Medical Applications of Enhanced Rule-Based Expert Systems

Lenka Lhotska, Vladimir Marik, Tomas Vlcek
The Gerstner Lab, Czech Technical University in Prague
Technická 2, 166 27 Prague 6, Czech Republic
e-mail: lhotska@feld.cvut.cz

Keywords: diagnostic expert system, blackboard control, taxonomy control structure, inheritance mechanism, multi-agent approach, medical applications

Summary:

The paper describes several types of efficiency enhancements of "classical" rule-based diagnostic expert systems: The blackboard control structure enables to explore more knowledge bases of the same syntax in parallel, the taxonomy structures make fast zooming of attention possible and provide additional inference mechanism based on inheritance principles. The applicability of the enhancing techniques is documented by four case studies exploring the extended FEL-EXPERT shell in different tasks of medical decision making. The authors consider the enhancing techniques as useful steps on the way from "classical" diagnostic expert systems towards more complex multi-agent decision tools.

1. INTRODUCTION

Even nowadays most of the practically applied expert systems are simple, rule-based diagnostic ones. They belong to the first generation of expert systems. From the programming technology point of view they represent quite a specific class of expert systems based typically on the inference net approach. Unfortunately the simple first generation expert systems are not sophisticated enough to solve complex tasks in a wider problem area, most of all because they are limited to exploitation of the surface (shallow) knowledge, and the knowledge representation (based on the inference net, only) leads to cumbersome behaviour of the system in the case the volume of relevant knowledge exceeds a certain threshold (if there are more than, let's say, 500 rules to be taken into account).

The advanced expert system technology explores – besides the original idea of propagating uncertainty in an inference network – additional AI techniques which significantly enhance the expert system capabilities. Specific methods and techniques have been developed to overcome the lacks of simple expert systems mentioned above, e.g.:

- Different <u>context/control</u> links (having a nature of meta-rules) have been introduced to ensure both the common sense ordering of questions as well as the efficient cutting-off of irrelevant parts of knowledge bases in the early stages of a consultation run.
- Various <u>blackboard</u> control structures which enable to split the knowledge into more knowledge bases (knowledge sources) and to cope with this set of knowledge bases in parallel. The WHEN-THEN rules are applied on the blackboard level to control the exploitation of the knowledge base set. Even a very simple blackboard control adds a lot of valuable features to the system. It enables, e.g., solving of monitoring tasks, problems requiring non-monotonic reasoning, zooming of attention etc.
- <u>Combination of different knowledge representation schemes</u> with the stress to enable representation of causal, deep knowledge. It is possible to combine the inference net with the semantic/taxonomy net, with a subsystem for qualitative simulation, etc.

The goal of this paper is to show the capabilities of rule-based diagnostic expert systems when certain extensions/modifications of the original inference net idea have been applied and implemented. The paper tracks – in some sense - the research which has been carried out under the FEL-EXPERT Research Project at the Gerstner Laboratory [1], [2], [3]. The applicability of different enhancements is documented in the form of experimental case studies.

The FEL-EXPERT is a family of diagnostic rule-based shells using a Prospector-like model for uncertainty handling. Basically, three types of knowledge representation elements are used in the knowledge base notation for the basic system versions:

- Production rules having the following $E \rightarrow H$ form: <evidence E> $\rightarrow$ <hypothesis H>, where both E and H are propositions. The sufficiency and necessity measures are assigned to each rule. The graphical representation of a set of production rules forms the inference net.
- Logical functions (enabling logical combinations of propositions).
- Context links: by means of a context link the investigation of the given node of the inference net (as well as of all the nodes belonging to the subtree rooting in the given node) may be blocked, if the actual probability/validity value of the context is outside the prespecified probability interval.

In the case the knowledge base covers a wider area and is huge enough (for instance, it contains several hundreds/thousands top hypotheses), the main problem is that of focusing attention on the proper part of the knowledge base according to the data at hand. There do exist several possibilities how to solve this problem. In the FEL-EXPERT research project we have concentrated at two different methods of knowledge organization and exploitation, namely at the blackboard architecture and exploitation of taxonomy structures.

Various techniques enhancing the original idea of rule-based expert systems are described in this paper. These techniques are aimed at creating different knowledge capturing structures behind a simple inference net which enable to combine more inference engines operating on different principles in parallel.

There were several similar, problem oriented techniques developed and applied already. Let's mention the combination of a classical rule based system with the Case-Based Reasoning (CBR) technique in [4]. Very often we can find an ad hoc integration of an expert systems and neural network emulator, e.g. in [5], but many other examples can be found in [6] or [7].

Our approach partially explores the idea of the blackboard [8] architecture, partially the idea of the inheritance principle firstly used in MYCIN [9]. The described techniques reflect our experience in uncertainty processing (e.g. we tackle the problem of inheritance with uncertainty). The proposed approach is comparatively modest as it handles quite particular problems we have been facing in particular medical applications. It seems to be suitable for smaller expert systems operating locally as a part of the more global, multi-agent based decision making tools.

2. BLACKBOARD CONTROL

In the case of a blackboard control structure, the knowledge is split into comparatively small knowledge bases and a (meta)knowledge control structure - blackboard [8] - is used for control of the exploitation of the set of the knowledge bases. Each partial knowledge base is considered as a separate knowledge source. The blackboard consists of two sections: a control section and a data structure. The control section contains (meta)knowledge used for accomplishing different "global" actions over the knowledge source set (switching-over the knowledge sources, changing the global hypotheses to be investigated, etc.). The data section serves as an information transfer medium among different knowledge sources.

The (meta)knowledge on the blackboard level is expressed in a form of WHEN-THEN rules [3]:

WHEN <situation> THEN <action> {<action>}*,

where situation is expressed as a logical combination of elementary conditions (a typical example of an elementary condition: the actual probability of a proposition U_i is greater/lower than a prespecified treshold t, $P(U_i)>t$ or $P(U_i)<t$ respectively).

Two control strategies are combined at the blackboard level: *demons* and *agenda.*

There are two (basic) stacks available to ensure the agenda strategy:

- the A-stack (an associative stack) can be used to create a list of partial hypotheses to be investigated, this stack is formed/changed by specific actions defined in the WHEN-THEN rules (associative links);
- the B-stack (a knowledge base stack) expresses the natural ordering of the knowledge base invoking; this ordering may be changed dynamically during the course of a consultation by means of the WHEN-THEN rules.

The possible actions on the right hand side of the WHEN-THEN rules have been specified with the aim to ensure both the strategies mentioned above. The following actions are considered:

- excluding of an arbitrary hypothesis/knowledge source from investigation,
- forming/changing the content of the A/B-stacks,
- focusing the attention to the specific hypothesis/knowledge source,
- quitting the investigation of the given hypothesis/knowledge source, etc.

In connection with the blackboard structure a very important feature - the ability of non-monotonic reasoning - has appeared by introducing so-called R-action. This action ensures the re-initiation of the pre-specified knowledge source(s) and the reactivation of the proper blackboard rules. If the intermediate results based on previous user's assumptions (expressed as the user's data) become later unacceptable, it is possible to re-initiate some of the knowledge sources and to explore them once more, new data (new assumptions) being taken into consideration. The corresponding part of the blackboard is re-initiated at the same time. The WHEN-THEN rules containing R-actions may influence the content of the control part of the blackboard (reactivate some of the WHEN-THEN rules). They can be viewed as rules expressing specific knowledge of the knowledge contained on the blackboard level. From this consideration follows, that the WHEN-THEN rules containing the R-actions on their right-hand sides are meta-rules (with respect to the blackboard level) or meta-meta-rules (from the viewpoint of the knowledge sources). Introduction of a new knowledge level enables to accomplish very complex and effective control strategies in a simple way.

3. TAXONOMY NET AND ITS EXPLORATION

In many applications, especially in the area of medical diagnostics, valuable knowledge may be contained in the taxonomy hierarchy. This idea has been used for the first time in MYCIN [9]. The taxonomy can be considered as a tree-like structure of IS-A hierarchical links. The idea behind implementing the taxonomy in the FEL-EXPERT shell is to make the focusing/zooming in a transparent as well as efficient way. After putting only several fundamental questions an expert is usually able to penetrate through several levels down in the taxonomy tree and to focus his/her attention to a smaller knowledge sub-area and to delete the majority of possible conclusions from consideration.

The situation might be more complicated in the cases when the taxonomy dependencies are not so strict. Vagueness in the taxonomy requires to introduce some uncertainty processing model into this structure.

The idea to combine the "classical" inference net with (even unstrict) taxonomy structures is our attempt to improve the efficiency of the diagnostic FEL-EXPERT shell in solving a specific class of diagnostic tasks - tasks with a high "number of top hypotheses to number of propositions" ratio. The taxonomies are used to reduce the number of top hypotheses to be investigated, but a substantial part of the "blocked" hypotheses is evaluated in parallel as a side effect. This fact enables to offer the user also promising alternative solutions. In comparison with the taxonomy structures used in MYCIN, the advantages of our solution are:

- the user has the possibility to use the taxonomy in an active way,
- the taxonomy may be an "unstrict" one, the exceptions in "pure" taxonomic rules being expressed in the form of meta-taxonomic rules,
- more taxonomy nets may be applied in parallel,
- alternative solutions produced as a side effect may be offered to the user.

3.1. TAXONOMY CONTROL STRUCTURE

Trying to develop a larger application, it has been proven that the first generation expert systems do not have tools powerful enough to process large knowledge bases. The reason is that their operation is based on representation of the shallow knowledge. For shallow knowledge, it is characteristic that it is **associative,** i.e. expressing relationships among various entities (objects, phenomena) without any expert reasoning, and **phenomenological,** i.e. expressing "exterior" of knowledge. This means that the corresponding rules usually have the form

<observable feature> → <possible conclusion>.

The knowledge hidden behind the shallow knowledge is regarded to be deep and is embedded into the system most frequently in the form of causal and hierarchical dependencies, or possibly in the form of more or less complete causal, qualitative or structural models. Let us consider the knowledge representation in the form of explicitly formulated production rules, as it is usual in the rule-based diagnostic expert systems (e.g. in the FEL-EXPERT systems). Sufficiently great number of rules formulated like this, if consistent and complete, gives the possibility to reach a correct problem solution through their application. If, in addition, they enable to specify, which propositions are more general and which are specialisation of the others, supplementary to each other and which exclude each other, alternative - then it is possible to structure and organise the knowledge into hierarchies. For evaluation of taxonomy dependencies, it is suitable to introduce certain formalism [1]. Three following basic operators can be defined, namely

- **TOG, a taxonomy operator of generalisation:**

  Let $M = \{E_1,...,E_n\}$ be a set of entities of a given taxonomy (see Fig. 1) and A an immediate generalisation (abstraction) of E_1 to E_n. Then $TOG(M) = A$ focuses attention to immediate generalisation of entities E_1 *to* E_n.

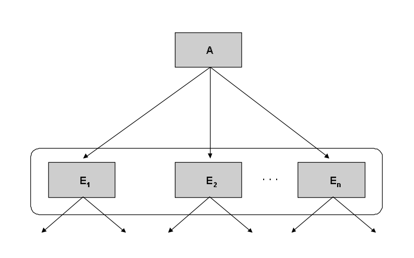

Figure 1

- **TOS, a taxonomy operator of specialisation:**

  Let A be an entity of a given taxonomy and $M = \{E_1,...,E_n\}$ (see Fig. 1) is a set of immediate specialisations A, i.e. A is generalisation (abstraction) of E_l to E_n. *TOS(A) = M* focuses attention to immediate specializations of the entity *A*.

- **TOP, a taxonomy operator of intersection:**
  Let $M = \{E_1,...,E_n\}$ be a set of entities of a given taxonomy, for whose elements it holds that they are not immediately taxonomically linked to each other (see simplified Fig.2). Let F_1 to F_n be sets of immediate specializations of E_1 to E_n, thus it holds that $F_i = TOS(E_i)$ *for* $i = 1..n$. Then $TOP(M) = E_1 \cap E_2 \cap ... \cap E_n$ focuses attention to common specialisation of the set *M*.

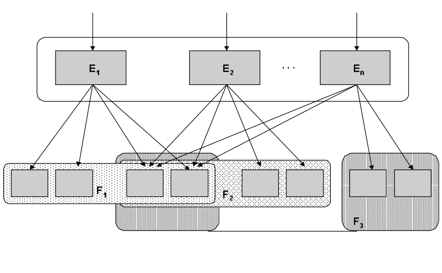

Figure 2

The suggested concept of focusing (zooming) mechanism has to be general enough to enable the exploration of more taxonomy structures constructed in parallel. The approach is based on the following methodology: at the knowledge base design phase, taxonomy structures are specified and each of them represents taxonomization of goal hypotheses from different points of view. Incorporation of these taxonomy structures into the early stage of the decision-making process enables to concentrate during the rest of the consultation only on those hypotheses belonging to a narrower set of hypotheses satisfying the taxonomy-based constraints. The situation is illustrated on the Fig. 3, where two taxonomy structures are defined for the goal hypotheses. From the point of view of global strategy of the expert system, two layers may be seen, namely the layer of the inference net and the layer of taxonomies for the attention zooming. At the very beginning of any consultation, the user specifies the consulted case by means of entities in the corresponding taxonomy classes.

In the case one taxonomy is chosen for inference control there is applied the TOS operator to this taxonomy class at the very beginning of a consultation. The TOS operator marks all hypotheses determined for examination. If the chosen taxonomy class contains further subclasses the TOS operator is applied recursively again. If more taxonomy classes are chosen for inference control then the TOS operator creates a set of all their subclasses and the TOP operator is applied to this set. In this way it is ensured that when the inference is controlled by more taxonomy structures at the same time only those hypotheses are chosen for examination that belong to all chosen classes.

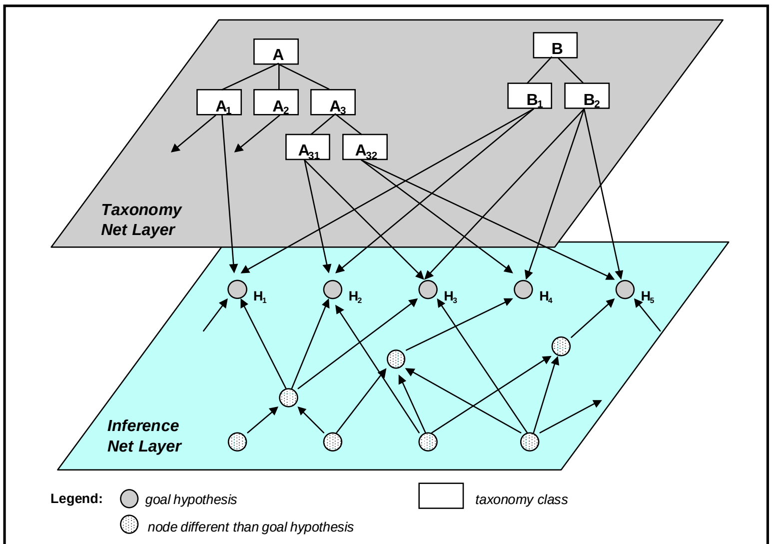

Figure 3

Let him/her specify, for example, that he/she is able to classify the case into the class A3 in the taxonomy A and into the class B2 in the taxonomy B (Fig.3). The zooming mechanism builds up a set of hypotheses on the level of inferencing - potential candidates for investigation, belonging to the given class, for each selected taxonomy class. In our case, the class A3 has been selected in the taxonomy A: by applying the taxonomy operator of specialisation TOS, the hypotheses of the subclass A31 (H2, H3) and the subclass A32 (H4, H5) are marked as adepts for investigation. Similarly, the class B2 has been selected in taxonomy B, i.e. the adepts for investigation are hypotheses H3, H4, H5. With respect to the statement that the consulted case satisfies criteria A and B in parallel, attention will be focused (by means of the taxonomy operator of intersection) on hypotheses belonging to the classes A31, A32 and B2, i.e. the hypotheses H3, H4, H5 will be investigated, only. Later on, the inference mechanism operates in the following way: in the phase of selection of a goal hypothesis for investigation, only those hypotheses are considered which were selected by the zooming mechanism. The taxonomy layer works - in principle - as a certain “mask” for the inference mechanism in the phase of goal hypothesis selection. However in the phase of inferencing it does not block the propagation of user's answer through the whole inference net. This approach has an advantage in keeping forward chaining strategy, but in a reduced state-space.

3.2. INHERITANCE MECHANISM FOR A RULE-BASED DIAGNOSTIC EXPERT SYSTEM

Motivation for the implementation of at least a simple inheritance mechanism exploring a certain taxonomy structure in rule based expert systems is obvious: in every non-trivial knowledge base dedicated for a rule based diagnostic expert system, a number of hierarchical dependencies among individual statements can be found. If these statements are represented by askable nodes, it is unpleasant - from the user's point of view - to be forced to provide data that are derivable from previously provided answers. The inheritance mechanism, as it is proposed, is understood as a second inference mechanism of the rule based diagnostic expert system closely co-operating with the main "classic" inference mechanism. Let us consider that the knowledge is represented by modified production rules in the knowledge base. Then the inheritance mechanism can be considered as an inference mechanism working, in principle, with two metarules, namely with the inheritance after a positive answer and with the inheritance after a negative answer. In the graph theory, each considered inheritance hierarchy is represented by a graph of a root tree type. Regarding that, the metarules can be expressed as follows:

- The metarule *inheritance after positive answer* is applied if the askable node, corresponding with the vertex v_i of inheritance hierarchy, is answered positively. In this case, the answer can be generalised: All the nodes corresponding with the vertexes on the oriented path from the vertex v_i to the root v_k can be considered as positively answered.
- The metarule *inheritance after a negative answer* is applied if the askable node, corresponding with the vertex v_i of the inheritance hierarchy, is answered negatively. In this case, the answer can be specialised and all the nodes, corresponding with vertexes which are predecessors of the vertex v_i in any depth under v_i can be considered as the negatively answered nodes.

3.3. CONDITIONS FOR CO-OPERATION OF INFERENCE AND INHERITANCE MECHANISMS

If the inference net on one hand and the inheritance hierarchies on the other hand are understood to be two separated, mutually co-operating layers, representing encoded knowledge in the knowledge base, the situation can be illustrated by Fig.4. With respect to the fact that a hierarchy of concepts and knowledge in the layer of inheritance hierarchies is defined explicitly, it can be regarded for superior to the inference net from the point of view of knowledge representation (it is a kind of metaknowledge) and utilised for checking correctness of the knowledge base. By correctness of the knowledge base, its non-contradiction with respect to the inheritance hierarchy is understood.

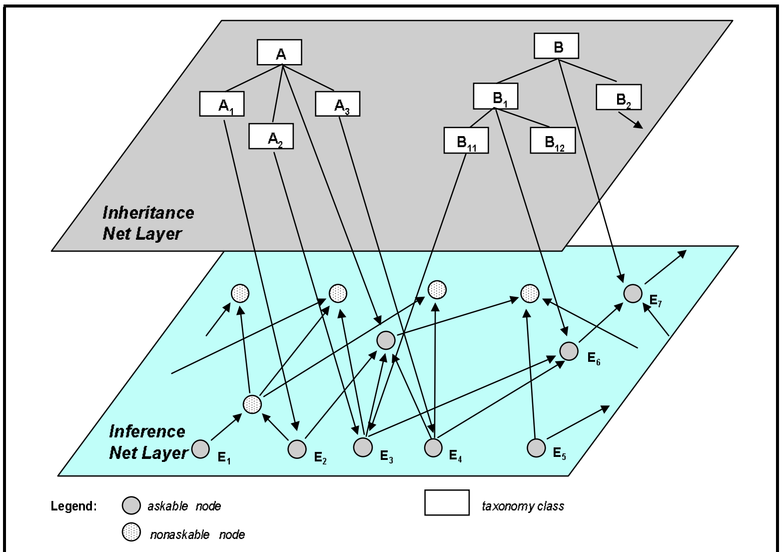

Figure 4

For the purpose of discussion concerning conditions of non-conflict co-operation of inference and inheritance mechanisms, it is advantageous to summarise premises that have been used: The knowledge base is built up from modified production rules. Its internal representation is implemented by an inference net that is represented by an oriented graph. The rule $E \rightarrow H$ is represented by an oriented arc, leading from the node E to the node H. The inheritance mechanism as it has been designed explores the inheritance net which is again an oriented graph. Every node in any inheritance hierarchy has unambiguously assigned the corresponding node in the inference net. The oriented arcs in the inheritance net represent IS-A links in the inheritance hierarchy (taxonomy). The arc orientation goes from more special to more general concepts. The cardinal problem of inheritance-inference co-operation is how to check whether the shallow knowledge encoded in the inference net is not in contradiction to the deep knowledge encoded in the inheritance net.

4. CASE STUDIES

4.1. CASE STUDY I - MYOPAT KNOWLEDGE BASE

This case study is oriented towards the genetic counselling area and results from the collaboration with the Genetic Counselling Center, 2nd Medical Faculty, Charles University, Prague.

The task was to diagnose three genetic syndromes: Duchenne's progressive muscular dystrophy, Becker's progressive muscular dystrophy and limb girdle dystrophy. All these syndromes are concerned with the muscular dystrophy of the lower and upper extremities. They are rare in the population (less than 0.02% in the Central Europe). Therefore, with the exception of several genetic laboratories, the experience of physicians in this area is rather limited. Usually, the clinical findings, as well as results of laboratory tests, are very similar. The principal difference is in the type of heredity and in the prognosis. In the early stage of defect manifestations (age 2-7) the symptoms are not distinct enough and it is difficult to produce the correct diagnosis, even for experienced geneticists. However, during these early stages the diagnosis is very important, since it is possible to estimate the risk of the birth of another affected child in the family and reconsider the consequences of another pregnancy. (At present, no therapy is known.) Later on, the syndromes are easier to be distinguished, as their symptoms become more evident and the progress of the disease provides additional information.

In the early stage, the clinical findings, results of laboratory tests and the pedigree information are the only data available for producing the diagnosis. Their mutual relations and their relations to the syndromes are described mostly by heuristic rules. The problem is well structured and, in our opinion, it was suitable for an expert system application.

For this purpose, the knowledge base MYOPAT has been designed, developed, implemented, repeatedly tested and permanently tuned. The final version of it contains 97 nodes, 153 rules and 24 contextual links . The depth of the inference net is 6.

The set of 35 cases has been used for debugging the MYOPAT knowledge base, another 113 cases were used for testing. In 112 cases the results achieved by exploring the FEL-EXPERT/MYOPAT system were identical with those provided by a panel of geneticists. In 9 cases which were not strictly decided by the physicians (the geneticists hesitated to express their final decision cathegorically), the final evaluations of the two "competing" syndromes were more or less similar. This fact demonstrates that the knowledge contained in the MYOPAT knowledge base emulates the physicians' knowledge in a reasonable way.

In the test set there were 7 pairs of brothers and/or sisters. They were investigated independently and the diagnosis was completely approved by the geneticists [11].

The possibility of explanation of the conclusions in terms of the most supporting and the most opposing symptoms has also been greatly appreciated. The FEL-EXPERT/MYOPAT system is currently used as a regular supporting tool for geneticists in the Prague Genetic Counselling Center.

4.2. CASE STUDY II - DIAGNOSTICS OF INHERITED METABOLIC DISORDERS

The first problem to be solved in collaboration with the Diagnostic Center for Inherited Metabolic Disorders, 1st Medical Faculty, Charles University of Prague, was the differential diagnostics of heterozygotes for hyperphenylalaninaemia I (classical phenylketonuria - PKU) [3]. This, with an incidence of about 1 in 6.000 is one of the most common inherited metabolic disorders in the Central Europe. High phenylalanine blood levels due to lack of hydroxylation lead, in untreated cases, to severe mental defects.

For genetic counselling purposes it is necessary to know which persons in the families with the occurrence of PKU are the heterozygotes. One of the most reasonable methods to detect the carrier status is the oral load test with L-phenylalanine. This test is based on an impaired hydroxylation of phenylalanine to tyrosine in heterozygotes as compared to healthy individuals. Carrying out this test is simple, but the evaluation needs in a deep experience. Many discriminants derived from the phenylalanine and tyrosine levels have been used all over the world. None of them is "strong" enough to distinguish completely the heterozygotes from healthy persons.

The knowledge base METABOL-AC which has been developed makes use of the experience of two experts and of the evaluation of 353 load tests which were previously performed. The rules express the interdependencies among the data (results of laboratory tests for instance), the values of the discriminants and the final decisions. To improve the efficiency of the decision making, the knowledge base has been improved by learning: The Fukunaga-Koontze method of feature extraction (originally developed in the area of pattern recognition) has been used for this purpose with advantage. As a result, new, far more efficient discriminants have been induced (they are of syntactic nature, of course) and, on the other hand, some of the rules have been cancelled.

The structure of the final version of the METABOL-AC knowledge base is presented in the Table 1. The 353 loading tests mentioned above have been evaluated by means of the METABOL-AC to test the reliability. The results were always thoroughly checked from different points of view by a panel of physicians. No substantial discrepancy was revealed. This system of evaluation of the loading tests with L-phenylalanine is routinely used, approximately 70 loading tests have been consulted on a routine basis.

The other and more serious and complicated problem can be found in deciding which laboratory tests should be used to establish the diagnosis of an inherited metabolic disease. The Diagnostic Center receives about 7.000 samples per year from the high risk population from all over the former Czechoslovakia. Only a small number of the samples arrive with a diagnosis designated; most of the samples are designated as being suspected of indicating metabolic disease without any precise specification. About 300 nosologic entities with primary metabolic defects are known.

The clinical findings and the course and time of manifestation are very different and are overlapping. It is usually impossible to establish the diagnosis of any inherited metabolic disorder from the clinical findings alone. The clinical findings enable the preliminary diagnosis (usually several diagnoses) to be determined. According to this preliminary diagnosis further investigations are performed to confirm or refute the diagnosis from the set of suspected ones. There are hundreds of methods, but the choice must be correct. Inherited metabolic diseases are not rare as a group (about 1-2 per 1.000 newborns). However, each of them is sufficiently rare that many clinicians may not have had considerable personal experience of dealing with it. Decision making is mainly based on the knowledge contained in very specialized handbooks.

There were two reasons for splitting the problem-oriented knowledge into several knowledge bases:

a) there are many nosologic entities and a vast volume of knowledge considered; and

b) the type of a disease can be fully confirmed only by biochemical tests, but the clinical findings have to be taken into account before these tests are recommended.

That is why the blackboard control structure has been applied and the structure of the knowledge base set had to be designed. The knowledge bases are arranged in two levels (see Fig. 5). The bases for evaluation of clinical findings have to be used first. The preliminary results obtained by them are exploited for the more precise diagnosis at the "deeper", second level. This level consists of knowledge bases for the evaluation of biochemical tests. Each of the "second level" knowledge bases has been constructed for precise decision making in a group of "related" diseases.

The total number of knowledge bases is not yet known (it might be around twenty), but only four of them plus the blackboard knowledge base METABOL have been fully implemented until now. There is one knowledge base on the first level (OB for the evaluation of biochemical findings) and three bases on the second one (MPK for the differential diagnosis of mucopolysachridoses, OK-2 for disorders of the metabolism of organic acids, and PK for disorders of the metabolism of purines and pyrimidines). Their structure is presented in the Table 1. Each partial knowledge base has been developed and tested separately.

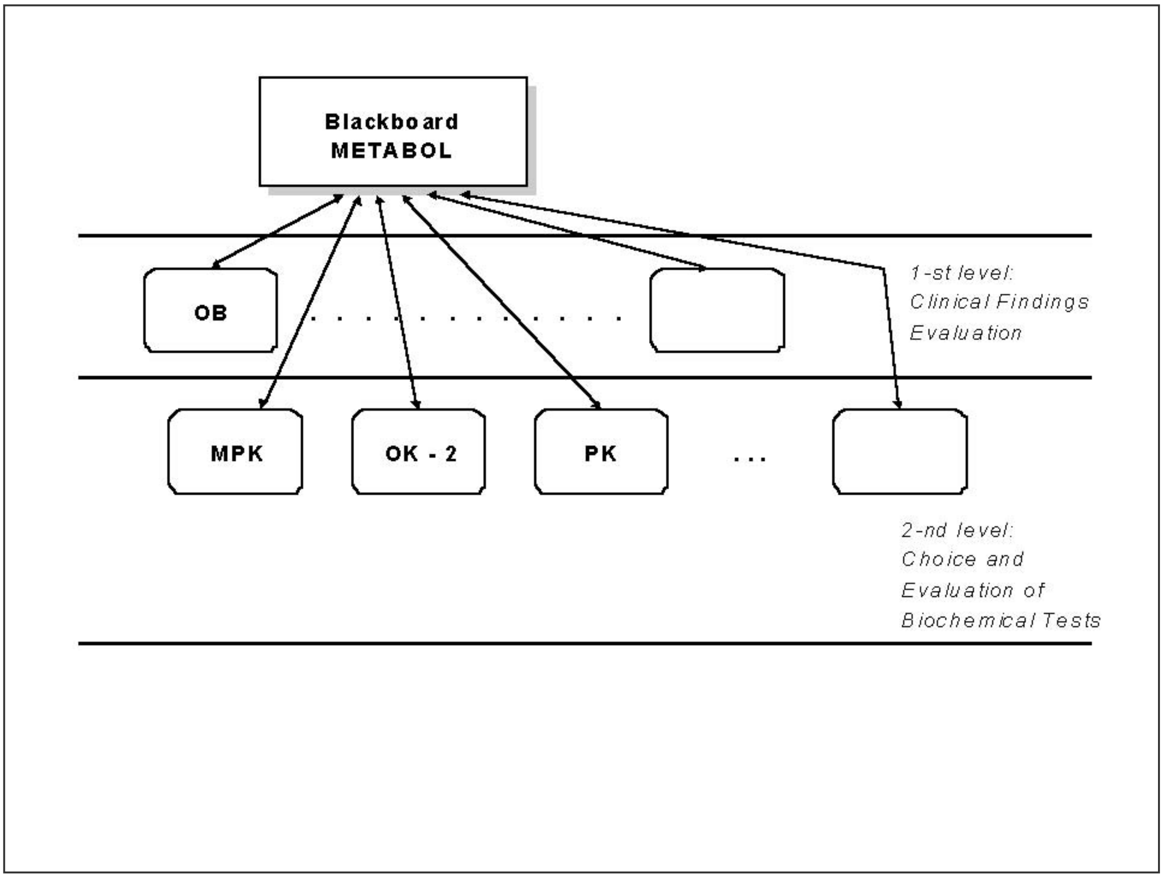

Figure 5

The METABOL blackboard knowledge base observes 26 partial hypotheses and contains 20 rules, for instance the following ones:

WHEN {The probability of the proposition "Changes of skeleton" < 0.1}
THEN {Start to investigate the knowledge base PK immediately},

WHEN {The probability of the proposition "Changes of skeleton" < 0.1
AND the probability of the proposition "It deals with the syndrome Sanfilippo" < 0.005}
THEN {The utilization of the MPK knowledge base is not relevant},

and so on.

The development of the other knowledge bases involved in the overall complex structure is expected to continue.

Table 1: The structure of the knowledge bases

Knowledge base	No. of nodes	No. of rules	No. of context links	No. of top hypothesis	No. of goals	Inference net depth
METABOL-AC	26	31	4	2	2	3
OB	87	135	31	15	15	6
MPK	82	130	27	8	8	5
OK-2	68	87	17	15	15	6
PK	78	82	22	17	17	5

4.3. CASE STUDY III – APPLICATION IN ONCOPATHOLOGY

If the knowledge covers a wide spectrum of diagnoses (several hundreds/thousands, typically), but the depth of the knowledge base is comparatively low and if some taxonomy in the diseases classification is available the taxonomy technique may be applied with advantage. Typically, in such cases, each feature influences a substantial number of top hypotheses (diagnoses).

A set of four knowledge bases covering all the diagnoses considered in the area of onco-pathology has been developed in cooperation with Hlava's Patho-Anatomical Institute, Charles University, Prague. The overview of the knowledge base set structure is presented in the following table (Table 2):

Table 2: The structure of the knowledge bases

Knowledge base	**T171.35**	**T171.40**	**PATOL.35**	**THEM.40**
No. of nodes	609	711	1421	160
No. of hypothesis	284	344	709	75
No. of rules	3740	357	716	85
No. of taxonomies	5	8	6	5
No. of taxonomic classes	152	130	226	89
Taxonomies	Topic Oncology Cytology Histology Organotype	Localization Morphology Bio-nature Makro-surface Makro-cut Description Organotype Growth-up	Localization Morphology Bio-nature Organotype Makro-surface Makro-cut	Localization Morphology Bio-nature Cyto-structure Histo-structure

Several taxonomies have been used to focus the attention inside the comparatively huge amount of knowledge. The system has been equipped with a multimedia module which enables the user to access the catalogue of pictures and the encyclopedia of more detailed description texts during the consultation. Thus, better and more qualified answers of the user can be achieved [1].

The application of more taxonomy structures in parallel has brought dramatical reduction in the number of the top hypotheses to be investigated during a consultation run (the system considers more than 1.800 top hypotheses – diagnoses). Also the number of questions put to the user was dramatically reduced (from 180-200 to 10-20 questions during one consultation run).

The system of knowledge bases is currently under testing in the Hlava's institute. It is expected that these bases in will be explored in a routine way.

4.4. CASE STUDY IV – DIAGNOSTICS OF THE VESTIBULAR ORGAN

In co-operation with the 2nd Medical Faculty, Department of Neurology, Charles University, Prague there is being developed a diagnostic expert system FELMOT for identification of vestibular organ diseases [2]. It is based on the FEL-EXPERT system shell. There has been designed and implemented the first version of a knowledge base containing available knowledge on the human vestibular system. The syndromes diagnosed by the system (till now) are the following ones: Menier disease, neuritis vestibularis (inflammation of the vestibular nerve), brainstem ischemia, vestibular schwannom, benign positional vertigo, perilymfatic fistula, sclerosis multiplex, and Cervicogen vertigo. These diseases may be accompanied by others. Their consequences can also be serious, e.g. unstability at forward or backward bend, danger of a fall and a subsequent injury or even unconsciousness.

There exist several methods for examination of the vestibular organ: audiometry, examination of the nervous system, cerebellum, vestibulospinal phenomena, stabilometry and electronystagmography. Besides that the medical doctor has to make personal and family anamnesis. Electronystagmography is one of the methods used for examination of the vestibular organ. It is based on monitoring reaction of the organ that is stimulated through various stimuli. There exist several kinds of electronystagmography, namely spontaneous, viewing, optokinetic, rotary and calorisation.

At present, the developed knowledge base contains 71 nodes, out of which 21 are leaf nodes, 42 intermediate ones, and 8 goal nodes representing individual diagnoses. There are 180 rules expressing relations between individual evidences and hypotheses representing state of the patient and his/her possible disease. The leaf nodes represent questions that can be answered by the doctor, e.g. subjective and objective complaints of the patient, found through various examination tests. There are statements concerning existence of vertigo and its kind (e.g. dependence on position), loss of hearing, its magnitude and manifestation (fluctuation), disorders of neocerebellum and paleocerebellum, affection of nerves (II., V. and VII.). Later on, the results of spontaneous, viewing and positional nystagmus are evaluated, positivity of fistula test and existence of tinnitus, migraine or imbalance. The goals represent individual diagnoses. Menier disease is influenced by vertigo, its fluctuation and spastic character. Neuritis is manifested by vertigo and imbalance. Brainstem ischemia is connected with vertigo and sudden loss of hearing. Schwannoma is connected with fluctuating loss of hearing, positional vertigo is mostly defined by vertigo and its dependence on the position. Perilymfatic fistula is strongly supported by a positive fistula test. Sclerosis multiplex is manifested through various disorders: nerves, cerebellum, C blockage of the spine and positive results of the viewing sys tem.

Experimental results: Recently the system has been tested on typical examples of individual diseases only. In all cases it has reached nearly 100 per cent correspondence with medical doctor's diagnosis. However, the testing of the expert system on more statistically significant sample of patients is expected.

5. CONCLUSIONS

The blackboard control architecture seems to be a very efficient tool suitable for highly sophisticated exploration of more knowledge bases in parallel. As a consequence of the two strategies which can be accomplished by the FEL-EXPERT blackboard, namely *demons* and *agenda,* it is possible to organize the sets of knowledge bases into hierarchical, heterarchical or sequential arrangements, to combine diagnostics with certain planning actions etc.

The FEL-EXPERT 3.5 system enables to construct arbitrary number of taxonomy structures over a single knowledge base. These structures may be freely used during a consultation run. The consultation may be run both with classical inference without using any of the taxonomy structures, as well as jointly with the inference controlled by one or more taxonomy structures in parallel

Introduction of taxonomy structures brings considerable decrease of number of investigated hypotheses. Therefore focusing attention to a smaller part of a knowledge base is made efficiently. As experiments in other application areas have shown, the decrease of number of investigated hypotheses can be really significant (from original 1000 hypotheses to 10-20 hypotheses).

The FEL-EXPERT system can be used not only as a diagnostic tool but also as an educational aid. For this purpose, it is advantageous that it is equipped with an extended explanation mechanism that enables to explain which symptoms support or reject the final diagnosis. In a certain way it may replace specialised libraries because accessibility of this information is surely higher. The possibility to execute external programs is another feature of this system and it is ideal for parallel interaction, for example, with a bibliographical database.

All the extensions and enhancements of "classical" rule-based diagnostic expert systems described in the paper represent steps towards more complex expert systems. The second-generation solutions consider separate knowledge-based modules of diverse nature, each being equipped by a specific inference algorithms. These modules are expected to be integrated by precisely specified, fixed links and relationships. The second generation expert systems are able e.g. to integrate "classical" expert systems with neural net emulators, genetic algorithm based modules and databases. The information transfer among the modules enabling the system integration is ensured by a *centralized integration program* running in the integration framework.

Both the first and second generation expert systems, even being significantly enhanced by very specific methods, do operate very efficiently, but usually only in a narrow field of expertise. It was demonstrated that the maintenance and up-grading modifications of these systems seem to be a hard-core, time-consuming problem. With the rapidly growing volume of available knowledge which can be stored in distributed databases accessible via Internet, the classical expert systems, including those described in this paper, seem to be too cumbersome and failing in solving the problems behind their scope of expertise.

Recently, the *multi-agent approach* [12] developed within the frame of the distributed artificial intelligence entered the scene. In this approach, each functional knowledge-based module operates autonomously and communicates with the other modules (agents) via message passing using the strategy "when needed". Each agent can be developed and maintained independently. There is *no central co-ordination unit* in the agents' community: all the control knowledge is distributed across the agents' community. The already existing and well running programs, including expert, neural networks, decision support systems etc. can be converted into the agents. During this "agentification process", the legacy software units are enriched by control and communication knowledge needed for their incorporation into the multi-agent community [13]. This knowledge is concentrated and organized in the agent's wrapper which is responsible for the behavior of an agent (encapsulated module) in the frame of the multi-agent community. The construction of wrappers is supported by the successful standardization efforts- currently they are already two standards used for the interagent communication, namely KQML and FIPA as a KQML extension. The idea of multi-agent systems enables to enrich the already existing distributed systems of decision making units whenever it is appropriated and by arbitrary modules needed. As all the knowledge maintenance is carried out locally. In practical applications it means that e.g. the knowledge needed for genetic counselling can be stored, maintained and explored locally, but offered for exploration globally, in the whole medical environment as well. Importance of the multi-agent approach is rather stressed by the possibility to explore it in the highly distributed Internet environment.

Acknowledgements
This article is based upon work sponsored by the Ministry of Education, Youth and Sports of the Czech Republic (under projects No. VS96047 "Gerstner Lab for Intelligent Decision Making and Control" and No. MSM210000012 "Transdisciplinary Biomedical Engineering Research").

References
[1] V. Campr, Z. Kouba, T. Vlcek, Information System for Onco-Pathology: From an Expert System to WWW, in: Proc. of the 6th International Conference DEXA'95, Springer Verlag, 1995, pp. 642-652
[2] L. Lhotska, M. Sorf, T. Vlcek, Application of the FEL-EXPERT System in Diagnostics of Vestibular Organ, in: Proc. of the 6th European Congress on Intelligent Techniques and Soft Computing EUFIT'98, Aachen: Verlag Mainz. 1998. pp. 1077-1081. - ISBN 3-89653-500-5
[3] V. Marik, L. Lhotska, Some Aspects of the Blackboard Control in the Diagnostic FEL-EXPERT Shell, in Proc. of the 5th International Conference on Artificial Intelligence and Information-Control Systems of Robots, Elsevier, Amsterdam, The Netherlands, 1989
[4] Y.M. Chae, S.H. Ho, C.S. Hong, C.V. Kim, Comparison of alternative knowledge model for the diagnosis of asthma, Expert Systems with Applications, 1996:11(4) 423-429
[5] D.V. Hillman, Integrating neural nets and expert systems. AI Expert, 1990:4;54-59
[6] U. Kraut, F. Mann, Current research on medical knowledge-based systems in Germany, in: Proc. of Medical Informatics Europe'96, 1996: 532-536
[7] J.I. Liebowitz (ed.), The handbook of applied expert systems. CRC Press LLC, 1998
[8] B. Hayes-Roth, A Blackboard Architecture for Control. Artificial Intelligence, Vol. 26, No. 3, July 1985, pp. 251-320
[9] B.C. Buchanan, E.H. Shortliffe, Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project. Addison-Wesley, Reading (Mass.), 1984
[10] G. Schreiber, B. Wielinga, R. de Hoog, H. Akkermans, W. van de Velde, CommonKADS: A comprehensive methodology for KBS development. IEEE Expert, 9, 6, 1994, p. 28-37
[11] T. Marikova, Current Ethical Problems in Genetic Counselling, in: Proc. 10th Int. Conference on System Research, Informatics and Cybernetics, Baden-Baden, 1998
[12] M.J. Wooldridge, N.R. Jennings (eds.), Intelligent Agents. LNAI No. 890, Springer Verlag, Heidelberg, 1995
[13] T. Hazdra, M. Pechoucek, ProPlanT: A Multi-Agent System for Operative Planning of Production. In: Proceedings: PAAM ´99. Vol. 1. London: The Practical Application Company. 1999. p. 315-324. - ISBN 1 902426 05 3