

Supporting decisions in medical applications: the knowledge management perspective

Stefania Montani ^{a,*}, Riccardo Bellazzi ^b

^a *Dipartimento di Informatica, Università del Piemonte Orientale 'A. Avogadro', Spalto Marengo 33, I-15100 Alessandria, Italy*

^b *Dipartimento di Informatica e Sistemistica, Università di Pavia, Pavia, Italy*

Abstract

In the medical domain, different knowledge types are typically available. Operative knowledge, collected during every day practice, and reporting expert's skills, is stored in the hospital information system (HIS). On the other hand, well-assessed, formalised medical knowledge is reported in textbooks and clinical guidelines. We claim that all this heterogeneous information should be secured and distributed, and made available to physicians in the right form, at the right time, in order to support decision making: in our view, therefore, a decision support system cannot be conceived as an independent tool, able to substitute the human expert on demand, but should be integrated with the knowledge management (KM) task. From the methodological viewpoint, case based reasoning (CBR) has proved to be a very well suited reasoning paradigm for managing knowledge of the operative type. On the other hand, rule based reasoning (RBR) is historically one of the most successful approaches to deal with formalised knowledge. To take advantage of all the available knowledge types, we propose a multi modal reasoning (MMR) methodology, that integrates CBR and RBR, for supporting context detection, information retrieval and decision support. Our methodology has been successfully tested on an application in the field of diabetic patients management.

© 2002 Elsevier Science Ireland Ltd. All rights reserved.

Keywords: Knowledge management; Case based reasoning; Multi modal reasoning

1. Introduction

Since 30 years ago, there has been a strong emphasis on the definition of the so-called

'expert systems' (ES), meant to help physicians solve particularly critical decision problems, and implemented in different medical domains. Several evaluation studies have proved the formal correctness of the suggestions obtained by applying these tools [1,2]; nevertheless, their on field exploitation has been relatively poor. Apart from usability,

* Corresponding author. Tel.: +39-0131-287-437; fax: +39-0131-287-440

E-mail address: stefania@mfn.unipmn.it (S. Montani).

accessibility and organisational factors, the main weakness of ES seems to be their attempt to apply medical knowledge out of its natural context [2]. As a matter of fact, ES were stand-alone computerised programs, meant to provide decision support in a certain domain, with automatic reasoning capabilities. They tried to take the role of the human expert, resorting to her/his knowledge, formalised in the system knowledge base. Moreover, they were usually not integrated with the hospital information system (HIS) (except the innovative system HELP [3]), and were unable to consider the complex interactions taking place among people within the medical environment. On the other hand, the role of different human agents, including patients themselves, has to be carefully considered, especially in the field of chronic diseases care.

Today the trend in decision support systems definition is the one of integrating the existing HIS with different methodologies and technologies, able to co-operate in a transparent way with respect to the user, and to promote information exchange among the people involved in the disease management process. Health care providers (and sometimes patients) can rely on a service (rather than on an isolated tool), able to assist them by supplying the relevant information at the right time, in the right form, but leaving them the responsibility of a contextual interpretation and evaluation of the information itself [4], and of course, of the final decision.

This goal is achieved by integrating the decision support functionality with the knowledge management (KM) task. Instead of being stand-alone solutions to afford particular problems, decision support systems can today be considered as building blocks of the HIS within a larger KM perspective. Information technology (IT) allows, therefore, to face simple, but large scale problems, by taking into consideration all the clinical practice

situations, thus overturning the former ES perspective, that was the one of addressing only few, particularly complex cases, on the physician's request.

The aim of this paper is the one of showing how KM can be implemented in practice. As a matter of fact, several theoretical approaches to the KM task have been described in the literature, but they have rarely been coupled to a computational model. Providing an operative solution, able to fulfil the needs of KM by implementing a real integration of different knowledge sources, has been our aim while developing this work. In detail, in [Section 2](#) we motivate the reasons for shifting the attention towards the KM view, by presenting the features of KM and of medical reasoning. Then, in [Section 3](#) we describe some methodological approaches to KM and decision support, and we introduce a Multi Modal Reasoning (MMR) paradigm, able to integrate two (or more) methodologies into a unique framework. In [Section 4](#) we present our experience in the definition of a MMR system, applied in the diabetes management domain; additional information about this system can be found elsewhere (see e.g. [5]). Finally [Section 5](#) summarises our conclusive observations.

2. Managing medical knowledge

2.1. Knowledge management features

Within a (medical) organisation, two knowledge types can be generally found: (i) explicit knowledge, corresponding to the already well established and formalised domain knowledge: it can be present in books and in written documentation, or can be represented by some formalisms for developing ES's knowledge bases, such as taxonomies or rules; (ii) implicit or tacit knowledge, consisting in

individual expertise, organisational practices, and past cases: it can be represented by heuristics that even experts are not fully aware of, or can be stored in data-bases and case libraries.

Such heterogeneous knowledge may be available only at certain times (especially if it's a tacit knowledge in the head of an expert) or in certain locations (in the library, in the expert's office, in a particular department). The knowledge quality may vary as well, being it incomplete, or for some aspects, out of date.

We could define KM as a discipline with the objectives of promoting knowledge growth, knowledge communication and knowledge preservation within an organisation [6], in order to achieve a better exploitation of this essential resource. Given the knowledge features listed above, KM aims, therefore, at [7]:

- 1) keeping the knowledge up to date and correct;
- 2) providing knowledge in the right location;
- 3) applying knowledge of the most suitable type;
- 4) providing knowledge at the time at which it is needed.

In a word, KM allows for securing and distributing knowledge, in order to assure and optimise its availability.

Various motivations for implementing a KM strategy within an organisation can be outlined:

- 1) to keep track of an expert's know-how, even if s/he moves or retires;
- 2) to enable an intelligent consultation of the available knowledge repositories, in order to exploit the experience acquired in the past, in terms of successes and mistakes;
- 3) to keep track of past cases and histories so that, even in presence of changes in the

staff, the quality of work will not decrease due to a lack of information;

- 4) to improve knowledge communication among personnel;
- 5) to improve the learning of inexperienced and personnel under training.

Obviously, such motivations are particularly relevant in the medical domain, where it is essential to keep the quality of care always at the highest level, and it is not admissible to miss this central goal due to a lack of information [8]. A proper KM strategy allows for a continuous quality improvement based on performance measurement and process data analysis [9]. A careful analysis of the data that are collected during the provision of clinical care may be a very effective way to identify problems in routine work and to provide directions on how to reorganise medical services. Moreover, it may motivate further investigations on the associations between clinically relevant variables that might have been neglected or unobserved before [10]. Finally, there is evidence that the combination of information from Electronic Patient Records together with patient registries and transactional data represent an invaluable source for planning organisational interventions [11] and, in synthesis, for supporting performance-based medicine.

2.2. Knowledge management in the medical domain: the nature of medical knowledge

In health care centres large amounts of information are generally present. In particular the introduction of HIS into clinical practice has led to the memorisation of a huge quantity of data, extracted from day by day activity, and reporting the unarticulated experience of individual workers. The HIS clearly represents a source of implicit knowledge; on the other hand, the increasing

distribution of evidence-based prescriptions, even in electronic format, makes the general domain knowledge more and more available for exploitation.

The need for KM and integration becomes very clear if we consider the nature of medical decision making. As a matter of fact, typically physicians use to reason by recalling past situations similar to the current one. The process is often biased by the tendency of recalling only more recent cases. A methodology enabling the retrieval of all the operative knowledge available, also related to more older examples, would thus be of help. Implicit KM comes out to be very useful also for sharing other clinicians' experiences: this is a key objective of the organisation, even though, in contrast to what happens in industrial settings, the individual ability of making inferences and of coping with the problem remains a central requisite to be respected. In particular, each physician has a personal way of organising the reasoning process, and each one can be particularly suitable in some situations. Explicit KM is very important as well: hypothetico-deductive reasoning is usually implemented by medical doctors. It consists in formulating a hypothesis about the correct diagnosis or therapy, and then in proving (or confuting it) on the basis of additional observations. The process goes on until the hypothesis cannot be confuted. In generating the hypothesis, the clinician does not rely only on data, but also on prior domain knowledge, which enables her/him making a pre-selection of the possible diagnostic explanations or of the therapeutic advice.

In a word, a physician has to be able to adapt her/his choices and her/his reasoning process to the current needs, using multiple strategies: an efficient KM approach is a valuable way for helping her/him exploiting all the available knowledge sources.

3. Methodologies for managing knowledge

One of the most suitable methodologies for managing knowledge of the implicit type [12], that has been successfully tested also in the medical field, is Case Based Reasoning (CBR) [13]. CBR is a reasoning paradigm able to exploit the information embedded into already solved instances of problems [13,14], called cases. Problem solving experience is explicitly taken into account by storing past cases in a library, and by suitably retrieving them when a new problem has to be tackled. A case consists of three basic information: the problem description, i.e. a collection of feature-value pairs able to summarise the problem at hand; the case solution, describing the solution adopted for solving the corresponding problem; the case outcome, reporting the result (positive or negative) obtained by applying the solution itself.

Case based problem solving can be summarised by the following four basic steps, known as the CBR cycle or as the four 'res' [14]:

- 1) retrieve the most similar case(s) from the case library;
- 2) reuse them, and more properly their solutions, to solve the new problem;
- 3) revise the proposed solution;
- 4) retain the current case in the library for future problem solving.

The retrieval and reuse of past data, and the possibility of retaining new information, closely resemble the KM objectives of keeping, increasing and reusing knowledge. The CBR paradigm hence comes out to be a really natural way of performing KM, when dealing with non-formalised information.

In the CBR cycle some steps may be missing or may be collapsed. Particularly in medical decision making applications it is very com-

mon to find examples of Case Based retrieval tools, able to extract relevant knowledge, but that leave to the user the responsibility of providing an interpretation of the current case and of proposing a solution.

In detail, Case Based retrieval involves three sub-steps [15]:

- 1) situation assessment, to elaborate case description and to clarify the relevant context to work with;
- 2) case memory search, to retrieve partially matching cases;
- 3) best cases selection.

Step 1 represents the identification of the contextual framework in which the input case has to be interpreted, and on which retrieval has to rely: the case search may be, therefore, restricted to relevant sub-parts of the library. Steps 2 and 3 are often considered as a whole: if a ranking procedure is implemented, cases are automatically ordered, and the best ones are identified during case memory search.

In conclusion, CBR is in some sense able to mimic human experts' analogical reasoning, by remembering solutions to similar problems adopted in the past, and by adapting them to the current situation. Applying CBR can also have a computational justification, especially when the domain knowledge is poor and difficult to explicit. The bottleneck of knowledge acquisition and of knowledge representation is reduced, as new implicit knowledge is automatically stored in the case base during the every day working process: no additional effort is required to the user.

CBR has been applied in several fields, mainly dealing with diagnostic problem solving [16] or planning [17]. Several tools for supporting medical diagnosis have been defined in different domains: from psychiatry [18], to audiology [19]. Pulmonary diseases [20], eating disorders [18], dysmorphic syn-

dromes [21] and acute abdominal pain [22] have been addressed as well. CBR has also been exploited for therapy planning: in diagnostic imaging procedure selection [23], in radiation therapy [24] and in antibiotics selection [25].

On the other hand, several methodologies for representing and managing explicit knowledge have been proposed in the Computer Science community, from taxonomies, to protocols, to mathematical models and production rules. In particular, when the problem can be described through a mathematical model, it is possible to rely on it for obtaining the optimal solution. Nevertheless, building a model requires that quantitatively (and qualitatively) sufficient data are available, for a correct parameter estimation, and this is not always the case in real contexts (e.g. in medical domains). On the other hand, also Rule Based Reasoning (RBR) is a successful methodology for managing knowledge of the explicit type [26], largely relied upon, since the first ES were developed (see e.g. [27]). In comparison to other formalisms, rules provide several advantages, i.e. the knowledge base modularity, the capability of representing also heuristic knowledge, and the integration with an inferential system that can be automatically applied to the information formalised by the domain experts and knowledge engineers. In detail, when building a RBR system, experts' knowledge has to be structured in a base of production rules, each one containing a piece of information. RBR consists in firing a set of rules through a chaining mechanism, exploiting the available data, to arrive at a decision.

A RBR system is, therefore, composed by three elements:

- 1) the set of rules (also called the knowledge base, or rule base);

- 2) a software program able to interpret and fire the rules, called inference engine;
- 3) a repository of information containing the available data during the overall reasoning process: this memory is thus updated every time a rule is applied.

The inference engine executes a cycle of three fundamental operations:

- 1) search of what rules may be fired on the basis of the information repository content;
- 2) selection of the most suitable rule if more than one could be fired (conflict resolution);
- 3) activation of the selected rule and update on the information repository.

Inference engines may implement two chaining mechanisms: forward chaining, that simulates data-driven reasoning, typically adopted by physicians in therapy planning, and backward chaining, that is able to simulate goal-directed reasoning, and is particularly well suited for the diagnostic task.

To properly integrate all the information available within a medical organisation, in order to mimic the real process of medical reasoning, the recent advances in IT have led to the design of a new generation of KM and decision support systems, able to integrate a set of knowledge sources, relying on different knowledge representation formalisms, and on a MMR strategy. MMR paradigms are meant to overcome the limitations of exploiting the single methodologies (e.g. CBR and RBR described above), and to better cope with the target domain problems.

Due to the success they have obtained in several applications, CBR and RBR have proved to be good candidates for realising a MMR approach. Several examples of integration between CBR and RBR have been proposed in the literature. In the majority of

them, RBR and CBR are applied in mutually exclusive ways: first RBR, meant to deal with standard problems, is applied; if it fails in proposing a suitable solution, CBR is exploited, in order to retrieve similar cases from a library of peculiar and non-standard situations [28]. Other authors [29,30] rely on CBR for contextualising rules, while RBR permits the extraction of more abstract and general concepts from cases. The methodology to be applied can be selected in a dynamic way, depending on the situation at hand: in particular, the rule base and the case memory can be searched in parallel for applicable entities; then the best entity (i.e. rule or case) to reuse (and, therefore, the reasoning paradigm to apply) can be selected, depending on its suitability for solving the current problem [30]. In other approaches, RBR is resorted to in the CBR solution revision step: if the memory does not contain suitable examples of adaptation of the retrieved solutions to situations similar to the current one, the system will employ some general adaptation rules [31].

The common basis of all the above MMR implementations (except perhaps the work in [29]) is that the various paradigms are used in a quite exclusive way. These approaches do not guarantee to overcome the intrinsic limitations of the two reasoning methodologies, which can be summarised as follows:

- 1) classical RBR systems don't have the capability of specialising the explicit knowledge embedded in the rules, by resorting to contextual knowledge (e.g. the single patient's features). To deal with as many peculiar situations as possible, it would be necessary to define a huge, and thus intractable, rule base (this is known as the qualification problem [32]);
- 2) CBR just relies on the implicit knowledge stored in the case library. If the library

presents competence gaps, i.e. it is too small, or biased by too specific examples, a misleading indication on how to solve the current problem may be provided.

Realising a tight integration between RBR and CBR to cope with these problems has been our aim while developing a MMR system, meant to perform KM in the context of type 1 diabetic patients management, and to provide decision support to physicians when dealing with therapy revision. The MMR system is, therefore, able to integrate population-based medical knowledge (knowledge of the explicit type) with case based information, that represents an individualisation of patients' problems and of responses to therapeutic prescriptions.

4. MMR for diabetes care

4.1. Type 1 diabetes mellitus

Diabetes mellitus is one of the major chronic diseases in the industrialised countries. In particular, patients affected by type 1 diabetes need to undergo intensive insulin therapy (IIT), consisting in three to four injections of exogenous insulin every day, in order to regulate blood glucose, and to reduce the risk of later life complications. Patients have to keep a hand-written diary, on which they record blood glucose level measurements taken before every injection, together with additional information about diet and life style. IIT also implies that patients are frequently visited (typically every 2–4 months); during visits, diary data are also reviewed, to assess the metabolic behaviour, and, if necessary, to change the insulin administration therapy. The major long-term intervention trial on type 1 diabetic patients, the DCCT [33], has clearly shown that the definition and

realisation of an appropriate individual therapeutic goal, customised on the single patient's needs, rather than the implementation of a specific therapy, is the key to an effective diabetes care.

4.2. The MMR system

Our system relies on CBR for managing knowledge of the implicit type, and embeds explicit knowledge in a set of production rules. Decision support is provided by exploiting Case Based retrieval results for specialising the RBR system behaviour, and tailoring it to the individual patient's needs.

In particular, we have mapped the concept of case to the one of periodical control visit. A new case is automatically stored in the case library every time a patient is visited. By cooperating with the pediatricians of Policlinico S. Matteo Hospital in Pavia, we were able to structure the case library resorting to a taxonomy of mutually exclusive prototypical classes, that express typical problems that may occur to type 1 diabetic patients in the age of infancy and puberty. Moreover, each case in the library is linked to the previous and to the following ones, in terms of time, by two chains of pointers [34].

Thanks to this case organisation, the CBR tool can be seen as an autonomous mean for managing implicit knowledge. Given a case, the user can easily retrieve the whole patient's history, and visualise the transitions from one class to another that took place, together with the therapeutic choices that made possible the transitions themselves. Moreover, before retrieving past cases similar to the current one, the system classifies the current case as belonging to one class. Classification implements context detection: the patient is categorised as being experiencing a particular clinical course condition or associated disease, and her/his data can be better interpreted, in

the light of the specific situation s/he is currently living. Finally, classification focuses the attention only on the relevant parts of the case library, by limiting retrieval to the most probable class, or to a subset of all the classes in the taxonomy; in this way retrieval time can be significantly reduced.

Moreover, as anticipated, in our system CBR is tightly integrated with RBR, to support therapy revision.

In the field of type 1 diabetes care, the therapy revision process is typically articulated in three consecutive tasks; in our implementation, the completion of the process is scheduled by the RBR system, within which each task is mapped onto a specific set of rules, fired through a forward chaining mechanism. In detail, the reasoning paradigm proceeds as follows (see Fig. 1):

4.2.1. Data analysis and problem identification

From the available monitoring data, the system extracts some indicators of the patient's metabolic condition, able to summarise the patient's response to the therapy s/he is following (see [34] for details on the data analysis techniques). The results trigger the identification of hyperglycaemia or hypoglycaemia problems in the different periods of the day. The RBR system can complete this

task independently, or its behaviour can be specialised resorting to the integration with Case Based retrieval. In particular, only classification results are exploited in this phase, to tune specific rule parameters, thus tailoring the identification of metabolic alterations to the peculiar patient's situation (see [34] for details).

4.2.2. Suggestion generation and selection

For each detected problem, a set of suggestions on how to modify the current insulin therapy are proposed; the most effective ones are selected resorting to the concept of insulin competence. The most competent insulin, that has the stronger effect on the moment of the day in which the problem has been found, is identified. Competence is evaluated relying on the pharmacokinetics of the different insulin types [35].

4.2.3. Therapy revision

The RBR system proposes an adjustment to the current insulin therapy, in accordance with the selected suggestions. It is meant to be general enough to be safely applicable in a variety of different situations, therefore, it typically proposes small variations to the current protocol insulin doses, quantitatively speaking. Even though the RBR behaviour was judged correct and quite satisfactory in a formal evaluation study [36], it came out to be sometimes not sharp enough to promptly face the patient's alterations. To overcome this weakness, we again resort to the integration with CBR. The CBR retrieval step is performed, restricted to the most probable class(es) identified during the problem identification phase. The physician is allowed to choose whether to retrieve only cases belonging to the most probable class, or to a set of very probable classes. In both situations, cases are retrieved by resorting to metrics able to cope with the problem of missing data, and to

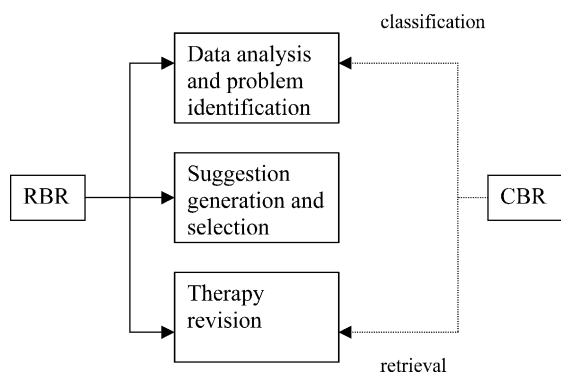


Fig. 1. Implementation of the integration between CBR and RBR in the automatic reasoning process for therapy revision.

treat both symbolic and numeric variables [34]. Some simple statistics are calculated on the retrieved cases, to set the insulin adjustments width that will then be applied to the current protocol.

The clinical correctness of the suggestions proposed by the MMR system have been tested through a formal evaluation procedure [5]. In particular, two diabetologists were asked to perform a fully-crossed, blind review of the therapies proposed by RBR, by MMR, and by two colleagues, to 30 real patient cases. About 98/118 therapies were judged as acceptable for the physicians, and 101/120 for the systems. The MMR and the RBR tools hence proved to be able to work at the physician level; on the other hand, no benefit could be found by exploiting MMR in comparison to RBR. The reason for such an unsatisfactory result was due to a case library affected by competence gaps. To cope with the problem of misleading retrieval information, we plan to implement a control strategy, that will enable the exploitation of retrieval results only if a sufficiently large number of similar cases have been retrieved, and if they are similar enough to the input case to justify their use. In this way, MMR will support decision making at various levels of complexity, as the case library grows: first, when the stored information is poor, RBR will come out to be the most competent methodology, and it will be applied without relying on CBR results. As far as new information is stored in the case library, retrieval results will become more reliable, and will be exploited by the MMR methodology. Note that the memorisation of new cases solved according to the system prescription is not made in an automatic fashion. As a matter of fact, the physician is required to analyse the system advice, and then s/he can decide if the new example should be added to the case base. In this way, the problem of circular inference,

which could lead to an overestimated self-confidence in the behaviour of the system, due to a large presence of cases solved relying to it, is avoided, since an independent expert's evaluation is always applied.

The tool will be integrated within the distributed architecture of a telemedicine service, whose implementation will be carried out in the context of the EU-funded project M²DM (IST-1999-10315).

4.3. MMR for knowledge creation

It has been postulated [37] that human knowledge creation and upgrade takes place through a process of continuous interaction between implicit and explicit knowledge, acting as mutually complementary entities, and converted one into the other. In this knowledge conversion model (Fig. 2), four different activities are implemented:

- 1) socialisation: it consists in implicit knowledge sharing and expansion, by putting together personal experience and by learning from experts;
- 2) externalisation: it is the process of formalisation of implicit knowledge into the explicit form;
- 3) combination: it consists in recombining and restructuring explicit knowledge, giving birth to new explicit knowledge;
- 4) internalisation: it is the process of acquiring explicit knowledge as a form of implicit one, by applying the explicit information in routine practice.

The MMR paradigm we have defined could also be seen as a means for implementing the knowledge conversion cycle presented in Fig. 2:

- 1) CBR is a natural way for performing socialisation, since, through past cases

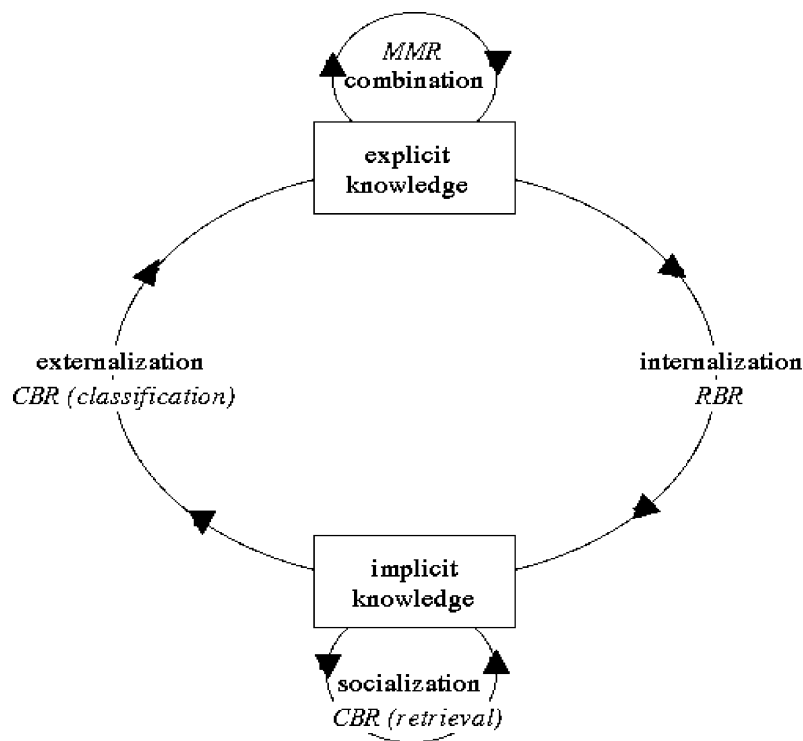


Fig. 2. MMR as a means for implementing the knowledge conversion cycle.

retrieval, it enables implicit knowledge sharing;

- 2) case classification supports externalisation: it is able to contextualise and categorise cases, transferring the implicit knowledge into an explicit form;
- 3) through a continuous application of MMR into clinical practice, combination could be supported as well: the acquisition and the analysis of new cases could be used to tune the classes definition; moreover, new production rules could be learnt for well identified patient situations, speeding up the therapy revision task;
- 4) finally, RBR naturally supports internalisation, by applying explicit knowledge to produce new cases solutions, that store in an implicit form the medical skills exploited to solve them.

5. Conclusions

In the medical domain, different knowledge types are typically available. Operative knowledge, collected during every day practice, and reporting expert's skills, is stored in the HIS. On the other hand, well-assessed, formalised medical knowledge, is reported in textbooks and clinical guidelines. All this heterogeneous information should be secured and distributed, and made available to physicians in the right form, at the right time, in order to support decision making, therefore, a decision support system cannot be conceived as an independent tool, able to substitute the human expert on demand, but should be integrated with the KM task.

KM is a discipline for keeping, upgrading and reusing the organisational intellectual

asset, that can be consequently exploited for reasoning and for therapy definition purposes. CBR is a very well suited methodology for performing KM when dealing with implicit knowledge: it is able to focus the attention on the problem at hand, retrieving past situations similar to the current one, to show a direction for solving the current problem. On the other hand, RBR is historically one of the most successful paradigms for managing knowledge of the explicit type.

To take advantage of all the available sources of information, in the field of diabetes care, we have defined a MMR methodology, that performs a tight integration of CBR and RBR, able to propose a therapy suggestion that is customised to the situation under examination. The large amount of patients data available when dealing with chronic care is also the typical application field for data mining techniques. Data can be processed through machine learning methods, in order to identify possible patterns in them, able to provide a suggestion for optimising patients management. The MMR methodology described in this paper could be interpreted in the data mining perspective as well: through the implementation of the classification step, it is able to provide an added value to the usefulness of information retrieval alone. The patient's data are analysed to identify meaningful behaviours in them, that allow us to categorise the case as belonging to a prototypical class. As a matter of fact, classification is a typical data mining technique, and in our system it provides a way for contextualising the situation at hand, for specialising the problem identification task, and for focusing the retrieval search space on the most relevant sub-parts of the case library.

In conclusion, our methodology seems to be a valuable way for exploiting all the knowledge available within the organisation, for mining and intelligently retrieving it, and for

distributing it among different health care providers, to optimise the processes of patients management and therapy revision.

Acknowledgements

Professor Mario Stefanelli and Professor Luigi Portinale are gratefully acknowledged their support in reviewing the methodological issues presented in this paper.

References

- [1] H.A. Heatefield, J. Wyatt, Philosophies for the design and development of clinical decision-support systems, *Methods of Information in Medicine* 32 (1993) 1–8.
- [2] M. Musen, Architectures for architects, *Methods of Information in Medicine* 32 (1993) 12–13.
- [3] H. Warner, *Computer-Assisted Medical Decision-Making*, Academic Press, New York, 1979.
- [4] O. Kuhn, A. Abecker, Corporate memories for knowledge management in industrial practice: prospects and challenges, *Journal of Universal Computer Science* 3 (1997) 929–954.
- [5] S. Montani, R. Bellazzi, Exploiting multi-modal reasoning for knowledge management and decision support: an evaluation study, *Proc. AMIA Symp.* 2000, 585–589.
- [6] L. Steels, Corporate knowledge management, *Proceedings of ISMICK '93*, Compiegne, France, 1993, pp. 9–30.
- [7] R. Van der Spek, A. Spijkervet, Knowledge management: dealing intelligently with knowledge, in: J. Liebowitz, L.C. Wilcox (Eds.), *Knowledge Mangement and its Integrative Elements*, CRC Press, 1997.
- [8] S.S.R. Abidi, Knowledge management in healthcare: towards 'knowledge-driven' decision support services, *International Journal of Medical Informatics* 63 (2001) 5–18.
- [9] R.M. Davis, E.G. Wagner, T. Groves, Advances in managing chronic disease, *British Medical Journal* 320 (2000) 525–526.
- [10] J. Van der Lei, Closing the loop between clinical practice, research, and education: the potential of electronic patient records, *Methods of Information in Medicine* 41 (1) (2002) 51–54.
- [11] S.M. Campbell, et al., Identifying predictors of high quality care in English general practice: observational study, *British Medical Journal* 323 (2001) 784.
- [12] G. Simon, M. Grandbastien, Corporate knowledge: a case study in the detection of metallurgical flaws, *Proceedings of ISMICK '95*, Compiegne, France, 1995, pp. 43–52.

- [13] J.L. Kolodner, *Case-Based Reasoning*, Morgan Kaufmann, 1993.
- [14] A. Aamodt, E. Plaza, Case-based reasoning: foundational issues, methodological variations and systems approaches, *AI Communications* 7 (1994) 39–59.
- [15] D.B. Leake, J.L. Kolodner, A tutorial introduction to CBR, in: *Case Based Reasoning: Experiences, Lessons and Future Directions*, AAAI Press, 1996, pp. 31–65.
- [16] J.L. Kolodner, R.M. Kolodner, Using experience in clinical problem solving: introduction and framework, *IEEE Transactions on Systems, Man and Cybernetics* 17 (1987) 420–431.
- [17] K.J. Hammond, *Case-Based Planning: Viewing Planning as a Memory Task*, Academic Press, 1989.
- [18] I. Bichindaritz, A case-based assistant for clinical psychiatry expertise, in: J.B. Ozbolt (Ed.), *Proceedings of the 18th Symposium on Computer Applications in Medical Care*, AMIA, 1994, pp. 673–677.
- [19] E. Bareiss, P. Porter, C. Wier, PROTOS: an exemplar-based learning apprentice, *International Journal of Man–Machine Studies* 29 (1988) 549–561.
- [20] R.M. Turner, Using schemas for diagnosis, *Computer Methods and Programs in Biomedicine* 30 (1989) 199–207.
- [21] C.D. Evans, A case-based assistant for diagnosis and analysis of dysmorphic syndromes, *Medical Informatics* 20 (1995) 121–131.
- [22] M. Fathi-Torbaghan, D. Meyer, MEDUSA: a fuzzy, expert system for medical diagnosis of acute abdominal pain, *Methods of Information in Medicine* 33 (1994) 522–529.
- [23] C.E. Kahn, G.M. Anderson, Case-based reasoning and imaging procedure selection, *Investigative Radiology* 29 (1994) 643–647.
- [24] J. Berger, Roentgen: a case-based approach to radiation therapy planning, in: *Proceedings of a Workshop on Case-Based Reasoning*, Morgan-Kaufmann, San Mateo, CA, 1989, pp. 218–223.
- [25] R. Schmidt, L. Boscher, B. Heindl, G. Schmid, B. Pollwein, L. Gierl, Adaptation and abstraction in a case-based antibiotic therapy adviser, in: P. Barahona, M. Stefanelli, J. Wyatt (Eds.), *Artificial Intelligence in Medicine*, *Lecture Notes in Artificial Intelligence*, 1995, p. 934.
- [26] K.M. Wiig, Roles of knowledge-based systems in support of knowledge management, in: J. Liebowitz, L.C. Wilcox (Eds.), *Knowledge Management and its Integrative Elements*, CRC Press, 1997.
- [27] E.H. Shortliffe, *Computer-Based Medical Consultations: MYCIN*, Elsevier/North Holland, New York, 1976.
- [28] J. Surma, K. Vanhoff, Integrating rules and cases for the classification task, in: *Lecture Notes in Artificial Intelligence 1010*, *Proceedings of 1st ICCBR*, Springer, 1995, pp. 325–334.
- [29] L.K. Branting, B.W. Porter, Rules and precedents as complementary warrants, *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI 91)*, Anaheim, 1991.
- [30] I. Bichindaritz, E. Kansu, K.M. Sullivan, Case-based reasoning in CARE-PARTNER: gathering evidence for evidence-based medical practice, in: *Lecture Notes in Artificial Intelligence 1488*, *Proceedings 4th EWCBR*, Springer, 1998, pp. 334–345.
- [31] D.B. Leake, Combining rules and cases to learn case adaptation, *Proceedings of the 17th International Conference of Cognitive Science Society*, Pittsburgh, 1995.
- [32] S. Russel, P. Norvig, *Artificial Intelligence, a Modern Approach*, Prentice Hall, 1995.
- [33] The Diabetes Control and Complication Trial Research Group, The effect of intensive treatment of diabetes on the development and progression of long-term complications in insulin-dependent diabetes mellitus, *The New England Journal of Medicine* 329 (1993) 977–986.
- [34] S. Montani, R. Bellazzi, L. Portinale, M. Stefanelli, A multi modal reasoning methodology for managing IDDM patients, *International Journal of Medical Informatics* 58–59 (2000) 243–256.
- [35] R. Hovorka, S. Svacina, E.R. Carson, C.D. Williams, P.H. Sonksen, A consultation system for insulin therapy, *Computer Methods and Programs in Biomedicine* 32 (1996) 303–310.
- [36] G. Nucci et al., Verification phase final report, T-IDDM deliverable 5.2, <http://aim.unipv.it/projects/tiddm/ftp.html>.
- [37] I. Nonaka, H. Takeuchi, *The Knowledge-Creating Company*, University Press, Oxford, 1995.