



0957-4174(94)00062-X

Knowledge Acquisition Tool for Case-Based Reasoning Systems

A. I. MECHITOV AND H. M. MOSHKOVICH

Department of Decision Sciences, Institute for Systems Analysis of the Russian Academy of Sciences, Prospekt 60 let Octiabrja, 9,
Moscow 117312, Russia

D. L. OLSON

Department of Business Analysis & Research, Texas A&M University, College Station, TX 77843-4217

B. KILLINGSWORTH

Department of Decision Sciences, East Carolina University, Greenville, NC

Abstract—*Knowledge acquisition is an important aspect of intelligent systems. The usefulness of such systems depends on system completeness and consistency. Case-based reasoning is a useful means of matching a natural human mode of dealing with complex problems through systematic recording of experience. However, this approach requires that representative case sets be considered in the knowledge acquisition phase, and once these case sets have been acquired, an efficient means of retrieving them is needed. M-CLASS is an ordinal classification model that gives structure to a knowledge acquisition environment where cases could be the source of knowledge. The M-CLASS system assures consistency and provides a means of checking for completeness. The system is demonstrated with a medical example.*

1. INTRODUCTION

KNOWLEDGE ACQUISITION is one of the most important problems in construction of knowledge-based systems. Tools and techniques have been developed to facilitate knowledge acquisition and aid the knowledge engineer in this process (Boose 1989, 1991; Gaines 1987; McGraw & Harbison-Briggs, 1989; Gaines and Boose, 1989; Reichgelt & Shadbolt, 1991; Scott, Clayton, & Gibson, 1991). However, knowledge acquisition can often become a bottleneck in the development of a knowledge base.

The acquisition of knowledge and its representation in a knowledge base generally is accomplished by means of production rules. Another way to approach the problem is through introduction of separate examples of specific objects. Rules generalize expert knowledge in a structured way (Hayes-Roth, Waterman, & Lenat, 1983). The

“IF . . . THEN . . .” production rules reflect expert logic in problem solving to some extent, making this approach highly popular in early knowledge-based systems. The disadvantage of production rules is that experts usually are unable to consider every possible situation that could arise. Experience shows that when people formulate rule bases for complicated situations, these rule bases tend to have gaps (Mechitov, Mashkovich, & Olson, 1994a, 1994b) and biases (Silverman, 1990), and when such gaps in the knowledge base occur, it is difficult to check such rule bases for consistency (Polat & Guvenir, 1991).

Another way to approach the problem is through introduction of separate examples of specific cases. The advantage of this approach is that the analysis and diagnosis of specific situations (examples) is more natural for experts, and does not require as much effort on their part. The experience of the expert consists of a lot of practical cases, situations, and other elements of his or her practical experience. These experiences can be used to formulate decision rules (Kolodner, 1991; Slade,

1991). It is much easier to remember, evaluate, and verbalize separate cases than to develop more abstract and generalized rules. Moreover, part of the expert's knowledge is maintained at the unconscious level (Larichev & Moshkovich, 1988; McGraw & Harbison-Briggs, 1989). In such cases experts know more than they are able to express. In many situations qualified experts can cope with complicated decision problems very well, but when asked to explain why they do or what they do, they face difficulties (Tversky & Kahneman, 1974; Kahneman, Slovic & Tversky, 1982).

Humans often reason by analogy, comparing new cases and situations with previous cases based on degrees of similarity (Helman, 1988). There has been widespread interest in case-based reasoning (CBR) systems where cases are used as a primary element of the knowledge base, and analogies between old cases and new cases are used as a tool for making decisions (Kolodner, Simpson, & Sycara-Cyransky, 1985; Hammond, 1986; Gentner, 1987; Ashley, 1991; Barletta, 1991; Rissland & Skalak, 1991). Humm, Schulz, Radtke, and Tversky (1991) applied CBR to a machining planning environment. Those authors noted that analyses of process planning show that specific planning tasks often repeat, either directly, or in some similar manner. They noted that human process planners usually rely on large collections of knowledge based on years of experience. These human planners often used reminders of old solutions for dealing with current tasks, and rarely found it necessary to generate entirely new plans. This approach is the basis of CBR. Mukhopadhyay, Vicinanza, & Prietula (1992) applied CBR to the difficult problem of estimating effort required to develop software for various applications. They concluded that the case-based approach is often used successfully in environments where no strong theoretical model for decision making exists, or where domain rules are incomplete, ill-defined, and inconsistent (Ashley & Rissland, 1987).

Two methodological problems have been highlighted in development of CBR systems (Kolodner, 1991). The first problem concerns development of a knowledge base consisting of a representative set of cases. The second problem concerns indexing cases, seeking a quick way to search among past relevant decisions.

If a large set of cases exists, representing the entire spectrum of anticipated situations, it would be reasonable to use case descriptions from some data base. Usually, however, it will be necessary to use experts as sources of case descriptions. The process of interviewing expert's for knowledge base construction can be done by having the expert form a set of cases relevant to the task under consideration, or by interviewing the expert using some knowledge elicitation system. There are clear limitations to the first approach having the expert select the cases. Usually experts can only remember part of the examples they have experienced, most reliably remem-

bering the most recent cases. Thus, there is a bias in the expert's judgment (Tversky & Kahneman, 1974, 1981; Kahneman, Slovic, & Tversky, 1982). It is difficult for experts to describe many of the cases they have experienced without help. Therefore, there is a need for a knowledge elicitation technique that would help organize the expert's interview in a systematic way, and also checking the consistency of the expert's judgments (Boose, 1989; Larichev, Moshkovich, & Rebrik, 1988; Larichev, 1992; Mechitov et al., 1994b).

In this paper we will discuss implementation of an ordinal classification technique (Ben-David, 1992; Larichev et al., 1989, 1991) for diagnostic case-based systems development. The technique is based on the interpretation of a diagnostic problem as a classification task, formulated in terms of a finite number of attributes describing objects, and classifying these objects into a finite number of diagnostic classes. The ordinal classification model used in this technique provides opportunities for systematic knowledge base construction, verification of expert knowledge, and effective indexing procedures for quick case search and reasoning.

2. SYSTEMATIC APPROACH TO CASE-BASED KNOWLEDGE ELICITATION

2.1. Diagnostic Task Description

Diagnostic tasks can be considered as classification tasks, where there are multiattribute descriptions of objects (cases) and a predetermined set of finite classes (or properties) that the objects need to be assigned to (Clancey, 1985). Such decision problems arise in a variety of tasks, for example, in medical diagnostics, machine repair diagnosis, loan applicant classification, personnel applicant review, and many other decision domains (Larichev & Moshkovich, 1990; Ben-David, 1992). Therefore, the diagnostic task can be described by a list of diagnostic properties (classes) and a discrete number of diagnostic signs or attributes used by the expert in a diagnostic process. The Cartesian product of attribute scales can be used to define all hypothetically possible multiattribute descriptions of an object.

Consider an example from medical diagnostics. Assume that the diagnostic system is designed to deal with several coma states, representing seven different decision classes (see Appendix 1). It is clear that an expert physician uses a set of symptoms (each with a finite number of values) to characterize a patient's condition. Assume that in order to make preliminary conclusions the expert-physician uses ten primary attributes (listed in the Appendix). For each attribute a scale of possible values is determined. The combination of all possible attribute values defines all possible patient states within this description. The diagnostic task is to classify each patient's state, defining appropriate disease or

diseases. In general, the patient may possess two or even more diseases at the same time. Therefore, in the process of classification it is necessary to decide for each possible object whether it must or must not belong to each class under consideration. As a result, objects are divided into two groups (membership or nonmembership) with respect to each class under consideration. As a result, objects are divided into two groups for each class: objects belonging to this class, and objects that do not belong to this class. We call these groups subclasses of the corresponding class.

It is also assumed that in the diagnostic process the expert can use a finite number of ordinal outcomes reflecting different levels of expert confidence that the particular object belongs to the assigned class. For example, in medical diagnostics classifying coma states, the physician may use three classes: high probability of a certain coma, moderate probability of a certain coma, and unlikely that this case is the certain coma. In the vehicle repair field, a mechanic could diagnose a transmission problem as being in one of the following states: transmission is in good condition, transmission should be operable but needs to be watched for deterioration, transmission is inoperable but can be made operable by changing fluid, or transmission is beyond salvage and needs replacement.

The diagnostic task we consider can be presented in the following formal way (Larichev et al., 1991).

GIVEN:

1. $P = \{P_1, P_2, \dots, P_L\}$ a set of classes that objects could possibly be assigned to;
2. $P_i = \{P_i^1, P_i^2, \dots, P_i^k\}$ a set of subclasses for class P_i ;
3. M , the number of attributes used to characterize objects;
4. $Q_m = \{q_{m1}, q_{m2}, \dots, q_{mn}\}$ a set of possible values on the m th attribute, where n is the number of possible values for attribute m , $n = n(m)$.
5. $A = Q_1 \times Q_2 \times \dots \times Q_M$ a set of all hypothetically possible objects;
6. $a_i \in A$ is a vector $a_i = (a_{i1}, a_{i2}, \dots, a_{iM})$ where $a_{im} \in Q_m$.

The task requirement is to identify the expert's definition of each object's performance level for all possible classes P . This means that it is necessary to identify the degree of membership of vectors a_i to the appropriate class P_i . For the medical task of diagnosing coma given in Appendix 1, there are seven possible outcome classes P_i , $M = 10$, n_4 and n_5 equal 3 whereas the other n_i equal 2. A consists of $2^8 \times 3^2 = 2304$ possible sets of symptoms a_i . The task is to classify possible patient states a_i into subclasses $\{P_i^1, P_i^2, \dots, P_i^k\}$ representing diagnoses (e.g. high probability of diabetic coma, unlikely that this is ichemic coma, etc.).

To further facilitate our description, we enumerate subclasses in decreasing order of membership (the lower the subclass number, the more of this characteristic the object has). In the coma diagnostic problem, for a set of

symptoms a_i we can rank order the possible diagnostic outcomes. Formally, this is equivalent to the definition of a binary relation:

$$s_l = \{(P_i^j, P_l^k) \in P_i \times P_l \mid i < j, i = 1 \text{ to } k_i, j = 1 \text{ to } k_l\}$$

(1)

The number of subclasses is usually not large (from two to five), because psychological research indicates that people do not usually consider more than this number of levels within a property (Miller, 1956; Larichev & Moshkovich, 1988). If there are only two subclasses, the expert wants to divide objects into those that have some property, and those that do not have that property. If more than two subclasses are used, the expert intends construction of a more fuzzy classification, introducing intermediate levels.

2.2. Problem Structuring

To capture the expert's decision-making process for the task at hand using the minimum number of questions, it is necessary to provide some structure for the task. The structure used in this paper is based on identifying the typical performance levels for an attribute considered by the expert (Larichev et al., 1991).

The attributes, which describe the possible states, can serve as a diagnostic tool when different attribute values have different typicalities for each class. It is common to categorize attribute performance into a finite number of ordered categories. For instance, in a medical diagnostic task involving coma, possible values on the attribute "skin state" are typically used. Dry skin is most likely to be due to diabetic coma, and moist skin is more typical for ichemic coma, and so on. In general, the rank ordering of the measure levels for different diseases (classes) would be different. In Appendix 1 the rank orderings for two types of coma—diabetic coma and hypoglycemic—are displayed. Number "1" corresponds to the most typical attribute value, "2" to the less typical value, and "3" to the least typical value. The first column corresponds to the typicality of each symptom for diabetic coma, and the second column to that of hypoglycemic coma.

Rank ordering of attribute values for each class defines a special binary relation over the set of attribute values. It is assumed that an expert, when comparing values q_{ij} and q_{ik} on attribute Q_i for class P_i , is able to select one of the following:

1. value q_{ij} is more typical of class P_i than is q_{ik} ;
2. value q_{ik} is more typical of class P_i than is q_{ij} ;
3. values q_{ij} and q_{ik} are equally typical of class P_i ;
4. values q_{ij} and q_{ik} are incomparable for class P_i .

For example, for diabetic coma, diminished pressure (less than 100) is more typical than normal pressure, and normal pressure is more typical than is high pressure. For other classes of coma (hypoglycemic, etc.) there can be a different ranking of these values.

Patient States					
## (PS1)		PS2		PS3	
Gradual coma development	=	Gradual coma development	=	Gradual coma development	
Skin is dry	--->	Moist skin	<---	Skin is dry	
Pathological respiration	=	Pathological respiration	--->	Normal/superficial respiration	
Rapid pulse	--->	Normal pulse	<---	Rapid pulse	
Diminished pressure	--->	Normal pressure	=	Normal pressure	
Diminished deep reflexes	=	Diminished deep reflexes	=	Diminished deep reflexes	
Diminished eyeball tension	=	Diminished eyeball tension	=	Diminished eyeball tension	
No convulsions	=	No convulsions	=	No convulsions	
No muscle rigidity	=	No muscle rigidity	--->	Muscle rigidity present	
Normal temperature	--->	Higher temperature	<---	Normal temperature	

FIGURE 1.

Therefore, we define a reflexive and transitive binary relation r_{il} on the set of attribute values Q_i for each attribute i and class l ($i=1$ to M ; $l=1$ to L). The statement $q_{ij}r_{il}q_{ik}$ means that value q_{ij} is not less typical for class P_l than is q_{ik} . Reflexiveness means that it is possible to compare equal attribute values. Transitivity means that $q_{ij}r_{il}q_{ik}$ and $q_{ik}r_{il}q_{im}$ imply $q_{ij}r_{il}q_{im}$ (if q_{ij} is more typical for class P_l than is q_{ik} , and q_{ik} is more typical for class P_l than is q_{im} , then q_{ij} is more typical for class P_l than is q_{im}).

Binary relations $r_{il}(i=1$ to M ; $l=1$ to L) define the binary relation R_l for each class P_l on set A :

$$R_l = \{(a_i, a_j) \in A \times A \mid a_{iq}r_{il}a_{jq} \text{ for all } m \in M; a_{iq}, a_{jq} \in Q_i\}.$$

Here a_{iq} is a vector over all attributes for a specific case. This relation allows comparison of multiattribute descriptions of objects. If each element of the object description is more (less) typical for a corresponding class than the corresponding elements of the second object, then the first object is certainly more (less) typical for this class than is the second object. Thus, for each class there is a quasiorder on the set of possible objects—comparing any two objects we can conclude that one of them is more typical than another; in this case we say that the first object dominates the second object, and the second object is dominated by the first object, or the objects are incomparable (if some elements of the first object are more typical than corresponding elements of the second object, and some elements of the second object are more typical than corresponding elements of the first object). For example (Figure 1), patient state PS1 is more typical for diabetic coma than PS2 because the symptom “Skin is dry” is more typical for diabetic coma than is “Moist skin,” “Rapid pulse” is more typical than

is “Normal pulse,” “Diminished pressure” is more typical than “Normal pressure,” and “Normal temperature” is more typical than “Higher temperature,” all following the situation given in Appendix 1. But patient states PS2 and PS3 are incomparable with respect to diabetic coma, because “Moist skin” is less typical than “Dry skin,” but “Pathological respiration” is more typical than “Normal respiration.”

In other words, binary relations r_{il} make some vectors from A comparable in each class. One vector is more typical for this class than another class if the vector has values that are not less typical on all components but one, and a value more typical on this component. Thus, the introduced structure allows comparison of some pair of vectors (objects) in each class. Vectors may be incomparable in this class if one vector has more typical values upon one subset of attributes and the other vector has more typical values upon another subset of attributes.

Let us emphasize that there are different quasiorders for each class corresponding to different rankings of attribute values for each class. That is why the same pair of objects can be comparable for one class (one is more typical than the other) and incomparable for another class. For example, object PS1 in Figure 1 is more typical for class “diabetic coma” than is object PS2. At the same time they are incomparable with respect to hypoglycemic coma, as for hypoglycemic coma “Dry skin” is less typical than “Moist skin” but “Normal temperature” is more typical than “Higher temperature.”

The structuring step results in two types of binary relations: R_l is the relation of dominance on set A for class P_l , and s_l is the linear order relation on the set of

subclasses for class P_i . Both of these relations reflect the typicality of different objects from set A for class P_i . That is why it is possible to discuss the mutual consistency of relations. Formally this is expressed in the following definition.

DEFINITION. Relations R_i and s_i are mutually consistent if (1) the expert says that a_i belongs to P_i at degree m (P_i^m) and there exists some a_j such that $a_j R_i a_i$ (a_j is at least as typical of class R_i as is a_i), then a_j should belong to class P_i^k , where the index of k is less than or equal to m .
 if (2) the expert says that a_i belongs to P_i^m and there exists a_j such that $a_i R_i a_j$ (a_i is at least as typical of class R_i as is a_j), then a_j should belong to P_i^k , where the index k is at least as large as m .

This definition infers that if object a_i is more typical than object a_j for class P_i , then object a_i is more likely than object a_j to belong to a class of P_i . In other words, the reflection of set A to set P is monotonic, as is shown in Figure 2.

Let us call classification of objects $a_i, a_j \in A$ for class P_i to be mutually inconsistent if their classification violates the mutual consistency of R_i and s_i . Therefore, in further analysis, we assume that conditions (1) and (2) should be satisfied in the process of knowledge base construction, and that violations of these conditions have to be due to inconsistency in expert judgments. Thus, in the task under consideration it is necessary for each class to construct ordinal classification. Such task formulation describes many diagnostic/decision-making problems, such as editorial decisions, housing selection, social worker judgments, medical diagnostics, and many others (Ben-David, 1992; Larichev et al., 1991).

The unique feature of this approach is that it provides a means to construct an effective algorithm for solution of some problems connected with CBR. In particular, 1) definition of a set of "effective" objects, whose classification is representative of the entire set of objects; 2) verification of expert judgment consistency; and 3) indexing (effective storage and location of cases).

2.3. Effective Procedures for Case Elicitation

First, let us note that classification of one object indirectly determines, to some extent, classification of

other objects from A as it limits feasible subclasses for other objects, connected with the classified object through the dominance relationship. For example, consider the medical diagnostic task given in the appendix. Assume the patient state described by the vector PS2, presented in Figure 1, has been classified by the expert-physician as a "high probability of a diabetic coma." We can then conclude that PS1 will also belong to the subclass "high probability of a diabetic coma," as this has more typical parameters for diabetic coma than PS2.

The same conclusion can be drawn for other cases that are connected by dominance relationships. Thus, classification of every object from A limits feasible subclasses for other objects. And if the set of feasible subclasses for some object from the set A contains only one subclass this object may be considered fully classified.

In the process of expert interview (while constructing the knowledge base) the expert can evaluate different objects, introduced by the expert or by the system, and their classification will lead to indirect classification of some other objects. The amount of information obtained as a result depends on both the object "position" in the corresponding quasiorders (corresponding to each class) and its concrete classification. For example, if the object is assigned to the first subclass of some class, and there are many objects that dominate it in this class, then all of them can be indirectly assigned to the same subclass, but if the number of such objects is small, then the amount of indirect information would be small too.

We therefore face a task of choosing the most informative objects, whose expert evaluation permits us to construct a representative knowledge base with a limited number of expert interviews. An algorithm based on a maximum criterion is proposed in Larichev et al. (1991). Initially, nonclassified objects are considered. The algorithm defines the number of indirectly classified objects potentially classified by an expert classification. Then the minimum of information needed is calculated. Thus, for each object not yet classified the minimum amount of direct and indirect information required is calculated. The last step of the algorithm chooses the object with the maximum guaranteed amount of this expected information.

Simulation of this algorithm proved that using the informative objects determined by the technique described above, a significant portion of objects in set A is classified by a relatively small number of expert classifications. For example, on average, classification of 70% of all objects requires evaluation of about 15% of these objects. Evaluation of all possible objects (i.e., construction of complete classification) requires classification of 25–30% of all objects (Larichev et al., 1991). As a result, even for large-scale diagnostic tasks, the approach gives a means to construct representative knowledge bases, providing classification of a considerable portion of possible objects.




FIGURE 2. Monotonic–nonmonotonic.

2.4. Consistency Checking

Knowledge base verification is one of the most important problems in building knowledge bases. It is clear that the quality of KBSs is mainly determined by the consistency of the KB. At the same time, there could be many reasons why expert judgments might be inconsistent, including problem complexity, fatigue, and casual expert error. Both field and lab studies indicate the possibility of having a large number of inconsistencies in expert judgments (Larichev & Moshkovich, 1988; Mechitov et al., 1994a, 1994b; Polat & Guvenir, 1991).

The task structuring method presented here provides a means to check the consistency of expert responses. The complete (all object subclasses defined) or partial (only some subclasses defined) object classification can be achieved by both direct expert classification of objects and indirect classification based on classification of other objects. If these two sources contradict each other, evidence has been uncovered of inconsistency in expert judgments.

Let us describe an example of such inconsistency. When evaluating object a_1 , the expert assigned it to the vector $(P_1^1 P_2^2 P_3^3)$, thus stating that there is a high probability of diabetic coma, moderate probability of hypoglycemic coma, and unlikely choropenic coma. Let object a_2 be more typical for the first class than a_1 . In this case we can conclude that object a_2 should also belong to subclass P_1^1 . But in general, classification of object a_2 cannot be completely defined as a result of previous classifications. The possible subclasses for hypoglycemic coma and choropenic coma cannot be determined. Let object a_2 be presented to the expert, who assigns it to subclasses $(P_1^2 P_2^1 P_3^3)$. Here the indirect classification into subclass P_1^1 does not coincide with the direct classification P_1^2 , which must be due to inconsistency in expert judgment for either a_1 or a_2 .

The next algorithm can be used to eliminate inconsistencies in expert judgments. Each object classification made by the expert is checked for consistency with previous object classifications. If a contradiction is identified, the contradictory knowledge fragments (the current object classification and the previous object classification) are presented to the expert for analysis with corresponding commentary. The expert can correct the last object classification and check its consistency with previous answers, or decide that one or more of the fragments of the current KB are wrong (i.e., prior classifications were erroneous). In the latter case the erroneous cases would be eliminated from the knowledge base. In general, these cases should be classified by the expert once again. This classification can be direct or indirect.

2.5. Indexing Search

The ordinal attribute scales give an opportunity to effectively store and search for objects "similar" to an




FIGURE 3. Binary relations between borders' elements.

object currently being analyzed. For this task the principle of Pareto sets from decision theory (Geoffrion, 1968) is used. The suggested approach is based on the implementation of upper and lower borders of subclasses. In this case it is not necessary to maintain data on all objects classified by an expert, but only on a key subset of these judgments.

Consider a set of all objects assigned to a certain subclass of some class. The upper border of this subclass contains all nondominated objects from this subclass according to the binary relation R , meaning that for these objects there are no objects from this subclass that dominate them. At the same time let us define the second subset of objects for which there is no object in the set dominated by members of the second subset. Let us call this second subset the lower border of this subclass (Figure 3).

In other words, the upper border of the subclass contains all nondominated (i.e., most typical) objects from this subclass according to the binary relation R . Analogously, the lower border of the subclass contains all objects from this subclass that do not dominate any other object from this subclass according to the binary relationship R (or the least typical objects of this subclass).

It is evident that keeping only these two subsets enables identification of all objects belonging to the subclass. To check the presence of an object, one has to check whether there is an object in an upper border that dominates the object, and whether there is an object in the lower border that is dominated by the object. If both conditions are satisfied, it can be concluded that the current object must belong to this subclass. Actually, as the current object is dominated by the object from the upper border, then it should belong to the same subclass or to a subclass with a higher number because otherwise its classification would be mutually inconsistent. Analogously, as the current object dominates even one object from the lower border, it should belong to the same subclass or subclass with a lower number. Thus, the current object can only belong to this subclass.

This means that it is sufficient to keep only the borders for each subclass, and that it is not necessary to record the classification of all objects classified, because the upper and lower borders of subclasses for each class

provide us all information on feasible subclasses for each object. To define these feasible subclasses for some object a_i in a certain class it is possible to use the next algorithm:

- Object a_i is compared sequentially with objects of the upper border of this class, beginning with the last subclass, until object a_j , belonging to subclass l and more typical than object a_i is identified. Then feasible subclasses for object a_i are $\{l, l + 1, l + 2, \dots, k\} = G1$ (where k is the number of subclasses for the given class).
- Object a_i is compared sequentially with objects of the lower border of this class, beginning with the first subclass, until object a_j , belonging to subclass t and less typical than object a_i is found. Then feasible subclasses for object a_i are $\{1, 2, 3, \dots, t\} = G2$.

If classification of an object is mutually consistent, then for each object intersection of sets $G1$ and $G2$ is not an empty set.

To facilitate search and storing of objects, it is possible to construct singular correspondence between multiattribute objects' descriptions (vectors) and numbers. Therefore, we can operate predominantly with numbers, converting them into vector description and vice versa when it is necessary.

3. SYSTEM IMPLEMENTATION

The knowledge acquisition system described has been implemented in an interactive system called M-CLASS (MultiCLASSification) consisting of three blocks:

- knowledge base,
- analog inference engine
- knowledge acquisition block.

The knowledge acquisition block directs the expert interview, defining the next most informative object for expert evaluation, checks consistency of different fragments of expert knowledge, and processes case classification, thus filling the systems knowledge base. This knowledge base consists of three main components: ranks of attribute values, classified objects, and an explanation subblock.

The logical structure of the system is outlined in Figure 4. The first step of the expert interview is to rank order all possible values on a separate attribute according to their typicality for each property (class). This means that for each property (class) P_i and attribute Q_i , the expert makes pairwise comparisons to determine which of two possible values on attribute Q_i is more typical for class P_i . Therefore, the rank ordering of all possible values q_{ij} , $j=1,2,\dots,n_i$ is constructed. These rank orderings are reflected in the binary relation r_{il} ($i=1,2,\dots,M; l=1,2,\dots,L$).

The second step consists of direct introduction of expert information into the knowledge base. The M-CLASS system provides two primary means of working with the expert: active and passive. In the active regime

the expert defines the fragment of his/her knowledge as a description of a case and its corresponding classification. In the passive regime the role of the communication initiator is transferred to the system. The system analyzes the current knowledge base, identifies unclassified objects, evaluates them from the point of view of information content, and presents it to the expert for classification (Figure 5).

In all regimes, the fragments of knowledge introduced into the system are checked for consistency with previous information. A newly classified object is sequentially checked for consistency with previously introduced cases. If the new case is totally consistent with the current knowledge base, it is added to the knowledge base. If a contradiction is detected, the contradictory knowledge fragments are presented to the expert for analysis (Figure 6). The expert can correct the last fragment and return to checking its consistency, or decide that one or more of the knowledge fragments in the current knowledge base (i.e., previously classified objects) are wrong. In the latter case, erroneous cases will be eliminated from the knowledge base, restoring it to a consistent state.

We note that in M-CLASS the process of analysis and elimination of contradictions between different case classifications is accomplished through a unified formal scheme. Elimination of contradictions is accomplished in the process of knowledge base construction in an on-line mode. After the introduction of each new fragment of expert information, this information is compared with previously built knowledge base elements. If the expert

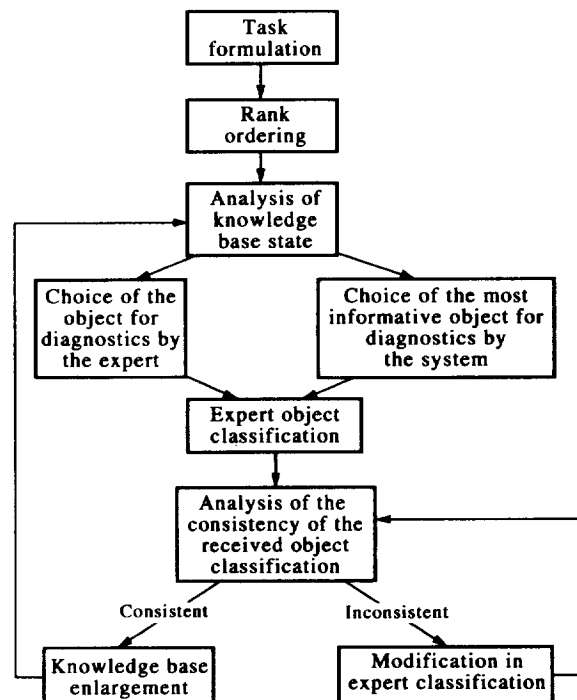


FIGURE 4. Logical structure of knowledge acquisition process.

oped diagnostic system was designed for rough preliminary diagnostics of coma type. For detection of seven types of coma in Appendix 1, ten attributes were used, such as palpation, body temperature, etc. For each attribute two or three of the most important values (from the perspective of differential diagnosis) were defined. All attribute values were formulated in a verbal form, two qualitative gradations being used for the quantitative attribute "body temperature."

In the first stage the expert ranked attribute values in terms of their typicality for corresponding comas. As the number of values for all attributes was small, the expert constructed strict linear orders for all attribute scales.

In the second stage the expert diagnosed different states, described in terms of the ten attribute values in Figure 5. The interaction between the system and the expert was held in a limited natural language: the presented situations were particular cases, and the expected answers were formulated as traditional diagnostic conclusions. For all comatose cases the expert

used three degrees of suspicion—high probability, moderate probability, and unlikely. Therefore, diagnosis was conducted as follows: "In this state there is a high probability of chloropenic coma, weak probability of eclamptic coma, and unlikely prospects of other types of coma."

The expert interview lasted 4 days. During that time about 300 states were considered directly by the expert physician, yielding construction of a complete knowledge base. The expert decided 24 states unlikely to be feasible. According to the attributes of these infeasible states, about 600 possible states of the 2304 possible were discarded.

Contradictory situations where expert diagnosis of current cases contradicted previous entries occurred 28 times. All of these contradictions were analyzed by an expert and reevaluated. As a rule, such errors were caused by the physician marking only one type of coma with a high probability while overlooking other possible types of coma that actually had a moderate probability.

<u>Last situation</u>		<u>Previous situation</u>
Gradual coma development	=	Gradual coma development
Skin is dry	→	Moist skin
Pathological respiration	=	Pathological respiration
Rapid pulse	→	Normal pulse
Diminished pressure	→	Normal pressure
Diminished deep reflexes	=	Diminished deep reflexes
Diminished eyeball tension	=	Diminished eyeball tension
No convulsions	=	No convulsions
No muscle rigidity	=	No muscle rigidity
Normal temperature	→	Higher temperature

moderate probability of a diabetic coma	high probability of a diabetic coma
---	-------------------------------------

According to your answers, in LAST situation You suspect MODERATE probability of a diabetic coma, and in this previous situation You suspect HIGH probability of a diabetic coma, THOUGH LAST situation is MORE TYPICAL for diabetic coma, than the PREVIOUS one.

- You can
1. Change your last answer (reevaluate last situation)
 2. Change your previous answers, contradicting to last one.
 3. To postpone the last situation evaluation.

Your choice > ...

FIGURE 6. Fragment of expert dialogue (analysis of contradiction).

The expert found the consistency analysis provided by the system to be very useful, helping him to be more attractive during the expert interview process. The 4-day expert interview session ended with construction of a complete knowledge base providing diagnosis for all possible patient states described in terms of the ten attributes.

Small medical diagnostic tasks as addressed in this example can be used as part of more complicated diagnostic systems, or as rough tools for preliminary analysis. In case of large tasks the system provides a quick and effective tool to construct a representative set of cases that guarantees classification (decision) of many possible objects (situations). On the other hand, different types of decomposition can be used to represent the initial problem as a set of such classification tasks of smaller dimensionality. The most general approach consists in implementation of hierarchical attribute structures. In this case the initial set of attributes is divided into several subsets with corresponding names and preliminary classifications are constructed for each subset. Then the names of the subsets with corresponding scales are considered as more general attributes, and the expert formulates decision rules (case classifications) that can be applied to them.

The M-CLASS system is a modification of the software system CLASS (Larichev & Moshkovich, 1990). CLASS was developed for selection problems under conditions of multiple criteria. M-CLASS software runs on a conventional PC framework.

5. CONCLUSIONS

Case-based approaches provide a natural means for experts to express their analysis and diagnosis of specific situations. The use of this approach has resulted in the development of many recent systems to aid in many decision supporting tasks. The case-based approach has especially been useful when there is not an underlying theoretical model for decision making, or where domain rules are incomplete, ill-defined, or inconsistent.

However, the case-based approach to knowledge acquisition suffers from two methodological problems. The knowledge base obtained through CBR may not be representative of all possible cases. And once the knowledge base of cases is built, there is a need for an efficient means to search among past relevant decisions.

This paper presents system M-CLASS, developed in the Institute for Systems Analysis of the Russian Academy of Science, that gives a means of addressing these two problems in case-based knowledge systems. M-CLASS treats the decision problem as a selection problem. This requires restating the problem in the form of a multicriteria selection problem with finite states for each criterion. This structure enables complete coverage of all possible combinations of states, thus enabling assurance of total representation of the decision problem.

The method carefully assures consistency of expert input during the knowledge acquisition phase. The problem structure also provides an efficient means of storing and retrieving cases within the system.

Future work is planned to test the system on subjects in both Russia and the U.S. The system is being applied to real problems in Russia. If readers wish to apply the system, please contact Dr. Mechitov or Dr. Moshkovich.

REFERENCES

- Ashley, K. (1991). Reasoning with cases and hypotheticals in hYPO. *International Journal of Man-Machine Studies*, **134**, 753–796.
- Ashley, K., & Rissland, E. (1987). Compare and contrast: A test of expertise. *Proceedings of AAAI-87*, Seattle, WA, pp. 273–278.
- Barletta, R. (1991). An introduction to case-based reasoning. *AI Expert*, **6**(8), 42.
- Ben-David, A. (1992). Automated generation of symbolic multi-attribute ordinal knowledge-based DSSs: Methodology and applications. *Decision Sciences*, **23**(6), 1357–1372.
- Boose, J. H. (1989). A survey of knowledge acquisition techniques and tools. *Knowledge Acquisition*, **1**, 3–38.
- Boose, J. H. (1991). Knowledge acquisition tools, methods, and mediating representations. In H. Motoda, R. Mizoguchi, J. H. Boose, & B. R. Gaines, (Eds.), *Knowledge Acquisition for Knowledge-Based Systems*, (pp. 25–62). Amsterdam: IOS.
- Clancey, W. (1985). Heuristic classification. *Artificial Intelligence*, **27**, 289–350.
- Gaines, B. R. (1987). An overview of knowledge acquisition and transfer. *International Journal of Man-Machine Studies*, **26**, 453–472.
- Gaines, B. R. and Boose, J. H. (1988). *Knowledge Acquisition for Knowledge-based Systems* (Vol. 1). New York: Academic Press.
- Gentner, D. (1987). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. Cambridge, MA: Cambridge University Press.
- Geoffrion, A. (1968). Proper efficiency and the theory of vector maximization. *Journal of Mathematical Analysis and Applications*, **22**, 618–630.
- Hammond, K. J. (1986). CHEF: A model of case-based planning. *Proceedings of the 5th National Conference on Artificial Intelligence—AAAI 86*, Menlo Park, CA, Vol. 1, pp. 267–271.
- Hayes-Roth, F., Waterman, D. A., & Lenat, D. B. (1983). *Building expert systems*. Reading, MA: Addison-Wesley.
- Helman, D. E. (Ed.). (1988). *Analogical Reasoning*. Dordrecht: Kluwer Academic Publishers.
- Humm, B., Schulz, C., Radtke, M., & Warnecke, G. (1991). A system for case-based process planning. *Computers in Industry*, **17**, 169–180.
- Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under Uncertainty: Heuristics and Biases*. Cambridge, MA: Cambridge University Press.
- Kolodner, J. L. (1991). Improving human decision making through case-based decision aiding. *AI Magazine*, 52–68.
- Kolodner, J. L., Simpson, R., & Sycara-Cyransky, K. (1985). A process model of case-based reasoning in problem solving. *Proceedings of IJCAI-85*, Los Angeles, CA, pp. 284–290.
- Larichev, O. I. (1992). Cognitive validity in design of decision aiding techniques. *Journal of Multicriteria Decision Analysis*, **1**, 127–138.
- Larichev, O. I., Mechitov, A. I., Moshkovich, H. M., & Furems, E. M. (1989). Expert knowledge acquisition system. In *Improving Decision Making in Organizations*, Lecture Notes in Economics and Mathematical Sciences. Berlin: Springer-Verlag. (pp. 517–522).
- Larichev, O. I., & Moshkovich, H. M. (1988). Limits to decision making ability in direct multiattribute alternative evaluation.

- Organizational Behavior and Human Decision Processes*, **42**, 217–233.
- Larichev, O. I., & Moshkovich, H. M. (1990). Decision support system "CLASS" for R&D planning. In *Proceedings of the First International Conference on Expert Planning Systems*, Brighton, England, pp. 227–232.
- Larichev, O. I., Moshkovich, H. M., Furems, E. M., Mechitov, A. I., & Morgoev, V. K. (1991). *Knowledge Acquisition for the Construction of Full and Contradiction Free Knowledge Bases*. Groningen, The Netherlands: iec ProGAMMA.
- Larichev, O. I., Moshkovich, H. M., & Rebrik, S. B. (1988). Systematic research into human behavior in multiattribute object classification problems. *Acta Psychologica*, **68**, 171–182.
- McGraw, K. L., & Harbison-Briggs, K. (1989). *Knowledge Acquisition: Principles and Guidelines*. London: Prentice-Hall.
- Mechitov, A. I., Moshkovich, H. M., & Olson, D. L. (1994a). Problems of decision rule elicitation in a classification task. *Decision Support Systems*, **12**, 115–126.
- Mechitov, A. I., Moshkovich, H. M., & Olson, D. L. (1994b). The role of rules and examples in the process of knowledge acquisition in direct classification tasks. *Expert Systems with Applications*, **8**, to appear.
- Miller, G. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, **63**, 81–97.
- Mukhopadhyay, T., Vicinanza, S. S., & Prietula, M. J. (1992). Examining the feasibility of a case-based reasoning model for software effort estimation. *MIS Quarterly*, **16**, 155–171.
- Polat, F., & Guvenir, H. A. (1991). A unification-based approach for knowledge base verification. *Expert Systems*, **8**(4), 251–259.
- Reichgelt, H., & Shadbolt, N. (1991). ProtoKEW: A knowledge-based system for knowledge acquisition. In D. Sleemans & O. Bernsen (Eds.), *Recent Advances in Cognitive Science*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Rissland, E. L., & Skalak, D. B. (1991). Combining case-based and rule-based reasoning: A heuristic approach. *Cognitive Models*, 524–530.
- Scott, A. C., Clayton, J. E., & Gibson, E. L. (1991). *A Practical Guide to Knowledge Acquisition*. Reading, MA: Addison-Wesley Publishing Company.
- Silverman, B. G. (1990). Critiquing human judgment via knowledge acquisition systems. *AI Magazine*, **11**(3), 60–79.
- Slade, S. (1991). Case-based reasoning: A research paradigm. *AI Magazine*, **12**(1), 42–55.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, **185**, 1124–1131.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, **211**, 453–458.

(Appendix 1 overleaf)

APPENDIX 1

MEDICAL TASK - diagnosis of coma states

Possible classes (comas)

- P₁ - diabetic coma
- P₂ - hypoglycemic coma
- P₃ - choropenic coma
- P₄ - oxycarbonic coma
- P₅ - apoplectical coma
- P₆ - eclamptic coma
- P₇ - ichemic insult

Possible subclasses for each class

- P₁¹ - high probability of diabetic coma
- P₁² - moderate probability of diabetic coma
- P₁³ - unlikely that this is diabetic coma

Possible attributes with scales $m = 10$

Rank-ordering of attribute values

	diabetic coma	hypoglycemic coma
Q ₁ - way in which coma began		
q ₁₁ - gradual coma development	1	2
q ₁₂ - sudden coma	2	1
Q ₂ - skin state		
q ₂₁ - skin is dry	1	2
q ₂₂ - skin is moist	2	1
Q ₃ - respiration		
q ₃₁ - pathological respiration	1	2
q ₃₂ - normal or superficial respiration	2	1
Q ₄ - pulse		
q ₄₁ - rapid pulse (greater than 80)	1	1
q ₄₂ - normal pulse	2	2
q ₄₃ - pulse rarus (less than 60)	3	3
Q ₅ - blood pressure		
q ₅₁ - diminished pressure (less than 100)	1	1
q ₅₂ - normal pressure	2	2
q ₅₃ - high pressure (greater than 145)	3	3
Q ₆ - deep reflexes		
q ₆₁ - diminished deep reflexes	1	2
q ₆₂ - high deep reflexes	2	1
Q ₇ - eyeball tension		
q ₇₁ - diminished eyeball tension	1	2
q ₇₂ - normal eyeball tension	2	1
Q ₈ - convulsions		
q ₈₁ - no convulsions	1	2
q ₈₂ - convulsions	2	1
Q ₉ - occipital muscle rigidity		
q ₉₁ - no muscle rigidity	1	2
q ₉₂ - muscle rigidity present	2	1
Q ₁₀ - temperature		
q ₁₀₁ - lowered or normal temperature	1	1
q ₁₀₂ - higher temperature	2	2

$$A = 2^8 \times 3^2 = 2304 \text{ possible combinations of values}$$