboard entries were split into five levels of abstraction for the physiological diagnosis panel, and four levels of abstraction for the clinical diagnosis panel. The physiological diagnosis panel was designed to function in a bottom-up manner, starting from raw patient data and proceeding up toward the root of the virtual PSM hierarchy, to produce a differential diagnosis of disorders of acid–base metabolism. The clinical diagnosis panel was designed to work in the opposite direction, starting with a clinical diagnosis entered by the user and proceeding down toward the leafs of the virtual PSM hierarchy to generate expected consequences, which were latter used to critique the results from the physiological diagnosis panel in order to refine complex interpretative hypotheses which could not be differentiated in the light of measurement data alone.

Following a series of tests performed in order to evaluate the translation of a theoretical design into a working prototype, the system was retrospectively evaluated with 60 cases of perturbed acid–base metabolism and was found to perform at the level of the expert who designed the knowledge-base. Not counting the cases of complex disorders, the evaluation study showed that BGAS was in agreement with either the expert or senior clinician involved in the study in 83% of the evaluation cases. In order to proceed with designing methods to resolve the diagnosis of complex disorders, it was necessary to integrate the KBS reasoning engine into the routine clinical information management environment of blood–gas analysis in the CCU. This has been achieved via the development of the ICIMS system.

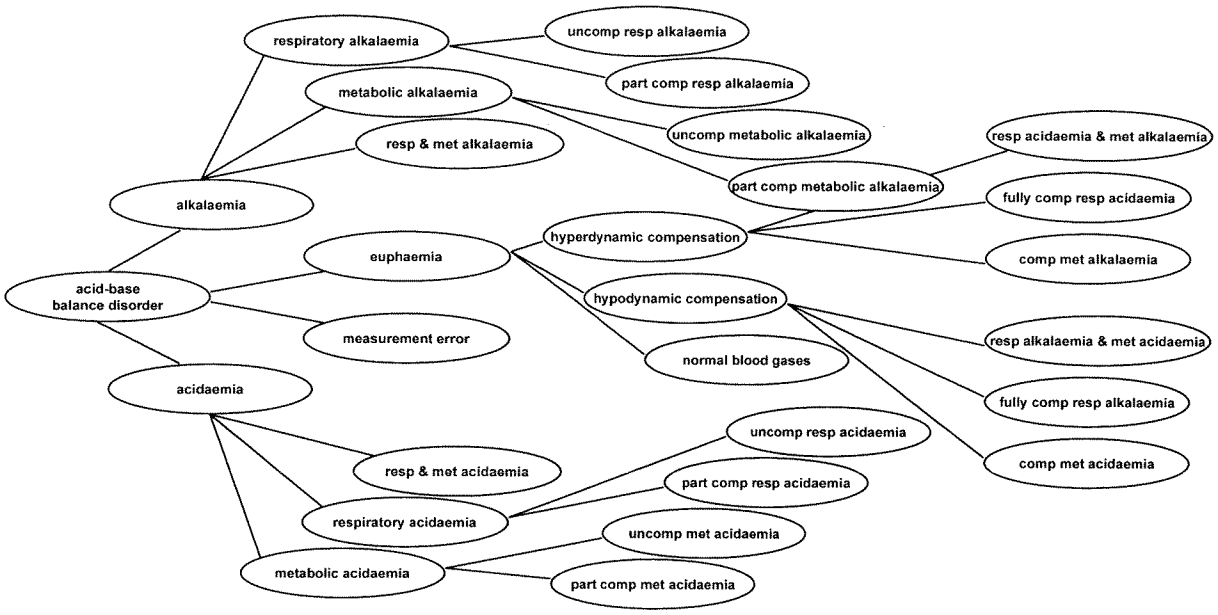

Fig. 1. The prototype knowledge-base for interpretative decision support in the domain of acid–base balance, organised in a strict hierarchy of disorder profile frames.

2.3. Methodology

The above stated three objectives underlying the development of the ICIMS system, have been pursued and accomplished via an iterative, incremental, and user-driven, object-oriented analysis, design, and implementation approach, which starts with the design of a clinical information object handling system, and proceeds upwards and closer to modelling the user's requirements for

Table 1
Acid–base balance disorders represented in the ICIMS system knowledge base with respective definitional features

Represented acid–base balance disorders	A priori belief	pCO_2	$[HCO_3^-]$	pH
Respiratory alkalosis, uncompensated	0.0833	L	N	H
Respiratory alkalosis, partially compensated	0.0416	L	L–N	H
Metabolic alkalosis, uncompensated	0.0833	N	H	H
Metabolic alkalosis, partially compensated	0.0416	H–N	H	H
Respiratory acidosis, uncompensated	0.0833	H	N	L
Respiratory acidosis, partially compensated	0.0416	H	H–N	L
Metabolic acidosis, uncompensated	0.0833	N	L	L
Metabolic acidosis, partially compensated	0.0416	L–N	L	L
Respiratory alkalosis+metabolic alkalosis	0.0833	L	H	H
Respiratory acidosis+metabolic acidosis	0.0833	H	L	L
Respiratory acidosis+metabolic alkalosis	0.0277	H	H	N
Respiratory acidosis, compensated	0.0277	H	H	N
Metabolic alkalosis, compensated	0.0277	H	H	N
Respiratory alkalosis+metabolic acidosis	0.0277	L	L	N
Respiratory alkalosis, compensated	0.0277	L	L	N
Metabolic acidosis, compensated	0.0277	L	L	N

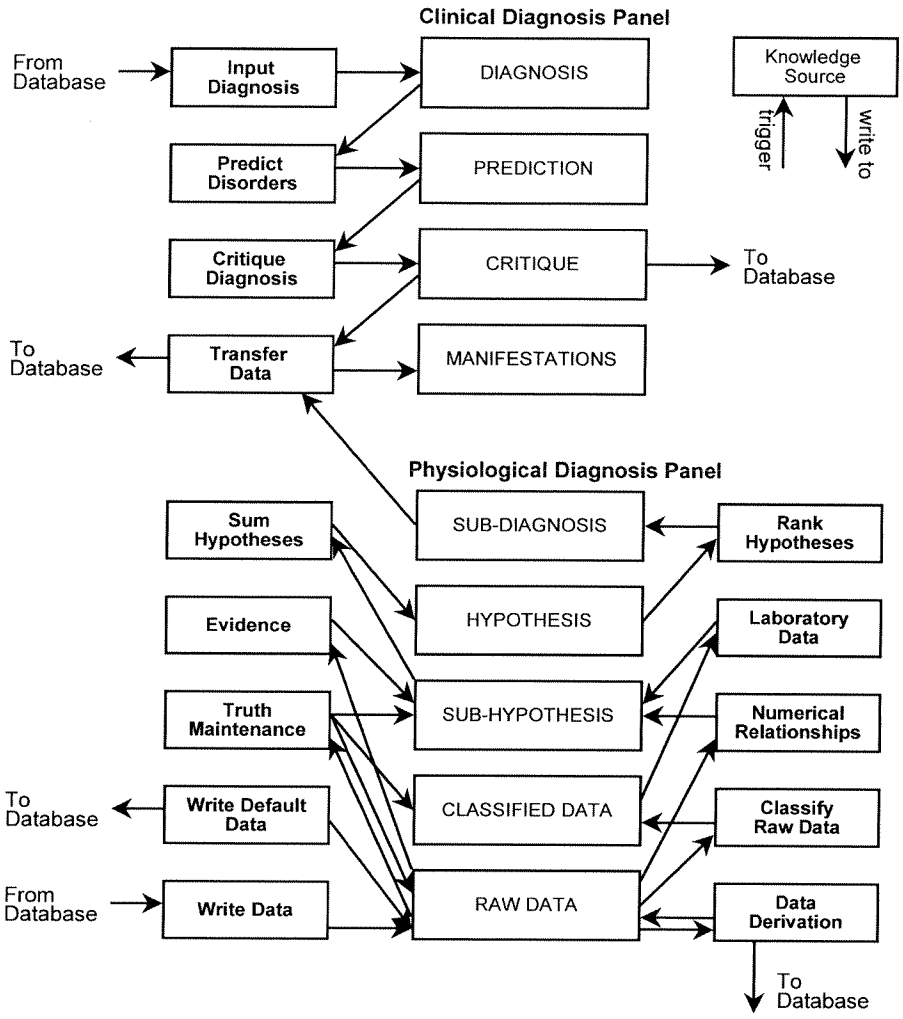

Fig. 2. Implementation of the blackboard model of problem solving in the KBS prototype (from Chelsom, 1990).

information processing support, by appending further object-oriented layers of clinical information management and decision support, to eventually include the BGAS interpretative task-domain model, and high-level interfaces for the review and update of the information acquired and generated in the process of monitoring CCU patients.

Object-oriented software engineering provides methods for the analysis, design, and implementation of IT systems, based on the principles of abstraction, encapsulation, modularity, hierarchy, typing, concurrency and persistence. What makes object-orientation particularly suited to the development of complex integrated systems is that the methodology brings these elements together into an incremental, unified decomposition, representation and implementation framework for modelling complex systems, with a structure-preserving transformation of concepts to maintainable implementations [34].

The ICIMS system has been developed within the Borland C++ environment, using the standard Windows Application Programmers Interface (API) for the construction of the users' high-level access interfaces. C++ is an object-oriented

software development environment, however, C++ does not provide mechanisms for creating and handling persistent objects, that is objects whose class instantiation inheritance, class structure inheritance, and state is saved, and transcends the lifetime of an individual program, thereby providing the ability to create and manipulate persistent world models. Introducing the concept of persistence to the object model, gives rise to object-oriented databases, and enables the development of integrated systems with superior performance in data handling both in terms of representational as well as reasoning power and efficiency. This property of the object-oriented support system has been added by means of the POET (Persistent Objects Extended Database Technology) pre-compiler, which reads class interfaces and creates persistent objects from classes and class structures, or models, which are declared persistent.

2.4. Top-level architectural view

Fig. 3 presents a top-level view of the layered, modular ICIMS system architecture. The clinical object base (COB) module shown in the middle, forms the kernel of the domain abstraction and support system integration process, by functioning as a global memory of persistent model-derived objects. There are three types of persistent model or persistent class structures in the system. The patient record model (PRM) class structure has been designed for the derivation of the persistent objects required to support the management of the clinical information generated in the process of monitoring the ICU patient. The domain knowledge model (DKM) class structure has been designed for the derivation of the persistent objects required to support the management of the knowledge-base utilised in the interpretation of the acquired clinical information. Finally, the pa-

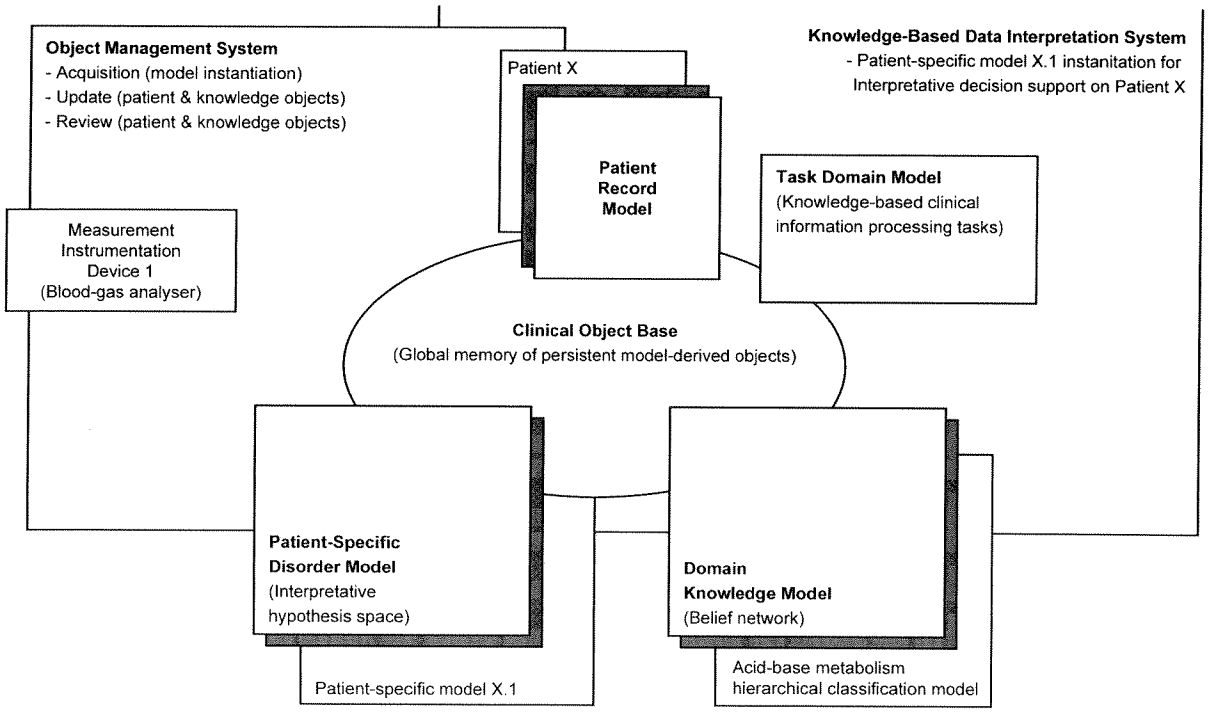

Fig. 3. Top-level view of the layered ICIMS modular architecture.

tient-specific model (PSM) class structure has been designed for the derivation of the blackboard objects required to support the application of the task-domain model (TDM) for evidence propagation in the hierarchical belief network contained in the COB.

The object management system (OMS), which includes an object communication system (OCS) module and the knowledge-based data interpretation system (DIS) module, has been designed to monitor the sources of patient data connected to the ICIMS system via serial data communication interfaces, integrate, store and organise the acquired patient data, in this case the results of the blood–gas analysis, by creating instances of the PRM class structure, interpret the data stored in the PRM model-derived object structures, and display their contents in a manner which converts it to information, thereby providing the clinical decision support required to avoid the misinterpretation and consequentially mismanagement of an observed clinical problem under conditions of information overload and contextual complexity.

Thus, for each set of evidence generated by the blood–gas analyser, the DIS module will instantiate a PSM class structure, by applying the reasoning operators that comprise the TDM of the KBS prototype (Fig. 2). So, as shown in Fig. 3, for patient [X], ICIMS will construct the patient-specific model [X.1] of the patient's interpretative hypothesis space for disorders of acid–base metabolism, given the set [E.1] of the patient's blood–gas measurements, and so on.

Finally, the OMS module may also be used in order to update and review the information contained in each patient record as well as the knowledge contained in the hierarchically connected frames that comprise the DKM, as depicted in Fig. 1. As in the case of the patient records, each knowledge frame is an instantiation of the DKM persistent class structure.

2.5. Entity models for the derivation of the COB

As stated above, the COB of the ICIMS system comprises a number of persistent objects which are derived from three models or class structures represented in the system and shown in Fig. 4. These are the PRM, DKM and PSM, which correspond to base requirements for patient data, domain knowledge and solution state acquisition, representation and manipulation, and which were used to support and facilitate the functional integration of the data interpretation KBS prototype, and the development of further clinical information management and decision support layers in the information processing task-domain of monitoring acid–base balance disorders.

As shown in Fig. 4, the persistent object structures derived from the three classes of persistent system model inherit the properties of the class of COB objects. This means that these objects are constructed and managed using the methods provided by the persistent class administration system (PCAS) of the POET pre-compiler. More specifically, amongst others, PCAS provides methods for constructing, opening and closing the COB, assigning objects to the COB, accessing the objects contained in the COB, creating, searching and manipulating sets of objects, retrieving objects from the COB, inserting and appending objects, storing, deleting, locking and watching objects, and querying the COB. These methods are inherited by each and every persistent class in the system.

2.6. Action models for the functional manipulation of the COB

Object-oriented systems are organised as co-operative collections of objects, each of which represents an instance of some class of objects, which in turn corresponds to some problem domain abstraction and which is a member of a hierarchy of classes united via inheritance relationships. Once instantiated, objects exist for some time, during which time they can act on other objects and be acted upon by other objects, thereby be changed, shared and destroyed. Thus, objects encompass two types of abstraction: entity and action abstractions [34].

In the ICIMS system, the persistent object structures contained in the COB constitute the majority of the entity abstractions required to support the integration-development process. This architectural feature confers the required ability to

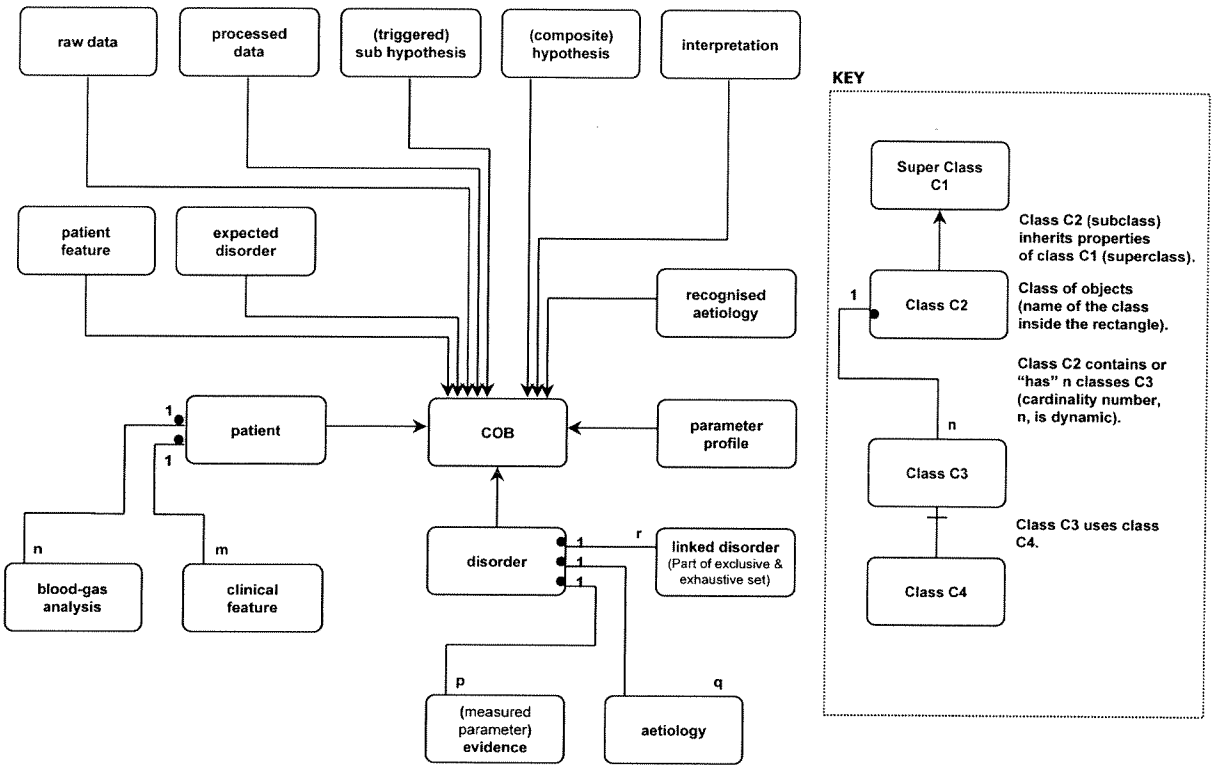

Fig. 4. Class diagram of the models used by the ICIMS system OMS module for the derivation of the persistent object structures required to support the management of the clinical information generated in the process of monitoring the ICU patient, the knowledge base utilised for the interpretation of the acquired clinical information, and the blackboard objects required to support the application of the prototype task-domain model.

develop, append, and constructively assess layers of action abstractions, which make use of the underlying COB entities in order to provide the required clinical information management and decision support functionality, in a manner which is ergonomically and cognitively compatible with the patient care activity of the user, without affecting the underlying object structures or their contents in the process.

Thus, as shown in Fig. 3, the first layer of the incremental ICIMS system development comprises the COB module, which uses the PCAS for handling the persistent object structures derived and maintained by means of the second layer, which comprises the object derivation models. Similarly, the OMS module is one level closer to the user since it provides the required functionality of patient data and domain knowledge acquisition, update and review, and a step closer to the integration of the KBS prototype, since it is at this stage that most of the clinical advisor's constructive assessment is translated into the evolutionary modifications pertaining to the KBS prototype integration process.

2.7. The object management system module

The first layer of action modelling which has been appended onto the basic object handling system is the OMS module. The OMS module, the class diagram of which is shown in Fig. 5, has been developed to generate Windows API dialogues with the user, which were designed to reflect the anatomy of the persistent object structures contained in the COB, thereby facilitating the review and update of their contents as described above.

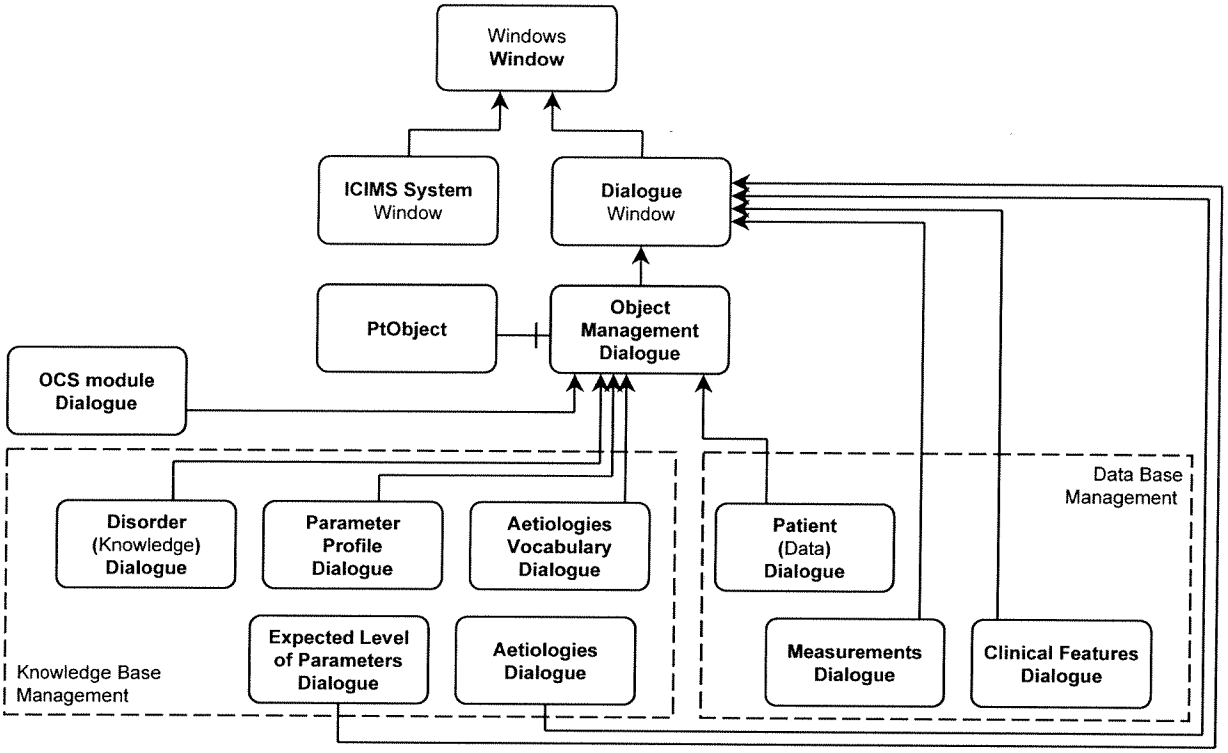

Fig. 5. Class structure diagram linking the ICIMS system COB to the OMS module, for the acquisition, review and update of the data and knowledge model derived objects.

An example of the type of dialogue generated by the system in order to review and update the contents of the patient record COB is given in Fig. 6. As shown in the figure, the information displayed in the patient window corresponds to instances of the elements of the class of patient objects (patient name, first name, hospital number etc.), instances of the elements of the class of blood–gas analysis objects, which are contained within patient objects, and instances of the elements of the class of clinical features, also contained within the class of patient objects. Furthermore, the class of patient dialogues uses the class of recognised aetiology objects shown in Fig. 4, in order to provide a vocabulary of terms for the construction of a patient description.

The actions represented in the class of dialogue windows, are performed by the buttons shown below the list boxes which display the contents of the object sets contained within patient objects, and by the buttons which are vertically arranged on the right hand side of the dialogue window. By contrast to the former, which are defined as member methods of the patient dialogue class and are therefore specific to that class of COB objects, the latter actions are inherited from the patient dialogue superclass of object management dialogues shown in Fig. 5, since the process of storing, deleting, and scanning objects is common to all classes of objects contained in the COB.

Fig. 7 shows the class structure diagram of patient dialogue objects such as the one shown in Fig. 6. As shown in Fig. 7, in order to perform the represented clinical information management and decision support function, patient dialogue objects inherit the properties and methods of persistent object management dialogues and use a number of other classes of objects. These comprise: (1) the class of persistent patient COB objects being visualised; (2) the class of the corresponding set

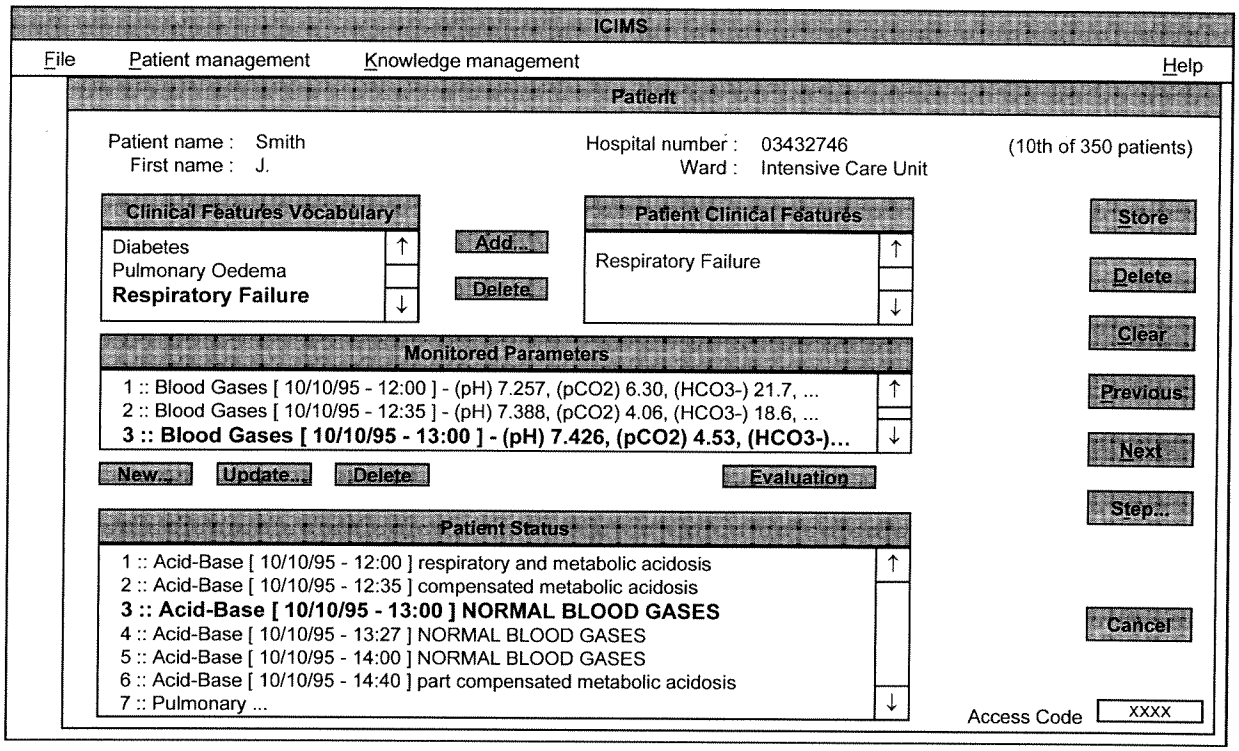

Fig. 6. Sample ICIMS system patient dialogue window.

of objects, which is required for searching the COB; (3) the class of recognised aetiology objects for editing patient descriptions; and (4) the class of recognised aetiology object-sets, which is required to search for available terms. The diagram also shows that the objects contained within patient objects, in this case blood–gas assay results, are manipulated by the methods of a separate class of dialogue windows, which is not shown here, but handled and stored automatically by the methods inherited from the class of object management dialogues, which makes use of appropriate PCAS methods. This dependency is denoted by the PCAS link shown in Fig. 7.

2.8. The data interpretation system module

As stated above, the OMS module layer has been appended onto the COB layer in order to satisfy the user's fundamental requirements for clinical information management and decision support. More specifically, the OMS module has been designed to generate ergonomic dialogues with the user, such as the one shown in Fig. 6, in order to create the persistent data and knowledge object structures contained in the COB, by deriving such object structures from the patient record and DKMs, and to thereby also support and facilitate the integration of the prototype interpretative TDM required for the contextual interpretation of the data acquired in the process of monitoring patients with disorders of acid–base balance. This corresponds to the third and final stage of the prototype KBS integration process.

The DIS module has thus been designed to be appended onto the OMS ICIMS system layer, in order to instantiate a number of interpretative dialogues with the user, which dialogues represent the cognitive information processing tasks comprising the prototype TDM [30], and to thereby derive the third class of persistent object structure, that of the PSM, which as described above has

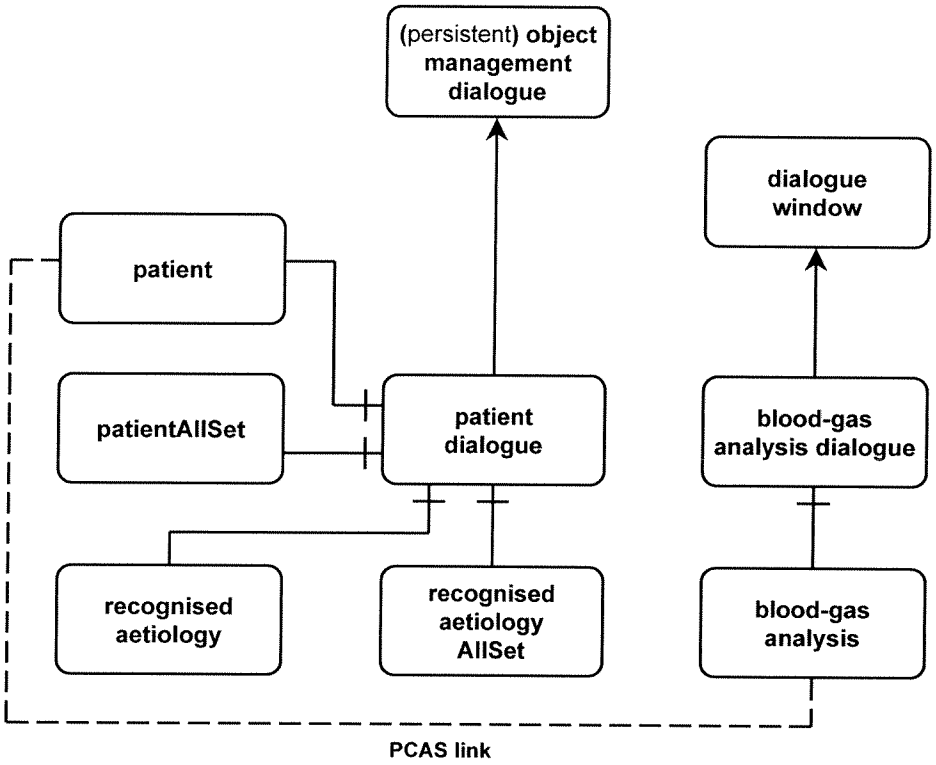

Fig. 7. Class diagram of the patient data dialogue object structure created by the OMS module.

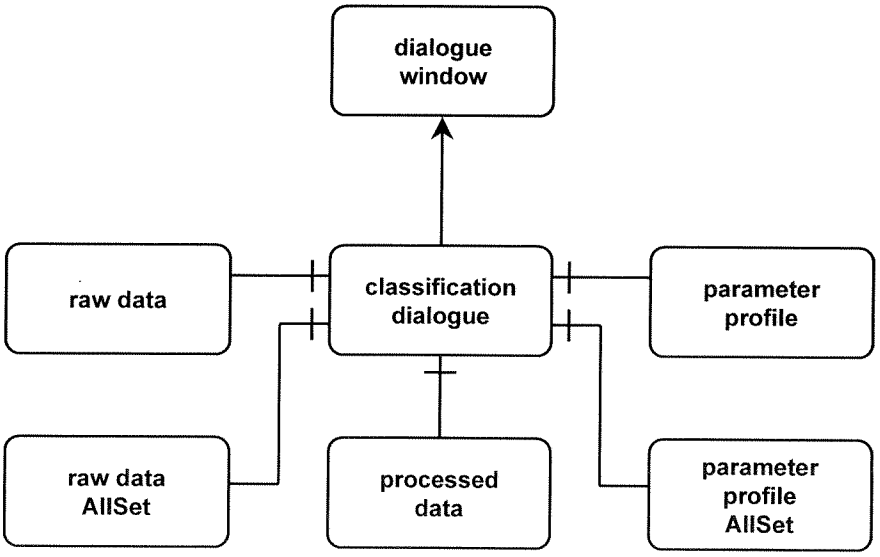

Fig. 8. Class diagram of data classification dialogues

been designed to take the place of the prototype KBS blackboard.

Fig. 8 depicts an example of the class structure of the dialogues generated by the DIS module for the case of the BGAS interpretative information processing knowledge source (KS) for the qualitative abstraction of the acquired patient measurements. The diagram shows that, following the original blackboard model depicted in Fig. 2, classification dialogue objects use: (1) raw data objects, which are constructed using the appropriate PSM model upon the creation of a DIS instance and filled-in with data from the blood–gas analysis object being interpreted; (2) parameter profile objects, in order to access the measurement classification information stored in the knowledge COB; and (3) processed data objects, which are constructed by the class of dialogues to save the state of the problem solution, following the application of the represented KS. As before, the corresponding sets of objects are used for searching the COB.

Again, following the principles of abstraction, encapsulation or information-hiding, and modularity, for the development of open, maintainable and re-usable systems, the classes of interpretative behaviour abstraction for raw data classification, and evidence impact, aggregation and propagation, were designed so that their client objects (the DIS module or other interpretative KS objects) are not required to know any of the implementation details of the represented behaviour. Thus, in the case of the example shown in Fig. 8, classification dialogue objects declare in their class interface: (1) as public (i.e. visible to client objects), only the object's constructor and destructor; and (2) as private (i.e. visible only to the objects constructed from the class), the parts of the COB accessed by the represented KS and the methods the class uses internally in order to perform the represented interpretative action.

Thus, the DIS is called to construct the required KS object using its class description, and to destroy the object following its application. In the meantime, the state of the system's problem-solving behaviour is encapsulated within the KS object. As stated above, this means that although implemented to support and facilitate the integration validation of the BGAS TDM within the ICIMS system and its environment, by progressively generating and consolidating interpretative hypotheses given the evidence available in a case, the classes of objects comprising the ICIMS TDM may be re-implemented without disturbing any parts of the system, and thus evolve into an integrated intelligent monitoring and control (IMC) TDM [35].

2.9. KBS integration validation

The ICIMS was developed within the routine CCU clinical information processing environment of the West Middlesex University Hospital and constructively assessed during the development process with regard to its information management contribution and its usability. The system was subsequently installed at the Mayday University Hospital intensive care unit where it automatically collected data from a blood–gas analyser via a serial connection to the measurement instrument over a period of 6 months. During this period the interpretations generated by means of applying the TDM to the online data were compared with those generated by the KBS prototype. The validation performed during this period confirmed that the integration process was successful. In the process, ICIMS functioned both as a valuable tool in the management of the vast volume of clinical information generated by the blood–gas analyser, as well as a platform for demonstrating the usefulness of reading descriptive and context-specific interpretations of numeric acid–base status measurements.

Fig. 9 shows an example combined parameter- and state-based trend display window, designed to augment the effective information yield generated by the patient record dialogues, and to provide the means for a preliminary assessment of the user's requirements for the integration of the BGAS TDM within the real-time information processing activity supported by the ICIMS system. The particular data set, which was selected from a total of approximately 1,800 blood–gas measurements taken from a total of 10 patients, corresponds to a ventilated patient with renal dysfunction during the period 12/2/96–7/3/96.

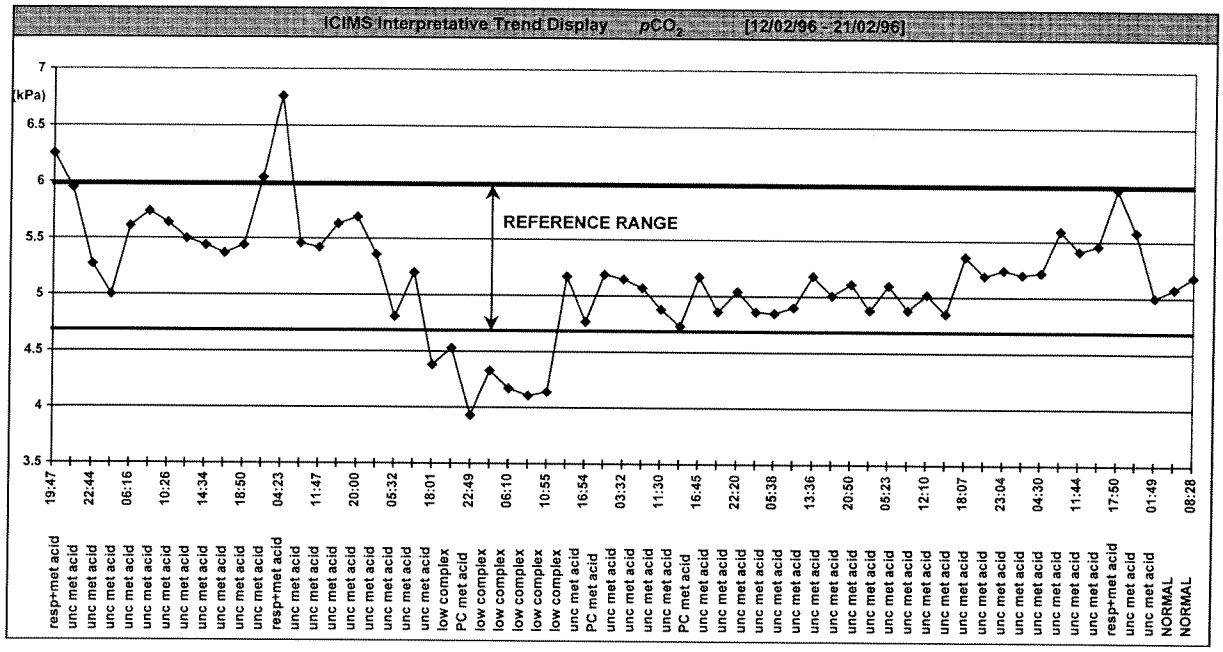

Fig. 9. Sample high-level interpretative trend display patient summary window.

In combination, the two trend displays were designed to provide further valuable clinical information management and decision support, in that they can be used to detect changes in patient state, and to distinguish those which merit patient-state control attention from insignificant or erroneous indications, due to measurement, transcription or execution errors. For example, the measurement taken at 04:23 on the 14th of February 1996 for the patient with acute renal failure, may either be the result of a mechanical ventilation execution error, a measurement or transcription error, or a secondary acid–base balance disorder. However, the measurement taken at 17:50 on the 20th of February appears to be a misclassification error which, being part of an on-going interpretative state–space trajectory, and in conjunction with the superimposed parameter trend, does not affect the validity of the decision support generated by the system in as much as it does in the case of consultation systems, such as the BGAS prototype, which are required to give the 'correct' answer to an interpretative problem.

3. Conclusions

Dealing with the increasing cost of health care delivery and the detrimental consequences on the quality and effectiveness of clinical outcome requires supportive and corrective action to be taken much closer to the source of the problem rather than in pursuit of compromising costs and the symptoms of an ailing financial control philosophy. Computer-based medical data processing has yielded methods and tools for managing the task away from the hospital management level and closer to the desired disease and patient management level. During the past decade a number of relevant methodologies have emerged, as dictated by key qualitative cost–benefit control protocols, and are rapidly disseminated into routine clinical practice. However, the underlying information and knowledge management technologies have failed to be integrated into routine clinical practice. This paper has described the development of a system which has successfully overcome the main obstacle encountered in the process of developing the

required technological infrastructure: a system that treats both information and knowledge as clinical objects with same modelling requirements and which, at the most fundamental level of clinical decision support, facilitates both the organised acquisition of clinical information and knowledge and provides a test-bed for the development and evaluation of clinically useful and effective knowledge-based decision support functions. Future development efforts must be directed towards the development and integration into the online intelligent clinical information management support system of the higher-level decision support functions embodied in the practice of evidence-based medicine.

References

[1] D.L. Sackett, W.M. Rosenberg, J.A. Gray, R.B. Haynes, W.S. Richardson, Evidence based medicinewhat it is and what it isn't, Br. Med. J. 13 (312) (1996) 71–72.

[2] A. Cowling, K. Newman, S. Leigh, Developing a competency framework to support training in evidence-based healthcare, Int. J. Health Care Qual. Assur. 12 (4) (1999) 149–159.

[3] J. Øvretveit, Evaluation informed management and clinical governance, Br. J. Clin. Governance 4 (3) (1999) 103–108.

[4] H. Sanderson, Information requirements for clinical governance, Br. J. Clin. Governance 5 (1) (2000) 52–57.

[5] G. Fairfield, A.F. Long, Measuring the outcomes of disease management, Int. J. Health Care Qual. Assur. 10 (4) (1997) 161–165.

[6] U. Svensk, Disease management at the system level—an effective way to improve health care, Int. J. Health Care Qual. Assur. 9 (7) (1996) 4–8.

[7] L. Hall, M. Eccles, Case study of an inter-professional and inter-organisational programme to adapt, implement and evaluate clinical guidelines in secondary care, Br. J. Clin. Governance 5 (2) (2000) 72–82.

[8] A. Millard, Evidence-based clinical guidelines–implementation plans in Scotland, Br. J. Clin. Governance 4 (3) (1999) 98–102.

[9] B.W. Ellis, S. Johnson, A clinical view of pathways of care in disease management, Int. J. Health Care Qual. Assur. 10 (2) (1997) 61–66.

[10] W.B. Schwartz, Medicine and the computer. The promise and problems of change, N. Eng. J. Med. 283 (1970) 1257–1264.

[11] E.H. Shortliffe, The science of biomedical computing, Med. Inform. 9 (3/4) (1984) 185–193.

[12] D.G. Cramp, E.R. Carson, The patient/clinician relationship, computing and the wider health care system, in: E.R. Carson, D.G. Cramp (Eds.), Computers and Control in Clinical Medicine, Plenum, New York, 1985, pp. 245–255.

[13] R.A. Greenes, E.H. Shortliffe, Medical Informatics: an emerging academic discipline and institutional priority, J. Am. Med. Assoc. 263 (8) (1990) 1114–1120.

[14] E.H. Shortliffe, The adolescence of AI in medicine: will the field come of age in the '90s?, Artif. Intell. Med. 5 (1993) 93–106.

[15] E.W. Coiera, Artificial intelligence in medicine: the challenges ahead, J. Am. Med. Assoc. 3 (6) (1996) 363–366.

[16] C. Ohmann, F.T. de Dombal, O. Winding, The International Evaluation Panel. Evaluation in the TELEGAS-TRO-Project, in: E.M.S.J. van Gennip, J.L. Talmon (Eds.), Assessment and Evaluation of Information Technologies in Medicine, IOS Press, Amsterdam, 1995, pp. 167–177.

[17] R.A. Miller, Medical diagnostic decision support systems—past, present, and future: a threaded bibliography and brief commentary, J. Am. Med. Inf. Assoc. 1 (1) (1994) 8–27.

[18] D.L. Hunt, R.B. Haynes, S.E. Hanna, K. Smith, Effects of computer-based clinical decision support systems on physician performance and patient outcomes: a systematic review, J. Am. Med. Assoc. 280 (15) (1998) 1339–1346.

[19] D.C. Classen, Clinical decision support systems to improve clinical practice and quality of care, J. Am. Med. Assoc. 280 (15) (1998) 1360–1361.

[20] D. Kalogeropoulos, An intelligent clinical information management support system for the critical care medical environment, Ph.D. Thesis, City University, Centre for Measurement and Information in Medicine, London, 1999.

[21] C.F. Aliferis, R.A. Miller, On the heuristic nature of medical decision-support systems, Meth. Inform. Med. 34 (1–2) (1995) 5–14.

[22] J.L. Talmon, R. Smeets, Case acquisition for knowledge-based decision support system validation, in: E.M.S.J. van Gennip, J.L. Talmon (Eds.), Assessment and Evaluation of Information Technologies in Medicine, IOS Press, Amsterdam, 1995, pp. 67–74.

[23] J. Brender, J. Talmon, P. Nykänen, P. McNair, M. Demeester, R. Beuscart, On the evaluation of system integration, in: E.M.S.J. van Gennip, J.L. Talmon (Eds.), Assessment and Evaluation of Information Technologies in Medicine, IOS Press, Amsterdam, 1995, pp. 189–208.

[24] T. Groth, P.O. Collinson, Strategies for decision support for fluid and electrolyte therapy in the intensive care unit-approaches and problems, Int. J. Clin. Monit. Comput. 10 (1993) 3–15.

[25] R.A. Bankowitz, M.A. McNeil, S.M. Challinor, R.C. Parker, W.N. Kapoor, R.A. Miller, A computer-assisted medical diagnostic consultation service. Implementation and prospective evaluation of a prototype, Ann. Intern. Med. 110 (10) (1989) 824–832.

[26] R.A. Bankowitz, M.A. McNeil, S.M. Challinor, R.A. Miller, Effect of a computer-assisted general medicine

diagnostic consultation service on housestaff diagnostic strategy, Meth. Inform. Med. 28 (4) (1994) 352–356.
[27] R.C. Parker, R.A. Miller, Creation of realistic appearing simulated patient cases using the INTERNIST-1/QMR knowledge base and interrelationship properties of manifestations, Meth. Inform. Med. 28 (4) (1989) 346–351.
[28] E.S. Berner, R.S. Maisiak, C.G. Cobbs, O.D. Taunton, Effects of a decision support system on physicians' diagnostic performance, J. Am. Med. Inform. Assoc. 6 (5) (1999) 420–427.
[29] C.P. Friedman, A.S. Elstein, F.M. Wolf, G.C. Murphy, T.M. Franz, P.S. Heckerling, P.L. Fine, T.M. Miller, V. Abraham, Enhancement of clinicians' diagnostic reasoning by computer-based consultation: a multisite study of 2 systems, J. Am. Med. Assoc. 282 (9) (1999) 1851–1856.
[30] J.J.L. Chelsom, The interpretation of data in intensive care medicine: an application of knowledge-based techniques, Ph.D. Thesis: City University, Department of Systems Science, London, 1990.
[31] E.R. Carson, J.J.L. Chelsom, R. Summers, Progress with measurement, information and decision-making in critical care medicine, Measurement 9 (3) (1991) 104–110.
[32] J. Pearl, On evidential reasoning in a hierarchy of hypotheses, Artif. Intell. 28 (1986) 9–15.
[33] W.J. Clancey, Model construction operators, Artif. Intell. 53 (1992) 1–115.
[34] G. Booch, Object oriented design, Benjamin/Cummings, Redwood, CA, 1991.
[35] S. Uckun, Intelligent systems in patient monitoring and therapy management: a survey of research projects, Int. J. Clin. Monit. Comput. 11 (1994) 241–253.