

## **SEARCHING FOR KNOWLEDGE IN A WORLD FLOODED WITH FACTS**

Ryszard S. Michalski  
Center for Artificial Intelligence  
George Mason University  
Fairfax, VA, 22030

### **ABSTRACT**

The wide availability of computer technology and large electronic storage media has led to an enormous proliferation of databases in almost every area of human endeavor. This naturally creates an intense demand for powerful methods and tools for data analysis. Current methods and tools are primarily oriented toward extracting numerical and statistical data characteristics. While such characteristics are very important and useful, they are often insufficient. A decision maker typically needs an interpretation of these findings, and this has to be done by a data analyst. With the growth of the amount and of the complexity of the data, making such interpretations is an increasingly difficult problem. As a potential solution, the paper advocates the development of methods for *conceptual data analysis*. Such methods aim at semi-automating the processes of determining high-level data interpretations, and discovering *qualitative* patterns in data. It is argued that these methods could be built on the basis of algorithms developed in the area of machine learning. An exemplary system utilizing such algorithms, INLEN, is discussed. The system integrates machine learning and statistical analysis techniques with database and expert system technologies. Selected capabilities of the system are illustrated by examples from implemented modules.

### **1. INTRODUCTION**

The current information age is characterized by an enormous expansion of data that are being generated and stored about all kinds of human activities. An increasing proportion of these data is recorded in the form of computer databases. This makes the data easy to access and easy to analyze by computer technology. The rapid growth of such databases has not been, however, matched with an equally intensive development of new type of methods and powerful new tools for analyzing and interpreting the data. As a result, we face today a growing problem of how to extract desirable knowledge from the large accumulations of data.

The existing data analysis tools are very useful and important for a whole range of data analysis tasks. They are, however, oriented primarily toward extraction of quantitative and statistical data

characteristics, and as such have certain inherent limitations. These tools include cluster analysis, numerical taxonomy, regression analysis, multivariate statistical methods, stochastic models, multidimensional analysis, times series analysis, nonlinear estimation techniques, and other. For a sample of research on data analysis see, e.g., papers by Tukey, (1986), Diday (1989), Daniel and Wood (1980), and Morgenthaler and Tukey (1989).

These traditional data analysis techniques are particularly useful for such tasks as producing statistical data summaries, fitting equations to data, revealing data organization on the basis of certain numerical measures, for developing mathematical data models, etc. Their results facilitate useful data interpretations, and can help to gain important insights into the processes that generated the data. These interpretations and insights are the ultimate knowledge sought for by those who build databases. Yet, this knowledge is not created by the current tools, but has to be derived by human data analysts. As the quantity of available data increases, the complexity of extracting knowledge from the data also increases, and may outstrip capabilities of data analysts.

Summarizing, traditional numerically-oriented techniques offer powerful tools and have important practical applications, but they can solve only a limited range of problems. For example, a statistical data analysis can discover a correlation between certain variables. It cannot, however, produce a conceptual characterization or a casual explanation why such a correlation exists. Nor it can develop a justification of this correlation in terms of higher-level concepts or analogies to known phenomena. A statistical analysis can determine a central tendency and variability of various properties, and a regression analysis can fit a complex curve to a set of datapoints. These techniques cannot, however, develop a qualitative characterization of the regularities in the datapoints and their dependence on various qualitative factors, nor can they draw an analogy between this characterization and some regularity or law in another domain, or relate it to similar past experiences.

A numerical taxonomy technique can create a classification of entities, and specify a numerical similarity among the entities assembled into the same or different classes. It will not, however, hypothesize reasons for the entities being in the same class, or build qualitative descriptions of the classes created.

Attributes that define the similarity, as well as the similarity measures, must be defined by a data analyst in advance. These techniques cannot generate relevant attributes and appropriate similarity measures by themselves. All the above processes seem to require complex symbolic reasoning that relates high level concepts and analysis goals to available quantitative measures and transformations relevant to these goals.

A question then arises as to the possibility for performing some of the above data analysis functions by an automatic or semi-automatic process. The paper tries to address this question, and discusses some novel data analysis tools.

## 2. CAN MACHINE LEARNING HELP DATA ANALYSIS?

The major premise of this paper is that various artificial intelligence methods, in particular, symbolic methods of machine learning and discovery, offer new and potentially useful tools for data analysis. These tools can perform new types of operations on data, and thus widen the scope of data analysis tasks that can be automated or semi-automated. Specifically, they can perform certain tasks of *conceptual data analysis*, that is, determine high-level symbolic data characterizations, and discover *qualitative* patterns in data. Let us try to briefly review some of these methods in the context of database applications.

### 2.1 Rule learning from examples

One class of machine learning methods that are potentially useful for data analysis is based on methods for inductive learning from examples. Given a set of examples of different classes (or concepts), and certain problem relevant knowledge ("background knowledge"), an inductive learning method hypothesizes a general description of each class. The description is usually expressed as a set of decision rules or as a decision tree.

A decision rule can have different forms; here we will assume the following form:

CLASS <:: CONDITION,

where CLASS denotes a class or a concept that is assigned to an entity if that entity satisfies the CONDITION. The CONDITION is often a conjunction of some elementary conditions on the values of attributes, or a disjunction of such conjunctions (a DNF form). We will assume here that if the CLASS needs a disjunctive description, then several conjunctive rules are linked to the same CLASS. For example, Figure 1 gives an example of a disjunctive description of *Class 1* in the form of two rules.

**Class 1 <:: Jacket Color is Red, Green or Blue &  
Head Shape is Round or Octagonal**

**Class 1 <:: Head Shape is Square and Jacket Color is Yellow**

*Figure 1.* A two-rule description of Class 1.

These rules characterize a class of robot-figures used in the EMERALD system for demonstrating machine learning capabilities. Paraphrasing, "A robot belongs to Class 1, if the color of its jacket is red, green or blue, and his head is round or octagonal; or, alternatively, if the color of its jacket is yellow and its head is square."

In a decision tree representation, nodes correspond to attributes, branches stemming from the nodes to attribute values, and the leaves to individual classes (see, e.g., Quinlan, 1986). A tree can be simply transformed into a set of decision rules (a ruleset) by traversing all paths from the root to individual leaves. The opposite process, that is, transforming a ruleset into a decision tree is not so direct. The reason for that is that a rule representation is generally more powerful than a tree representation. The word "powerful" means here that, e.g., for some simple rulesets, the equivalent decision tree may be much more complex (e.g., Michalski, 1990).

The EMERALD system, mentioned above, combines five programs that display different kinds of learning capabilities (Kaufman, Michalski, and Schultz, 1989; and Kaufman, Schultz and Michalski, 1991). These capabilities include rule learning from examples, learning distinctions between structures, conceptual clustering, prediction of object sequences, and derivation of equations and rules characterizing data about physical processes. The rules in Figure 1 were generated by the rule learning program (version AQ-15) from a set of "positive" and "negative" examples of robot-figures (Michalski, Hong and Mozetic, 1986). This paper concentrates on the applicability of the two of the above capabilities to data analysis, namely, rule learning and conceptual clustering. For a description of other capabilities see, e.g., (Kaufman, Michalski and Kershberg, 1991).

Most inductive rule learning methods learn *attributional* descriptions of entities in a class, i.e., descriptions that involve only binary or multiple-valued attributes. Some methods learn *structural descriptions*, which characterize entities in terms of both, attribute values, as well as relationships that hold among components of the entities. Such relationships are represented by multiplace predicates (Michalski, 1983). For data analysis, the most interesting programs seem to be ones for learning attributional descriptions, because typical databases characterize entities in terms of attributes.

The input to attributional learning programs are sets of examples of individual classes, and some "background knowledge" relevant to the learning problem. The examples are in the form of vectors of attribute-value pairs associated with some decision class. In many cases, background knowledge (BK) is limited to the information about the legal values of the attributes, their type (the scale of measurement), and the *preference criterion* for choosing among candidate hypotheses. Such a criterion is defined by the user in advance. In addition to BK, a learning

method may have a *representational bias*, i.e., it constrains the form of descriptions to only a certain type of expressions, e.g., single conjunctions, decision trees, sets of conjunctive rules, DNF expressions, etc. In some methods, BK may include more information, e.g., constraints on the interrelationship between various attributes, rules for generating higher level concepts or attributes, and/or some initial hypothesis (e.g., Michalski, 1983). Learned rules are usually *consistent* and *complete* with regard to the input data. This means that they completely and correctly classify all the original "training" examples. Section 4 presents example solutions from the inductive concept learning program AQ15. In some applications, especially those involving learning rules from noisy data or learning *flexible* concepts (Michalski, 1990), it may be advantageous to learn descriptions that may be incomplete or inconsistent (Bergadano et al, 1990).

Attributional descriptions can be easily visualized by mapping them into a set of cells in a certain diagram. For example, Figure 2, shows a diagrammatic visualization of the rules from Figure 1.

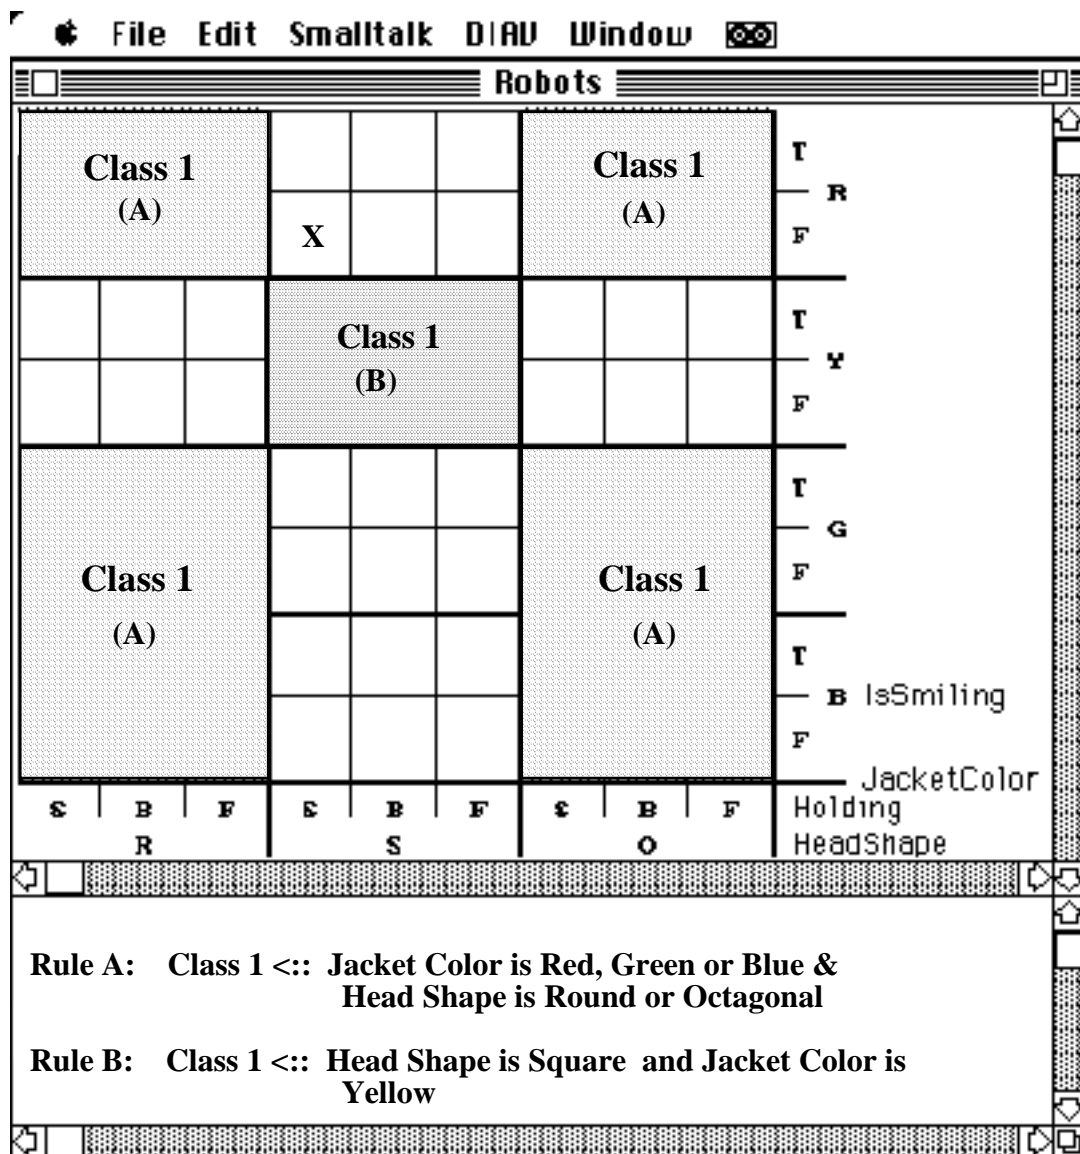


Figure 2. A diagrammatic visualization of rules from Figure 1.

Such a diagram is a planar representation of a multidimensional space spanned over the set of attributes (Michalski, 1978; Wnek et al., 1990). The diagram in Figure 2 was generated by the visualization program DIAV (Wnek, et al., 1990). Each cell in the diagram represents one combination of values of the attributes. For example, the cell marked by an X represents the vector: (HeadShape=S, Holding=S, Jacket Color=R, IsSmiling=F). The four shaded areas marked Class1 (A) represent rule A, and the shaded area marked Class 1 (B) represents rule B. In such a diagram, conjunctive rules correspond to certain regular arrangements of cells and can be easily recognized (Michalski, 1978).

The diagrammatic visualization can be used for displaying the *target concept* (i.e., the concept to be learned), the training examples (the examples and counter-examples of the concept), and the actual concept learned by a method. By comparing the target concept with the learned concept, one can determine the *error area*, i.e., the area containing all examples that would be incorrectly classified by the learned concept. Such a diagrammatic visualization method can illustrate any kind of attributional learning process (Wnek, 1990). Because a data table used in the data analysis can be viewed as a set of points in a multidimensional space, the visualization technique can be a useful tool for representing the data and the learned symbolic descriptions (when the space is not too large).

Two types of data analysis operators can be based on methods for learning concept descriptions from examples :

- A. An operator for determining a general symbolic description of any designated group or groups of entities in a data set. Such a description expresses the common properties of the entities in the group. It can use abstract concepts in the description that are not present in the original set via the mechanism of constructive induction (see later). This operator is based on a program for learning the so-called *characteristic concept descriptions*.
- B. Determining differences between different groups of entities. Such differences are expressed in the form of rules that define the properties that characterize one group but not the other. This operator is based on a program for learning the so-called *discriminant concept descriptions*.

Section 3 will illustrate these two types of descriptions. For their definitions see (Michalski, 1983). Basic methods for concept learning assume that examples do not have errors, that all attributes have a specified value in them and that concepts to be learned have a precise ("crisp") description. In many situations one or more of these assumptions may not hold. Therefore, some more advanced research problems in the area of concept learning from examples include:

- *Learning from incorrect data*, i.e., learning from examples that may contain certain number of errors or noise.
- *Learning from incomplete data*, i.e., learning from examples in which the values of some attributes are unknown.
- *Learning flexible concepts*, i.e., concepts that lack precise definition and whose meaning is context-dependent.

All these problems are relevant to data analysis, and thus the methods for solving them being developed in machine learning can be potentially useful. For some of these methods see, e.g., (Quinlan, 1990; Michalski, 1990; Bergadano et al., 1990).

## 2.2 Conceptual clustering

Another class of machine learning methods that are important to data analysis are those concerned with developing "conceptual" classifications of a given set of entities. The problem is similar to that considered in traditional cluster analysis, but is defined in a different way. Given a set of attributional descriptions of some entities, a description language for characterizing classes of such entities, and a "classification quality" criterion, the problem is to split the entities into classes that maximize this criterion, and simultaneously determine general (extensional) descriptions of these classes in the assumed language. Thus, a conceptual clustering program seeks not only a classification (a dendrogram), but also a symbolic description of the proposed classes (clusters). Also, unlike in cluster analysis, the process of determining the classes is dependent on the properties of descriptions of the generated classes.

To clarify the difference between conceptual clustering and conventional clustering, notice that a conventional method typically determines clusters on the basis of a similarity measure that is a function solely of the properties (attribute values) of the entities being compared, and not of any other factors:


$$\text{Similarity}(A, B) = f(\text{properties}(A), \text{properties}(B))$$

In contrast, a conceptual clustering program clusters entities on the basis of a *conceptual cohesiveness* that is not only a function of properties of the entities, but also of two other factors: the description language  $L$ , which the system knows a priori and uses for describing the classes of entities, and of the set of neighboring examples, the *environment*,  $E$ :

$$\text{Conceptual cohesiveness}(A, B) = f(\text{properties}(A), \text{properties}(B), L, E)$$

Thus, two objects may be similar, i.e., close according to some distance measure, they may have a low conceptual cohesiveness. An example of such a situation is shown in Figure 3.





*Figure 3.* An illustration of the difference between closeness and conceptual cohesiveness.

According to conceptual clustering, the points (black dots) in Figure 3 would be clustered into two "ellipses." The points A and B, which are "close" to each other, have small conceptual cohesiveness, because they belong to configurations representing different concepts.

A classification quality criterion may take into consideration several factors, such as the "fit" of a *cluster description* to the data, the simplicity of the description, and/or other properties of the entities or of the concepts that describe them (Michalski, Stepp and Diday, 1981.). Section 4 gives an example of conceptual clustering.

### 2.3 Other related symbolic operations on data

Methods for learning rules from examples usually assume that attributes used for describing the examples are given a priori. These attributes must be sufficiently relevant to the learning problem, otherwise, the learned rules will be poor. An important advantage of symbolic methods over, e.g., statistical methods, is that they can relatively easily handle *nonessential* attributes. An attribute is *nonessential*, if there is a complete and consistent description of the class(s) under consideration that does not use this attribute. Thus, a nonessential attribute is either just irrelevant, or relevant but dispensable. Inductive learning programs such as the rule-learning program AQ, or the decision tree learning ID3, can relatively easily cope with a large number of nonessential attributes in data descriptions.

If there are, however, very many nonessential attributes in the initial description of the examples, the efficiency of a learning program would decrease. Such situations call for a method that would determine the most relevant attributes for a given classification from among those given initially. Only attributes that are most relevant would then be used in the description learning process. Determining the most relevant attributes can therefore be a useful data analysis operator. Such an

operator can also be useful on its own merit, as it may be of interest to know which attributes are, e.g., most discriminatory for a given set of classes.

Another related operator is to determine the most representative or, generally, the most important examples in a given set. Such an operator would be needed when there are very many examples of a given class. The most important examples would be those that are either most typical or most extreme. A method for determining the latter ones, the so-called "outstanding representatives," is described in (Michalski and Larson, 1978).

In many applications, it is not easy to determine a priori what attributes will be most relevant to the classification problem at hand. Often, attributes are dictated by a given application or already supplied in the database. Yet, these attributes may not be the most relevant for the task. Such situations require a method for generating new attributes, and selecting from them the most relevant ones (e.g., Bongard, 1970). This type of problems is considered in *constructive induction* (Michalski, 1983). A recently developed program for constructive induction, AQ17, can generate new attributes by combining initially given attributes in many different ways, by creating complex relationships among them, or by obtaining advice from an expert (Michalski, Bleodorn and Wnek, 1991).

### 3. LEARNING PROBLEMS AND A GENERAL DATA TABLE

The above described learning problems can be simply illustrated by means of a *general data table* (GDT). Such a table is a generalization of a typical data table used in data analysis (Figure 4). The columns in a GDT correspond to attributes. These may be some initial attributes, given a priori, or additional ones generated through a process of constructive induction. Each attribute is assigned a *domain* and a *type*. The domain specifies the set of all legal values that the attribute can take on in the table, which includes "?" ("unknown") and N/A ("not applicable"). The type defines the order of the values in the domain (the scale). For example, the AQ15 learning program (Michalski et al. 1986) allows three types of attributes: nominal (no order), linear (total order), and structured (a hierarchical order). The attribute type determines the kinds of operations that are allowed on this attribute during a learning process.

Rows in the table correspond to individual entities that are being characterized by the attributes. An entry in the table can be a specific value of the corresponding attribute, the symbol "?," meaning that that value is unknown for the given entity, or the symbol N/A, if the given attribute does not apply to the given entity. For example, the color attribute usually applies to physical objects, but does not apply to relations or abstract entities.

An important problem of data analysis is to determine whether some attribute in a table depends on some other attributes. A related problem is to determine a specific form of this relationship. A case of the latter problem is when it is known that a nominal-scale attribute depends on other attributes, and the problem is to hypothesize a general description that characterizes this relationship. The latter problem is equivalent to a typical problem of concept learning from examples. In such a case, one column in the data table is designated to represent the *output attribute*. The values of that variable are classes whose descriptions are to be learned. In Figure 4, it is the first column that represents values of the output variable (in this case, the class of the examples). In conceptual clustering, there is no such column, as there are no a priori classes which entities belong to (as in any other "unsupervised learning" problem).

Using a GDT, one can simply characterize the problems described above (Figure 4):

- Learning from examples:* Classes of examples are sets of rows in the table which have the same value of the output attribute (variable). The problem is to determine a set of general rules characterizing these classes of examples.
- Example selection:* The problem is to select from the table rows that correspond to the most representative examples (rows) of different classes.
- Attribute selection:* The problem is to select columns that correspond to the most relevant attributes for characterizing given classes or the differences between classes.
- Generating new attributes:* The problem is to generate additional columns that correspond to new attributes generated by a constructive induction procedure. These new attributes are created by using the problem background knowledge, and/or special heuristic procedures (Michalski, Bloedorn and Wnek, 1991).
- Conceptual clustering:* The problem is to split the rows of the table into groups of rows that correspond to "conceptual clusters," that is sets of entities with high conceptual cohesiveness. An additional column is added to the table that corresponds to a new "output attribute." The values of this attribute in the table denote the proposed class of each entity. Rules that describe clusters are stored separately (see section 5).

*Learning from imperfect data:* In this case, some entries of the table are missing, or be incorrect. The problem is to determine the best (e.g., the most plausible) hypothesis that accounts for the data or for the most of the data.

Methods for performing the above operations have been developed and implemented in various machine learning programs (see, e.g., Michalski, Carbonell and Mitchell, 1983 and 1986; Forsyth and Rada, 1986; Kodratoff, 1988, and Kodratoff and Michalski, 1990).

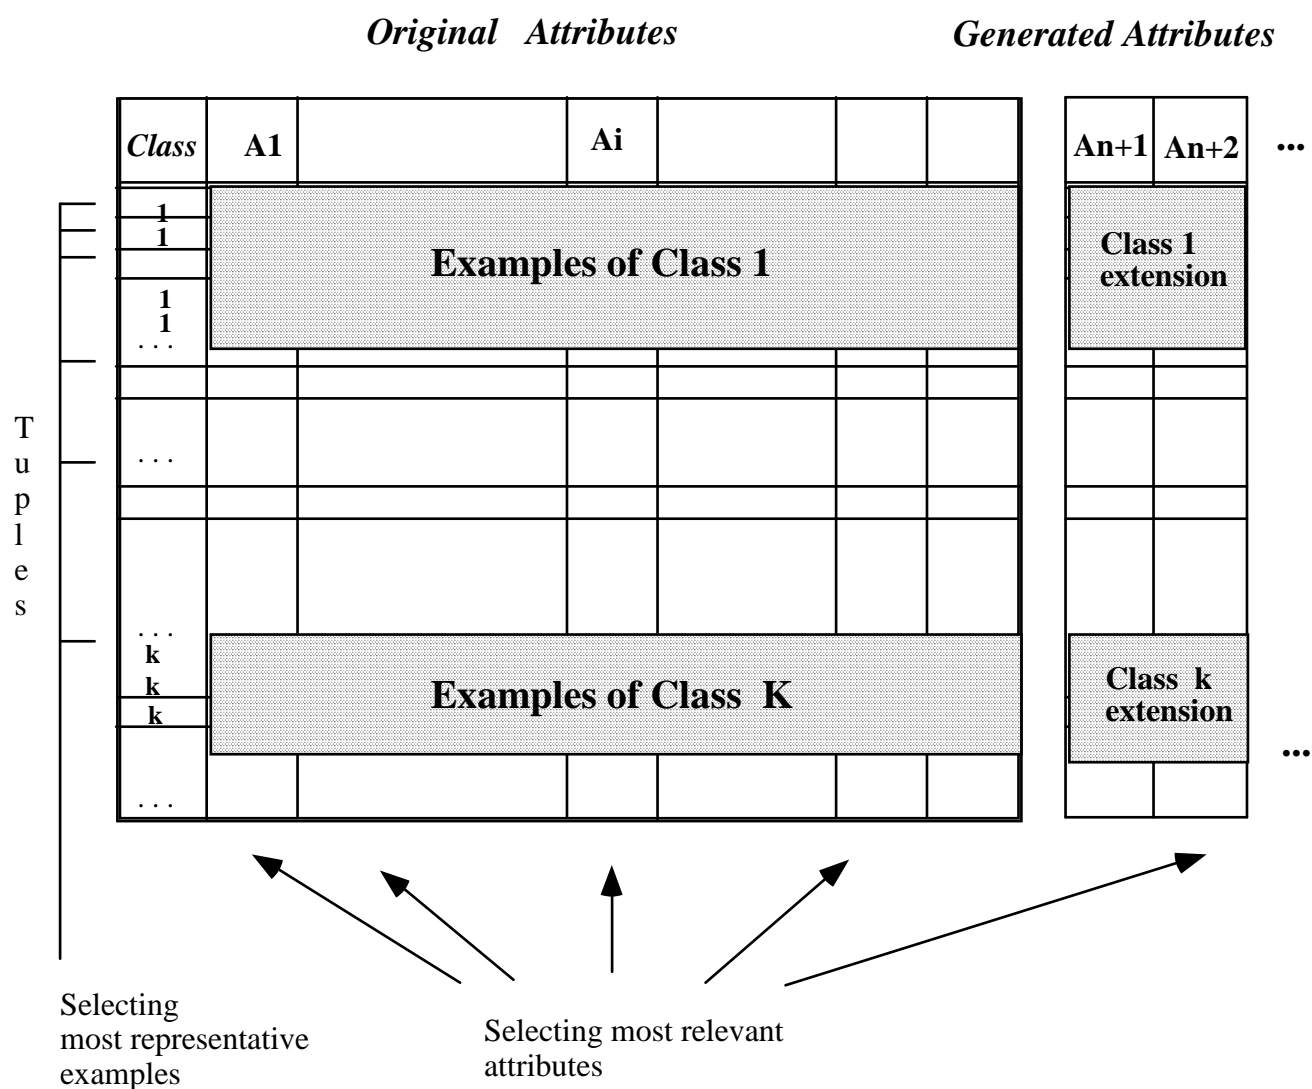


Figure 4. An illustration of the role of different symbolic operators

#### 4. EXAMPLES OF KNOWLEDGE GENERATION OPERATORS

As mentioned above, the programs for different forms of symbolic learning and knowledge transformation can be used as data analysis operators. These operators are applied to a data table and produce certain data characteristics. To illustrate this idea, we will use examples of how a conceptual clustering program and a learning from examples program can be used as data analysis operators (based on the material in Kaufman, Michalski and Kershberg, 1991).

Suppose we are given a data table as shown in Figure 5.

Microcomputer	Display	RAM	ROM	Processor	No_Keys
Apple II	Color_TV	48K	10K	6502	52
Atari 800	Color_TV	48K	10K	6502	57-63
Comm. VIC 20	Color_TV	32K	11-16K	6502A	64-73
Exidi Sorceror	B/W_TV	48K	4K	Z80	57-63
Zenith 118	Built_in	64K	1K	8080A	64-73
Zenith 1189	Built_in	64K	8K	Z80	64-73
HP 85	Built_in	32K	80K	HP	92
Horizon	Terminal	64K	8K	Z80	57-63
Challenger	B/W_TV	32K	10K	6502	53-56
O-S 11 Series	B/W_TV	48K	10K	6502C	53-56
TRS-80 I	B/W_TV	48K	12K	Z80	53-56
TRS-80 III	Built_in	48K	14K	Z80	64-73

*Figure 5.* Illustrating learning problems by a general data table.

Suppose now that the task is to produce the "best" conceptual classification of the data into two and three classes. The table in Figure 5 is an input to a conceptual clustering program, acting as a CLUSTER operator.

The results of applying this operator have two components: a new extended data table, and a set of rules. The new table has two additional columns; the first column indicates the class assigned to the tuples in the generated two-class clustering, and the second column indicates the class assigned to the tuples in the generated three-class clustering (Figure 6).

The second component are two sets of rules; the first ruleset describing classes in the two-class clustering, and the second set describing classes in the three-class clustering (Figure 7).

INPUT						OUTPUT (added to input table)	
Microcomputer	Display	RAM	ROM	Processor	No_Keys	2-Group	3-Group
Apple II	Color_TV	48K	10K	6502	52	1	1
Atari 800	Color_TV	48K	10K	6502	57-63	1	1
Comm. VIC 20	Color_TV	32K	11-16K	6502A	64-73	1	2
Exidi Sorceror	B/W_TV	48K	4K	Z80	57-63	1	2
Zenith 118	Built_in	64K	1K	8080A	64-73	2	3
Zenith 1189	Built_in	64K	8K	Z80	64-73	2	3
HP 85	Built_in	32K	80K	HP	92	1	2
Horizon	Terminal	64K	8K	Z80	57-63	1	2
Challenger	B/W_TV	32K	10K	6502	53-56	1	1
O-S 11 Series	B/W_TV	48K	10K	6502C	53-56	1	2
TRS-80 I	B/W_TV	48K	12K	Z80	53-56	1	1
TRS-80 III	Built_in	48K	14K	Z80	64-73	1	1

Figure 6. An extended table generated as result of the CLUSTER operator.

Rules characterizing the generated 2-class clustering:

[Class 1] <== [RAM = 16K..48K]

[Class 1] <== [No\_Keys ≤ 63]

[Class 2] <== [RAM = 64K] & [No\_Keys > 63]

Rules characterizing the generated 3-class clustering:

[Class 1] <== [Processor = 6502 v Z80] & [ROM = 10K..14K]

[Class 2] <== [Processor = 6502A v 6502C v HP]

[Class 2] <== [ROM = 1K..8K] & [Display ° Built\_in]

[Class 3] <== [Processor = 8080A v Z80] & [ROM = 1K..8K] & [Display = Built\_in]

Figure 7. Rules characterizing classes created by the CLUSTER operator

Suppose now that we use the extended data table in Figure 6 as an input to a program for learning concept learning from example. Suppose that the parameters of the program are set so that the program should try to differentiate among the classes, i.e., to determine the simplest discriminant class descriptions (Michalski, 1983). In this case, the program plays the role of the so-called DIFF operator. The results are shown in Figure 8.

*Rules for 2-class clustering obtained by the DIFF operator*

[Class 1] <== [Display or Built\_in]  
 [Class 1] <== [ROM ≥ 14K]  
 [Class 2] <== [RAM = 64K] & [No\_Keys = 64-73]

*Rules for 3-class clustering obtained by the DIFF operator*

[Class 1] <== [Processor = Z80 v 6502] & [ROM = 10K..14K]  
 [Class 2] <== [Processor = 6502C v 6502A v HP]  
 [Class 2] <== [ROM = 4K..8K] & [Display = B/W\_TV v Term]  
 [Class 3] <== [ROM = 1K..8K] & [Display = Built\_in]

*Figure 8. Discriminant rules generated by DIFF operator*

Comparing rules in Figure 7 with those on Figure 8 (note that the latter were generated without knowledge of the former), one can see that they are similar but not identical. The reason for the difference is that given a set of examples, there are usually many different ways of creating a general description of them. The rules in Figures 7 and 8 are both complete and consistent with all the examples in Figure 6, i.e., they classify these examples in the same way.

## **5. INTEGRATING DIFFERENT OPERATORS IN ONE SYSTEM**

To make symbolic data analysis operators, such as CLUSTER or DIFF, easily available to a data analyst, it is desirable to integrate them into one system. Such an idea was implemented in the INLEN system (Kaufman, Michalski and Kershberg, 1989, 1991). INLEN incorporates a whole spectrum of machine learning operators, and also conventional statistical data analysis operators. To facilitate the application of all the operators, INLEN combines relational *database* technology with *knowledge base* technology. The database technology is used for storing and updating data tables, and the knowledge base technology is used for storing, updating and applying rules.

A general diagram of INLEN is presented in Figure 9. The name is an acronym from **in**ference and **l**earning.




Figure 9. A general diagram of the INLEN system for conceptual data analysis

INLEN offers to a data analyst three classes of operators:

- **DMO:** Data Management Operators. These operators are conventional data management operators that are used for creating, modifying and displaying relational tables.
- **KMO:** Knowledge Management Operators. These operators play a similar role to the DMOs, but they apply to the rules and other structures in the knowledge base.
- **KGO:** Knowledge Generation Operators. These operators perform symbolic and numerical data analysis operations on data. They are based on various machine learning and inference programs, on conventional data analysis techniques, and visualization operators for displaying graphically the results of analysis. The diagrammatic visualization method, described briefly above, is used for displaying the effects of symbolic learning operations on data.

The KGOs are the heart of the INLEN system. To facilitate their use, the concept of a *knowledge segment* was introduced. A knowledge segment is a structure that links one or more tables from the database with one or more structures from the knowledge base. Such knowledge segments are both, the inputs to and outputs from KGOs. Thus, KGOs can be viewed as modules for performing complex inferences on knowledge segments in order to create new knowledge segments.



The execution of a KGO usually requires some background knowledge (BK) from the knowledge base, and is guided by various parameters. The BK specifies facts about the application domain, and provides information about attributes, the constraints and relationships among attributes, etc. The parameters specify how to choose an output description from among multiple candidates. KGOs can usually work in either incremental or batch mode. In the incremental mode, they try to improve or refine the existing knowledge; while in the batch mode, they try to create entirely new knowledge from facts in the database, using knowledge in the knowledge base.

KGOs in INLEN can be classified into several groups, based on the type of the output they generate. Each group includes a number of specific operators that are instantiated by a combination of the parameters.

- GENRULE operators generate various kinds of rules from given facts. They include operators that generate symbolic descriptions of data, e.g., generate rules characterizing a set of facts, discriminate between groups of facts, build decision trees, characterize a sequence of events, and determine differences between sequences. They also include operators generating equations qualitatively and quantitatively characterizing numerical data sets, and build conceptual hierarchies.
- TRANSRULE operators perform various transformations of the rules, e.g., generalize or specialize, abstract or concretize given rules.
- GENATR operators generate new attributes, or select the most representative attributes from a given set.
- GENEVE operators generate events, facts or examples that satisfy given rules, select the most representative events from a given set., determine an example that is similar to a given example, or predict a value of a given variable
- ANAREL operators analyze various relationships that exist in the data, e.g., determine the degree of similarity between two examples, check if there is an implicative relationship between two variables, determine various statistical properties of the data.
- TEST operators test the performance of a given set of rules on an assumed set of facts. The primary output from these operators is a confusion matrix, i.e., a table whose (i,j)th element shows how many examples from the class i were classified by the rules to be in class j.

For more details about these operators the reader can consult the paper (Kaufman, Michalski and Kerschberg, 1991). Summarizing, INLEN integrates a large range of operators for creating data, creating rules, and for performing many different symbolic and numeric operations on the data and/or rules.

## 6. CONCLUSION

An enormous growth of databases has created an intense need for developing a new type of data analysis tools that can not only numerically but also *conceptually* characterize data. Such conceptual characterizations include symbolic descriptions, logical relationships and qualitative evaluations, as well as causal dependencies among the entities in the data. These characterizations represent knowledge that may be more directly usable in human decision-making than numerical characterizations.

To determine such knowledge, a data analysis system has to be able to represent, and take into consideration, various kinds of background knowledge about the data and the domain of discourse. This background knowledge may include, for example, a specification of the domain and the type of the attributes, the relationships among them, causal dependencies, theories about the objects or processes that generated the data, and other high-level knowledge.

This main idea of the paper is that modern methods developed in symbolic machine learning can be a basis for the development of such conceptual data analysis tools. As examples, we described methods for learning general concepts from examples and conceptual clustering. We have also described an architecture of a large-scale system, INLEN, which integrates machine learning and statistical analysis techniques with database and knowledge base technologies.

The machine learning techniques allow the system to perform a whole range of symbolic data manipulation and knowledge extraction operations. These operations include, e.g., the generation of rules characterizing or discriminating between groups of facts, the creation of symbolic descriptions of sequences, the determination of new attributes, the building of conceptual hierarchies, and the construction of equations and logical preconditions for their application among others.

INLEN is a very complex system, and its development requires a significant amount of effort. Its experimental version is currently being implemented in the Center for Artificial Intelligence at George Mason University. Some of the INLEN components are based on programs already developed, and some require new research. Even those that are based on existing programs need to be properly modified and/or improved to be adequate, or sufficiently efficient, for the data analysis tasks. Further research is also needed to determine what other types of symbolic learning techniques might be applicable to data analysis.

## ACKNOWLEDGEMENTS

The author thanks Michael Hieb, Ken Kaufman, and Elizabeth Marchut-Michalski for very useful comments and criticism. This research was done in the Artificial Intelligence Center of George Mason University. The activities of the Center are supported in part by the Defense Advanced Research Projects Agency under grant, administered by the Office of Naval Research, No. N00014-87-K-0874, in part by the Office of Naval Research under grant No. N00014-88-K-0226, and in part by the Office of Naval Research under grant No. N00014-88-K-0397.

## REFERENCES

- Bergadano, F., Matwin, S., Michalski, R.S., and Zhang, J., "Learning Two-Tiered Descriptions of Flexible Concepts: The POSEIDON System," *Reports of Machine Learning and Inference Laboratory*, MLI 90-9, Center for Artificial Intelligence, George Mason University, Fairfax, VA September 1990.
- Bongard, N., *Pattern Recognition*, Spartan Books, New York, 1970 (a translation from Russian).
- Daniel, C. and Wood, F.S., *Fitting Equations to Data*, Wiley, New York, 1980.
- Diday, E. (ed.), *Proceedings of the Conference on Data Analysis, Learning Symbolic and Numeric Knowledge*, Nova Science Publishers, Inc., Antibes, September 11-14, 1989.
- Forsyth, R. and Rada, Roy, *Machine Learning: applications in expert systems and information retrieval*, Pitman, 1986.
- Kaufman, K.A., Michalski, R.S. and Schultz, A.C., "EMERALD 1: An Integrated System of Machine Learning and Discovery Programs for Education and Research, User's Guide", *Reports of the Machine Learning and Inference Laboratory*, MLI 89-11, Center for Artificial Intelligence, George Mason University, Fairfax, VA, 1989.
- Kaufman, K.A., Michalski, R.S., and Kerschberg, L., "Mining for Knowledge in Databases: Goals and General Description of the INLEN System," *Proceedings of IJCAI-89 Workshop on Knowledge Discovery in Databases*, Detroit, MI, August 1989.
- Kaufman, K.A., Michalski, R.S., and Kerschberg, L., AN ARCHITECTURE FOR Knowledge Discovery from Facts: Integrating Database, Knowledge Base and Machine Learning in INLEN, in *Machine Learning and Inference Reports*, to appear in 1991.
- Kodratoff, Y., *Introduction to Machine Learning*, Pittman, 1988.

Michalski, R. S. , Mozetic, I, Hong, J. and Lavrac, N., "The AQ15 Inductive Learning System: An Overview and Experiments," *ISG Reports*, 86-20, UIUCDCS-R-86-1260, Department of Computer Science, University of Illinois, Urbana, 1986.

Michalski, R. S., Stepp , R. and Diday, E., "A Recent Advance in Data Analysis: Clustering Objects into Classes Characterized by Conjunctive Concepts," Invited chapter in *Progress in Pattern Recognition*, Vol. 1, L. Kanal and A. Rosenfeld (Editors), pp. 33-55., 1981.

Michalski, R.S., "Learning Flexible Concepts: Fundamental Ideas and a Method Based on Two-tiered Representation", in *Machine Learning: An Artificial Intelligence Approach, Vol. III*, Y. Kodratoff and R.S. Michalski (Eds.), Morgan Kaufmann Publishers, 1990.

Michalski, R.S., "A Planar Geometrical Model for Representing Multi-Dimensional Discrete Spaces and Multiple-Valued Logic Functions," *ISG Report No. 897*, Department of Computer Science, University of Illinois, Urbana, January 1978.

Michalski, R.S., "A Theory and Methodology of Inductive Learning," *Artificial Intelligence*, 1983., pp. 111-161, 1983.

Michalski, R.S., Bleodorn, E. and Wnek, J., Inventing Relevant Attributes and Relations in the AQ17 Program for Constructive Induction, to appear in the *Reports of Machine Learning and Inference*, 1991.

Morgenthaler, S. and Tukey, J.W., The Next Future of Data Analysis, in *Proceedings of the Conference on Data Analysis, Learning Symbolic and Numeric Knowledge*, Diday, E. (ed.), Nova Science Publishers, Inc., Antibes, September 11-14, 1989.

Quinlan, J.R., "Induction of Decision Trees," *Machine Learning Journal*, Vol. 1, No. 1, pp. 81-106, 1986.

Quinlan, J.R., "Probabilistic Decision Trees," *Machine Learning: An Artificial Intelligence Approach, Vol. III*, Y. Kodratoff and R.S. Michalski (Eds.), pp.140-152, Morgan Kaufmann Publishers, 1990.

Tukey, J.W., *The Collected Works of John W. Tukey, Vol. V*, Philosophy and Principles of Data Analysis: 1965-1986, L.V. Jones (ed.), Wadsworth & Brooks/Cole, Monterey, CA. 1986.

Wnek, J., Sarma, J., Wahab, A., Michalski, R.S., "Comparing Learning Paradigms via Diagrammatic Visualization: A Case Study in Single Concept Learning using Symbolic, Neural Net and Genetic Algorithm Methods," *Reports of the Machine Learning and Inference*

*Laboratory*, MLI 90-2, Center for Artificial Intelligence, George Mason University, Fairfax, VA, January 1990.