



ELSEVIER

International Journal of Medical Informatics 64 (2001) 355–367

International Journal of
**Medical
Informatics**

www.elsevier.com/locate/ijmedinf

Cased-Based Reasoning for medical knowledge-based systems

Rainer Schmidt ^{a,*}, Stefania Montani ^b, Riccardo Bellazzi ^b, Luigi Portinale ^c,
Lothar Gierl ^a

^a *Institute for Medical Informatics and Biometry, University of Rostock, Rembrandtstrasse 16/17, 18055 Rostock, Germany*

^b *Dipartimento di Informatica e Sistemistica, University of Pavia, Pavia, Italy*

^c *Dipartimenti di Scienze e Tecnologie Avanzate, Universita del Piemonte Orientale, Alessandria, Italy*

Abstract

In this paper we present the results of the MIE/GMDS-2000 Workshop ‘Case-Based Reasoning for Medical Knowledge-based Systems’. While in many domains Cased-Based Reasoning (CBR) has become a successful technique for knowledge-based systems, in the medical field attempts to apply the complete CBR cycle are rather exceptional. Some systems have recently been developed, which on the one hand use only parts of the CBR method, mainly the retrieval, and on the other hand enrich the method by a generalisation step to fill the knowledge gap between the specificity of single cases and general rules. And some systems rely on integrating CBR and other problem solving methodologies. In this paper we discuss the appropriateness of CBR for medical knowledge-based systems, point out problems, limitations and possible ways to cope with them. © 2001 Elsevier Science Ireland Ltd. All rights reserved.

Keywords: Cased-Based Reasoning (CBR); Medical knowledge-based systems; Multi Modal Reasoning

1. Introduction

Cased-Based Reasoning (CBR) has become a successful technique for knowledge-based systems in many domains, while in the medical field some more problems still arise. In this paper, we are going to discuss the appropriateness of CBR for medical knowledge-

based systems and to point out its problems, limitations and possible ways how they can partly be overcome.

Case-Based Reasoning means to retrieve former, already solved problems similar to the current one and to attempt to modify their solutions to fit for the current problem (Fig. 1 shows the Cased-Based Reasoning cycle developed by Aamodt [1]). The underlying idea is the assumption that similar problems have similar solutions. Though this assumption is not always true, it holds for many practical domains.

* Corresponding author. Tel.: +49-381-394-7307; fax: +49-381-394-7203.

E-mail address: rainer.schmidt@medizin.uni-rostock.de (R. Schmidt).

CBR fulfils two main tasks [1,2]: the first is the retrieval, which means to search for or to calculate the most similar cases. If the case base is rather small, a sequential calculation is possible, otherwise faster non-sequential indexing [2,3] or classification algorithms (e.g. ID3 [4] or Nearest Neighbour match [5]) should be applied. For this task much research has been undertaken in the recent years and for nearly every sort of application problem it has actually become correspondingly easy to find a suitable, sophisticated CBR retrieval algorithm. The second task, the adaptation (reuse and revision), means a modification of solutions of former similar cases to fit for a current one. If there are no important differences between a current and a similar case, a simple solution transfer is sufficient. Sometimes only few substitutions are required, but in other situations the adaptation is a very complicated process. So far, no general adaptation methods or algorithms have been developed. The adaptation is still absolutely domain dependent.

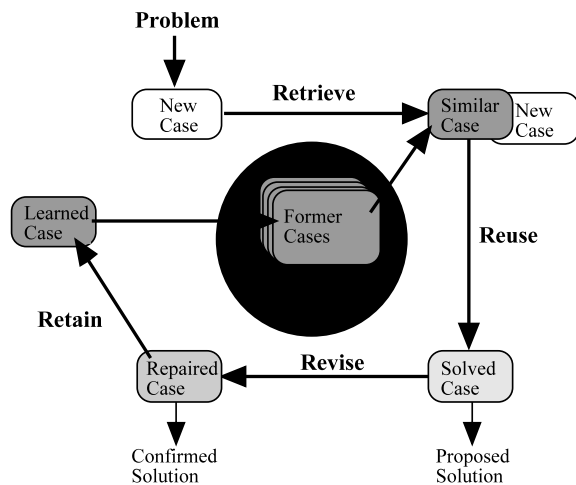


Fig. 1. The Cased-Based Reasoning cycle developed by Aamodt.

1.1. Why Cased-Based Reasoning for medical decision making?

Especially in medicine, the knowledge of experts does not only consist of rules, but of a mixture of textbook knowledge and experience. The latter consists of cases, typical and exceptional ones, and the reasoning of physicians takes them into account [6]. Medical knowledge based systems therefore contain two knowledge types: objective knowledge, which can be found in textbooks, and subjective knowledge, which is limited in space and time and changes frequently.

Both sorts of knowledge can clearly be separated: objective textbook knowledge can be represented in forms of rules or functions, while subjective knowledge is contained in cases. The problem of updating the changeable subjective knowledge can partly be solved by incrementally incorporating new up-to-date cases [6].

So, the arguments for the use of case-oriented methods can be summarised as follows:

1. Reasoning with cases corresponds with the typical decision making process of physicians.
2. Incorporating new cases means automatically updating parts of the changeable knowledge.
3. Objective and subjective knowledge can be clearly separated (of course they can be used together in one system).
4. As cases are routinely stored, integration into a Hospital Information System (HIS) is easy.

2. Medical Cased-Based Reasoning systems

In medicine, CBR has mainly been applied for diagnostic and partly for therapeutic tasks. Related methods have been used in further fields: case-oriented methods for tu-

toring (e.g. D3 [7]) and retrieval methods to search for similar images (e.g. MACRAD [8]). In this paragraph just three medical case-based decision support systems are mentioned. For further systems we refer to Ref. [9].

One of the earliest medical decision support systems that applies CBR is CASEY [10]. It deals with heart failure diagnosis. The systems reasoning functionality follows three steps: a search for similar cases, a determination process concerning differences and their evidences between a current and a similar case, and a transfer of the diagnosis of the similar to the current case or—if the differences between both cases are too important—an attempt to explain and modify the diagnosis. If no similar case can be found or if all modification attempts fail, CASEY uses a rule-based domain theory. The most interesting aspect of CASEY is the ambitious attempt to solve the adaptation task by general adaptation operators. However, as many features have to be considered in the heart failure domain and as consequently many differences between cases can occur, not all differences between former similar and current cases can be handled by the developed general adaptation operators.

The FLORENCE system [11] deals with health care planning in a broader sense, for nursing, which is a less specialised field. It fulfils all three basic planning tasks: diagnosis, prognosis, prescription. Diagnosis is not used in the common medical sense as the identification of a disease, but it seeks to answer the question: “What is the current health status of this patient?” Rules concerning weighted health indicators are applied. The health status is determined as the score of the indicator weights. Prognosis seeks to answer the question: “How may the health status of this patient change in the future?” Here a Cased-Based approach is used. The

current patient is compared to a similar previous patient for whom the progression of the health status is known. Similar patients are searched for first concerning the overall status and subsequently concerning the individual health indicators. As the further development of a patient not only depends on his situation (current health status, basic and present diseases), but additionally on further treatments, several individual projections for different treatments are generated. Prescription seeks to answer the question: “How may the health status of this patient be improved?” The answer is given by using general knowledge about likely effects of treatments and also by considering the outcome of using particular treatments in similar patients. That means it is a combination of a rule-based and a Cased-Based approach.

The most interesting aspect of medic [12] is its memory organisation. MEDIC is a schema-based diagnostic reasoner on the domain of pulmonology. Schemata represent the problem solvers knowledge. These are packets of procedural knowledge about how to achieve a goal or a set of goals. The memory does not only consist of schemata, but additionally of diagnostic memory organisation packets of individual cases of diagnosis and of scenes. A scene represents an instantiation of a schema in a particular case. This memory organisation and retrieval allows a reasoner to find the most specific problem-solving procedures available.

3. Problems of Cased-Based Reasoning for medical applications

To use Cased-Based Reasoning a few problems have to be solved: a representation form for cases has to be determined, and an appropriate retrieval algorithm has to be selected. Moreover, an infinite growth of the case base

should be avoided e.g. by clustering cases into prototypes and removing redundant ones, or by restricting the case base to a fixed number of cases and updating it during expert consultation sessions [8]. However, the main problem of the CBR method is the adaptation task. Little research has been undertaken on this topic and only formal adaptation models [13], but no general methods have been developed so far. The adaptation still depends on domain and application characteristics. Sometimes no adaptation is necessary, because e.g. the field and the cases are as unspecialised as in FLORENCE. Sometimes the adaptation is a simple solution transfer or only a little bit more, sometimes just a few constraints have to be checked (e.g. in GS.52 [14]), but in other situations many differences between current and former similar cases have to be considered (e.g. in CASEY). Adaptation is not only a problem for medical applications. However, in medicine it increases, because cases often consist of an extremely large number of features. In non-medical CBR applications, the adaptation is usually solved by a set of specific adaptation rules, which usually have to be acquired during expert consultation sessions. As these rule sets have to consider all possible important differences between current and former similar cases, for medical applications it is mostly impossible to generate such sets. So, some adaptation solutions have been developed that are not limited to, but are rather typical for medical domains.

3.1. Focusing on retrieval

An idea to avoid the adaptation problem is to build retrieval-only systems. These are programs that only retrieve similar cases and present them as information to the user. Some of them additionally point out important differences between current and similar

cases. The justification for giving up the adaptation task is that in some application domains it is much too complicated or even impossible to acquire adaptation knowledge [15] and that physicians are interested to get information about former similar cases, but wish to reason about current patients themselves [8]. Examples of successful retrieval-only systems are mainly in the fields of images [8] and of organ function courses [16].

3.2. Multi Modal Reasoning

Multi Modal Reasoning represents another way to avoid the adaptation problem, mainly by combining CBR retrieval with other reasoning methodologies, to provide decision support. The interest in multi modal approaches involving CBR is recently increasing in different application areas [17,18] including the medical one [19,20]. Different reasoning methods can be combined in the same application, or one form of reasoning can be used to support another, or it can be possible to switch among alternative reasoning paradigms. CBR is well suited for integration with Rule Based Reasoning (RBR) or Model Based systems. Particular attention has received the combination of CBR with RBR, since rules are the most common explicit knowledge representation formalism for intelligent systems.

Different levels of integration between RBR and CBR are possible. Usually RBR and CBR are applied in mutually exclusive situations, where RBR deals with knowledge on standard or typical problems, while CBR faces exceptions. In this view, RBR is usually applied first; when it fails to provide the user with a reliable solution, CBR allows one to retrieve similar cases from a library of peculiar and non-standard situations [20–22]. Sometimes, RBR can be applied to routine problems, MBR to more complex ones, and

CBR on a few remaining cases, to improve system performances [23]. Other approaches suggest to making use of the differences in generality between rules and cases. Rules are used as an ‘abstract’ description of a situation, while cases represent a further ‘specialisation’. Cases assist RBR by instantiating and by providing suitable contexts to rules, while rules assist CBR by permitting the extraction of more general concepts from concrete examples [24]. The resulting architecture may be more flexible than previously, as it is possible to decide ‘a priori’ which method should be applied first, or to select the most convenient one in a dynamic way, depending on the situation at hand [19,24]. 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 [19]. Finally, RBR can support CBR just in the adaptation phase, by providing some general adaptation rules [25].

3.3. Generalised cases

As one reason for the adaptation problem is the extreme specificity of single cases, an idea is to generalise from single cases into abstracted prototypes or classes [19]. Though the main ideas for this generalisation are to structure the case base, to decrease the storage amount by erasing redundant cases, to speed-up the retrieval and sometimes to learn more general knowledge, additionally it can at least partly help to solve the adaptation problem. An example is the diagnostic system for dysmorphic syndromes GS.52 [14] described in Section 4.3.

The idea to partly solve the adaptation task by generalising can only work for diagnostic tasks where abstracted typical cases

represent diagnoses and additional specific features of former single cases can be neglected. Abstracted cases fill the gap between general rules and specific cases. If a hierarchy of abstracted cases exists (as in MEDIC), adaptation can be seen as a top down search to find the most specific case that fits for the current problem [13].

4. Examples

4.1. Retrieval-only: time course prognoses of the kidney function

As intensive care patients are often no longer able to maintain adequate fluid and electrolyte balances due to impaired organ functions or because they are ventilated, physicians need objective criteria for the monitoring of the kidney function and to diagnose therapeutic interventions as necessary. At our intensive care unit the renal function monitoring system NIMON [26] was developed that daily prints a renal report that consists of 13 measured and 33 calculated parameter values. However, the interpretation of all reported parameters is quite complex and needs special knowledge of the renal physiology. Our aim was to develop a system, called ICONS [16], that gives an automatic interpretation of the renal state to elicit impairments of the kidney function on time. In the domain of fluid and electrolyte balance, neither prototypical courses in ICU settings are known nor exists complete knowledge about the kidney function. So we had to design our own method to deal with course analyses of multiple parameters.

The method consists of three main steps: two data abstractions plus CBR retrieval. We have got the idea of abstracting many parameters into one single parameter from RÉSUMÉ [27] where the course of this single

parameter is analysed by means of a complete domain theory. The comparison of parameter courses with well-known course pattern is performed in some medical knowledge based systems ([28] and in VIEVENT [29]). As no such pattern are yet known for the kidney function, we use single courses and incrementally learned prototypes instead of well-known course pattern to compare with. We attempt to learn course pattern by structuring the case base by prototypes.

As the interpretation of all NIMON parameters is too complex, we decided to abstract them. For this data abstraction we have defined states of the renal function which determine states of increasing severity starting with a normal kidney function and ending with a renal failure. Based on these definitions, we ascertain the appropriate state of the kidney function per day. Based on the sequence of assessments of transitions of the state of a day to the state of the, respectively next day, we generate four different trends. These trends describe courses of states. Subsequently, we use Case-Based Reasoning retrieval to search for similar courses. We present the current course in comparison to similar ones to the user, the course continuations of the similar courses serve as prognoses (Fig. 2). As there may be too many different aspects between both patients, the adaptation of a similar to the current course is not done automatically. ICONS [16] offers only diagnostic and prognostic support, the user has to decide about the relevance of all displayed information (e.g. additional renal syndromes and courses of single kidney function parameter values).

4.1.1. Retrieval

The parameters of the trend descriptions are used to search for similar courses. Since the aim is to develop an early warning sys-

tem, a prognosis is needed. Since there are many different possible continuations for the same previous course, it is necessary to search for similar courses and different projections. Therefore, we have divided the search space into nine parts corresponding to the possible continuation directions within 3 days. Each direction forms an own part of the search space. During the retrieval these parts are searched separately and each part may provide at most one similar course. The retrieval consists of two steps for each projection part. First we search with an activation algorithm [30] concerning qualitative features. Subsequently, we check the retrieved cases with an adaptability criterion that looks for sufficient similarity, since even the most similar course may differ from the current one significantly. If several courses are selected in the same projection part, in a second step a sequential similarity measure concerning the quantitative features is used. It is a variation of TSCALE [31] and goes back to Tversky [32].

4.2. Multi Modal Reasoning: managing diabetic patients

At the University of Pavia we have developed a Multi Modal Reasoning (MMR) methodology, that performs a tight integration of CBR, RBR and MBR, with the aim of suggesting a therapy properly tailored on the single patient's needs, in the field of type 1 diabetic patients management. This methodology allows the exploitation of the implicit knowledge embedded in patients' visits (past cases) and in monitoring data, respectively through Case-Based retrieval and model identification. On the other hand the explicit domain knowledge is formalised in a set of production rules, and in the resulting model itself.




Fig. 2. Comparative presentation of a current and a similar course. In the lower part of each course the (abbreviated) kidney function states are depicted. The upper part of each course shows the deduced trend descriptions.

4.2.1. Application domain

Type 1 diabetic patients suffer from a impaired functionality of the pancreatic beta cells, and need to inject themselves exogenous insulin 3–4 times a day to regulate blood glucose metabolism. Such an intensive therapy may lead to hypoglycemic episodes: Blood Glucose Level (BGL) has therefore to be frequently tested and logged. In order to improve the quality of care, the implicit knowledge about patients' histories (and physicians' expertise) needs to be kept, man-

aged and distributed across the institution, and to be integrated with the other available knowledge sources, i.e. the explicit domain knowledge, formalised in knowledge bases or rule sets. It seems therefore of interest to provide instruments for managing all available knowledge types, and for supporting decisions in therapy planning, since revising insulin administration is a complex task, that can be correctly afforded only by customising general indications on the basis of the single patient's features [33].

4.2.2. Multi Modal Reasoning paradigm

The backbone structure of the decision support procedure we have implemented is based on the following reasoning tasks (Fig. 3):

1. Identification of metabolic problems,
2. Generation of a set of suggestions, able to cope with the identified metabolic problems, and selection of the most suitable ones,
3. Application of the selected suggestions to the current insulin protocol and selection of additional library protocols that could also fit the situation at hand.

The reasoning paradigm described above is completed resorting to the combination of a rule system, a Case Based retrieval system, and a model of the glucose–insulin interaction. In particular, the RBR system is able to schedule the tasks execution; in each task different methods are used, and the results are then deployed in the following steps. The methods used for the different tasks are outlined below.

4.2.2.1. Identification of metabolic problems.

The RBR fires some specialised procedures for data analysis and metabolic indicators extraction. The raw data are first analysed through a Temporal Abstractions (TA) technique [34]: in particular, state abstractions (e.g. low, normal, high values) are extracted and aggregated into intervals called episodes. From the most relevant episodes of hypoglycemia, hyperglycemia and normoglycemia, the so-called BGL modal day is derived [34]. The BGL modal day is an indicator able to summarise the average response of the patient to a certain therapy: it consists of the probability distributions of each BGL state abstraction in the different periods of the day (i.e. before breakfast, before lunch, etc.). Problems are identified when the probability of an undesired BGL state (hypo/hyperglycemia) in a certain time period is over a suitable threshold. Such threshold is computed relying on the Case Based retrieval tool. Case Based retrieval is implemented as a two-step procedure: a classification step, and an actual retrieval step. In the problem iden-




Fig. 3. The Multi Modal Reasoning methodology developed for managing type 1 diabetic patients.

tification task, only classification is exploited. The input case is classified relying on a taxonomy of prototypical classes, that describe the most common situations a paediatric diabetic patient may incur in. Through a Naive Bayes technique [35], the most probable classes are found. Classification results are used to choose the thresholds associated to each class.

4.2.2.2. Suggestion, generation and selection.

A number of rules are applied to propose suitable solutions to the problems derived in task 1. Such a reasoning task is therefore performed by the RBR tool (see [36] for further details).

4.2.2.3. Application of the selected suggestions to the current insulin protocol.

The RBR system typically suggests default solutions that consist in small variations of the current protocol insulin doses, as it is meant to be conservative enough to be safely applicable in a variety of different situations. However, when possible, a model is used to calculate the optimal insulin doses related to the daily insulin schedule proposed at step 2. Such model is a stochastic version of the model proposed by Deutsch et al. [37]. It is able to predict the steady-state BGL, obtained in response to a certain therapeutic protocol, on the basis of the current patient steady-state BGL and of the former insulin protocol. In our implementation, the steady-state BGL is represented through the modal day, i.e. through a set of probability distributions. The resulting dynamic model is therefore a Markov chain that is used to compute the optimal insulin doses according to decision theory. Having defined a utility function for the BGL, that is maximum for normoglycemia and minimum for hypos and hypers, the doses that maximise the expected utility function are chosen. Unfortunately,

not always the model turns out to give reliable predictions. This problem may be easily detected during the model parameters identification. When this situation holds, the MMR system also performs the CBR retrieval step, restricted to the most probable classes. Retrieval resorts to suitable Nearest Neighbour techniques [35]. Some simple statistics are calculated on the retrieved cases, to set the insulin adjustments width that will then be applied to the current protocol. Therefore, MBR and Case Based retrieval are used in a mutual exclusive way to specialise the rules behaviour.

After having adapted the current protocol to the problem at hand, similar protocols can be retrieved from a library of past protocols.

4.2.3. First evaluation results

A first evaluation procedure of the MMR methodology described above was carried out resorting to a patient's data set, provided by the paediatric department at Policlinico S. Matteo. The results obtained may be summarised as follows:

1. It is usually possible to obtain a model that leads to reliable predictions when dealing with 'simple' situations (i.e. cases in which the correct therapeutic strategy can be easily identified). Obviously, when the model can be exploited, it provides the optimal insulin doses adjustments. In particular, simple situations correspond to all situations in which a clear causal effect of insulin dosages on the BGL can be detected in the data.
2. On more complex situations (such as 'brittle control' or 'Somogyi effects' [38]), the model cannot be effectively used; in these examples, the possibility of exploiting past cases similar to the current one, retrieved through the CBR methodology, is very helpful for the definition of a proper therapy. In comparison to the application of

RBR with no integration, the exploitation of retrieval results leads to a sharper and more suitable insulin doses adjustment, customised for the patient at hand.

3. On the other hand, when the case library content is poor, the retrieval results may lead to an unfit rule specialisation. In this condition, only RBR can provide a reliable (even if maybe too conservative) solution. Nevertheless, the CBR methodology enables an easy knowledge storing and upgrading. The overall system will automatically improve its competence during routine clinical practice, as new cases will be stored in the HIS without requiring an additional work load to physicians, and will contribute to reduce the competence gaps. Through the memorisation of new information, the system is therefore able to learn how to cope with more and more complex situations.

4.3. Generalised cases: diagnosis of dysmorphic syndromes

GS.52 [14] is a prototype-based expert system which is routinely used in the children hospital of the University of Munich for many years. It is a diagnostic support system for dysmorphic syndromes. Such a syndrome means a non-random combination of different disorders. The major problems are the high variability of the syndromes (hundreds), the high number of case features (between 40 and 130) and the continuous knowledge modifications of dysmorphic syndromes. This means there are so many differences between a current and a similar case that an adaptation that takes all of them into account is impossible. So, for all cases with the same dysmorphic syndrome a prototype is created, which contains the most frequent observed features of these cases (Table 1). Such an abstracted prototypical case represents a dys-

Table 1
Portion of an example of a generated prototype

Diminished postnatal growth rate	77%
Hypercalcaemia	30%
Prenatal onset	75%
Mild microcephaly	67%
Full cheeks	46%
Anteverted nares	63%
Prominent lips	17%
Long philtrum	17%
Fullness of peri-o. region	75%
Medial eyebrow flare	25%

The numbers are the relative frequency in percentages the features occurred in the cases of the prototype.

morphic syndrome and usually contains only up to 20 features. For a current case the most similar prototypes are calculated. Subsequently, for the adaptation only few constraints have to be checked.

The prototypes are acquired by an expert consultation session. An experienced physician selects a new or an existing syndrome and typical cases for this syndrome. Subsequently, GS.52 determines the relevant features and their relative frequencies.

Diagnostic support occurs by searching for the most adequate prototypes for a current case. A similarity value between each prototype and the current case is calculated and the prototypes are ranked according to these values.

We evaluated the similarity measure of Tversky and the measure of Rosch and Mervis. Tversky [32] determines the similarity between a case and a prototype by adding up the number of shared features and subtracting the number of features of the prototype which the case does not share with the prototype. In contrast to him Rosch and Mervis [39] ignore those case features which the prototype does not share. Our experiment (Fig. 4) shows that their measure performed better than Tversky's, which indicates to ignore

those features of the current case the prototype under consideration does not share.

The result additionally indicates to present the most probable syndromes rather than to produce the one and only diagnosis. For both measures the correct diagnosis was always among the first 10, mostly among the first five and in majority, the first position. GS.52 contains about 230 diagnoses and more than 800 symptoms.

GS.52 differs from typical CBR systems, because cases are clustered into prototypes, which represent diagnoses, and the retrieval searches only among these prototypes. The sequential retrieval considers every prototype, calculates a similarity value for each prototype and ranks them according to these values. The adaptation consists of two examinations of the probable prototypes: a plausibility check with general rules (constraints) and a check of evidences for specific syndromes (some syndromes are nearly a proof for or against some diagnoses).

5. Conclusion

Cased-Based Reasoning seems to be a suitable technique for medical knowledge based systems. However, the adaptation task is the

bottleneck that has to be solved. Though adaptation is sometimes a rather easy task (as in FLORENCE), in many medical applications it may become an insurmountable difficulty. In this paper we have presented three possible solutions, all of them are justified for specific applications and none of them is an ultimate solution. Retrieval-only systems are especially useful for visualisation tasks, e.g. of images or organ function courses, because the users wish to see and interpret all specific details themselves [8]. Solving the adaptation by generalising is restricted to diagnostic problems where the following condition holds: the more abstracted a case the more typical are its features. This means to adapt by searching top down in a hierarchy of abstracted cases: the further down cases are placed in the hierarchy, the more specific and less typical are their additional features [13]. Combining CBR with rule- and model-based components should not really be seen as a solution for the adaptation problem, but as an opportunity to incorporate CBR subtasks (mainly the retrieval) into more complex methodologies instead of applying the complete CBR cycle, with the further advantage of making use of all available knowledge types.

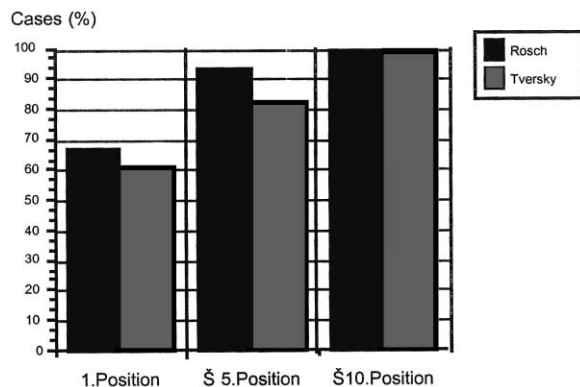


Fig. 4. Sensivity of GS.52 using cases of trisomy-21.

References

- [1] A. Aamodt, Cased-Based reasoning: foundation issues, *AICOM* 7 (1994) 39–59.
- [2] J. Kolodner, *Cased-Based Reasoning*, Morgan Kaufmann, San Mateo, 1993 ISBN: 1-55860-237-2.
- [3] R. Stottler, et al., Rapid retrieval algorithms for Cased-Based reasoning, in: *Proceedings of 11th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Mateo, 1989, pp. 233–237.
- [4] J. Quinlan, *C4.5, Programs for Machine Learning*, Morgan Kaufmann, San Mateo, 1993.

- [5] A. Broder, Strategies for efficient incremental nearest neighbor search, *Pattern Recognition* 23 (1990) 171–178.
- [6] L. Gierl, *Klassifikationsverfahren in Expertensystemen für die Medizin*, Mellen University Press, Lewiston, 1992 ISBN: 0-7734-4054-2.
- [7] F. Puppe, et al., Generated critic in the knowledge based neurology trainer, in: P. Barahona, et al. (Eds.), *Proceedings of Artificial Intelligence in Medicine*, Springer, Berlin, 1995, pp. 427–428.
- [8] R. Macura, K. Macura, MacRad: radiology image resource with a Cased-Based retrieval system, in: M. Veloso, A. Aamodt (Eds.), *Proceedings of 1st International Conference on CBR*, Springer, Berlin, 1995, pp. 43–54.
- [9] L. Gierl, et al., CBR in medicine, in: M. Lenz, et al. (Eds.), *Cased-Based Reasoning Technology, from Foundations to Applications*, Springer, Berlin, 1998, pp. 273–297 ISBN: 3-540-64572-1.
- [10] P. Koton, Reasoning about evidence in causal explanations, in: J. Kolodner (Ed.), *First Workshop on CBR*, Morgan Kaufmann, San Mateo, 1988, pp. 260–270 ISBN: 0-934613-93-1.
- [11] C. Bradburn, J. Zeleznikow, The application of Cased-Based reasoning to the tasks of health care planning, in: S. Wess, et al. (Eds.), *Proceedings of European Workshop on CBR*, Springer, Berlin, 1993, pp. 365–378.
- [12] R. Turner, Organizing and using schematic knowledge for medical diagnosis, in: J. Kolodner (Ed.), *First Workshop on CBR*, Morgan Kaufmann, San Mateo, 1988, pp. 435–446 ISBN: 0-934613-93-1.
- [13] R. Bergmann, W. Wilke, Towards a new formal model of transformational adaptation in Cased-Based reasoning, in: L. Gierl, M. Lenz (Eds.), *Proceedings of 6th German Workshop on CBR*, University of Rostock, 1998, pp. 43–52.
- [14] L. Gierl, S. Stengel-Rutkowski, Integrating consultation and semi-automatic knowledge acquisition in a prototype-based architecture: experiences with Dymorphic Syndromes, *ArtifIntell. in Med.* 6 (1994) 29–49.
- [15] B. Fuchs, A. Mille, A knowledge-level task model of adaptation in Cased-Based reasoning, in: K.-D. Althoff, et al. (Eds.), *3rd International Conference on CBR*, Springer, Berlin, 1999, pp. 118–131.
- [16] R. Schmidt, et al., Medical multiparametric time course prognoses applied to kidney function assessments, *International Journal in Medical Informatics* 53 (2–3) (1999) 253–264.
- [17] D. Aha Daniels, *Proceedings of AAAI Workshop on CBR Integrations*, AAAI Press, 1998.
- [18] E. Freuder (Ed.), *AAAI Spring Symposium on Multi-modal Reasoning*, AAAI Press, 1998.
- [19] I. Bichindaritz, et al., Cased-Based reasoning in CARE-PARTNER: gathering evidence for evidence-based medical practice, in: B. Smyth, P. Cunningham (Eds.), *Proceedings of 4th European Workshop on CBR*, Springer, Berlin, 1998, pp. 334–345.
- [20] L.D. Xu, An integrated rule- and Cased-Based approach to AIDS initial assessment, *International Journal of Biomedical Computing* 40 (1996) 197–207.
- [21] J. Surma, K. Vanhoff, Integrating rules and cases for the classification task, in: *Lecture Notes in Artificial Intelligence 1010*, *Proceedings of 1st IC-CBR*, Springer-Verlag, 1995, pp. 325–334.
- [22] M. Fathi-Torbaghan, D. Meyer, MEDUSA: a fuzzy expert system for medical diagnosis of acute abdominal pain, *Methods of Information in Medicine* 33 (5) (1994) 522–529.
- [23] C. Sary, et al., Trend analysis for spacecraft systems using multimodal reasoning, in: *Proceedings AAAI Spring Symposium on Multi-modal Reasoning*, AAAI Press, 1998, pp. 157–162.
- [24] L.K. Branting, B.W. Porter, Rules and precedents as complementary warrants, in: *Proceedings 9th National Conference on Artificial Intelligence (AAAI 91)*, Anaheim, 1991.
- [25] D.B. Leake, Combining rules and cases to learn case adaptation, in: *Proceedings of 17th International Conference of Cognitive Science Society*, Pittsburgh, 1995.
- [26] U. Wenkebach, et al., Visualization of large datasets in intensive care, *Proceedings of Annual Symposium Computing Applied Medical Care* (1992) 18–22.
- [27] Y. Shahar, M.A. Musen, RÉSUMÉ: a temporal-abstraction system for patient monitoring, *Computers and Biomedical Research* 26 (1993) 255–273.
- [28] I.J. Haimowitz, I.S. Kohane, Automated trend detection with alternate temporal hypotheses, in: R. Bajcsy (Ed.), *13th International Joint Conference on Artificial Intelligence (IJCAI-93)*, Morgan Kaufmann Publishers, San Mateo, 1993, pp. 146–151.
- [29] S. Miksch, et al., Therapy planning using qualitative trend descriptions, in: P. Barahona, et al. (Eds.), *AIME'95, Lecture Notes in Artificial Intelligence*, vol. 934, Springer-Verlag, Berlin, Heidelberg, New York, 1995, pp. 197–208.

- [30] J.R. Anderson, A theory of the origins of human knowledge, *Artificial Intelligence* 40, Special Volume on Machine Learning (1989) 313–351.
- [31] W.S. DeSarbo, et al., TSCALE: a new multidimensional scaling procedure based on Tversky's contrast model, *Psychometrika* 57 (1992) 43–69.
- [32] A. Tversky, Features of similarity, *Psychological Review* 84 (1977) 327–352.
- [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] C. Larizza, R. Bellazzi, A. Riva, Temporal abstractions for diabetic patients management, in: E. Keravnou, C. Garbay, R. Baud, J. Wyatt (Eds.), *Lecture Notes in Artificial Intelligence*, Springer-Verlag, Berlin, 1997, pp. 319–330.
- [35] S. Montani, R. Bellazzi, L. Portinale, G. d'Annunzio, S. Fiocchi, M. Stefanelli, Diabetic patients management exploiting Cased-Based reasoning techniques, *Computer Methods and Programs in Biomedicine* 62 (2000) 205–218.
- [36] S. Montani, R. Bellazzi, C. Larizza, A. Riva, G. d'Annunzio, S. Fiocchi, R. Lorini, M. Stefanelli, Protocol-based reasoning in diabetic patient management, *International Journal of Medical Informatics* 53 (1999) 61–77.
- [37] T. Deutsch, A.V. Roudsari, H.J. Leicester, T. Theodorou, E.R. Carson, P.H. Sonksen, UTOPIA: a consultation system for visit by visit diabetes management, *Medical Informatics* 21 (1996) 345–358.
- [38] *The Merck Manual*, Merck Sharp and Dohme, 1987.
- [39] E. Rosch, C.B. Mervis, Family resemblance: studies in the internal structures of categories, *Cognitive Psychology* 7 (1975) 573–605.