# Artificial Intelligence Techniques for Diabetes Management: the T-IDDM Project

**Stefania Montani**[1]**, Riccardo Bellazzi**[1]**, Alberto Riva**[1]**, Cristiana Larizza**[1]**,**
**Luigi Portinale**[2]**,** and **Mario Stefanelli**[1]

**Abstract.** We present a successful application of Artificial Intelligence (AI) methodologies in the context of a telemedicine service for diabetic patients management, developed within the EU-funded T-IDDM project. The system architecture is distributed, and composed by a Patient Unit and by a Medical Unit, connected through a telecommunication link. Several AI methods have been exploited to implement the T-IDDM functionality. The data base relies on an explicit representation of the domain ontology. Temporal Abstractions and other Intelligent Data Analysis techniques are used to analyse the patient's monitoring data; the Case Based Reasoning (CBR) methodology is applied to perform the Knowledge Management task. Finally, CBR is integrated with Rule Based Reasoning to provide physicians with a multi-modal reasoning decision support tool. The T-IDDM service is being tested through a small on field trial in Pavia; the first results, though preliminary, seem to substantiate the hypothesis that the use of an AI-based telemedicine system could present an advantage in the management of type 1 diabetic patients, leading to a more tight control of the patients' metabolic situation, in a cost-effective way.

## 1 INTRODUCTION

Diabetes Mellitus is a major chronic disease, affecting up to 3% of the population in the industrialised countries. In particular, Insulin Dependent Diabetes Mellitus (IDDM) patients need exogenous insulin injections to regulate blood glucose metabolism, in order to prevent ketoacidosis and coma, and to reduce the risk of later life invalidating complications. It has been proved [1] that Intensive Insulin Therapy (IIT), consisting in 3 to 4 injections every day, or in the use of sub-cutaneous insulin pumps, is the most effective way to stabilise blood glucose, and therefore to reduce or delay IDDM complications; the increase in therapy planning complexity and in costs is the obvious drawback. IDDM management normally consist in visiting patients every 2/4 months; during these visits the data coming from home monitoring are analysed, in order to assess the metabolic control achieved by the patients. Laboratory results and historical and/or anamnestic data are verified as well, to finally revise the patient's therapeutic protocol. Diabetic patients management is obviously a complex task; how to balance the advantages coming from IIT with its disadvantages is a matter of discussion, that involves social and ethical considerations [2]. It has been advocated that the use of current advances of information technologies and decision-support systems may improve cost-effectiveness of IIT, by reducing the number of periodical control visits, while increasing the patient/physician communication rate. Several tools and advisory systems for therapeutic plan assessment are now available, both on a day-by-day and on a visit-by-visit basis [3], and for some of them the capability of providing proper decisions has been shown experimentally [4]. The exponential growth in the availability and in the use of telecommunication services pushes towards the integration of such tools in a networking environment, in order to provide long-distance assistance to patient, as well as long-distance monitoring capability to the physician [5]. The use of appropriate Artificial Intelligence (AI) techniques, such as knowledge based systems, Intelligent Data Analysis and Case Based Reasoning, may enhance the design of the overall service: it should be possible to allow the users exploiting an intelligent desk for periodic therapy assessment and revision. For these reasons, we have worked at the development of a telemedicine system, that, reinforced by previous experiences, is able to offer a new integrated solution to the IDDM management problem.

[1] Dipartimento di Informatica e Sistemistica, Università di Pavia, via Ferrata 1, I-27100 Pavia, Italy, email: stefania@aim.unipv.it
[2] Dipartimento di Scienze e Tecnologie Avanzate, Università del Piemonte Orientale, Alessandria, Italy

## 2 THE T-IDDM PROJECT

The motivations described in section 1 led to the definition of the EU funded project T-IDDM (Telematic Management of Insulin-Dependent Diabetes Mellitus) [6], concerned with the design, implementation and testing of an intelligent telemedicine service to assist IDDM patients, able to provide physicians with a collection of AI techniques for improving management of patients according to the best current medical practice. The project goal was to accomplish the following specific aims: (i) to provide patients with an effective treatment leading to good glycemic control, and to achieve a careful balance between insulin therapy, diet and physical activity, thus delaying the onset and/or slowing the progression of chronic complications; (ii) to provide patients at home or in other non-clinical environments with an appropriate level of continuous and intensive care through telemonitoring and teleconsultation services, taking into account the needs of remote or isolated individuals that are unable to reach frequently the hospital institutions; (iii) to allow for a cost-effective monitoring of a large number of patients, automatizing data collection and the management of a large set of therapeutic protocols; (iv) to support a continuous education of patients through teleconsultation services; (v) to allow the patient to customise the insulin therapy within the bounds established by the physicians. The T-IDDM service exploits two main components: a Patient Unit (PU) and a Medical Unit (MU), connected through a telecommunication system (Internet or the Public Switched Telephone Network). Patients collect metabolic data together with insulin and food intake information every day, and store them in the PU data base. Whenever a problem occurs, they can send an alarm and the data to the MU to get a quick response from the physician, otherwise all the information will be sent every 7-10 days. The MU is a web-based workstation (a demo is available at http://aim.unipv.it/projects/tiddm) in which several distributed servers cooperate in a transparent way to help the physician in managing diabetes mellitus through: consultation and analysis of the patients' data, communication with patients' home, revision of the therapeutic protocol and information repositories consultation. T-IDDM was a three-year project, which ended in October 1999. The T-IDDM service has been implemented and preliminary tested through an on field trial that started in June 1998. It is routinely used by six pediatric patients and by two diabetologists at the Policlinico S. Matteo hospital in Pavia.

## 3 ROLE OF AI IN T-IDDM

Within the T-IDDM project, several AI techniques have been exploited:

- the structure of the MU data base tables has been automatically generated on the basis of the domain ontology, stored in the knowledge base;
- Intelligent Data Analysis (IDA) is applied to analyse the patient's monitoring data, sent by the PU; some indicators of the patient's metabolic behavior can then be extracted. A Rule Based Reasoning (RBR) system, meant to provide the physician with therapy planning suggestions, relies on such indicators to identify possible metabolic alterations;
- in order to take advantage from the amount of information contained in the MU data base, and from the "operative" knowledge of experts, embedded in past cases, we have implemented a Knowledge Management (KM) tool, based on the Case Based Reasoning (CBR) methodology;
- finally, a multi-modal reasoning system, that specialises the RBR behavior with the results of Case Based retrieval, has been defined and made available to physicians.

### 3.1 Domain ontology

In the case of a distributed system, the existence of a global, shared ontology is essential to ensure the possibility of communication between the architecture components: the ontology acts as the common terminology to which all the modules refer when exchanging information, and is used to determine the behavior of the whole system. The T-IDDM ontology is stored in a knowledge base built using a *frame* system [7, 8], that supports multiple inheritance and typed slots. It is organised into taxonomies, which describe entities (e.g. patients, laboratory values), events (e.g. monitoring data measurements), abstractions (e.g. hyperglycemia), drugs, therapeutic protocols, and so on. Since slots are typed, we also needed to define a hierarchy of classes representing types (e.g. the class of numerical values, and the one of fixed-length-string values). Types are *abstract* classes: no instances are created from them, but they are used to store information needed by the instance slots.

We use the knowledge base mainly as a declarative tool, to describe the domain ontology, while we exploit a relational data base to manage the actual data. The structure of the data base tables, as well as the commands to store and retrieve the data, are generated automatically on the basis of the ontology information. A specialised SQL interface, able to receive the queries addressed to the data base, and to return the resulting data, ensures the communication between the user and the data repository [9].

### 3.2 Intelligent Data Analysis

The data collected by patients during home-monitoring, and sent through the PU-MU connection, are time-stamped, and acquired several times (from three to four) a day. In order to allow a proper interpretation of the data, we have subdivided the 24-hour daily period into a set of consecutive non-overlapping time slices, centered on the time of meals or insulin injections; each datum is hence associated to a given time slice.

The MU exploits the home monitoring information through a set of tools to visualise and to analyse collected data. Data analysis ranges from a set of statistical methods, such as the extraction of the daily average value of Blood Glucose Level (BGL), the daily insulin requirement, and the number of serious hypoglycemic events in a given period of time, to more complex forms of abstractions. In particular an abstract description of the course of longitudinal data is obtained through the Temporal Abstractions (TA) technique.

The basic principle of TA methods is to move from a time-point to an interval-based representation of the data. Given a sequence of time stamped data (*events*), the adjacent observations which share a common feature are aggregated into intervals (*episodes*). In more detail, two main classes of abstractions can be described: *basic* abstractions for detecting predefined patterns in a univariate time series, and *complex* abstractions, for discovering specific temporal relationships between episodes as well as for analysing multivariate patterns [10]. Basic abstractions extract *states* (e.g. low, normal, high values) or *trends* (increase, decrease or stationarity patterns) from a uni-dimensional time series. Complex abstractions search for specific temporal relationships between episodes which can be generated from a basic abstraction or from other complex abstractions. This kind of TA can be exploited to extract multi-dimensional patterns or to detect uni-dimensional patterns of complex shapes. Exhaustive examples of the application of TA methods to diabetic patients monitoring data can be found in [11].

The decision support activity of the MU is based on the extraction of state abstractions, dealing in particular with BGL measurements. When detecting state patterns in time series of numerical variables, a preliminary qualitative abstraction is carried out. The mapping between the qualitative abstractions and the quantitative levels of each numerical variable depends on the time slice and on the specific patient's characteristics. For example, the BGL normal range is wider in the morning than around lunch, and it is wider in pediatric patients than in adult ones. Then, the BGL state abstractions are derived, moving from the original time scale to a new scale obtained from the sequence of relevant patterns detected in the data.

After having identified the most significant episodes, the BGL Modal Day can be extracted. The Modal Day is an indicator, well known in the literature, able to summarise the mean response of the patient to the therapy followed in the period under examination. It is usually derived as the collection of the most probable BGL qualitative level in each time slice. Several approaches for extracting the Modal Day have been presented, from simple statistics to time series analysis [12, 13]. Our choice has been the one of applying a Bayesian method described in [14] that is able to explicitly take into account the presence of missing data, i.e. missing measurements in one or more time slices.

Five BGL state TAs are considered, *Severe Hypoglycemia, Hypoglycemia, Normoglycemia, Hyperglycemia, Severe Hyperglycemia.*

Before starting data collection in a given period we may assign a prior probability to the occurrence of each state TA equal to $\frac{1}{5}$. After a certain monitoring period of $N$ days, we collect $D$ measurements, while the remaining $M = N - D$ data are missing. The posterior probability bounds of the occurrence of a generic $k$-th of the 5 levels, given by difference between the upper ($p_{sup}$) and the lower($p_{inf}$) probability bound, 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, for each time slice, an interval probability distribution over the BGL state abstractions. The modal day is extracted by taking the BGL states with the highest $p_{inf}$ in each time slice.

## 3.3 Knowledge Management

Managing all the data piling in the hospital data base, and extracting from them reliable information about the patients' status, is a complex problem, difficult to be satisfactorily solved. When dealing with chronic diseases management, one of the most effective instruments for the Knowledge Management (KM) task is Case Based Reasoning (CBR). CBR [15] is a problem solving paradigm that uses the "operative" knowledge of previously experienced situations, called cases. Past cases similar to the current one are retrieved and shown to the user; past solutions may be reused and applied (after, if necessary, an adaptation step) to the present situation. Case Based retrieval, in particular, by enabling physicians performing an intelligent consultation of the available case library, keeps track of the "problem/solution" patterns that occurred in the past, maintains a specific patient's history, and makes the overall health care organisation expertise available even in presence of changes in the physicians' staff.

Within the T-IDDM MU architecture, we have implemented a KM tool that exploits Case Based retrieval. We define a case as a set of feature-value pairs, together with a solution and an outcome. The case features are the data collected during a periodical visit. The solution is the therapeutic protocol assigned by the physician after the features examination, and the outcome of such therapy is given by the number of hypoglycemic episodes and by the value of HbA1c collected at the following visit (i.e. the following case). The case library structure strongly influences the case search; to make retrieval more flexible, through the collaboration with the diabetologists of the Pediatric Department of Policlinico S. Matteo in Pavia, we have been able to structure the library by resorting to a taxonomy of mutually exclusive prototypical classes, that express typical problems that may occur to patients in the age of infancy and puberty (see figure 1).

Retrieval is hence implemented as a two step procedure: a *classification* step, able to identify the class to which the current case could belong, and a proper *retrieval* step, meant to extract the "closest" cases. Classification relies on a Naive Bayes strategy, a method that assumes conditional independence among the features given a

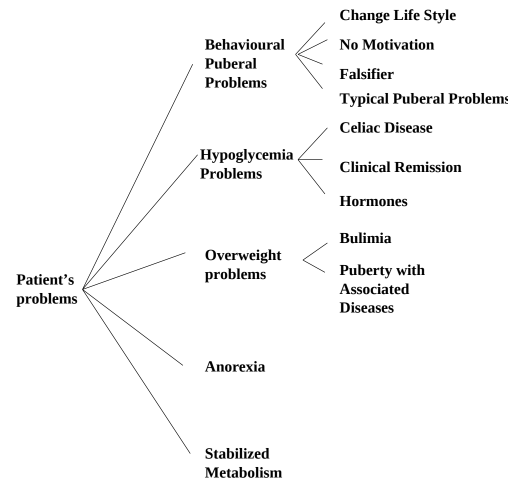


**Figure 1.** Taxonomy of classes of prototypical situations in pediatric IDDM patients

certain class, but that is known to be robust in a variety of situations [16, 17] even in the presence of conditional dependencies. For applying Naive Bayes, we calculate the probability that a case belongs to class $c_i$, given that the set of its features $f = \{f_1, \dots, f_M\}$ is $\hat{f}$, through the following formula:

$$P(c_i \mid f = \hat{f}) \quad \propto \quad \prod_{j=1}^{M} p(c_i) p(f_j = \hat{f}_j \mid c_i)$$

The method classifies a case as belonging to the class that maximises $P(c_i \mid f = \hat{f})$.
The conditional probabilities $p(f_j = \hat{f}_j \mid c_i)$ are obtained through the Bayesian update formula for discrete distributions [18, 19]; in particular, we use a re-parameterised version of the update formula known as *m-estimate* of probability [20], that modifies the prior knowledge with the information coming from the cases of the case memory as follows:

$$p(f_j = \hat{f}_j \mid c_i) \quad = \quad \frac{m\, \hat{p}_{ij} + \hat{N}_{ij}}{m + D_i}$$

where $\hat{N}_{ij}$ is the number of cases in the case memory of class $i$ whose feature $f_j$ assumes the value $\hat{f}_j$, while $D_i$ is the total number of cases in class $i$. The medical knowledge is synthesised by the prior probability distribution ($\hat{p}_{ij}$), whose reliability is expressed by the implicit number of samples $m$. In other words, the larger is $m$, the larger is the confidence of the expert on the prior. Our initial case library was composed by 145 real cases from the histories of 29 pediatric patients. In our application, the prior probability value ($p_{ij}$) was derived from experts opinion through a technique described in [21].

Retrieval may be performed just on the most probable class identified by the classification step, or on a subset of the most probable classes. In both situations the system relies on a Nearest-Neighbor (NN) technique and classical metrics, able to treat numeric and symbolic variables, and to cope with the problem of missing data, are applied to calculate distances [22]. When dealing with a large case-base, our application implements a non exhaustive search procedure that exploits an anytime algorithm called Pivoting-Based Retrieval (PBR) [23], whose efficacy has been proved on a 10000 cases library.

### 3.4 Decision Support

The CBR methodology has been exploited not only for the KM task, but also for decision support: we have provided the T-IDDM MU with a multi-modal reasoning system that relies on the integration between CBR and the RBR methodology. Such a tool is able to overcome the weakness of the two reasoning paradigms when applied independently. When developing a RBR system, in fact, it is possible to incur in the so-called *qualification problem*: trying to deal with as much peculiar situations as possible leads to the definition of a huge rule base, reducing the system performances. On the other hand, the case library of a CBR system could be biased on too specific examples, and may present some regions not covered by a sufficient number of cases (*competence gaps*).

Our approach is strongly innovative from a methodological point of view: we have realised a tighter integration in respect to other examples that can be found in the literature [24, 25, 26], as it takes place within the general problem solving cycle (see figure 2). Our tool is able to exploit Case Based retrieval results to specialise the rules behavior, and to dynamically adapt it to the specific patients characteristics.

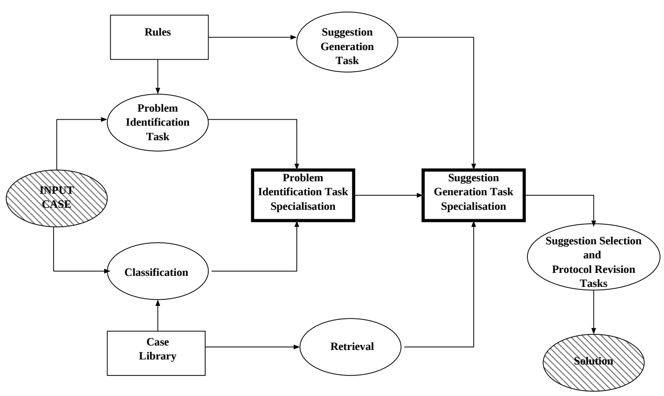


**Figure 2.** CBR-RBR integration within the T-IDDM Medical Unit

The RBR component embeds the domain knowledge into a taxonomy of production rule classes, fired through a forward chaining mechanism. Each rule class performs an action [27]:
- **data analysis and problem identification**: after having computed the BGL Modal Day (see section 3.2), the system identifies hypoglycemia or hyperglycemia problems. In particular, when $p_{inf}$ is higher than a given threshold *alpha*, and the *ignorance* is smaller of a given threshold *beta*, a problem is recognised;
- **suggestion generation**: for each detected problem, a set of alternative suggestions, dealing with insulin therapy, diet or physical exercise, is generated;
- **suggestion selection**: the most suitable and effective suggestions are selected and applied to the current therapeutic protocol;
- **protocol revision**: the adjusted protocol, together with other library protocols suitable for the situation at hand, is listed to physician for her/his final judgment.

CBR results are integrated in the RBR framework by means of a rule refinement process involving the change of suitable rule parameters on the basis of information obtained from the case library (i.e. classification and retrieval). In particular, the rule classes dealing with **problem identification** and **suggestion generation** are affected by the integration procedure, following a series of steps:
- the Bayesian classifier is invoked on the patient's visit data;
- if the physician chooses to rely on the classification results, the most probable class information is used to specialise the **problem identification** rules parameters *alpha* and *beta*;
- the physician may want to complete the classification information with retrieval results. A test on the retrieved cases is performed; only cases whose protocol has the same injection number of the input case will be considered. If no such a case is retrieved, RBR is applied without CBR integration;
- among the remaining cases, the tool just exploits the ones with a positive outcome (i.e. cases for which the applied protocol has resulted in a low number of hypoglycemic events and in a HbA1c decreasing trend). On them, it computes some descriptive statistics, to set parameters such as the number of insulin injections, the variation in the daily insulin requirement and the variation of a single insulin dose, thus specializing the **suggestion generation** rules class.

The result of the previous rule specialisation is a list of refined suggestions, which can be used to complete the RBR cycle, for providing the physician with a final outcome. On the other hand, CBR can take advantage from the results of RBR as well. When no suitable case is retrieved by the CBR component (either because no positive outcome is found or because the retrieved therapeutic protocols are significantly different from the current one), we can infer that the input case belongs to a competence gap region. In such a situation, our tool performs RBR without integration, in order to avoid wrong specialisation due to misleading cases. As soon as the outcome of the proposed protocol is available (normally at the next periodical visit), a new case is learnt, and stored in the memory, to fill the competence gap.

## 4 RESULTS OF THE T-IDDM SERVICE IN PAVIA

The T-IDDM demonstration phase, conducted in Pavia from June 1998 to October 1999, involved six pediatric patients and a diabetologist from Policlinico S. Matteo hospital. Although this small number does not allow to obtain statistically significant results, there are several interesting issues that arise from the collected data [28]. First, the high frequency of logging on the system, and the high number of messages exchanged (56 messages sent from the physician and 35 messages sent from patients in an average period of 415 days) demonstrates that both patients and the physician have found helpful to rely on the system's functionality, and that the system is quite easy to use, having a sufficiently user-friendly interface. Moreover, the average number of therapy changes has been augmented with respect to clinical practice; while in normal clinical practice such changes would be up to 6, during the demonstration phase it reached the value of 8. In fact, when working without the system support, the physician just evaluates the patient's condition on the visit day (every 2/4 months), while T-IDDM has allowed him to verify the metabolic state more frequently. In such a way, if a problem arose, the physician was immediately able to make a therapy revision to cope with it, helped by the AI-based tools, and in particular by the IDA techniques embedded in the MU, that he largely exploited. We believe that such an increased frequency in therapy adjustments has led to a reduction of the mean HbA1c level in the patients under monitoring (HbA1c variation range: [-6.7 +1.1]; average: -1.23): helped by the system, the physician has been able to better and sooner understand the patients' peculiar needs, and to define insulin treatments properly tailored on the specific situation under examination; additionally a statistically significant reduction in insulin requirement has taken place (variation range: [-0.15 +0.08]; average: -0.03; $p < 0.03$ with Wilcoxon's test for paired data).

To test the patients compliance, some questionnaires about the use of T-IDDM were distributed. The general technical judgment on the PU software was positive for all patients. The usability questionnaire answers presented more heterogeneous results: while the PU was recognised by all patients to be easy to use, nice and efficient enough, the help provided by the system for performing self monitoring has been considered acceptable only by some patients. Such results, although obtained on a very small number of subjects, give the suggestion that a better compliance could be obtained by designing a PU configurable on the single user's needs.

Finally, the reliability and the soundness of the RBR system suggestions have been positively tested within the T-IDDM verification phase; the system, even if a bit too conservative, generally mimics the expert's behavior.

## 5 CONCLUSIONS

The encouraging results we obtained from the T-IDDM project verification phase in Pavia, though preliminary, seem to substantiate the hypothesis that the use of the telemedicine systems in association with decision support systems could present an advantage in the management of type 1 diabetic patients, leading to a more tight control of patients' metabolic situation, in a cost-effective way. As the system is currently used at the pediatric clinic, we will be able to collect additional data which, in our opinion, will probably enforce such conclusions. As a final step of the T-IDDM evaluation procedure, we plan to verify the reliability and the correctness of the multi-modal reasoning tool, by implementing the following verification protocol: i) two physicians, who never used the T-IDDM MU decision support functionality, will analyse a set of real patients cases; ii) a panel of expert physicians, again unfamiliar with T-IDDM, for each case, will compare the three available therapeutic prescriptions (two coming from the colleagues, and the other from the multi-modal reasoning system), without knowing who is the author of the various answers. By applying such methodology, we foresee to understand if the T-IDDM multi-modal reasoning system mirrors the reasoning of an expert diabetologist, other than those who provided the knowledge. At the same time we will test whether there are conflicting opinions among the physicians, and we will find out in how many cases there is a complete inter-expert consensus. Finally, concepts of the T-IDDM project will be applied within a project funded by the EU in the Fifth Framework Programme, called M$^2$DM (Multi-Access Services for the management of Diabetes Mellitus).


## ACKNOWLEDGEMENTS

Dr. Giuseppe d'Annunzio is gratefully acknowledged his support in the CBR and in the RBR systems definition, and in the EU TAP project T-IDDM (HC 1047) on field testing.


## REFERENCES


[1] 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**, 977–986, (1993).

[2] R.D. Lasker, 'The diabetes control and complication trial. Implications for policy and practice', *The New England Journal of Medicine*, **329**, 1035–1036, (1993).

[3] E.D. Lehmann and T. Deutsch, 'Application of computers in diabetes care - a review (I and II)', *Med. Inform.*, **20**, 281–329, (1995).

[4] S. Andreassen et al., 'A probabilistic approach to glucose prediction and insulin dose adjustment: description of metabolic model and pilot evaluation study', *Computer Methods and Programs in Biomedicine*, **41**, 153–165, (1994).

[5] E.J. Gomez et al., *A Telemedicine Distributed Decision-Support System for Diabetes Management*, 1238–1239, IEEE- 14th Ann. Int. Conf. of the IEEE Eng.in Med. and Biol. Soc., 1992-CH3207-8, 1992.

[6] R. Bellazzi, C. Cobelli, E. Gomez and M. Stefanelli, *The T-IDDM project: Telematic management of Insulin Dependent Diabetes Mellitus*, 271–276, Proceedings of Health Telematics '95, M. Bracale, F. Denoth eds., Ischia, 1995.

[7] M.L. Minsky, *A framework for representing knowledge*, The psychology of computer vision, P.H. Winston ed., pp. 211-277, Mc Graw Hill, 1975.

[8] S.G. Pauker, G.A. Gorry, J.P. Kassirer, W.B. Schwartz, 'Toward the simulation of clinical cognition: taking a present illness by computer', *Americal Journal of Medicine*, **60**, 981–995, (1976).

[9] A. Riva, R. Bellazzi, G. Lanzola, M. Stefanelli, 'A development environment for knowledge-based medical applications on the world-wide web', *Artificial Intelligence in Medicine*, **14**, 279–293, (1998).

[10] R. Bellazzi, C. Larizza, A. Riva, *Cooperative intelligent data analysis: an application to diabetic patients management*, Intelligent data analysis in medicine and pharmacology, N. Lavrac, E. Keravnou, B. Zupan eds., Kluwer, 1997.

[11] R. Bellazzi, C. Larizza, A. Riva, 'Temporal abstractions for interpreting chronic patients monitoring data', *Intelligent Data Analysis - an International Journal*, **2**, (1998), http://www.elsevier.com/locate/ida.

[12] M.G. Kahn et al., 'Intelligent computer-based interpretation and graphical presentation of self-monitored blood glucose and insulin data', *Diab. Nutr. Metab.*, **4**, 99–107, (1991).

[13] T. Deutsch et al., 'Time series analysis and control of blood glucose levels in diabetic patients', *Computer Methods and Programs in Biomedicine*, **41**, 167–182, (1994).

[14] M. Ramoni and P. Sebastiani, *The use of exogenous knowledge to learn Bayesian Networks for incomplete databases*, 537–548, Advances In Intelligent Data Analysis, X. Liu, P. Cohen and M. Berthold eds., Springer Verlag, Berlin, 1997.

[15] J.L. Kolodner, *Case-Based Reasoning*, Morgan Kaufmann, 1993.

[16] I. Kononenko, 'Inductive and Bayesian learning in medical diagnosis', *Applied Artificial Intelligence*, **7**, 317–337, (1993).

[17] I. Zelic, I. Kononenko, N. Lavrac, V. Vuga, *Induction of decision trees and Bayesian classification applied to diagnosis of sport injuries*, 61–67, Proceedings of IDAMAP 97 workshop, IJCAI 97, Nagoya, 1997.

[18] D. Spiegelhalter, A. Dawid, S. Lauritzen, R. Cowell, 'Bayesian Analysis in Expert Systems', *Statistical Science*, **8**, 219–283, (1993).

[19] A. Riva, R. Bellazzi, 'Learning temporal probabilistic causal models from longitudinal data', *Artificial Intelligence in Medicine*, **8**, 217–234, (1996).

[20] T. Mitchell, *Machine Learning*, Mc Graw Hill, 1997.

[21] S. Montani et al., *A case-based retrieval system for diabetic patients therapy*, 64–70, Proceedings of IDAMAP 98 workshop, ECAI 98, Brighton, 1998.

[22] D.R. Wilson, T.R. Martinez, 'Improved heterogeneous distance functions', *Journal of Artificial Intelligence Research*, **6**, 1–34, (1997).

[23] L. Portinale, P. Torasso, D. Magro, *Selecting most adaptable diagnostic solutions through pivoting-based retrieval*, 277–288, Lecture Notes in Artificial Intelligence 1266, Springer Verlag, 1997.

[24] J. Surma, K. Vanhoff, *Integrating rules and cases for the classification task*, 325–334, Proceedings of the first ICCBR, Lecture Notes in Artificial Intelligence 1010, Sesimbra, 1995.

[25] I. Bichindaritz et al., 'CARE-PARTNER: a computerized knowledge-support system for stem-cell-post-transplant long-term-follow-up on the World-Wide-Web', *JAMIA, Symposium supplement 1998*, 386–390, (1998).

[26] D.B. Leake, *Combining rules and cases to learn case adaptation*, Proceedings of the 17th Annual Conference of the Cognitive Science Society, Pittsburgh, 1995.

[27] S. Montani et al., 'Protocol-based reasoning in diabetic patient management', *International Journal of Medical Informatics*, **53**, 61–77, (1999).

[28] S. Montani et al. Demonstration Phase and Demonstration of the Extensions Final Report. T-IDDM Deliverable 11.1, http://aim.unipv.it/projects/tiddm, 1999.