Appl Intell (2008) 28: 275–285
DOI 10.1007/s10489-007-0046-2

Exploring new roles for case-based reasoning in heterogeneous AI systems for medical decision support

Stefania Montani

Received: 11 October 2006 / Accepted: 29 March 2007 / Published online: 31 May 2007

Abstract *Background* Supporting medical decision making is a complex task, that offers challenging research issues to Artificial Intelligence (AI) scientists. The Case-based Reasoning (CBR) methodology has been proposed as a possible means for supporting decision making in this domain since the 1980s. Nevertheless, despite the variety of efforts produced by the CBR research community, and the number of issues properly handled by means of this methodology, the success of CBR systems in medicine is somehow limited, and almost no research product has been fully tested and commercialized; one of the main reasons for this may be found in the nature of the problem domain, which is extremely complex and multi-faceted.

Materials and methods In this environment, we propose to design a modular architecture, in which several AI methodologies cooperate, to provide decision support. In the resulting context CBR, originally conceived as a well suited reasoning paradigm for medical applications, can extend its original roles, and cover a set of additional tasks.

Results and conclusions As an example, in the paper we will show how CBR can be exploited for configuring the parameters relied upon by other (reasoning) modules. Other possible ways of deploying CBR in this domain will be the object of our future investigations, and, in our opinion, a possible research direction for people working on CBR in the health sciences.

S. Montani (✉)
Dipartimento di Informatica, Università del Piemonte Orientale, Alessandria, Italy
e-mail: stefania.montani@unipmn.it

1 Introduction

The first papers analyzing the possibility of providing decision support to physicians by exploiting computers appeared in the 1950s [1], and the first accurate experimental prototypes were soon implemented [2]. As a matter of fact, decision making is a central challenge in medical practice, and looking at the still partially unexplored potentialities of machines seemed a natural objective to the earliest medical informatics scientists. The research field of computer assisted decision making has remained very active over the years, and has led to the definition of several different solutions.

Among the exploited reasoning methodologies, Case-based Reasoning (CBR) [3, 4] soon appeared to be a very well suited choice, with the earliest CBR systems in medicine dating back to twenty years ago (see e.g. [5]). CBR is a problem solving paradigm that utilizes the specific knowledge of previously experienced situations, called cases. It basically consists in *retrieving* past cases that are similar to the current one and in *reusing* (by, if necessary, *revising*) past successful solutions; the current solved case can then be *retained* and put into the system knowledge base, called the *case base* or the *case library*. The retrieve, reuse-revise and retain procedures are known as the steps of the *CBR cycle* [4].

Actually, several arguments for the exploitation of CBR in the medical domain can be recognized:

- *CBR resembles human reasoning* in general, and medical decision making in particular. As a matter of fact, physicians are used to reason by recalling past situations similar to the current one. The process is often biased by the tendency to recall only more recent or difficult cases, or only the positively solved ones. CBR seems therefore to be a natural solution for the diagnostic and/or therapeutic goals in this field, by enabling the retrieval of older, simpler or negative examples as well. Storing and recalling practical cases comes out to be very useful also for sharing other clinicians' skills, and for training un-experienced personnel, which are key objectives of every health care organization.

- *CBR allows operative knowledge management*. Managing knowledge is a relevant issue to be addressed in the medical domain, where large amounts of information are generally available. The introduction of Hospital Information Systems (HIS) into clinical practice has led to the memorization of a huge quantity of data, extracted from the day by day activity, thus providing a new type of operative knowledge, which can be exploited together with the general domain knowledge. Actually this kind of data may be effectively used to change organizational settings [6], and to improve the overall quality of care. The role of operative knowledge is even more central for those diseases that are still not well understood, or for which generalized rules or models do not apply. CBR is one of the most suitable methodologies for managing operative knowledge. Actually, retrieving and reusing past data, and retaining new information, which are the basic steps of the CBR cycle, fit very well the Knowledge Management discipline [7] objectives of keeping, increasing and reusing knowledge. Moreover, representing a real world situation as a case is often straightforward: given a set of meaningful features for the application, it can be sufficient to identify the value they assume in the situation at hand; the case can also store information about the solution applied and sometimes about the outcome obtained. Every past case then *implicitly* embeds a bunch of domain knowledge (i.e. the problem-action pattern adopted on that occasion), which can be memorized without the need of making it *explicit* in a more generalized form, thus mitigating the well known knowledge acquisition bottleneck that affects other methodologies (e.g. Rule or Model-based Reasoning).
- *CBR allows to integrate different knowledge types*. When domain knowledge is available, extracted from textbooks or physicians committees expertise, and formalized by means of rules, ontologies, or computerized guidelines, its integration with operative knowledge may represent a significant advantage. Cases can be quite naturally adopted to complement guidelines and to make them operational in a real setting, or can be of help in situations where multiple diagnoses interact, and some of them cannot be deducted by an existing model representing the disease process [8]. Actually, CBR is well suited for integration with other reasoning methodologies, and has been largely deployed in multi-modal reasoning systems (see Section 2.2).

As a matter of fact, CBR has successfully proven to solve or handle issues that previously had been too difficult to manage with other methods and techniques. The results of CBR research in healthcare are so far valuable, and they show a significant potential for the future.

Nevertheless, some limitations can be still outlined. In particular, despite the fact that several CBR systems have been proposed in the literature (see Section 2), a few of them have been fully tested or commercialized, a few are autonomous with respect to human intervention (for example, many authors have implemented pure retrieval systems, that leave to the physician the responsibility of providing the current case solution), and several problems remain open (see Section 3).

The next section investigates the possible reasons for such limitations, and the possible improvement directions.

1.1 Reasons for failure and improvement directions

The limitations outlined in the introduction are essentially due to two categories of reasons: reasons related to the intrinsic "weakness" of the CBR methodology, and reasons related to the (increasing) complexity of the healthcare domain. In the first category, we may highlight:

- *Difficulties in feature mining*. Despite the above comments about the simplicity in the acquisition of a case, in several applications case representation is becoming more complex than in the past [8]: case data can partly come in the form of time series [9, 10], or images [11, 12], or free text [13], and can be intrinsically high-dimensional.
- *Presence of competence gaps*. Retrieval results may be affected by a low quality of the case base content. Actually, a misleading indication on how to solve the current problem may emerge when the number of cases is too small, or when the retrieved information is polarized on too specific examples. The capability of filling such competence gaps is crucial especially in medical applications, since final decisions should be always based on well established knowledge.
- *Challenge in the definition of a suitable adaptation strategy*. No general framework for performing adaptation in medical CBR has been proposed so far. In non-medical fields, adaptation is often solved by defining some adaptation rules, elicited by experts. The complexity of many medical domains makes this effort (applied e.g. in [14]) very hard and time consuming, and leads back to the knowledge acquisition bottleneck. Moreover, case solutions are frequently associated with risks in clinical domains: risk analysis should thus be conducted when defining adaptation strategies, making the process too slow if completed at run time.

Other reasons for the limited success of CBR systems are related to the increasing complexity of medical domains. Actually, many healthcare applications are simply too complex and multi-faceted to be handled using CBR [15]. The choice to rely on a combination of different Artificial Intelligence (AI) reasoning methodologies appears a possible way to tackle the problem. In this direction, several multi-modal reasoning systems, in which CBR is just one of the available reasoning tools, have been reported in the literature (see Section 2.2).

Nevertheless, deploying CBR just for reasoning is a limitation with respect to the potentialities that this paradigm can offer to medical decision making. In this paper, we propose an alternative approach, in which CBR is not exploited (only) as a reasoning methodology any more, but fulfills different tasks, such as parameter configuration or classification, providing its output as an input to other modules in a heterogeneous environment, where several AI techniques cooperate to generate the final result. In this way, additional issues, that remain as open problems in multi-modal reasoning systems, might be correctly addressed as well (see Section 3).

The paper is organized as follows: Section 2 briefly introduces past work about CBR and multi-modal reasoning systems in healthcare. Section 3 motivates and describes our approach, and presents two already implemented systems in which the approach itself has been realized. Finally Section 4 concludes the paper.

2 CBR as a reasoning methodology: a classical role

The observations stated in the introduction justify the significant number of CBR systems developed to support medical applications.[1]

In particular, some of them were pure CBR systems, i.e. they resorted to CBR alone as a reasoning methodology. Among these ones, some were tools for supporting medical diagnosis in different domains: from psychiatry [5], to cardiology [17], to oncology [18]. CBR has also been exploited for therapy planning: for instance in diagnostic imaging procedure selection [19], in radiation therapy [20] and in antibiotics selection [21].

But pure CBR systems did not successfully tackle the open issues analyzed in Section 1.1. In more recent years, therefore, new tendencies appeared, giving birth to an extremely active research panorama (see e.g. [22, 23]). In the following subsections, we will categorize the resulting systems in relation to the issue they were meant to cope with.

2.1 Dealing with feature mining

Some medical applications are intrinsically data intensive. Recent technological achievements allow to automatically sample and record biological signals, e.g. in Intensive Care Unit, haemodialysis, or instrumental diagnostic procedures. In these fields, data typically come in the form of time series, and when they represent the features of a case, they cannot be simply stored "as they are". Similar problems emerge when features have to be extracted from diagnostic images, again a high dimensional data type.

Data pre-processing and dimensionality reduction techniques are required in these situations. Several recent systems address this point. Time series data pre-processing is pursued by means of different methodologies, such as Temporal Abstractions [24] in [25], filtering and distortion reduction techniques in [10]. Data dimensionality is reduced e.g. by applying the Discrete Fourier Transform [26]. In image pre-processing, on the other hand, suitable algorithms have been designed in order to abstract features such as color and shape from the image itself; methodologies to learn generalized cases from images and similarity measures suited for this field have been introduced as well [12].

Taking advantage of this kind of techniques, also textual data can be mined for features extraction (see e.g. [13]), facilitating the development of CBR systems from the medical literature.

2.2 Dealing with competence gaps

The quality of the case library content might be unsatisfactory, especially when the system is initially put into operation, and may remain unsatisfactory for some particular situations, not represented by a sufficient number of past examples. The presence of such competence gaps makes retrieval results highly unreliable. A reasonable way to tackle this problem is to resort also to other knowledge sources (e.g. different kinds of formalized background knowledge), when available, in order to deploy all the information collected at the health care organization, when supporting decision making.

This observation has led to the design of many multimodal reasoning systems, implemented (at least in the form of prototypes) in different domains. The interest in multimodal approaches involving CBR dates back to more than ten years ago, with famous systems such as CASEY [27] and FLORENCE [28], and is recently increasing. In particular, CBR has proved to be well suited for integration with Rule-based Reasoning (RBR) or Model-based Reasoning (MBR). In the literature, the combination of CBR with RBR has received particular attention, since rules are truly the most successful explicit knowledge representation formalism for intelligent systems. As it is well known, RBR consists in firing a set of rules through a chaining mechanism (forward or backward), exploiting the available data to arrive at a decision. Each rule is a conditional statement, which formalizes domain knowledge by relating observations to an associated inference that can be drawn from them.

[1] The following review is intentionally non exhaustive; interesting surveys can be found in [8, 15, 16].

Different levels of integration between the two paradigms are described. Usually RBR and CBR are applied in mutually exclusive ways (see e.g. [29]), where RBR deals with knowledge on standard or typical problems, while CBR faces exceptions. In this view, CBR is exploited to retrieve similar cases from a library of peculiar and non-standard situations only when RBR has failed to produce a solution [30, 31]. Other approaches rely on CBR for instantiating and providing suitable contexts to rules, while rules are used to assist CBR by permitting the extraction of more general concepts from concrete examples [32]. It is possible to select which methodology to apply first in a dynamic way, depending on the situation at hand [32, 33]. 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 on the basis of its suitability for solving the current problem [33].

2.3 Dealing with adaptation

Different strategies to cope with the difficulty in the definition of case-based adaptation in the medical domain have been devised. In the simplest case, the implemented systems just avoid the adaptation step (i.e. they are pure retrieval systems) [21, 26, 34]. Multi-modal reasoning methodologies (see Section 2.2) represent an alternative solution: in some of these contributions adaptation is not dealt with because a different reasoning paradigm provides the final solution to the input problem. Finally, it has been proposed to rely on the use of prototypes [35], i.e. generalized cases on which the definition of an adaptation strategy becomes easier, since the specific details of ground cases leave space to a more general kind of knowledge [36]. Anyway, as already observed, a theoretical and generalizable paradigm to implement adaptation in medical problems has not been devised so far. As a matter of fact, adaptation remains as a weak point in medical CBR applications.

3 Devising new roles for CBR in medical decision making

The relatively simple procedure that often allows case acquisition, and the natural way in which operative knowledge is managed by the CBR cycle, makes CBR potential applications in medicine very different. Deploying it as (the main) reasoning paradigm in a system for supporting medical decisions, as in Section 2, is actually not the only way in which it is possible to take advantage of this methodology.

Moreover, despite the fact that some recent approaches (see Section 2.2 in particular) succeed in addressing part of the problems that affect CBR systems in healthcare, some issues still remain open, when CBR is deployed to cover the reasoning task. In particular:

1. in many real settings, data memorization into the HIS is still incomplete. This situation often holds true for instance in Italy, where the adoption of computerized devices in routine clinical practice is still opposed by many doctors and nurses, who feel unconfident with them. Moreover, sometimes the HIS tables schemas cannot be directly converted into the case structure defined in the CBR system. The lack of information in the electronic format makes case mining very hard;
2. in medical practice, past diagnostic or therapeutic recommendations tend to become quickly obsolete, due to the development of new technologies and to the availability of new scientific evidences. Therefore, despite memorizing operative knowledge is still a desirable task, the case base should be maintained and updated in a life-long learning perspective [8], where recent developments and findings are integrated, while old ones are carefully evaluated, and possibly discarded. This issue, together with issue 1, in a way leads back to the knowledge acquisition bottleneck, a problem that CBR should be able to reduce (see Section 1);
3. despite the fact that the efforts in implementing multi-modal reasoning systems have led to the possibility of exploiting different knowledge sources, at least to some extent, current CBR systems do not clearly address interoperability issues with other knowledge bases and case bases. This would be a desirable result [37], since case mining can be a hard task, as outlined above;
4. CBR systems do not associate their output with probabilities and statistics, which, on the other hand, add a scientific dimension to clinical research, and would make results acceptable to physicians [8]. Actually, at the moment most of the implemented CBR systems require the physician's intervention to accept/validate the obtained results ([15]; as a counter-example see [29]). This limitation is particularly critical in emergency-treating systems, in diagnostic systems, and in systems meant to provide advice directly to patients (although CBR might not always be the most suited methodology for this kind of applications);
5. finally, CBR may have to face restrictions to medical data access, due to legal issues. Moreover, currently CBR systems do not address data protection [8].

We therefore propose to enrich the panorama of the possible ways of applying CBR, by studying new roles the paradigm could cover, within a modular architecture, in which different AI methodologies interact among them to provide the final result (i.e. medical decision support). In particular, CBR would be extremely well suited for tasks such as classification, or parameter configuration for other AI modules.

In [10], for example, a CBR system is deployed to classify biological signals in order to diagnose Respiratory

Sinus Arrhythmia. The classified signals are thereafter sent to a second subsystem, the pattern-identifier. The pattern-identifier analyses the classified signals and searches relevant sequences in them. The identified sequences give clinicians a more complete insight of the measurements, providing them with a better basis for diagnosis. Classified signals manually diagnosed by experts can then be provided as an input to a knowledge discovery procedure [38]: the newly learnt knowledge (i.e. which sequences are the most important for which diagnosis) can later be used by a reasoning methodology (maybe a second CBR system, but not necessarily) with the aim of automating the diagnostic task.

Parameter configuration, on the other hand, is a critical issue in many AI processes (e.g. Rule-based Reasoning, or Temporal Abstractions), especially when they are applied to complex domains like medical ones. Configuration of reasoning modules usually needs domain knowledge. When general domain knowledge is not available, we can resort to the case base (i.e. to the data), which provides specific knowledge hidden in cases, to solve configuration problems. In several applications, in particular, the use of contextual knowledge represents an appropriate means for parameter setting. Defining a (prototypical) case as a set of feature/value pairs keeping contextual information, and storing the suggested parameter configuration as the corresponding solution, CBR can be resorted to in order to fulfill this task. The main advantage of a CBR approach obviously stands in the fact that the knowledge acquisition process for configuring parameters is made easier by the use of already configured cases, retrieved because similar to the current input situation. In the next subsections, by introducing two already (partially) implemented systems, we provide an overview on the use of CBR for parameter configuration, and we demonstrate how the approach can provide an added value to a more complex decision support architecture.

Note that the newly proposed CBR tasks defined in this section may overcome (some of) the open problems listed above. As a matter of fact, in these situations cases do not store instances of solutions to the decision making process, but address the maybe less ambitious—but still relevant—goal of keeping contextual information, or of supporting a kind of data pre-processing. Issues 2, 3 and 4 appear less critical in this light, since the case solution is not a diagnosis/therapy to be accepted by a physician (i.e. based on well-established knowledge and scientifically reliable). This observation also mitigates issue 1, because the patient's problem description, still required in the case structure when CBR is used for classification or parameter configuration, is usually automatically stored in the HIS, while the health care procedures executed by nurses/physicians sometimes are not. Legal issues and data protection (issue 5), on the other hand, still need to be correctly considered.

3.1 CBR for parameter configuration in a Rule-based Reasoning system

The first application we present refers to supporting therapy modification in young type 1 diabetic patients management.

Diabetes mellitus is one of the major chronic diseases in the industrialized countries. Patients affected by type 1 diabetes need to undergo Intensive Insulin Therapy (IIT), consisting in 3 to 4 injections of exogenous insulin every day, in order to regulate blood glucose, and to reduce the risk of later life complications. Before every injection, patients have to measure and record their blood glucose level. Blood glucose time series will then be inspected by physicians to assess therapy efficacy and to revise the therapy itself in case of need.

In this domain, we have developed a Rule-based Reasoning (RBR) system that implements the following reasoning tasks:

1. *Identification of metabolic problems*. We extract an abstract description of the course of longitudinal blood glucose data collected by the patient through the Temporal Abstractions (TA) technique [24]. TA are an AI methodology able to solve a data interpretation task, the goal of which is to derive high level concepts from time stamped data. The basic principle of TA methods is to move from a *point-based* to an *interval-based* representation of the data, where: (i) the input points (*events* henceforth) are the elements of a (discretized) time series; (ii) the output intervals (*episodes* henceforth) aggregate adjacent events sharing a common behavior, persistent over time.

These abstractions can be further subdivided into *state* TA and *trend* TA. State TA are used to extract episodes associated with *qualitative levels* of the monitored feature, e.g. *low*, *normal*, *high* values; trend TA are exploited to detect specific *patterns*, such as *increase*, *decrease* or *stationarity*, from the time series.

In the RBR system, the identification of metabolic problems task is based on a careful analysis of state abstractions on the blood glucose time series. In more detail, the episodes of five state abstractions (namely *very low*, *low*, *normal*, *high*, *very high* values) are searched for in the blood glucose measurements.

After having identified the state abstractions episodes in every time slice of the day (e.g. at breakfast, lunch, dinner time), the blood glucose *modal day* is extracted. The modal day represents the characteristic daily blood glucose pattern that summarizes the patient's response to the therapy in a specific monitoring period; it is used to evaluate the insulin protocol performance over the selected time interval, even when the information is poor.

In our approach we derive the modal day by calculating the marginal probability distribution of the state abstractions

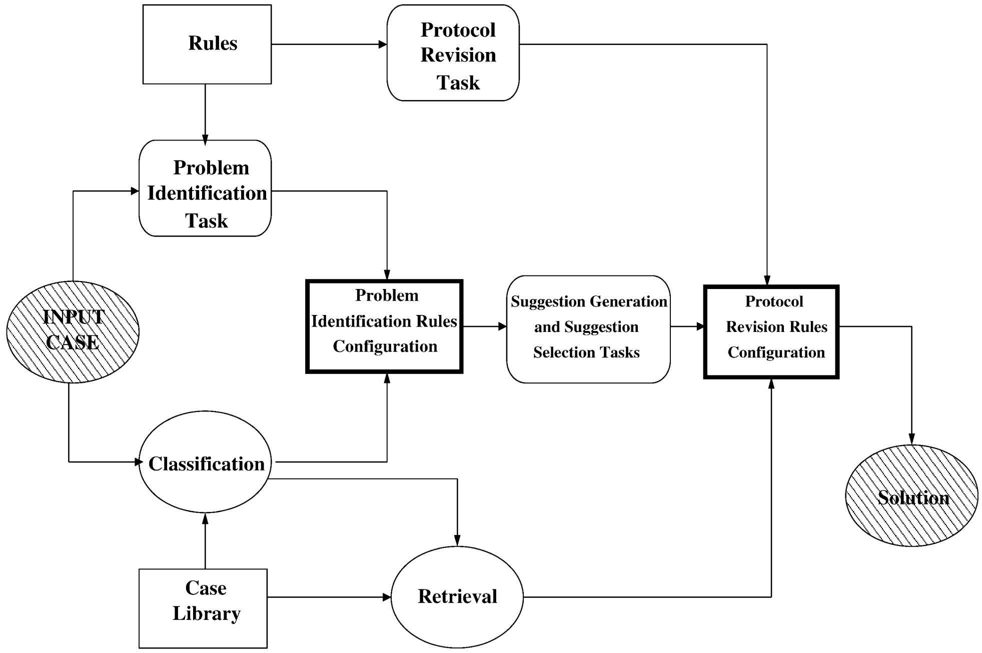

Fig. 1 Case-based configuration of rules parameters

listed above. In particular, we apply a Bayesian method described in [39] that is able to explicitly take into account the presence of missing data: since we considered 5 state abstractions, before starting data collection in a given period we assign a prior probability to the occurrence of each state abstraction equal to $\frac{1}{5}$ in every time slice. After a certain monitoring period of N days, we collect D measurements, while the remaining $M = N - D$ data are missing. The posterior probability lower ($p_{\inf}$) and upper ($p_{\sup}$) bounds of the occurrence of a generic k-th of the 5 levels can be derived as:

$$p_{\inf} = \frac{1 + d_k}{5 + N},$$

$$p_{\sup} = \frac{1 + d_k + M}{5 + N}$$

where d_k is the number of occurrences of the k-th level in the monitoring period.

The difference between $p_{\sup}$ and $p_{\inf}$ is proportional to the number of missing data and is denoted as the *ignorance* in the monitoring period.

As the monitoring process proceeds, the bounds on the probabilities are updated. At any time we obtain an interval probability distribution over the blood glucose state abstractions. The modal day is extracted by taking the blood glucose states with the highest $p_{\inf}$ in each time slice of the day.

When $p_{\inf}$ for (very) low/high blood glucose values in a certain time slice of the modal day is higher than a given α threshold, and when the number of missing data is sufficiently small to rely on such information (i.e. the ignorance is smaller than a given β threshold), a *problem* is identified. α and β are percentages obtained from medical knowledge.

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 then selected resorting to the concept of *insulin competence*. The most competent insulin, that has the strongest effect on the time slice of the day in which the problem has been found, is identified. Competence is evaluated relying on the pharmacokinetics of the different insulin types [40].

3. *Insulin protocol 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 (i.e. $+1, -1$ insulin units) to the current protocol insulin doses.

Even though the RBR system behavior was judged correct and quite satisfactory in a formal evaluation study [41], it came out to be sometimes not sharp enough to promptly face the patients alterations. To overcome this weakness, we have set up an architecture (developed in the context of the EU-funded project M2DM IST-1999-10315) in which a case retrieval system properly tunes rule parameters, thus tailoring them to the specific patient's needs (see Fig. 1). Observe that the major long-term intervention trial on type 1 diabetic patients, the DCCT [42], has clearly shown that the definition and realization of an appropriate individual therapeutic goal, customized on the single patients characteristics, is the key to an effective diabetes care, rather than the implementation of a specific therapy.

Exploiting CBR for rule parameter setting is, to our knowledge, an innovative approach, in which retrieval results can really modify and properly tune the rules behavior. As a matter of fact, in classical integrations between CBR and RBR (see Section 2.2), the two paradigms are often exploited in a mutually exclusive way, where RBR deals with knowledge on standard or typical problems, while CBR faces exceptions, and does not really "touch" the rules definition.

In our approach, the concept of case has been mapped to the data collected during a periodical control visit. Since in routine clinical practice patients are visited every 2–4 months, on those occasions a new case can be automatically stored in the case library.

By co-operating with the paediatricians of Policlinico S. Matteo Hospital in Pavia, Italy, we have been able to collect 145 cases from the histories of 29 pediatric patients, and to structure the case library resorting to a taxonomy of mutually exclusive classes, which express the typical problems that may occur to type 1 diabetic patients in the age of infancy and puberty. Then, it has been possible to design a multi-step retrieval strategy, in which:

- first, a new case is classified as belonging to one of the classes. Classification is crucial in making efficient the retrieval by restricting search only to relevant parts of the whole case library. Moreover, this step implements *contextualization*: the patient is categorized 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 she/he is currently living. The identified class (a sort of prototypical case which abstracts the information of the subsumed ground cases) thus embeds the contextual knowledge to be reused, and is relied upon to tune rule parameters in task 1 (identification of metabolic problems). In particular, classification results are used to tune the α and β thresholds (see task 1) to more proper values, thus tailoring the identification of problems to the particular patient's condition.

  As an example, by default the RBR system considers low glycaemia episodes as severe (i.e. a problem generation rule is triggered) when $p_{\text{inf}} >= 30\%$ and $(p_{\text{sup}} - p_{\text{inf}}) <= 10\%$ (i.e. $\alpha = 30$ and $\beta = 10$). But, if the patient is classified as suffering from anorexia, even when $20\% <= p_{\text{inf}} <= 30\%$ and $10\% <= (p_{\text{sup}} - p_{\text{inf}}) <= 50\%$, the problem is raised. The solution of the anorexia prototypical case (i.e. class) thus sets $\alpha = 20$ and $\beta = 50$. This is motivated by the fact that such patients run a higher risk of hypoglycaemia, since they habitually do not have a sufficient food intake. A prompt reaction is therefore necessary in this case, even when a few episodes have been identified, or in presence of many missing data;
- secondly, cases in the selected class are retrieved. Some sufficient statistics on the insulin protocols (i.e. the solutions) prescribed in the retrieved cases (when available) are used to tune the number of insulin doses to be added/subtracted from the current therapy in task 3 (protocol revision).

The clinical correctness of the suggestions proposed by the RBR system configured by case retrieval results (RBR-config henceforth) has been tested through a formal evaluation procedure [43]. In particular, on both retrospective and simulated data, RBR-config suggested sharper adjustments to the insulin administration in comparison to RBR, reacting more promptly to hypoglycaemia or hyperglycaemia, depending on what was the most relevant problem at hand. As a second validation step, two diabetologists were asked to perform a fully-crossed, blind review of the therapies proposed by RBR, by RBR-config, and by two colleagues, to 30 real patient cases. 65/118 therapies were globally judged as acceptable for the physicians, and 62/120 for the two systems (RBR and RBR-config were initially grouped together for this study, in order to draw a comparison with respect to humans). These small differences support the hypothesis that the RBR-config and the RBR tools are able to perform at an expert level. On the other hand, by deeply inspecting the differences between RBR and RBR-config answers, the benefits obtained by exploiting RBR-config in comparison to RBR appeared to be quite limited, but the reason was easily identified in a case library affected by competence gaps. To cope with the problem of misleading retrieval information, we foresee the implementation of a control strategy, that enables the exploitation of retrieval results only if a sufficiently large number of similar cases have been retrieved, if they are similar enough to the input case to justify their use, and, of course, if the details about the prescribed insulin protocols are available. In this way, RBR-config will support decision making at various levels of complexity, as the case library grows: first, when the stored information is poor, RBR will be applied without exploiting CBR results. As far as new information is stored in the case library, retrieval results will become more reliable, and will be resorted to by the RBR-config methodology.

3.2 CBR for parameter configuration in a Temporal Abstraction server

Temporal Abstractions (TA) (see the definition in Section 3.1) can be adopted as a means for reducing data dimensionality (see also Section 2.1) in time series processing [9], and for subsequently interpreting the information carried by the time series themselves (see e.g. Section 3.1), in particular when:

- a *qualitative* abstraction of the time series values is sufficient;
- a clear mapping between raw and transformed data has to be made available;
- the mapping itself needs to be easily understood by end users as well.

We are currently applying TA as a pre-processing step to the End Stage Renal Disease (ESRD) domain, within a project funded by the Italian Ministry of Education (grant PRIN 2004 number 2004094558). ESRD is a severe chronic condition that corresponds to the final stage of kidney failure. Without medical intervention, ESRD leads to death. Haemodialysis is the most widely used treatment method for ESRD; it relies on an electromechanical device, called haemodialyzer, which, thanks to an extra-corporal blood circuit, is able to clear the patient's blood from catabolites, to re-establish acid-base equilibrium and to remove water in excess. On average, haemodialysis patients are treated for four hours three times a week. Each single treatment is called a haemodialysis session. Haemodialyzers typically allow to collect several variables during a session, most of which are in the form of time series. In our system, time series pre-processed through TA can then be manually analyzed by physicians, or provided as an input to an automatic reasoner (whose description is outside the scope of this paper, see [9]).

The TA output results depend on the value assigned to specific parameters. In particular, for trend TA (see Section 3.1), the following parameters need to be set [44]: Minimum/Maximum Rate (i.e. the minimum and maximum slope allowed for the trend episode); Minimum/Maximum Duration (i.e. the minimum and the maximum duration in time for the trend episode). As regards state TA, on the other hand, we need to specify [44]: Lower/Upper Bound (i.e. the lower and upper bounds of data values allowed for the state episode); Minimum/Maximum Duration (defined as above).

In the domain of haemodialysis, contextual knowledge plays a fundamental role in the time series pre-processing and interpretation (and thus in TA parameters configuration). Since defining the right configuration for each possible contextual situation may be impractical, we have proposed to adopt a case-based approach, where the suitable configuration can be obtained by looking at the most similar already configured case.

In our system, a case stored in the case base is defined as follows: (i) problem description: the *context description*, composed by patient and session characteristics which tend to be stable in the long/medium run (such as patient's age and session duration); (ii) case solution: the *configuration* of the various signals (i.e. of the time series variables collected by the haemodialyzer). In turn, the configuration of each signal consists of a list of state and trend TA symbols to be searched for in the time series to which the configuration refers, together with the corresponding parameter values (and optionally, of a list of suitable combinations of the obtained states and trends, known as *joint TA*).

As a concrete example, let us consider an ESRD patient suffering from hypotension. A case supporting parameter configuration for the hypotension context can be defined as follows:

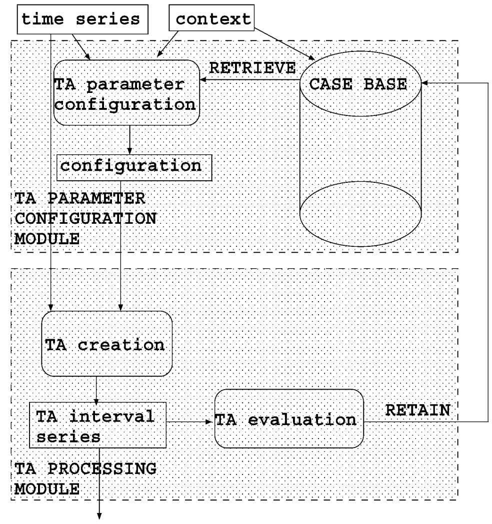

Fig. 2 Case-based configuration of TA parameters, and TA processing

- *context description*: (i) presence of hypotension disease (i.e. systolic pressure below 110 mmHg or diastolic pressure below 60 mmHg [45]); (ii) short session duration (typically 3 hours and 30 minutes); (iii) age in a range in which the collapse is most probable (i.e. 64.4 years or more [46]); (iv) nurse intervention with proper drug (i.e. mannithol).
- *configuration* (for the sake of brevity, we just take into account the diastolic pressure signal, which is one of the most interesting, and a couple of abstractions to be searched for):
  - TA type = Decreasing *Trend*
    - Minimum Rate = 85 degrees
    - Maximum Rate = 88 degrees
    - Minimum Duration = 10 min
    - Maximum Duration = no bound
  - TA type = Low *State* (Hypotension)
    - Lower Bound = 23 mmHg
    - Upper Bound = 59 mmHg
    - Minimum Duration = 6 min
    - Maximum Duration = no bound

The system we are implementing is conceived as a two-module architecture [47], composed as follows (see Fig. 2):

- a case-based module for TA parameter *configuration*;
- a module for TA *processing*.

An input case contains, together with the *context description*, a set of raw time series, instances of the signals on which TA must be extracted. Given an input case, the TA parameter configuration module retrieves the less distant (i.e. most similar) case, with respect to the input case context

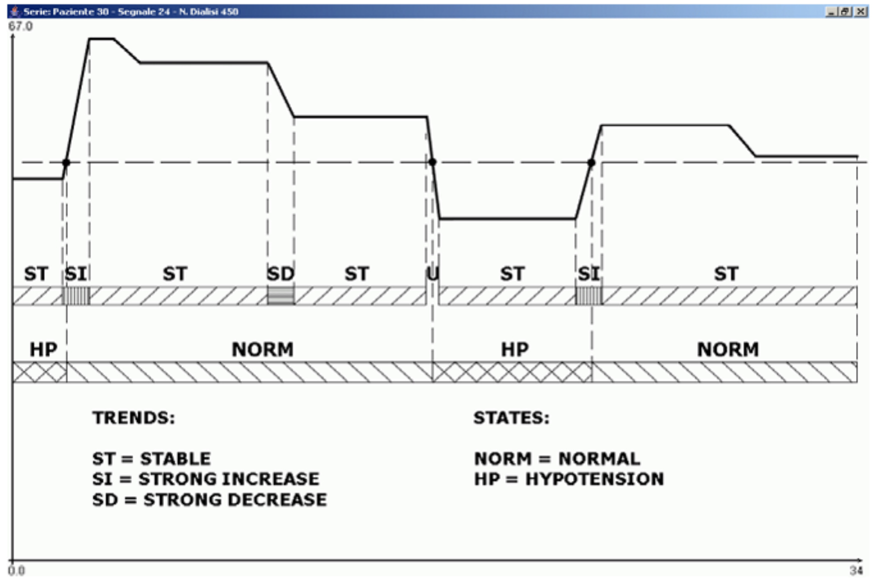

Fig. 3 The output of the TA processing module on the diastolic pressure time series for a patient suffering from hypotension. The U symbol in the middle of the trend time series represents an interval in which no trend could be identified (due to an insufficient duration of the interval itself)

description. The retrieved *configuration* information, corresponding to the signals present in the input case, is extracted and passed to the TA processing module, together with the raw data. The TA processing module, implemented as a TA web service and described in [44], provides a set of qualitative states, trends and suitable combinations of both as a result.

As an example, Fig. 3 shows the output of the TA processing module on the diastolic pressure time series of a patient suffering from hypotension. In the beginning, the pressure value is too low, and the nurse normalizes it by means of a proper drug. A second intervention is required in the middle of the session, and again the nurse succeeds in adjusting the parameter, helping the patient reaching a (more or less) stable diastolic pressure behavior until the end of the session itself. In the figure, trend and states are highlighted, and are extracted relying on the parameter definition for the hypotension context listed above.

The processing of a time series exploiting a given configuration may provide a significant feedback for the possible revision of the used configuration itself. In particular, the resulting TA series can be relied upon to guide case-based maintenance, by suggesting the elimination of non-representative cases, and by keeping the knowledge embedded in the case library always up to date. As described in Fig. 2, the two modules thus give birth to a closed loop architecture, where parameter configurations suggested by the CBR module are adopted for TA processing, while the obtained TA series are evaluated to support case-based maintenance. Case-based maintenance is conceived as a semi-automatic procedure, to be always supervised by a domain expert. However, such a partially data-driven approach is very appealing in this domain, where a well established knowledge about (context, parameter configuration) pairs does not exist.

In particular, cases which are candidate for revisions can be identified by investigating:

1. the frequency of the each suggested abstraction in the input situations;
2. the quality of the extracted TA series

where the last issue is in turn evaluated by considering: (i) the presence of *unknown* abstractions (i.e. of intervals in which the process could not recognize any abstraction); and (ii) the presence of overlaps of (conflicting) TA episodes, over significant time spans (i.e. time spans longer than a given threshold, defined on the basis of medical knowledge). As an example, in Fig. 3 the *U* symbol in the middle of the trend time series represents an *unknown* abstraction, i.e. an interval in which no trend could be identified (due to an insufficient duration of the interval itself, shorter than 10 minutes in that specific situation—see the parameter definition for the hypotension context).

As regards issue 1, an infrequently identified TA can be simply deleted from the case solution. Issue 2, on the other hand, emerges in front of a competence gap region, or of a weakly defined case. In both situations, the expert's intervention is required, in order to acquire a new case from scratch, or to carefully revise the existent ones, in search of conflicting or improper parameter settings.

The system implementation is still on the way, and will have to be followed by a testing phase on real patients' data. To this end, we plan to provide the service to the physicians of the Nephrology and Dialysis Unit of the Vigevano Hospital in Italy, which are involved in the project.

As a final observation, despite the fact that the architecture is being implemented and tested in the haemodialysis domain, it appears to be sufficiently general to be adopted in other (medical) applications as well.

4 Conclusions and future works

In the increasing complexity of medical domains, it makes sense to devise a heterogeneous and modular architecture, in which several AI methodologies cooperate to provide decision support. In this context CBR, originally conceived as a well suited reasoning paradigm for medical applications, can extend its original roles, and cover a set of additional tasks. As an example, in this paper, we have shown how CBR can be exploited for configuring the parameters relied upon by other (reasoning) modules. Other possible ways of deploying CBR in such a multi-faceted domain will be the object of our future investigations, and, in our opinion, a possible research direction for people working on CBR in the health sciences. We believe that a modular and heterogeneous decision support environment, developed along these lines, would be able to provide a valuable support to physicians, thus reinforcing the claim that AI techniques can favor the actual adoption of computer science tools within the medical community.

References

1. Ledley R, Lusted L (1959) Reasoning foundations in medical diagnosis. Science 130:9–21
2. Warner H, Toronto A, Veasy L (1964) Experience with Bayes' theorem for computer diagnosis of congenital heart disease. Ann New York Acad Sci 115:2–16
3. Kolodner JL (1993) Case-based reasoning. Kaufmann, San Mateo
4. Aamodt A, Plaza E (1994) Case-based reasoning: foundational issues, methodological variations and systems approaches. AI Commun 7:39–59
5. Kolodner JL, Kolodner RM (1987) Using experience in clinical problem solving:introduction and framework. IEEE Trans Syst Man Cybern 17:420–431
6. Stefanelli M (2001) The socio-organizational age of artificial intelligence in medicine. Artif Intell Med 23(1):25–47
7. VanderSpek R, Spijkervet A (1997) Knowledge management: dealing intelligently with knowledge. In: Liebowitz J, Wilcox LC (eds) Knowledge management and its integrative elements. CRC, Boca Raton
8. Bichindaritz I, Marling C (2006) Case-based reasoning in the health sciences: what's next? Artif Intell Med 36:127–135
9. Montani S, Portinale L (2006) Accounting for the temporal dimension in case-based retrieval: a framework for medical applications. Comput Intell 22:208–223
10. Nilsson M, Funk P, Olsson E, vonScheele B, Xiong N (2006) Clinical decision-support for diagnosing stress-related disorders by applying psychophysiological medical knowledge to an instance-based learning system. Artif Intell Med 36:159–176
11. ElBalaa Z, Strauss A, Uziel P, Maximini K, Traphoner R (2003) Fm-ultranet: a decision support system using case-based reasoning applied to ultrasonography. In: Proceedings of the case-based reasoning in the health sciences workshop, international conference on case based reasoning (ICCBR), Trondheim, pp 37–44
12. Perner P, Janichen S, Perner H (2006) Case-based object recognition for airborne fungi recognition. Artif Intell Med 36:137–157
13. Wu D, Weber R, Abramson FC (2004) A case-based framework for leveraging nutrigenomics knowledge and personalized nutrition counseling. In: Proceedings of the case-based reasoning in the health sciences workshop, European conference on case based reasoning (ECCBR), Madrid, pp 73–82
14. Leake DB (1995) Combining rules and cases to learn case adaptation. In: Proceedings of the 17th international conference of cognitive science society, Pittsburgh
15. Nilsson M, Sollenborn M (2004) Advancements and trends in medical case-based reasoning: an overview of systems and system development. In: Barr V, Markov Z (eds) Proceedings of the 17th international Florida artificial intelligence research society conference–special track on case-based reasoning. AAAI, Menlo Park, pp 9–18
16. Gierl L, Bull M, Schmidt R (1998) Cbr in medicine. In: Lenz M, Bartsch-Sporl B, Burkhard HD, Wess S (eds) Case-based reasoning technology: from foundations to applications. Springer, Berlin, pp 273–297
17. Jang Y (1993) HYDI: a hybrid system with feedback for diagnosing multiple disorders. PhD thesis (TR-576), Computer Science, Massachusetts Institute of Technology, MA
18. Hunter L (1989) Knowledge acquisition planning: gaining expertise through experience. PhD thesis (DCS-TR-678), Computer Science, Massachusetts Institute of Technology, MA
19. Kahn CE, Anderson GM (1994) Case-based reasoning and imaging procedure selection. Invest Radiol 29:643–647
20. Berger J (1989) Roentgen: a case-based approach to radiation therapy planning. In: Proceedings of the workshop on case-based reasoning. Kaufman, San Mateo, pp 218–223
21. Schmidt R, Pollwein B, Gierl L (1999) Case-based reasoning for antibiotics therapy advice. In: Althoff KD, Bergmann R, Branting LK (eds) Proceedings of the third international conference on case-based reasoning. Springer, Berlin, pp 550–559
22. Bichindaritz I, Marling C (2004) In: Second workshop on case based reasoning in the health sciences, ECCBR 2004. Technical Report 142-04, Departamento de Sistemas Informaticos y Programacion, Univesidad Complutense de Madrid, Madrid
23. Bichindaritz I, Marling C (2005) In: Third workshop on case based reasoning in the health sciences, ICCBR 2005. DePaul University, Chicago
24. Shahar Y (1997) A framework for knowledge-based temporal abstractions. Artif Intell 90:79–133
25. Schmidt R, Gierl L (2001) Temporal abstractions and case-based reasoning for medical course data. Two prognostic applications. In: Perner P (ed) Proceedings of the machine learning and data mining in pattern recognition: second international workshop. Lecture notes in computer science, vol 2123. Springer, Berlin, pp 23–34
26. Montani S, Portinale L, Leonardi G, Bellazzi R, Bellazzi R (2006) Case-based retrieval to support the treatment of end stage renal failure patients. Artif Intell Med 37:31–42
27. Koton PA (1989) Using experience in learning and problem solving. PhD thesis (TR-441), Computer Science, Massachusetts Institute of Technology, MA
28. Bradburn C, Zeleznikow J (1994) The application of case-based reasoning to the tasks of health care planning. In: Wess S, Althoff KD, Richter MM (eds) Proceedings topics in case-based reasoning: first European workshop, EWCBR-93. Springer, Berlin, pp 365–378
29. Marling C, Whitehouse D (2001) Case-based reasoning in the care of Alzheimer's disease patients. In: Aha DW, Watson I (eds) Proceedings of the 4th international conference on casebased reasoning. Springer, Berlin, pp 702–715
30. Surma J, Vanhoof K (1995) Integration rules and cases for the classification task. In: Veloso M, Aamodt A (eds) Proceedings of the 1st international conference on case-based reasoning, Sesimbra, Portugal, October 1995. Lecture notes in computer science. Springer, Berlin, pp 325–334
31. Xu LD (1996) An integrated rule- and case-based approach to AIDS initial assessment. Int J Biomed Comput 40:197–207

32. Branting LK, Porter BW (1991) Rules and precedents as complementary warrants. In: Proceedings of the 9th national conference on artificial intelligence, Anaheim, CA, July 1991. AAAI, Menlo Park
33. Bichindaritz I, Kansu E, Sullivan K (1998) Case-based reasoning in care-partner: Gathering evidence for evidence-based medical practice. In: Smyth B, Cunningham P (eds) Proceedings of the 4th European workshop on case-based reasoning, Dublin, Ireland, September 1998. Lecture notes in computer science, vol 1488. Springer, Berlin, pp 334–345
34. Macura R, Macura K (1995) Macrad: Radiology image resource with a cased-based retrieval system. In: Veloso M, Aamodt A (eds) Proceedings of the 1st international conference on CBR. Springer, Berlin, pp 43–54
35. Gierl L, Stengel-Rutkowski S (1994) Integrating consultation and semi-automatic knowledge acquisition in a prototype-based architecture: experiences with dysmorphic syndromes. Artif Intell Med 6:29–49
36. Schmidt R, Montani S, Bellazzi R, Portinale L, Gierl L (2001) Case-based reasoning for medical knowledge-based systems. Int J Med Inf 64(2–3):355–367
37. Bichindaritz I (2006) Memoire: A framework for semantic interoperability of case-based reasoning systems in biology and medicine. Artif Intell Med 36:177–192
38. Funk P, Nilsson M, Xiong N (2005) Knowledge discovery and case based reasoning in medical applications with time series. In: Proceedings of the of the third workshop on case based reasoning in the health sciences, ICCBR 2005. DePaul University, Chicago, pp 42–51
39. Ramoni M, Sebastiani P (1997) The use of exogenous knowledge to learn Bayesian Networks for incomplete databases. In: Liu X, Cohen P, Berthold M (eds) Advances in intelligent data analysis. Springer, Berlin, pp 537–548
40. Hovorka R, Svacina S, Carson ER, Williams CD, Sönksen PH (1996) A consultation system for insulin therapy. Comput Methods Programs Biomed 32:303–310
41. Nucci G et al. (1999) Verification phase final report, T-IDDM deliverable 5.2. http://aim.unipv.it/projects/tiddm/ftp.html
42. The Diabetes Control and Complication Trial Research Group (1993) The effect of intensive treatment of diabetes on the development and progression of long-term complications in insulin-dependent diabetes mellitus. New Engl J Med 329:977–986
43. Montani S, Bellazzi R (2000) Exploiting multi-modal reasoning for knowledge management and decision support: an evaluation study. In: Journal of the American medical informatics association (JAMIA) symposium supplement, pp 585–589
44. Bellazzi R, Larizza C, Lanzola G (1999) An http-based server for temporal abstractions. In: Proceedings of the IDAMAP'99, pp 52–62
45. Cases A, Coll E (2002) Chronic hypotension in the dialysis patient. J Nephrol 15:331–335
46. Tisler A, et al (2002) Comparison of dialysis and clinical characteristics of patients with frequent and occasional hemodialysis-associated hypotension. Kidney Blood Press Res 25:97–102
47. Portinale L, Montani S, Bottrighi A, Leonardi G, Juarez J (2006) A case-based architecture for temporal abstraction configuration and processing. Technical Report TR-INF-2006-05-02, Dipartimento di Informatica, Universita' del Piemonte Orientale, Alessandria, Italy