AN ADCSP-BASED NON-MONOTONIC FRAMEWORK FOR MEDICAL DIAGNOSIS

Sabina Costache

"Dunarea de Jos" University, Faculty of Automatics, Computers and Electrical Engineering

Abstract: The ability to reason within a dynamical environment is of crucial importance in Artificial Intelligence. The present paper models nonmonotonic reasoning by means of a DCSP (Dynamic Constraint Satisfaction Problems) framework, taking advantage of the representation facilities of direct argumentation systems. The algorithm presented below applies dynamic backtracking for the approximate computation of the admissible semantics, which was used to define the concept of multiple diagnosis. The final application of our work is a system for medical diagnosis, that models its search space efficiently and dynamically, while confronted with sequential tests. It asserts and rejects beliefs in different component elements of the diagnosed domain following a nonmonotonic schema which is very close to a human expert's reasoning model.

Keywords: clinical decision making, hybrid intelligent systems, direct argumentation systems, nonmonotonic reasoning.

1. INTRODUCTION

Medical diagnosis represents a major challenge to the world of AI, due to its complexity and non-typicality.

A major drawback of the first-generation artificial intelligence programs in medicine (INTERNIST/CADUCEUS, MYCIN, PIP) comes from the fact they do not use a deep causal structure for the relationship between disorders and their symptoms, while for a human expert an explanation is seen as a deduction inferred on basis of a cause-effect chain. An important consequence is that interaction among multiple disorders is impossible to approach, only by associations between phenomena, with no causal details.

The problem of complex interactions occurs when multiple disorders are present in one and the same patient, and their symptoms unexpectedly interact. Even CASNET, with all its causal representation, has serious problems with interacting or overlapping symptoms, and therefore resumes its utility to single-disorders cases, because of the difficulties with the probabilistic treatment of uncertainty and inference.

The probabilistic approach to uncertainty is also to blame for the unappropiate tackling of contradictions. When two rules are in conflict, this is treated –

This paper was recommended for publication by Cornelia Tudorie

likewise concordance- by adjusting the degree of trust in some related hypotheses. But in the real world reasoning, human experts have a much deeper and complex reaction to the detection of a contradiction: they reconsider previously accepted data, and/or add new possible hypotheses to the active set (i.e. those currently taken into consideration). The conclusion is that a probabilistic model is inherently inadequate to deal with contradictions, and a categorical approach is needed.

Patil has completely replaced probabilistic measures with structural criteria in Abel (Patil, 1981), trying to surpass the difficulties described above. His system uses links of a special kind to model competition/ contradiction, and only categorical decisions are allowed. The system is based on a hierarchically partitioned data representation defined by Lynch (Lynch, 1960) (conceptual maps[1]).

Yet, the medical field is far too complex to completely give up probabilities, like Abel does. As structural and probabilistic measures complement each other, they should both be used in diagnosis. Moreover, the applicability of Abel in large fields is restricted, because general strategies are needed to initially pre-process extended medical contexts. Probabilistic / associative efficient types of reasoning would be useful exactly during this phase of pre-processing, in order to focus the search. The DiaMed system presented below combines probabilistic/ categorical reasoning, taking advantage of the qualities of both of them, and leading to a combinative hybridization (the same general architecture was used by CHECK).

The present paper subsequently presents an original hybrid system for medical diagnosis (DiaMed). Its originality consists mainly of an original modeling and reasoning framework for diagnosis, although DiaMed is, at the same time, a possible continuation of the ideas in CASNET and CHECK systems.

Section 2 is dedicated to the detailed presentation of the DiaMed system. Regarding the nonmonotonocity of diagnostic reasoning and the necessity of asymmetric attack relations, Section 2.2 introduces the modern formalism of direct argumentation systems, and suggests a new definition for multiple diagnosis making use of the argumentative admissible semantics. The section also presents the basis of an efficient computation for the admissible semantics using dynamic backtracking for dynamic constraint satisfaction problems (DCSP). An important contribution of this chapter is the translation of the dynamic testing mechanism for (medical) diagnosis into a dynamic constraint satisfaction problem. Moreover, we use an original arguments-based knowledge representation model for the medical field. We believe our paper is a pledge for direct argumentation systems and their applicability in difficult real-world problems.

[1] "In approaching large maps it is recommended to work with partial images sets, which can be more or less inter-related or overlapped" (instead of a large, global image).

2. THE DIAMED HYBRID SYSTEM

2.1. Introduction

DiaMed (**Dia**gnosis in **Med**icine) is a hybrid-combinative system with two levels. Combinative hybridization of the type chosen here seemed the best choice, not only for the reasons resumed above and further detailed in this section, but also because it was a good option when compared to, for instance, neuro-fuzzy or neuro-symbolic hybrids, with their curse of dimensionality and difficulties related to modeling interactive, dynamic problems (like medical diagnosis is).

In DiaMed, uncertainty is modeled logically, by nonmonotonic reasoning. The problem of complex interactions is approached in a generative manner: composite hypotheses are built based upon *admissible* solutions to a dynamic constraint satisfaction problem- DCSP- (instead of an explicit codification of all possible composite hypotheses and their effects) (Minh Dung, 2006; Chesnevar, 2000). Admissibility is a theoretic-argumentative view of consistency, appropriate for a diagnosis problem (as we shall see following) (Caminada, 2007). This generative approach needs a causal model, in order to better understand possible interactions among different elements of the medical model as we have already emphasized.

Therefore, composite hypotheses (i.e. multiple disorders at the same patient) are defined as covering admissible sets. Admissibility is defined through individual attack relationships, and allows us to dynamically compose hypotheses, dependently on a given context of manifested symptoms. Like in CASNET, CHECK or Abel, DiaMed is built around a causal knowledge representation. Complete causal models are not necessary, but only their restrictions to the nodes relevant for the decision process.

Although the architecture of DiaMed resembles CHECK, the implementation differs. The first level implements hypotheses' selection with an efficient associative method (it uses fuzzy decision functions to rank disorders) (Munteanu, 2005). The main advantage of these decision functions, compared to the law of evidence combination in CHECK resides in the fact they can accurately express a great variety of vague criteria (for instance, *the majority, at least x out of n, a significant part of*, etc.)

The second level uses a deep causal model, restricted to the context of hypotheses selected at the first level, in order to discriminate and refine the final diagnostic, and to solve the conflicts generated – if any. A DCSP algorithm controls which constraints are active at a given moment, having the role to focus on interesting sub-parts of the model. Moreover, this phase considers a complete and precise model (to the maximum possible extent), which represents exceptions in a natural and efficient way, and the reasoning scheme suits the nonmonotonicity of diagnosis. To this purpose, the second level uses the logical and symbolic methods of direct argumentation systems and CSP algorithms in order to refine and explain diagnostic results.

2.2. Knowledge representation

The knowledge model of DiaMed contains causal associations between classes and their characteristics. Its components are described following.

The *diagnostic classes* (i.e. the diseases) are modeled by a special type of *causal nets*, with different kinds of nodes and arcs, which describe the deep causal model. There exist three types of nodes:

- root-nodes corresponding to classes (diagnostic hypotheses); they are primary deep causes of observed manifestations;
- nodes related to deep manifestations (inaccessible or accessible only through expensive/ time-consuming/ invasive tests);
- nodes related to shallow manifestations (easy to access or direct observations).

The nodes of the net (either deep or shallow) can be of two kinds: **necessary** or **supplemental**. If a necessary node is infirmed by tests, the diagnostic hypothesis which contains it is eliminated.

Arcs linking the nodes can also be of various types:

- <u>Necessary implications:</u> the cause always determines the effect;
- <u>Possible implications:</u> the cause *may* determine the occurrence of the consequence, but it is not compulsory; (this uncertainty comes from the model's incompleteness: there exist certain elements/ conditions that influence the validity of the implication but which were not explicitly modeled);
- <u>Attacks</u> (either bi- or unidirectional): these relations connect elements that cannot be simultaneously assumed "in" (i.e. *true*) in the case of one and the same system (i.e. patient, in the medical field).

Each diagnostic class is defined by such a causal net that contains all possible elements related to the class, and these elements are organized in progressively shallow (i.e. accessible to direct observation) levels. Intermediary nodes between the root and the leaves are usually inaccessible or difficult to access (only through expensive, invasive, time-consuming tests). The degree of accessibility grows as we approach the leaves (Figure 2.1).

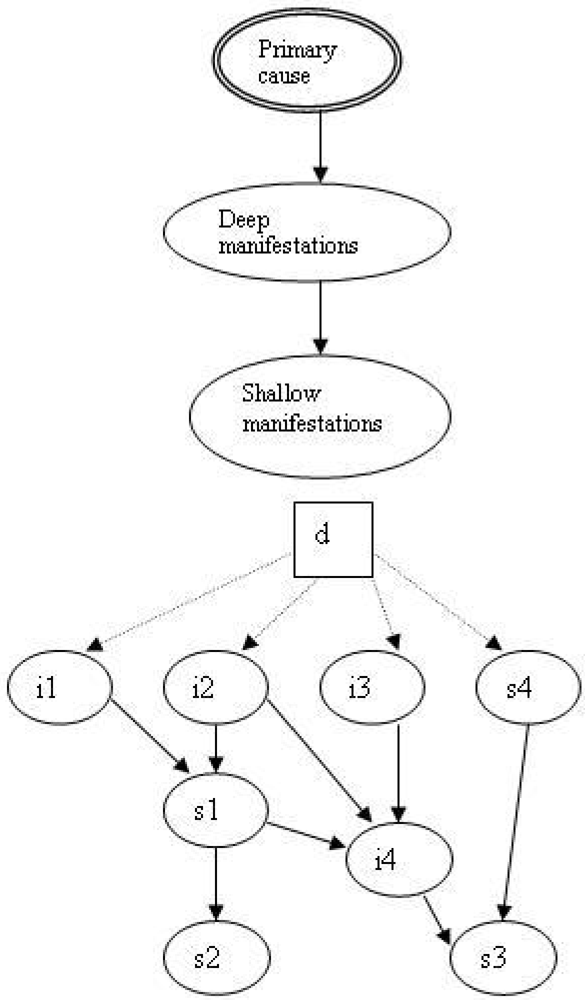

Fig. 2.1.The general structure of a causal net and a sub-net example (d-disorder, i-intermediary inaccessible nodes, s –accessible manifestations/ symptoms)

Definition. An *argument* associated to a class is an instantiation of the causal net that defines the class. An instantiation of a causal net is a subset of its nodes that contains at least an observed manifestation (the rest of the nodes being assumed true).

A natural and original definition for a multiple diagnosis is given below.

Definition. A *multiple diagnosis* (i.e. a non-empty set of possible diseases for a given patient) is an admissible[2] hypotheses[3] set that covers all observations and is minimal with this property.

Example. Let A={leukemia}, B={normal PT, PTT level}, C={medication affecting the PT, PTT level} be three arguments. The attack relation is described by the graph: A←B←C. Then A is acceptable with respect to C, C being a defence of A against attack B. Consequently, {A,C} is an admissible set, and “leukemia” is a possible diagnostic within the context defined by C.

Remarks!

- P.M. Dung has made the remark that attack relations among arguments in defeasible reasoning only depend upon the assumptions (i.e. hypotheses, that have not been either confirmed or infirmed yet) the arguments are based on, and suggested a definition for argument as a deduction whose premises are all assumptions; moreover, argument ***a*** attacks argument ***b*** if ***a*** attacks an assumption ***b*** is based on. (An argument attacks an assumption ***α*** when the conclusion of ***a*** is the contrary ***not α*** of ***α***).

- An argument corresponds to the peculiar configuration of symptoms observed at a given patient, through which the disease manifests itself, and which can vary from patient to patient. Each argument can be viewed as an instantiation of a class. It can dynamically be added/ deleted from the current context, depending on the related attack relations.

Definition. *The abstracted attack graph* is, simply, the attack graph considered between the arguments that contain the attacking nodes, ignoring their internal structure of nodes (we suppose there are no attacking nodes within one and the same argument).

The present model uses three types of attacks. Firstly, two alternative causes for the same effect are attacking each other (see Figure 2.2.) (the idea being that only one of them actually produced the effect, and it is useless to consider them both as “in”). The situations that might contradict this assumption are rare, if ever. We haven’t met such an exception, within a context of 30 disorders chosen for testing different examples.

[2] Argument A is *acceptable* to the set S of arguments iff any argument defeating A is defeated by an argument in S. A set S of arguments is *admissible* iff any argument in S is acceptable to S.

[3] A hypothesis is any active disease, which can be, in particular, associated to the argument that sustains it.

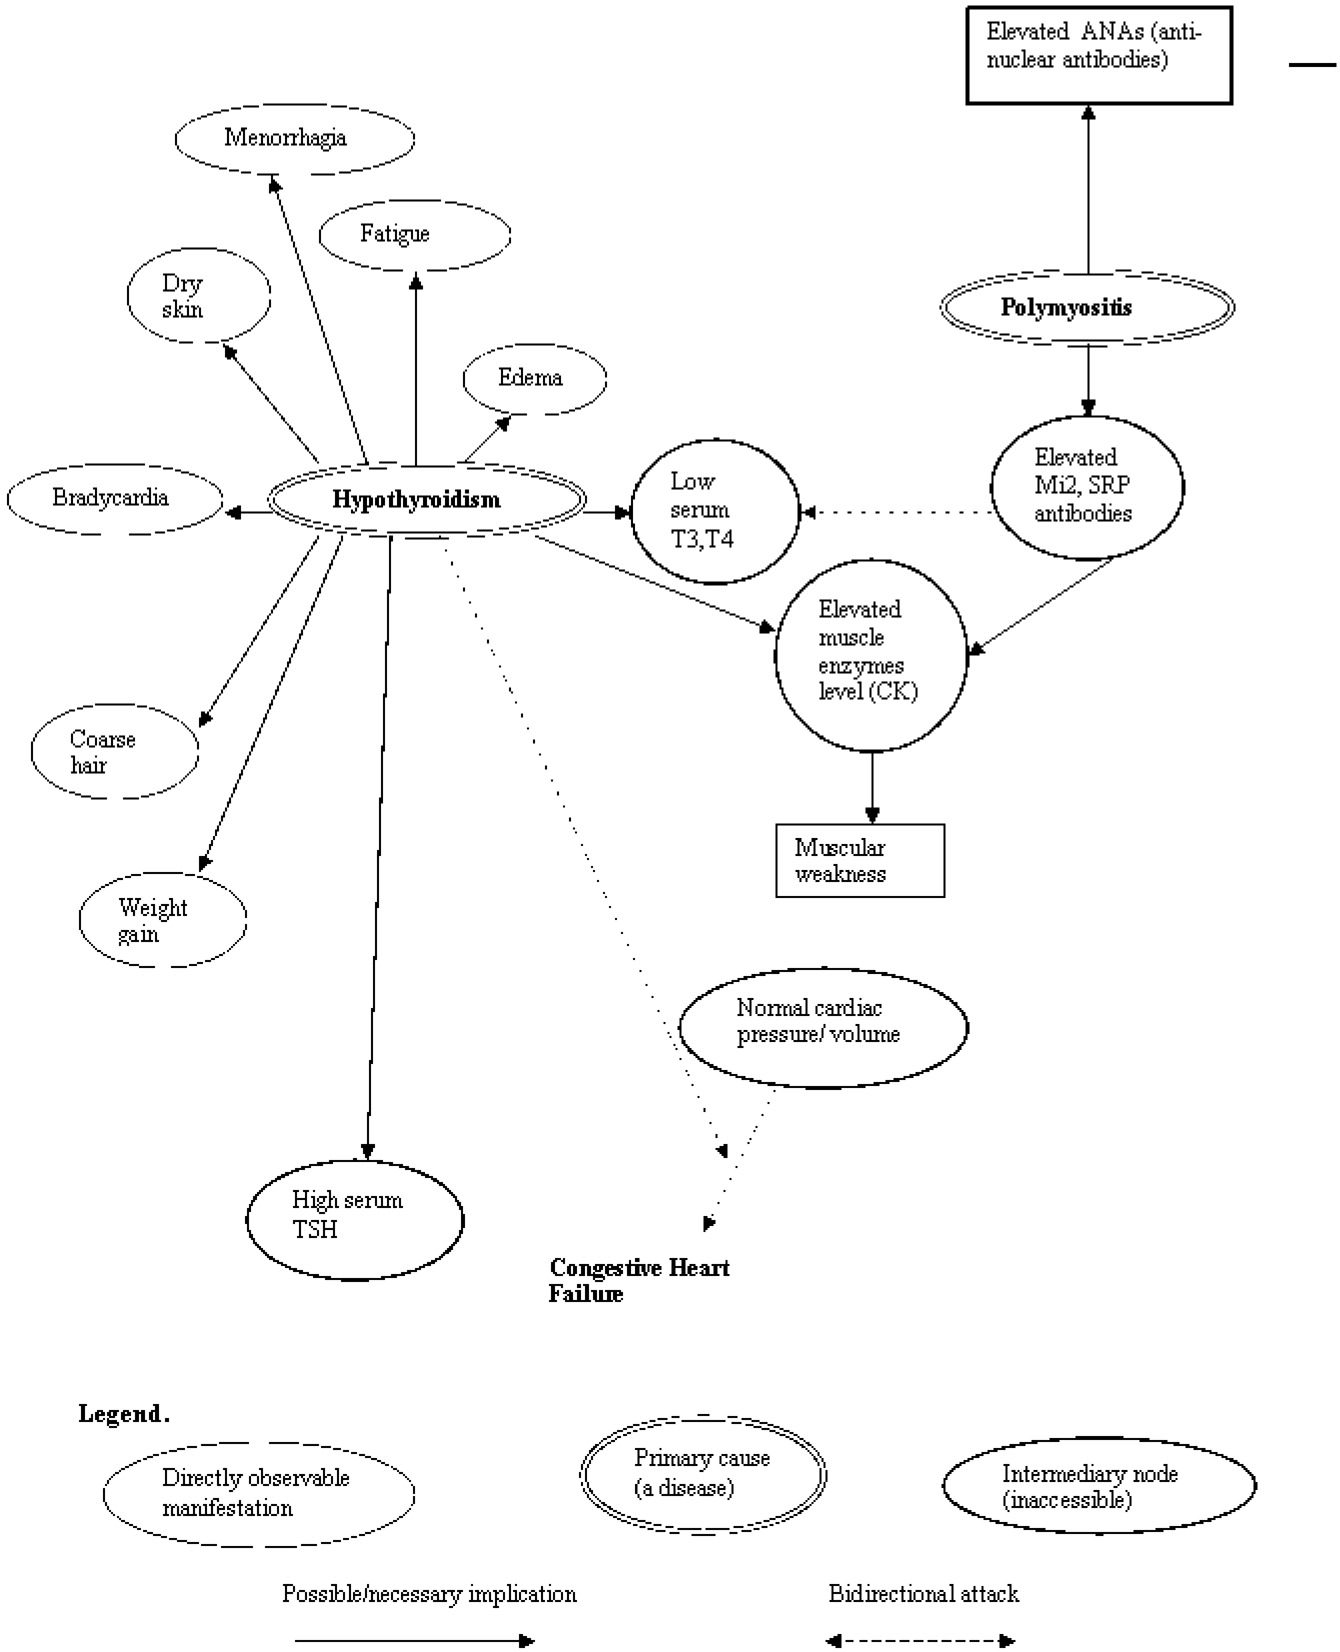

Fig. 2.2. Example of causal dependencies

Secondly, attacks model the following situation: if a disorder ***b*** can manifest itself under the mask of disorder ***a***, such that ***b*** is often confused with ***a***, then we consider that ***b*** attacks ***a***. The idea here is that ***a*** is initially assumed true, but supplementary tests infirm ***a*** and confirm ***b***. This attack comprises experts' experience about some possible temporal order for the occurrence of hypotheses within a reasoning chain. For instance, "congestive heart failure" attacks "cirrhosis", because sometimes "heart failure" is confused with "cirrhosis".

Finally, the third kind of attack is given by logical inconsistency: if it is necessary to have an affected heart pressure/ volume in order to conclude heart failure, then we consider "normal heart pressure/ volume" as a direct attack to heart failure (abstracting other internal network elements). But there are exceptions to this situations that make things even more difficult; for instance, in hypothyroidism we have symptoms of heart failure, although pressure or volume were not affected (and this is because the generating mechanism is different, originating in hypocalcemia), hence we consider that hypothyroidism directly attacks "normal heart pressure/ volume", within the context "heart failure". In other words, hypothyroidism is a defence of "heart failure" against "normal heart pressure/ volume" (see the attack graph in Figure 2.3). This exception is, as a matter of fact, an attack of the first type - "alternative cause" (and therefore will not receive a distinct representation in our model).

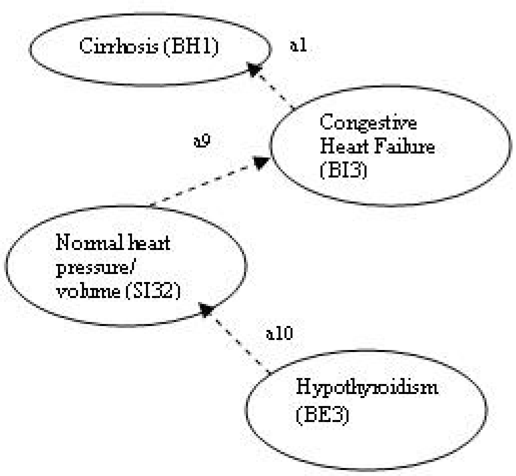

Fig. 2.3. The attack graph for the firstly selected context

2.3. Problem modeling within a DCSP framework

Notations.

- The diagnostic hypotheses: $d_1,\ldots,d_N$ (N classes/ diseases);
- Their related causal nets: $C_1,\ldots,C_N$;
- Arguments: $A_1,\ldots A_L$ (where L represents the number of possible instantiations of causal nets);
- **Ip**=$\{d_1,\ldots,d_t\}$is the set of selected hypotheses;
- *Confirmed_Manifestations*=$\{mc_1,\ldots,mc_k\}$ – the set of confirmed manifestations (that were proved being true);
- *Infirmed_Manifestations*=$\{mi_1,\ldots,mi_p\}$ –the set of infirmed manifestations (that were proved being false);
- *Context*= *Confirmed_Manifestations* $\cup$ *Infirmed_Manifestations* –the complete set of symptoms either confirmed or infirmed during the testing process;
- G_a – the attack graph restricted to the set **Ip**;
- The set **Ip** is partitioned into: *Infirmed_hypotheses, Considered_hypotheses, New_hypotheses.*

The algorithm starts with an initial set *Confirmed_Manifestations,* selects hypotheses through fuzzy decision (the abductive phase of hypotheses' generation), and it ends up with a set of active hypotheses:$\{d_1,\ldots,d_t\}$, and their associated causal nets $C_1,\ldots,C_t$. Firstly, we compute the state of every node linked through deterministic relations to the confirmed/ infirmed nodes (for instance, a node attacked by a "true" node–when the attack is not defensible-, will surely be "false"). The rest of the nodes from the activated causal nets represent the set of active variables *Active_Var*=$\{V_1,\ldots,V_p\}$, on which an adapted version of dynamic backtracking for DCSP (Verfaille, 1994) is applied (during the consistency-check phase). Active variables are usually associated to defensible nodes: their status is "in" (assumed true) or "out" (assumed false) at a certain moment, but they can be contradicted/reinstated later during the algorithm. The active constraints of the CSP algorithm – *Active_Constr*- are those that contain at least one active variable. As already defined, a multiple diagnosis is a minimal solution- i.e. a minimal, admissible, completely covering set – to the CSP problem. We have made things easier in practice by selecting into *Active_Var* only the diagnostic hypotheses themselves.

The evidence of tested nodes is kept in *Confirmed_Manifestations* set (which contains diagnostic hypotheses and their supporting symptoms and is used for the final explanation) and *Infirmed_Manifestations* - used for keeping track of already infirmed nodes, in order not to reconsider them again.

The three possible types of implications in a causal net are represented as binary constraints (this representation was inspired from BCP- Boolean Constraint Propagation- problems):

- Necessary implications: $a \rightarrow b$: C_{ab}={00,11,01};
- Attacks: $a \mapsto b$: C_{attack}={10, 01,00};
- Possible implications do not restrict the domain of values of any variable ({00,01,10,11}).

Let:
- *Cons(v)*= constraints which contain variable v;
- *Var(c)*= variables occuring in constraint c.

The set of current constraints *Active_ Constr* is naturally defined as:

$$Active\_Constr = \bigcup_{v \in Active\_Var} Cons(v) \quad (2.1)$$

Definition. A *solution* to a diagnostic problem is a complete and consistent (admissible) assignment of truth values to all the active variables (i.e. activated through the selection of some particular hypotheses), which covers all confirmed manifestations. A solution is *minimal* if it has a minimum number of nodes, while still respecting the previous conditions. This definition corresponds to the definition of multiple diagnosis above.

2.4. Fuzzy decision –based selection of hypotheses in DiaMed

The phase of selection of relevant hypotheses from a large context should use efficient techniques (rather than precise and transparent ones), in order to quickly reduce the search space. The selection of hypotheses in DiaMed uses a model based upon fuzzy decision functions (Munteanu, 2005).

2.5. Discrimination of hypotheses and argumentation of diagnostic decision in DiaMed

The present section describes the algorithm used by the discrimination/ explanation level in DiaMed, which is an original adaptation of the algorithm in (Verfaillie, 1994), such that the dynamics of constraints is governed by the results from the medical tests. A relevant example ends the section.

Algorithm
Dynamic Backtracking_DCSP _for_medical_diagnosis (V_n, V_a, C)

1.Order the *Active_Var* set starting with the most constrained variables (this is a well-known heuristic to make CSP algorithms perform better). As attack is the most restrictive constraint, place variables which take part in attacks first;

2.We have adapted the dynamic backtracking algorithm for DCSP from (Verfaillie, 1994) to compute the diagnostic as follows below (*Active_Var* is the set of variables, *Active_Constr* are their associated constraints). (The differences to the original version are marked in bolded letters).

(In, out nodes =defensible nodes; i.e. they can change state subsequently during computation).

Initializations
Let $V_a = \phi$ be the set of assigned variables, V_n= *Active_Var* the set of unassigned variables, C= *Active_Constr*, the set of constraints that contain at least one variable from *Active_Var*; initialize variables' *Domains* with the set *{in,out}*.

V_{total}= the set of decision-relevant hypotheses and nodes;

1. If (V_n=Ǿ) return V_a (assigned variables)
if V_a is an admissible complete covering of present manifestations/ symptoms
then print the solution;
if V_a is not admissible
Inadmissible_Solutions= Inadmissible_Solutions $\cup$ V_a;
If V_a admissible but incomplete Partial_Solutions=Partial_Solutions $\cup$ V_a;
(memorized for a potential subsequent extension)

Else
Pick a variable v from V_n;
Current_Domain(v)= Domain(v)- (Values inconsistent with V_a, with respect to C)
If (Current_Domain(v)= Ǿ)
Let V_e be an explanatory explanation of the conflict (i.e. a set of variables in conflict with v)
If (V_e= Ǿ) (*we do not have a solution*) (*)
New_Data=Test;
Repeat the Selection of hypotheses (through fuzzy decision) using the New_Data set;
Change the working Context:
CA = constraints that define the new context of hypotheses,
CS = constraints related to those specific contexts in which a necessary node was attacked by some node associated to one of the symptoms recently tested;

Let V be the set of variables that contradict at least one constraint from CA
Un-assign variables from V:

$V_n = V_n \cup V$; $V_a = V_a - V$;

Delete eliminatory explanations which contain variables from V;

Delete eliminatory explanations which contain variables that contradict at least one constraint from CS;

GOTO1; (Restart from the begining)

OR: try to extend inadmissible and partial solutions, keeping V_a fixed and working only with V_n.

Else

Un-assign v' –the most recent variable from V_e

$V_n = V_n \cup \{v'\}$; $V_a = V_a - \{v'\}$;

Create an explanation for the elimination of value val' = $\{V_e - v'\}$;

Delete eliminatory explanations which contain v'; (→ refresh the domains of all variables)

GOTO 1.

Else

For each (val$\in$ Current_domain(v))

v←val; $V_a = V_a \cup v$; $V_n = V_n - \{v\}$; GOTO 1.

We delete a node from the Context when it appears as *false* in a (partial) solution. Only the solutions whose true nodes' total weight exceeds a given threshold are kept in a list and further taken into consideration, this heuristic being very useful in making the algorithm perform better faster.

Example

1.Let's suppose we initially observe the following symptoms (see Figure 2.2 and the Appendix):

Confirmed_Manifestations ={anorexia, arrhythmias, ascites, dyspnea, edema, fatigue, muscular weakness, anemia of chronic disease}. DiaMed selects, through fuzzy decision (Munteanu, 2005), the following disorders (hypotheses), which have accumulated a score greater or equal to the chosen threshold (i.e. 0.2):

- Angina pectoris 0.51 (BI1)
- Cirrhosis 0.22 (BH1)
- Hypothyroidism 0.25 (BE3)
- Myocardial infarction 0.26 (BI2)
- Congestive heart failure 0.29 (BI3)
- Myocarditis 0.70 (BI4)
- Pericarditis 0.34 (BI5)

(The codification of diseases used by the implementation is listed between the brackets). A special-status node is also automatedly selected by our algorithm into the list of hypotheses: “normal heart pressure/ volume” (SI32). Its special status is given by the fact that the node occurs in the attack relations that imply selected hypoptheses, without being a hypothesis itself, or a symptom of a disorder. (Figure 2.3 presents the attack graph associated to the selected context).

The algorithm computes three minimal admissible covering solutions:

1. Hypothyroidism, congestive heart failure, myocarditis (BE3, BI3, BI4)
2. Hypothyroidism, congestive heart failure, myocardial infarction (BE3, BI3, BI2)
3. Hypothyroidism, congestive heart failure, angina pectoris (BI1, BI3, BE3)

Suppose the user (physician) is not satisfied about the results and wishes to further complete the investigations with new tests. At this moment, a list of possible necessary symptoms for the selected disorders is displayed, to be eventually infirmed and narrow down the context. To this purpose, a list of necessary symptoms possible for the current context is built by the program (this being the only automatic clue to help the testing process). Necessary symptoms/ manifestations have a preferential status over supplementary ones because they can restrict the search space, if denied (Figure 2.3).

The new test results lead to *Infirmed_Manifestations* ={“chest pain sensitive to nitroglycerin”, “chest pain insensitive to nitroglycerin”, “gallop rhythm”, “pericardial friction rub”}, which excludes from the list of possible nodes the set {BI1, BI2, BI4, BI5}-those disorders whose certain necessary manifestations are missing at the given patient (see Appendix). The infirmed manifestations support “normal heart pressure/ volume”, which defends the hypothesis *„cirrhosis”*. The presence of “elevated Mi2, SRP antibodies” is furthermore discovered. This symptom adds “Polymyositis” (BREUM3) to the list of possible hypotheses (its score being now 0.45). The attack graph becomes the one in Figure 2.4.

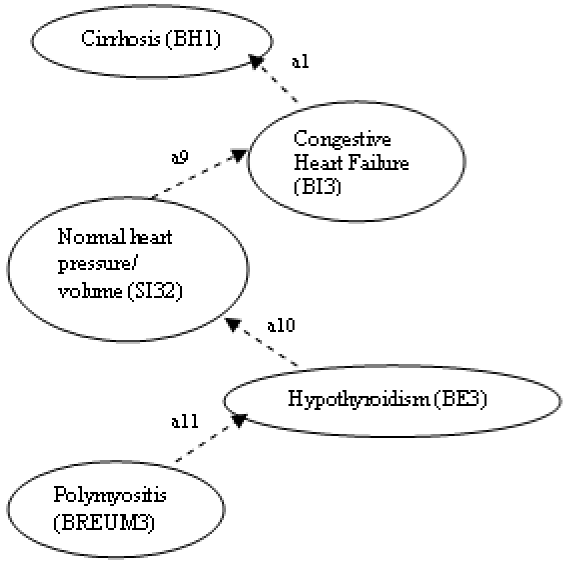

Fig. 2.4. The attack graph for the secondly selected context

The following phase of the algorithm tries to extend inadmissible assignments, generated at the previous step, with defences. For instance, in the present example, it starts with the inadmissible assignment {"cirrhosis" "in", "normal heart pressure/ volume" "in"} and gets an admissible assignment within the new context (although not completely covering): {"cirrhosis" "in", "normal heart pressure/ volume" "in", "polymyositis" "in"}. This means that the node BREUM3 is a defence of the set {SI32, BH1}, against {BI3, BE3} –the latter have been activated by the new evidences. In conclusion, the system manages to follow very closely the human expert's non-monotonic steps of reasoning, through a dynamic simulation of the search space, which is controlled by sequential phases of testing and denial of previously assumed hypotheses.

3. CONCLUSIONS AND FUTURE WORK

The DCSP-based approach from Section 2.5 represents an efficient translation of the dynamic re-modelling of the working context, which is directed by the evidences resulted from tests. This re-modelling focuses reasoning on limited sections of the medical domain. The activity constraints add or delete variables to/from the problem according to the context of selected hypotheses, which is dynamically tuned through testing and through the application of domain-dependent rules. These activity constraints are implicitly defined by the fuzzy decision functions that perform the selection and by the arguments (i.e. active instances of causal nets).

Dynamic backtracking for DCSP is an incremental method, which keeps the valid part of a solution when moving to different parts of the model (efficiency). Moreover, the system is flexible: it generates hypotheses even when provided only with insufficient information (they shall be retracted if further information is contradictory). Our algorithm also shows how DCSP can approximate the admissible semantics in a tractable manner. The choice for the dynamic version facilitates a contextual computation of admissibility, which is naturally context-dependent.

Multiple diagnosis is originally defined in terms of arguments (using the admissible semantics), and arguments are adapted to match the medical field, by structuring information and grouping disorders according to possible interactions. Because arguments were especially created to model human reasoning confronted with uncertainty and incremental evidence gathering, they are appropriate for iterative belief revision which is a main characteristic of medical diagnostic reasoning, and they can handle the interactivity of sequential testing which interleaves with hypotheses' generation (see also (Caminada, 2006), (Rahwan, 2003)).

The nonmonotonic mechanism of belief generation and cancellation is reflected in the addition and deletion of constraints within DCSP. The main advantage of this method over CHECK, for instance, resides in its tractability, as compared to the computational approaches of indirect abduction.

The original approach of DiaMed, which uses argumentative non-formal logic and DCSP algorithms, can be very useful during the phase of discriminating among alternative diagnoses. "Further research in nonmonotonic reasoning should focus on computational aspects, because it is only so that nonmonotonicity can have an impact on Artificial Intelligence and an utility for real-world problems" (Brewka et al., 1997).

The system has to be further improved. A great part of the decisions associated to testing are still delegated to the user (which maybe is not a drawback after all). Also, the medical model needs to be completed by a team of human experts, in order to test the system on a significant amount of real data. It would also be worth to study the impact intelligent techniques can have on propositional inference in general.

REFERENCES

Brewka, Dix, Konolige (1997) *"Nonmonotonic Reasoning: an overview"*, *CLSI Lecture Notes* ***73***, Stanford, California;

Caminada M.W.A.(2006), „On the issue of reinstatement in argumentation" . In M. Fischer, W. van der Hoek, B. Konev, and A. Lisitsa, editors, *Logics in Artificial Intelligence; 10th European Conference, JELIA 2006*, pages 111–123. Springer, LNAI 4160;

Caminada M., Amgoud L. (2007), On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171(5-6):286–310;

Ches˜nevar C. I., Maguitman A., Loui R. P. (2000), Logical models of arguments. *ACM Computing Surveys*, 32(4):337–383;.

Lynch K. (1960): "*The image of the city*", MIT Press, Cambridge, Massachussets;

Minh Dung P., Mancarella P., and Toni F. (2006) , A dialectical procedure for sceptical, assumption-based argumentation. In P.E. Dunne and TJ.M. Bench-Capon, editors, *Computational Models of Argument; Proceedings of COMMA 2006*, pages 145–156. IOS;

Munteanu, S.(2005), "A Hybrid Model for Diagnosing Multiple Disorders", *International Journal of Hybrid Intelligent Systems*, Vol **2.** No.1 IOSPress, :35-55;

Patil R., (1981): "Causal representation of patient illness for electrolyte and acid-base diagnosis", MIT Lab for Computer Science, *Technical Report* **TR-267**;

Rahwan I., Ramchurn S. D., Jennings N. R., McBurney P., Parsons S., Sonenberg L. (2003) Argumentation-based negotiation. *Knowledge engineering review*, 18(4):343–375;

Verfaillie G., Schiex T. (1994), "Dynamic Backtracking for Dynamic Constraint Satisfaction Problems", *Proceedings of the ECAI '94 Workshop on Constraint Satisfaction Issues Raised by Practical Applications*, Amsterdam, the Netherlands.

APPENDIX

The weights of symptoms within disorders' definitions

ANGINA_PECTORIS	
Symptom	weight
ARRHYTHMIAS	0.7
DYSPNEA	0.8
CHEST_PAIN_RELIEVED_BY_NITROGLYCERIN	0.9
NAUSEA/VOMITING	0.5
CIRRHOSIS	
CRONIC_DISEASE_ANEMIA	0.6
ANOREXIA	0.6
ASCITES	1
ENCEFALOPATHY	0.7
ELEVATED_FIBRINOGEN	0.7
LEUCOPENIA	0.7
MELENA	0.7
FATIGUE	0.5
PROTEINURIA	0.8
VARICEAL_BLEEDING	1
LOW_ALBUMINE_SYNTHESIS	0.8
SPLENOMEGALIA	0.9
HIGH_SERUM_TRANSAMINASES(TGP,TGO)	1
LIQUID_INSIDE_ABDOMEN	1
THROMBOCYTOPENIA	1
HYPOTHYROIDISM	
BRADYCARDIA	0.6
EDEMA	0.7
ELEVATED_MUSCULAR_ENZYMES(CK…)	0.8
WEIGHT_GAIN	0.8
MENORRHAGIA	0.8
FATIGUE	0.8
COARSE_HAIR	0.6
DRY_SKIN	0.6
MUSCULAR_WEAKNESS	0.6
TACHYCARDIA	0.6
LOW_SERUM_T3,T4	1
HIGH_SERUM_TSH	1
MIOCARDIAL_INFARCTION	
ARRHYTHMIAS	0.8

DYSPNEA	0.9
CHEST_PAIN_NOT_RELIEVED_BY_NITROGLYCERIN	1
ELEVATED_FIBRINOGEN	0.6
NAUSEA/VOMTING	0.5
PANIC	1
LOW_BLOOD_PRESSURE	0.7
HIGH_SERUM_TRANSAMINASES(SGOT)	0.8
CONGESTIVE_HEART_FAILURE	
CRONIC_DISEASE_ANEMIA	0.5
ANOREXIA	0.5
ASCITES	0.7
PULMONARY_VASCULAR_CONGESTION	0.7
ELEVATED_VOLUME_IN_CENTRAL_VESSELS_WHEN_LAYING	0.6
DYSPNEA	0.9
PAROXYSTIC_NOCTURNAL_DYSPNEA	1
EDEMA	0.7
HEMOPTYSIS	0.5
BIG_HEART	1
FATIGUE	0.5
PULMONARY_RALES	1
GALOP_RHYTHM	1
TACHYCARDIA	0.7
LIQUID_INSIDE_PULMONARY_ALVEOLA	0.7
JUGULARY_TURGESCENCE	0.9
NOCTURNAL_COUGHING	0.8
MYOCARDITIS	
ARRHYTHMIAS	0.7
DYSPNEA	0.8
FATIGUE	0.7
GALOP_RHYTHM	0.9
PERICARDITIS	
ASCITES	0.6
DYSPNEA	0.8
CHEST_PAIN_AT_INSPIRATION	0.8
EDEMA	0.7
HIGH_FEVER	0.7
PERICARDIAL_RUBBER	1
HEPATOMEGALIA	0.6
JUGULARY_TURGESCENCE	0.9