

Concepts as Flexible and Context-dependent Sets: The Two-tiered View*

R. S. MICHALSKI

George Mason University

The paper proposes to view concepts as sets with flexible and context-dependent boundaries. An efficient representation of such concepts consists of two components (tiers), an explicit one, called the *Base Concept Representation* (BCR), and an implicit one, called *Inferential Concept Interpretation* (ICI). The BCR represents typical, simple to describe and relatively stable concept properties, and the ICI represents admissible transformations of the BCR and context-dependency of the concept. The admissible transformations are performed by *flexible matching* procedures and inference rules reflecting the reasoner's background knowledge. The presented *two-tiered* concept representation is justified by cognitive economy, and illustrated by several psychological studies and computer experiments.

I. INTRODUCTION

Cognitive scientists have been searching for years to find essential ingredients of intelligence, yet this problem remains unsolved. Two abilities, however, clearly stand out as associated with

*

This research was supported in part by the Defense Advanced Research Projects Agency under the grants administered by the Office of Naval Research, No. N00014-87-K-0874 and No. N00014-91-J-1854, in part by the Office of Naval Research under grants No. N00014-88-K-0397, No. N00014-88-K-0226, No. N00014-90-J-4059, and No. N00014-91-J-1351, and in part by the National Science Foundation under grant No. IRI-9020266. The author expresses his thanks to Mike Hieb and the members of Machine Learning Doctoral Seminar at George Mason University for helpful discussions and suggestions. Correspondence and requests for reprints should be sent to R. S. Michalski, Center for Artificial Intelligence, George Mason University, Fairfax, VA 22032 (*email*: michalski@aic.gmu.edu).

intelligent behavior. One is the ability to acquire knowledge or skill through experience, i.e., the ability to learn. The second is the ability to apply the knowledge or skill possessed to solve new problems, i.e., the ability to reason. A precondition for the above abilities is the capability to represent diverse forms of knowledge. As our knowledge is built of individual concepts, to represent knowledge one needs to represent concepts. Consequently, it is generally agreed that concept representation is a fundamental problem underlying efforts in the quest to understand intelligence. Although cognitive scientists, psychologists, linguists, philosophers and artificial intelligence researchers have given considerable attention to this problem, no conclusive solution of it exists.

Among most widely known approaches to concept representation there are the classical view, the prototype view, and exemplar view (e.g., Medin & Smith, 1984; Rosch & Mervis, 1975, Barsalou, 1987). Although these views, and various extensions or new proposals (e.g., Barsalou, & Medin, 1986; Keil, F., 1989) explain many aspects of concept representation, they do not yet provide a totally satisfactory and comprehensive solution.

This paper discusses a new approach to this problem that overcomes some of the weaknesses of other representations. This new approach, called the *two-tiered (TT) concept representation*, considers human concepts not as fixed, well-defined structures, but as flexible, context-modifiable and background-knowledge-dependent units of our knowledge. To represent concepts so understood, the TT view recommends splitting the total concept representation into two components (“tiers”), an explicit one that defines the most stable concept properties, and an implicit one, that handles rare or exceptional cases and context-dependency of a concept. The latter component consists of flexible procedures for matching the concepts with instances, and a body of rules for characterizing concept variability in different contexts and situations.

As a background for the new approach, the following section outlines our notion of concepts and categories. Later sections provide details about the two-tiered representation, relate it to other views, and finally presents some experimental results from a computer model.

2. THE MEANING OF CONCEPTS AND CATEGORIES

Before defining the two-tiered concept representation, we would like to clarify the terminology and basic notions used in this paper, in particular, those of *concepts* and *categories*. While everybody agrees that concepts are fundamental building blocks for our knowledge and thinking, there is no established definition of what concepts are. For example, Barsalou (1990) states: “Whereas most

theorists would probably reserve 'concept' for categorization rules, I will also use 'concept' for people's conceptualization of categories." Smith (1989) views a concept as a "mental representation of a class or individuals and deals with *what* is being represented and *how* that information is typically used during the categorization." Medin (1989) holds the view that "a concept is an idea that includes all that is characteristically associated with it." He also suggests that "a category is a partitioning or a class to which some assertion or set of assertions might apply."

These characterizations of concepts are not all equivalent and raise various questions. A frequently used definition of concepts that identifies them with categorization rules (or a concept description) is not satisfactory, because for any given concept one can usually think of many different categorization rules (or concept descriptions) that specify how to recognize this concept, i.e., how to distinguish between concept instances and non-instances. To see this point, suppose that a concept was defined as a certain categorization rule. If one would replace this rule by a syntactically different one, but which produces exactly the same classification, then we would have a contradiction: since the initial and the newly created rule are different from each other, then they would constitute different concepts, yet—since they generate the same classification—they would represent the same concepts. For example, there can be many different medical procedures ("categorization rules") for distinguishing between a specific type of cancer from other types of cancer. Although these medical procedures are different, the cancer is all the same. In short, concepts should not be identified with their descriptions.

Characterizations of a concept as a "mental representation," a "conceptualization of a category" or as an "idea" are also not very satisfactory. These characterizations use complex terms that themselves require definition or at least clear explanation. If notions of "mental representation," "conceptualization" or "idea" are not properly defined, then these definitions suffer the same problem as those that identify concepts with "categorization rules." These definitions also leave the concept of "category" unexplained and unclearly related to the notion of "concept."

In our view, concepts should not be identified with "categorization rules" or "mental representations" However, to recognize a concept, one always needs to use some knowledge (explicit or implicit) about it, and that knowledge can be viewed as a "categorization rule." So a basic problem is, then, what is the relationship between a concept and a categorization (or classification) rule used for its recognition. Another problem is what is the relationship between

concepts and categories. This section attempts to develop a new approach to the solution of these problems.

Our stand is that *concepts* can be viewed formally as sets of entities, physical or abstract, that are grouped together because of some reason or according to some apriori given principle. The set constituting a concept may consist of just one individual (e.g., the sun in our solar system), a countable number of individuals, or an uncountable number of individuals (e.g., real numbers). The reason for grouping entities to a concept may be to put together those that have the same function or purpose, the similar appearance, similar physical, spatial, or temporal characteristics, the same role in a structure of other entities or concepts, or because of some combination of these and/or other factors. The reason for creating a concept may be incorporated in a body of knowledge that is encoded in the structure of an agent's (concept bearer's) mind as a result of some learning or creative process. Such a body of knowledge may, or may not, be precisely translatable into an effective classification rule. It may be externalized only by the classifications one makes, e.g., one may not be able to tell why a given painting is beautiful, but may be quite confident in this evaluation. In sum, the reason (or principle) underlying a concept is not necessarily expressible linguistically, and a concept (a set) should be distinguished from a classification rule for that concept.

A classification rule is an expression in some explicit language that constitutes an effective procedure for concept recognition, i.e., for telling which instances belong to a concept and which do not. Such a procedure uses various attributes and/or relations (generally, *descriptors*) which are dimensions for characterizing concept instances, and measurable either directly or indirectly. As mentioned earlier, there may be many different classification rules for recognizing a given concept.

In many cases, the reason for grouping entities to a concept may involve abstract notions expressed in terms of other abstract notions, and these notions may depend on the individual experience of the agent holding the concept. Consider, e.g., the concept of "friendship" or "revolution." The meaning of such concepts changes with an agent's experience, and thus these concepts are living personal constructs. Such concepts may be manifested primarily by instances the agent may generate. To define a precise classification rule for such concepts, even only for a specific agent and for a specific time period, one would have to inspect all possible concept instances generated by the agent. But this may be impossible because the agent may not be able to classify some instances, and/or because the number of instances may be too large. Consequently, some human concepts are intrinsically imprecise, and there are, in principle, no precise classification rules for

them. For such concepts there can only be approximate classification rules derived from abstract linguistic or other descriptions, or from subsets of instances generated by a reasoning agent.

Since concepts are viewed as sets, then set-theoretic operations apply to them. Consequently, two concepts are identical, if the corresponding sets of instances can be shown to be the same. The difference between concepts can be defined by the set difference, or by a characterization of such a set difference. A similarity between concepts can be measured by a similarity between sets. A concept may have a name, or may have no name. For example, a collection of entities such as crowds, tanks, speeches, white-red-and-blue flag, and broken monuments of formerly admired communist leaders may constitute an unnamed concept of “things that come to mind when thinking about the August 1991 revolution in USSR.” Different people may associate different objects with this revolution, therefore this concept is a personal construct.

Concepts that were created by defining an a priori principle or classification rule (e.g., the concept of prime numbers), or for which an effective classification rule has been developed (e.g., mammals) are called *categories*. Categories are thus special kinds of concepts. An effective procedure should be expressed in terms that are either directly observable or measurable, or themselves have an effective recognition procedure. From this viewpoint, concepts such as “peace” or “freedom” are not categories, because no effective recognition procedure for them have been developed. On the other hand, concepts such as “students that received an A in my class” or “a set of blue jays” are categories, because statements describing them define an effective recognition procedure (assuming that there exists an effective procedure for identifying students in my class and blue jays).

Thus, concepts have the logical status of sets. Some concepts are sets that are relatively stable, and thus permanently useful in describing our world or our actions in it. Therefore, we develop effective classification rules for them, and they become categories. For example, animal species, plants, diseases, astronomical objects, etc. are such concepts. We call such concepts *natural categories*. Other concepts represent sets that may have only temporary or theoretical utility, like, for example, the concept of “things to carry out of a burning house” (Barsalou, 1987), or a “set of pins in my drawer.” Such concepts are not natural categories, although they may be categories. As stated earlier, for a given concept, there may be many different classification rules. Some of them may be simple but only approximate (e.g., those defining a type of animal based on its physical appearance), and some may be complex but precise (e.g., those defining the type of animal by the genetic code analysis). Depending on the need, different types of rules are employed.

Concepts that are sets with precise boundaries that do not change with the context or situation in which the concepts are used, are called *crisp*. Crisp (“well-defined”) concepts occur primarily in science, where—whenever it is useful—they are given a name (e.g., the concepts “even number” or “vertebrates”). These concepts remain crisp usually and only if used within the scope of a given scientific discipline, e.g., “triangle” or “rectangle” in geometry, a “group” in mathematics, “water” in chemistry. Outside of this discipline, they may represent different sets of entities (e.g., a triangle may be an arrangement of streets or cities; a group may mean a set of people, etc.). There are relatively few crisp concepts outside of science (e.g., kin relationships, such as “father” or “grandmother,” are among the few).

Concepts that intrinsically do not have precise boundaries, either because they are open-ended and/or because they are context-dependent, are called *flexible*. The boundaries of flexible concepts may dynamically evolve with time, and/or change with the context in which they are used. For example, the concept “computer” (i.e., the set of entities we call “computers”) does not have a well-defined and context-independent boundary in principle. The entities called computers have been rapidly evolving in time, and it would be unwise to a priori limit the meaning of this concept.

In order to make the above notions more clear, one needs to explain what is meant by context. By *context* of a concept we mean a set of concepts relevant to the intended meaning of the given concept. What concepts are relevant to the intended meaning is specified by the agent’s knowledge (sometimes also called agent’s theory). A context can be explicitly specified by evoking the name of a more general concept that includes the concept discussed, or by providing a description characterizing it.

In different contexts, a flexible concept may stand for similar or quite different sets of entities. For example, in the context of “old-fashioned office equipment” the concept “printer” would have a different meaning than in the context of “modern office equipment.” As another example, consider the concept “key.” In the context of doors, rooms, houses, buildings, etc., “key” would mean a set of entities that serve for opening door locks. In the context of “travel,” “suitcases,” “briefcases,” etc., the concept so named would stand for a different set of entities, related only at an abstract functionality level to the entities in the previous set. Finally, in the context of “problems,” “solutions,” “methods,” etc., “key” would stand for a very different set, related only by analogy to the previous sets, and thus could be considered as being a different concept. This example suggests that sets associated with a given name of a flexible concept consist of elements that may be bounded to each other by different relationships, such as identity, surface similarity,

analogy, functionality, or just by having certain relationships to other concepts. The context determines which sets are referred to in a given discourse.

Crisp concepts, by their very nature, can be represented by logical-style descriptions that evaluate to “true” or “false” for any given entity. Such descriptions may be in the form of disjunctive normal expressions, sets of rules, decision trees, grammars, etc. These descriptions can always be transformed, by the use of intermediate terms or relations, to a single conjunction of single conditions. Consequently, crisp concepts can be defined by a set of jointly sufficient and individually necessary conditions. One therefore can say that the classical view (e.g., Smith and Medin, 1981; Smith, 1989) applies to crisp concepts.

The above precise logic-style classification rules do not directly apply to flexible concepts. This is because boundaries of such concepts—as mentioned earlier—are intrinsically imprecise and/or dependent on the current and future contexts in which the concept will be used. For these reasons, only approximate classification rules may be attainable for flexible concepts.

It should be noted, however, that instances of concepts, both crisp and flexible, may occur with different frequencies. One meaning of “typicality” of a concept instance is associated with such a frequency of occurrence (the “frequency-based” typicality). For example, there could be typical or atypical geometrical triangles, in the sense that they occur frequently or infrequently. Since a triangle in geometry is a crisp concept (with sharp boundaries), every instance is a “prototype” triangle. In other words, all geometrical figures satisfying the definition of a geometrical triangle are “equally good” triangles.

A different situation exists regarding flexible concepts. Such concepts do not have precisely defined boundaries, or possess different boundaries in different contexts. Therefore, to create a classification rule, one needs to rely on the distributions of instances over a space of the observable descriptors. If there are some instances that belong to the concept frequently and in many situations and many contexts, then they have “high typicality.” Such typicality can be measured by the degree of similarity (defined in terms of the chosen observable descriptors) between an instance and an “average” instance in the space spanned over these descriptors. Which descriptors are chosen to characterize an instance depends on how relevant they are to the original principle underlying the concept. A judgment of relevance of different descriptors depends on the knowledge or experience that an agent has with the concept.

The prototype view of concepts (see, e.g., Rosh and Mervis, 1975) concerns flexible concepts whose instances have a “well-behaved” distribution (just one or few modalities) in the space spanned over the chosen descriptors. However, for many concepts, e.g., chair, game, house, music, mechanism, etc., it is very difficult to determine a set of measurable descriptors that would produce a space in which all their instances would be distributed in such a “well-behaved” fashion. The key problem in representing such concepts lies in the difficulty of describing all their possible manifestations and context-dependency in terms of measurable properties. It is a remarkable challenge to understand how people can communicate with such concepts without having any explicit classification rules or explicit descriptions. If specifically asked, different people usually produce different descriptions of such concepts. (To see this point clearly, the Reader may ask different people to define such common concepts like “body,” “force” or “house” to see what a wide range of definitions they may produce). Although the proposed below two-tiered representation can handle both crisp and flexible concepts, its primary aim is to represent the latter ones.

3. THE TWO-TIERED CONCEPT REPRESENTATION

The fundamental assumption underlying the two-tiered concept representation is that total classificatory information about a concept should reside not only in its explicit representation, but also in the inference and matching procedures applicable to it. Through the application of these inference and matching procedures the meaning of a concept can be easily extended, modified, or adapted to various contexts. The two-tiered representation thus constitutes a significant departure from the previous approaches, which attempted to represent the complete meaning of a concept within one explicit structure. Nevertheless, as explained later, the TT view builds upon and relates in various ways to other views.

The TT view was originally proposed in Michalski (1986), and subsequently applied to various problems in concept learning (e.g., Michalski, 1990; Bergadano et al., 1988b,c; 1991). The original formulation of this approach was based on an observation that although most human concepts lack precise definition when used outside of a specific context, they acquire a precise meaning when used in a combination with other concepts in a specific context. Consider, for example, the statement: “This blonde woman near the center of the room is a student at GMU.” If there is only one woman with hair of lighter color than any other woman near the center of the room, the statement above precisely defines the person of interest (although the concepts “blonde,” “near,” and “room” are imprecise by themselves, outside of a specific context). Thus, statements made of imprecise concepts can convey precise meaning. In other words, the fact that individual

concepts are imprecise and context-dependent (i.e., flexible) does not prevent us from communicating precise meaning.

The above example illustrates the basic supposition for the TT approach (Michalski, 1986) that the imprecision of human concepts stems not from an undesirable vagueness of our concept definitions, but rather from the universal need for cognitive economy. By allowing individual concepts to be imprecise and context-modifiable, the expressive power of concepts is greatly enhanced. The latter means that one can employ fewer concepts for expressing a greater set of meanings. This leads to a simplification of descriptions of our immensely complex universe. Various experiments reported here and elsewhere (e.g., Michalski, 1990; Bergadano et al., 1988b,c; 1991) have confirmed this idea in a microworld to which it was applied.

The two-tiered approach assumes that (flexible) concepts have a certain central tendency and basic usage, which should be described explicitly, as the “first approximation” of the concept meaning. The complete concept meaning, however, may vary significantly beyond the basic (or typical) usage, especially in different contexts. The TT view postulates that this variability and context-dependency is best represented implicitly, by appropriate matching methods and context- and background knowledge-dependent rules of inference. Accordingly, a two-tiered concept representation consists of two components (tiers):

- the BCR, the *base concept representation*, that represents the basic concept meaning explicitly. This basic meaning may be characterized by specifying the typical function of the concept instances, their usual physical properties and appearance, their structure and the components, their role in a system of concepts by indicating ancestor and descendent concepts, etc.
- the ICI, the *inferential concept interpretation*, that consists of inference rules and matching procedures, that modify the basic meaning accordingly to any given situation or context. In matching an instance with the concept representation, these inference rules and matching procedures involve the agent's background knowledge (BK) and the context of discourse. The ICI may involve any type of inference—deductive, analogical or inductive.

Thus, the basic concept meaning is handled by the BCR, and imprecision, exceptions and the context-dependency are handled by ICI. The ICI captures the allowable transformations applicable to a concept in different contexts. Entities that satisfy the base concept representation are called typical, and those that do not are characterized by a degree of typicality. Factors affecting the degree of typicality are the type and the amount of transformation needed for matching the entity

with the BCR, and the frequency of the occurrence of an entity. Highly atypical entities are called *exceptions*.

Figure 1 illustrates the relationship between the BCR and the ICI in a TT concept representation. It shows that the ICI can, in general, extend the concept meaning beyond the BCR in one area of the description space, and reduce the meaning in another area.




Figure 1. An illustration of the relationship between the Base Concept Representation (BCR) and the Inferential Concept Interpretation (ICI).

Note that BCR is not a “prototype,” but a characteristic description of a class of entities that are viewed as typical members of a given concept. Such a description may be a disjunction of conjunctions (i.e., a DNF expression). A “prototype” instance should satisfy the BCR, but there could be a number of distinct “prototype” instances that also satisfy it. For example, a chair, a desk, and a bookshelf are all typical examples of the concept “office furniture.” The BCR can be expressed at different levels of abstraction. For example, a BCR of the concept “table” can include statements stating their typical function, typical physical properties, the typical material they are made of, etc. In different situations different components of BCR can be used for concept recognition. A method of “dynamic recognition” that uses various subsets of properties for recognizing a concept in different situations (a situation is defined by the set of properties that have been observed or can be observed) has been described in Michalski (1989).

Here is an example of a TT representation of the concept “chair”:

BCR: *Function:* To seat one person.

Structure: A seat supported by legs and has an attached backrest from one side.

Physical properties: Made of wood or plastic. The seat is flat and is located about 14-18 inches above the ground. There are usually four legs which are located in the vertices of an imaginary square, etc.

(BCR may also include a picture or a 3D model of typical chairs)

Ancestor concepts (parent concept only): Furniture

Descendent concepts (children concepts only): straight chair, armchair, chaise lounge, folding chair, highchair, lawn chair, rocking chair, barber chair, throne, wheelchair.

ICI: *Possible variations of the properties in BCR:* The number of legs can vary from one to four. The legs may be replaced by any support. The shape and the material of the seat, the legs and the backrest are irrelevant, as long as the function is preserved. The backrest may be very small or missing; the seat may be closer to the ground or much higher above the ground than specified in BCR, etc.

Situation or Context-dependent variations:

Situation: If no chair is available ---> any object on which a person can sit can be called a chair.

Context = museum exhibit ---> chair is not used for seating persons any more, although may have all the physical characteristics of a typical chair.

Context = {toys, doll's house, dolls, children playing,...} ---> the size can be much smaller than stated in BCR. The chair does not serve for seating persons, but correspondingly small dolls.

To recognize a concept instance, the instance is matched against the BCR of candidate concepts. This process first attempts to find a BCR that directly matches the instance. A match is direct if the BCR's conditions are directly satisfied by the instance. If there is no direct match, or there are several direct matches, then the process employs the ICI rules of inference and matching procedures. If there is no direct match with any BCR, the system searches for a deductive match, i.e., tries to determine if the instance is a logical consequence of some BCR. For example, if there is a BCR condition "X is a very good student" and ICI has a rule that "a student who has a high grade average and gets along well with most of the peers and teachers then that student is a very good student," then an instance "John has high grade average and gets along with almost

everyone” deductively matches the BCR condition. If there is no deductive match or there is more than one direct match, a *plausible match* procedure is employed. Such a procedure determines the degree to which an instance matches different concept candidates by using matching procedures and rules contained in the ICI. The best matching description defines the concept.

Most of the known concept representations, such as semantic networks, decision trees, and frame representation usually employ a direct match. Plausibility matching employs usually a method for determining some degree of similarity or partial match. In a general case, a plausibility matching may employ any type of plausible inference (e.g., approximate deduction, analogy or induction; Collins and Michalski, 1989). Illustrative examples of plausible matching are given in the next section. Performing plausible inferences may involve concept meta-knowledge, e.g., the importance of concepts and their frequencies of occurrence, the relation to other concepts and other relevant domain knowledge. Such meta-knowledge is assumed to be a part of the second tier.

As mentioned earlier, the ICI includes the background knowledge relevant to the proper interpretation of the concept. Since such background knowledge (called by some authors “theory,” e.g., see \$\$\$ by Murphy) grows with a person's experience, the same concept may have a different meaning for different people. For example, the concept of a computer is quite different for a novice than for a computer expert who witnessed the evolution of computers over many years. Thus, the two-tiered view provides a simple model for explaining the observable personal differences in concept meaning. In sum, by dividing the concept representation into an explicit and implicit part, the TT view combines elements of the classical representation, prototypical representation, and other elements not present in other representations.

In the general theory of two-tiered representation, the “distribution” of the meaning between the BCR and the ICI is not assumed to be fixed for a given concept, but modifiable, according to a *criterion of description quality for a given reasoning agent*. One extreme of such a distribution is when the BCR represents explicitly all possible concept variations in different contexts, and the ICI is just a direct matching procedure. Another extreme is when the BCR is empty, and whole meaning of the concept resides in context-dependent inference rules. The description quality criterion reflects computational properties of the reasoning agent (e.g., the relative costs of directly remembering concept properties vs. deriving them through inference).

In the cognitively-oriented form of the TT representation, the distribution of the concept meaning between BCR and ICI is more fixed. The BCR expresses the general unifying idea, the typical

function, and/or measurable properties implied by or correlated with this idea or function. Such a BCR can be viewed as representing the “first approximation of the concept.” The ICI, in this case, defines the matching procedures and inference rules for handling less typical instances and context-dependency.

An advantage of distributing the concept meaning between the BCR and the ICI is that it permits a learner to flexibly modify or extend the concept meaning by varying matching procedures and inference rules, and/or by changing the context of discourse. The concept meaning can thus be changed without having to alter the base concept representation. As mentioned earlier, by evaluating the type and the amount of inference involved in matching the BCR with an instance, one may produce a qualitative or quantitative estimation of the degree of typicality.

The ability to produce a measure of the degree of match indicates one principal difference between this approach and the fuzzy set approach (e.g., Zadeh, 1965, 1976). In the fuzzy set approach, a set defining a concept is associated with a membership function, which needs to be *defined* to the learner by a person. The influence of the context is hidden in the definition of this membership function. In the proposed approach, a concept is associated with interpretation procedures and context-dependent rules, which implicitly define the membership of an instance in a concept. These rules and procedures can be used to *compute* the membership function in different contexts.

4. EXAMPLES JUSTIFYING THE TWO-TIERED APPROACH

There are many examples that show that human concepts are not characterized solely by explicitly stated properties, relations, or prototypes. A full description of most concepts involves general and domain-specific rules and background knowledge, in order to capture the concepts’ variations, and context-dependency. These rules tell us what transformations of instances preserve their concept membership and which do not, which are potentially feasible, and what transformations move them outside of a concept. For example, repainting a chair preserves its membership in the category of chairs, but shredding it does not. Consequently, to learn a concept one needs to learn not only the base concept representation, but also inferential concept interpretation. Let us consider a few examples illustrating and confirming the ideas of the two-tiered representation.

Example 1. A pizza or a coin?

This example is based on recent experiments by Rips (1989), which involved concepts that can be characterized, using the TT terminology, as having significantly different ICIs, and somewhat comparable values of one attribute in their BCRs. Specifically, a subject was presented with a partial description of an object that expressed only the value of one of its attributes (e.g., its diameter). Based on this description, the subject was supposed to classify the object to one of the two categories. The attribute value given in the description was somewhere between the subject's average values of this attribute for the two categories. For example, an object was described as "three inches in diameter," and the categories of choice were pizzas and quarters. The results of experiments were that the subjects consistently classified such an object as likely to be a pizza, although the specified size was closer to the quarter than to a typical pizza.

The two-tiered model explains this behavior simply. The subjects not only possess the description of the typical values of objects in each class (the BCR), but also have an understanding of the possible variability of the objects in each category (the ICI). This understanding comes either from a generalization of their observations, or from their relevant background knowledge, e.g., about the process that produces objects of the two categories. Consequently, they could admit an object to the category of pizzas, although unusually small for a pizza, but not to the category of quarters, since quarters do not have a variable diameter.

Example 2. Concept of sugar maple

Our prototypical image of a sugar maple is that it is a tree with three- to five-lobbed leaves that have V-shaped clefts. Some of us may also remember that the teeth on the leaves are coarser than those of red maple, that slender twigs are brown, and that the buds are brown and sharp-pointed. As a tree, of course, a maple has a root, trunk and branches.

Suppose that while strolling on a nice winter day someone tells us that a particular tree is a sugar maple. A simple introspection tells us that even though the tree does not have leaves, this is not a contradiction of what we know about typical sugar maples. Yet, clearly, the presence of leaves of a particular shape is deeply embedded in our typical image of a maple tree. The two-tiered theory explains this phenomenon simply: the inferential concept interpretation associated with the general concept of deciduous trees evokes a rule "in winter deciduous trees lose leaves." Since a maple is a deciduous tree, the rule would apply to the maple tree. The result of this inference would override the stored standard information about maple trees, and the inconsistency would be resolved.

Matching an instance with the concept requires in this case deductive reasoning from the knowledge associated with a more general concept.

Example 3. A coffee pot used as a bird feeder

This example is taken from a study of childrens' concepts performed by Keil (1989). In this study, children were shown objects representing some known concepts. These objects then underwent a sequence of transformations that produced an object that resembled an instance of a different concept. In one experiment, the transformations were realistic, i.e., corresponded to changes that can be obtained by some simple operations. In the other experiments, the transformations were purely imaginary (they did not correspond to changes that would be allowed by the ICI).

For example, in one case the initial object was a coffee pot. Subsequently, the pot was punched full of holes, and its spout was removed. It was then filled with bird feed and hung from a branch of a tree. The children were then asked if the object was a coffee pot or a bird feeder. In another case, the initial object (presented as a picture) was a raccoon. It was then presented as having been transformed to black with a white tail, and able to produce a super smelly stuff from beneath its tail, like a skunk. The children were then asked if the end result of these transformations was a raccoon or a skunk. The experiments thus can be interpreted as testing the children's understanding of the feasibility of various transformations of instances of different concepts.

The results of the first experiment were that all children (5 to 9 years old) indicated that the object was a coffee pot that was turned into a bird feeder. In the second experiment, however, all children experienced difficulty in accepting the change of the raccoon into a skunk. These experiments indicate that children knew not only physical characteristics of the objects in the involved categories (here, the coffee pots and raccoons), but had also the knowledge of transformations that are applicable to objects in these categories.

Example 4. Concept of a triangle

Let us consider the concept of "triangle." In geometry, a triangle is a plane figure consisting of three non-colinear points connected by straight lines. Using the notation of *annotated predicate calculus* (APC), which is like predicate calculus, but employs more compact forms of some logical expressions (Michalski, 1983), one can write:

Triangle(T, P1, P2, P3) \Leftarrow Consists(T, P1 & P2 & P3) & Type(P1 & P2 & P3, point) &
 Connected_by_straight-line(P1,P2 & P1,P3 & P2,P3) & Non-colinear(P1, P2, P3)

In the above expression, symbol “&” is used in two related meanings: first, to denote an ordinary (*external*) conjunction connecting predicates, and second, to denote an *internal conjunction*, i.e., a conjunction of terms, which is treated as a *compound* argument of a predicate. For example, the predicate “Type(P1 & P2 & P3, point) states that P1 and P2 and P3 are points, thus is equivalent to a conjunction of three predicates: Type(P1, point) & Type(P2, point) & Type(P3, point).

Suppose that someone tells us that the tall towers in his/her hometown form a big triangle. Obviously, the meaning of the triangle in this statement differs from that in the formal geometrical description. To match the statement with the concept of triangle, the following assumptions and transformations need to be made:

a. In the context of describing a configuration of physical objects such as towers, the individual objects play the role of nodes. Thus, the statement implies that towers correspond to nodes, and that there are three towers in the town. The matching operation involves drawing an analogy between the abstract nodes and the towers, which can be characterized as consisting of one step of generalization (GEN):

Point ----GEN----> Object

and one step of specialization (SPEC):

Object ----SPEC----> Tower

b. In the context of towers, the presence of a “straight line” is imaginary, i.e., there is no physical connection, but one could imagine a straight line between the objects (towers). The condition “ConnectedBy” is then satisfied in such an abstract sense. This is an operation of abstraction. Thus, matching the statement about a triangular arrangement of towers with the definition of a triangle involves here a generalization, specialization and abstraction. [A reader interested in an analysis of the differences between generalization and abstraction may consult Michalski (1991).]

The examples above show that relating a concept instance to a concept representation is not just a straightforward comparison of attribute values in an instance with those in the concept

representation, as done in various mechanized decision processes. They show that such a process may involve different forms of inference, as postulated in the two-tiered representation.

5. LEARNING TWO-TIERED CONCEPT REPRESENTATIONS

The essence of the two-tiered representation is a recognition that a concept may have explicit and implicit properties. The explicit properties are represented by logical-style expressions, and the implicit ones by a set of rules and/or matching procedures. Consequently, to acquire a TT concept representation, one needs to learn both, the explicit (BCR), and the implicit (ICI) components of a TT representation.

In general, each concept has a different BCR that has to be learned. On the other hand, the ICI is often the same or very similar for various classes of concepts. Consequently, it can be learned only once for a given class of concepts. All concepts of this class, and/or descendent concepts in the type hierarchy will share the same ICI. For example, all diseases in the class of liver diseases may share the same flexible method for matching the BCR descriptions with symptoms. By sharing the ICI among concepts, a significant economy of the concept representation can be achieved.

The two-tiered approach does not impose restrictions on what type of language is used to represent BCR and ICI parts of a concept. Any type of knowledge representation formalism could potentially be employed. It is assumed that the language employed should depend on the goal of the concept representation. The goal implies the degree of abstraction with which the concept needs to be represented, and this sets the criteria for the representation language. For example, in a diagnostic situation, it may be sufficient to represent a disease at a high level of abstraction, e.g., by an attributional description in the form of propositional calculus rules. In a medical research situation, the representation may have to be much more complex, and involve not only structural representation (e.g., using a predicate calculus language), but also images and drawings.

In most of the experiments done so far with the TT representation, the representation language was an attributional calculus, called VL₁. In this calculus, descriptions are expressed in the form of decision rules. A decision rule links a logical product of elementary conditions with a decision (e.g., a class membership). The elementary conditions relate a multiple-valued attribute to a value or a set of values. For example, elementary conditions may be [sex = male], or [color = blue or red] or [weight > 2 kg]. [See, e.g., (Michalski, 1974; Bergadano et al., 1991)] for details on this representation.). In some experiments, an extension of predicate calculus (called *annotated predicate calculus*) was used (Bergadano et al., 1988b).

The rest of this section outlines an implemented method for learning two-tiered concept representations from concept examples using the above mentioned attributional calculus. It also presents results from comparing the method with several other methods that use different concept representations. Full details of the method and the experiments performed are described in Bergadano et al. (1988b&c; 1991).

In the method, the *Base Concept Representation* (BCR) is created in two phases. In Phase I, a complete and consistent concept description is induced from the supplied examples. This is done by an inductive learning program AQ-15 (Michalski et al., 1986). Such a description explains (covers) all positive examples of the concept being learned, and none of its counter-examples (negative examples). The description is in the form of rules that link single conjunctive conditions with a concept name.

In Phase II, the obtained description is optimized according to a *description quality criterion*. The criterion prefers the descriptions that are computationally simple and at the same time have a high predictive accuracy, i.e., correctly classify new concept instances. The optimization process involves making various changes in the original description, such as removing a rule, removing a condition, or modifying a condition in a rule. These changes correspond to either generalization or specialization of the original rules. For example, rule removal is a specialization operation, and condition removal is a generalization operation. A condition modification can be either a specialization or generalization operation depending on the type of modification. The rules obtained as a result of the optimization process may be significantly simpler than the original rules.

The obtained rules, however, may not directly match all the positive examples, i.e., they may constitute an *incomplete* concept description, and/or may directly match some negative examples, i.e., may constitute an inconsistent concept description. It turns out, however, that the incompleteness and/or inconsistency is often well compensated by the inferential concept interpretation (ICI), so that the resulting rules may actually have a higher predictive accuracy than the original rules.

Depending on the type of description learned, the ICI may consist of just a procedure for *flexible matching*, or the procedure plus a set of rules that capture the concept exceptions. In the latter case, the rules are acquired through an interaction with a teacher.

The method has been implemented in POSEIDON (AQ16), and experimentally tested on two real-world problems: learning the concept of an acceptable union contract, and learning voting patterns

of Republicans and Democrats in the U.S. Congress. For comparison, a few other learning methods were also applied to the same problems. These methods included simple variants of exemplar-based learning, an ID3-type decision tree learning (ASSISTANT program; Chestnik et al., 1987) a method that utilizes complete and consistent descriptions, and a method that uses a very simple BCR (the “Top rule” description) and flexible matching.

To give the reader some sense of the results obtained, Table 1 presents a summary of the results from testing these methods (based on Bergadano et al., 1988b,c; 1991).

The “Simple exemplar-based method” illustrates the exemplar view of concepts, i.e., represents concepts by sets of their examples. An unidentified concept instance is recognized by determining examples that are most “similar” or “close” to it. In the “k-nearest” procedure, k most similar examples are determined, and the concept associated with the majority of them is assigned to the instance (in the experiments reported, k was 1, 3 and 5).

The predictive accuracy of each description expresses the percent of the correct decisions (as evaluated by an expert) when the description was applied to a set of testing examples. The complexity of the description was measured by the number of rules (#Rules) in it, and the total number of elementary conditions in all the rules (#Conds).

The “Complete and consistent” concept description is a disjunctive normal expression (equivalent to a set of rules) that describes all positive examples, and none of the negative examples. The “Top rule” description represents a very simple variant of two-tiered representation. The BCR is just a single rule that covers the largest number of examples among all the rules in the consistent and complete description. Such a rule can be viewed as representing the most typical concept examples.

The ICI is a procedure for *flexible matching* that determines a similarity between an instance and the BCR of candidate concepts. The concept associated with the most similar BCR is assigned to the unidentified instance. It is interesting to notice although this form of two-tiered representation is extremely simple, it still has a relatively high predictive accuracy.

The “Optimized two-tiered description” is a more advanced variant of the TT representation. The BCR is a set of rules generated by removing some rules and/or conditions from the consistent and complete description. This is done by the so-called TRUNC/SG optimization procedure (Bergadano 1988b,c; 1991).

<u>Simple exemplar-based method</u>	<u>Labor Contract</u>	<u>Congress Voting</u>
	Predictive accuracy	Predictive accuracy
<i>Performance</i>		
(%Correct / % Incorrect)		
<i>1-nearest neighbor</i>	77% / 23%	86% / 14%
<i>3-nearest neighbor</i>	83% / 17%	84% / 16%
<i>5-nearest neighbor</i>	80% / 20%	84% / 16%
<i>Complexity</i>		
(#Rules / #Conds)	27 / 432	51 / 969

<u>Pruned decision tree</u>		
(ASSISTANT + PRUNING)		
<i>Performance</i>		
(%Correct / % Incorrect)	86% / 14%	86% / 14%
<i>Complexity</i>		
(#Leaves/ #Nodes)	29 / 53	19 / 28

<u>Complete and consistent description</u>		
(AQ15 without rule truncation)		
<i>Performance</i>		
(%correct / % incorrect)	80% / 17%	86% / 14%
<i>Complexity</i>		
(#Rules / #Conds)	11 / 29	10 / 32

<u>"Top rule" description</u>		
(AQ15 with rule truncation)		
<i>Performance</i>		
(%Correct / % Incorrect)	83% / 17%	85% / 15%
<i>Complexity</i>		
(#Rules / #Conds)	2 / 6	2 / 6

<u>Optimized two-tiered description</u>		
(AQ16 or POSEIDON with rule truncation and condition optimization)		
<i>Performance</i>		
(%Correct / % Incorrect)	90% / 10%	92% / 8%
<i>Complexity</i>		
(#Rules / #Conds)	9 / 12	10 / 21

Table 1. Summary of the results for testing descriptions learned by different methods.

The procedure generalizes some parts and specializes other parts of the consistent and complete description in order to maximize a *general description quality criterion*. The ICI of the concept was

a set of expert-generated rules for covering exceptional cases identified by the system. As shown in Table 1, the “Optimized TT description” had the highest predictive accuracy, and was also the next simplest, after the “Top rule” description.

6. THE TIERED VIEW VS. OTHER VIEWS

As mentioned earlier, cognitive science literature distinguishes among several *views* of concept representation. The major ones are the classical view, the prototypical (or probabilistic) view, an exemplar view, and the theory view (Smith and Medin, 1981; Medin and Smith, 1984; Nosofsky, 1987, Allen et al., 1988; Murphy and Medin, 1985; and \$\$\$\$s by Barsalou and Hale, by Hampton, and by Murphy in this book). Some authors also distinguish between different *types of formalisms* (or representation languages) employed in the concept representation, e.g., feature lists, decision trees, connectionist models, predicate logic expressions, simple and complex frames (e.g., Barsalou and Hale, this book). The distinction between a “view” and a “representational formalism” is usually not well explained in the literature. Our interpretation is that a view of a concept concerns inherent aspects of concepts as sets, while the representational formalism is a language in which those aspects are expressed. In principle, a sufficiently rich representational formalism, e.g., predicate calculus, with a proper interpretation, could be used for expressing any of the concept views.

The classical view assumes that concepts have well-defined borders, and can be represented by singly necessary and jointly sufficient conditions. This implies that they can be represented by logic-style conjunctive definitions. The inadequacy of this view has been widely recognized, e.g., Wittgenstein, 1922; McCloskey and Glucksberg, 1978; Barsalou and Medin, 1986; and Lakoff, 1987. To overcome the limitations of this view, other views have been advanced, such as the *prototype view* and the *exemplar view* (e.g., Smith and Medin, 1981; Medin and Smith, 1984; Nosofsky, 1987, Allen et al., 1988).

The prototype view represents concepts by prototypes, and uses the so-called *family resemblance* principle (e.g., Rosch and Mervis, 1975). Unlike the classical view, this view does not require that all concept instances share the same properties. It can, therefore, explain the observed phenomenon that instances of the same concept may differ in their *typicality*. The exemplar view claims that concepts can be adequately represented by their exemplars (e.g., Smith and Medin, 1981; Bareiss, Porter and Craig, 1990). Both these views can also be criticized on various grounds. The prototype view, which formally is based on the idea of linear separability, disregards the existence of correlations between the attributes, the context-dependency, and other information that has been

shown to be relevant to human concept understanding (e.g., Estes, 1986; Flannagan, Fried and Holyoak, 1986).

The exemplar view employs similarity-based and context-sensitive matching, which has received support in the cognitive science literature. It is also supported by an observation that people do store examples of concepts, especially exceptional examples. This view disregards, however, the importance of general concept descriptions, that clearly play a role in human concept representations. Such general descriptions are useful in many ways, for example, for quickly identifying the differences between concepts, for recognizing them from partial information, for handling context-dependency, for efficiently storing invariant information about concepts, etc. The above operations would be difficult to perform, if concepts were represented only by examples. In some work using the exemplar view, general aspects of concepts are captured under the idea of “category structure,” which is a network of domain knowledge that specifies the relevance of exemplars to the concept they define (Bareiss, Porter and Craig, 1990).

Some recent work has advocated a theory view (or knowledge-based view) which emphasizes the need to define concepts through their role in theories in which they exist as interrelated components (Carey, 1985; Hofstadter, 1985; Murphy and Medin, 1985; Schank, Collins and Hunter, 1986; Murphy - this book).

The two-tiered (TT) view constitutes a significant departure from the existing approaches, although it has a close relationship to most of them. It assumes that concepts have a certain central tendency, and proposes to describe this tendency explicitly, in the form of a logic style description. Such a description is the “first approximation” of the concept. On the other hand, it also assumes that concepts may have a great variability and be context-dependent, and that these aspects are best handled not explicitly, but implicitly, by a flexible matching method and context- and background knowledge-dependent rules of inference.

Since the TT view proposes to represent basic concept properties (BCR) in the form of explicit logic-style expressions, it is related to the classical view (it admits, however, that the BCR can be not only conjunctive descriptions, but also disjunctive descriptions). It also recognizes that most concepts have a central tendency, and therefore one needs to distinguish between more or less typical concept examples. In this sense, the TT view is similar to the probabilistic view and the fuzzy set concept representation (e.g., Zadeh, 1965 and 1976). Unlike these views, it proposes to handle this aspect by the inferential concept interpretation (ICI) that employs a flexible concept matching procedure (Bergadano et al., 1991).

The TT view has also a relationship to the exemplar view, as it postulates the use of sophisticated matching procedures and inference rules in classifying new instances. Also, the TT view recognizes the usefulness of storing individual examples, especially in the case of exceptions (e.g. Zhang, 1990a, b). The TT approach is also closely related to theory view, as it stresses the role of background knowledge and inference in concept representation, especially, in the case of non-typical or borderline instances. The main differences between the TT view and the theory view are that the former makes a distinction between the explicit and implicit representation, and postulates that the agent's knowledge or "theory" plays a role primarily in the inferential concept representation, rather than in the basic concept representation.

7. SUMMARY AND OPEN PROBLEMS

The most significant aspect of the presented two-tiered view is that it represents concepts as a combination of an explicit and implicit representation, in contrast to other views that attempt to compact concept representation into a single monolithical structure. In the two-tiered representation, the first tier, the base concept representation (BCR), captures the explicit and common concept meaning, and the second tier, the inferential concept interpretation (ICI), defines allowable modifications of the base meaning and exceptions. The typical concept instances strictly match the BCR, and thus can be recognized efficiently. The exceptional and/or context-dependent cases involve the inferential concept interpretation, which takes more time, but is done relatively less often.

Such a two-tiered representation is particularly suitable for learning flexible concepts, i.e., concepts that inherently lack precise definition and are context-dependent. The TT view asserts that it is impossible, in principle, to give such concepts a precise meaning in an explicit and context-independent form. Therefore, it proposes to describe precisely only their central tendency, and to use flexible matching procedures and inference rules to characterize the concept's variability, to resolve borderline cases, and to express their context-dependency.

The TT view seems to be confirmed by an observation that people usually can recognize rapidly and without any difficulty typical concept instances, e.g., that chair is an example of furniture. When, however, they need to decide about special borderline cases (e.g., if a lamp or painting is an example of furniture), they try to resolve the problem by reasoning and/or an evaluation of the similarity to the cases with clear membership. This seems to contradict approaches that try to resolve such cases by weighting features or statistical evaluation. In the TT view, the reasoning

and similarity eventuation is conducted by matching procedures, inference rules and agent's background knowledge residing in the second tier of the concept representation.

The TT view also assumes that the imprecision of human concepts stems not from an undesirable vagueness of our concept definitions, but rather from the universal need for cognitive economy. By allowing concepts to have a context-modifiable meaning, and making them precise only to the extent to which a given situation and/or context requires them to be precise, the expressive power of concepts is greatly enhanced. This means that one can employ fewer concepts for expressing more meanings, and this helps us to simplify our descriptions of our immensely complex universe.

Thus, by compacting the variability of concepts into rules and matching procedures, rather than trying to express it explicitly, one can reduce the overall complexity and improve the effectiveness of concept representation. The experiments on learning TT representations have demonstrated that the obtained TT representations tend to be considerably simpler and have better predictive accuracy than other representations.

The ideas of TT representation are recent, and many problems remain unsolved. Among especially interesting and important problems is how to represent not just one, but a complex system of concepts using the TT representation. As indicated in the paper, the ICI usually can be shared or inherited from higher level concepts. This problem is then how to allocate different parts of the ICI within a structure of TT representation. There is also a need for a clearer analysis of the relationship between the TT view and other views, and experimental studies with human subjects to determine the adequacy of the TT model as a cognitively viable concept representation. There are also many problems related to the development of systems capable of learning TT representations, and efficiently using them for concept recognition processes. In particular, an important problem is how people learn individual tiers of different concepts.

Finally, a theoretical analysis and more experimental investigations are needed to confirm or disconfirm the initial finding that the TT representation can lead to significant memory savings and a simultaneous improvement of the predictive accuracy of concept representations over other representations.

Original Submission Date: September 26, 1991.

REFERENCES

- Allen, S.W., Brooks, L.R. and Norman, G.R., "Effect of Prior Examples on Rule-based Diagnostic Performance," *Proceedings of the 29th Annual Meeting of the Psychological Society*, November 1988.
- Bareiss, E.R., Porter, B.W. and Craig, C.W., "PROTOS: An Exemplar-Based Learning Apprentice," \$\$\$ in *Machine Learning: An Artificial Intelligence Approach Volume III*, Y. Kodratoff and R.S. Michalski (eds.), Morgan Kaufmann Publishers, 1990.
- Barsalou, L.W., "The Instability of Graded Structure in Concepts," in *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, U. Neisser (Ed.), pp. 101-40, Cambridge University Press, 1987.
- Barsalou, L. W. *Cognitive Psychology: An Overview for Cognitive Scientists*, D.A. Norman and A. Ortony's (Eds.), Tutorial series on Cognitive Science, Lawrence Erlbaum Associates, September 1990.
- Barsalou, L.W. and Medin, D.L., Concepts: Fixed Definitions or Dynamic Context-Dependent Representations, *Cahiers de Psychologie Cognitive*, Vol. 6, pp.187-202, 1986.
- Bergadano, F., Matwin, S., Michalski, R.S. and Zhang, J., "Measuring Quality of Concept Descriptions," *Proceedings of the Third European Working Session on Learning*, D. Sleeman, Pitman, London, pp. 1-14, 1988a [An extended version was published under the title "A Measure of Quality of Descriptions," *Reports of the Machine Learning and Inference Laboratory*, Vol. 88-3, Artificial Intelligence Center, George Mason University, 1988a.]
- Bergadano, F., Matwin, S., Michalski, R.S., and Zhang, J., "Learning Flexible Concept Descriptions Using a Two-tiered Knowledge Representation: Implementation and Experiments," *Reports of the Machine Learning and Inference Laboratory*, Vol. 88-35, Artificial Intelligence Center, George Mason University, 1988c.
- Bergadano, F., Matwin, S., Michalski R. S. and Zhang, J., "Learning Two-tiered Descriptions of Flexible Concepts: The POSEIDON System," *Machine Learning Journal and Reports of the*

Machine Learning and Inference Laboratory, MLI 91-3. Center for Artificial Intelligence, George Mason University, Fairfax, VA, May 1991. (Accepted for publication in *Machine Learning Journal*, 1992.)

Bergadano, F., Matwin, S., Michalski, R.S., and Zhang, J., "Learning Flexible Concept Descriptions Using a Two-tiered Knowledge Representation: Ideas and a Method," *Reports of the Machine Learning and Inference Laboratory*, Vol. 88-4, Artificial Intelligence Center, George Mason University, 1988b.

Bratko, I., Mozetic, I. and Lavrac, L., *Kardio: A Study in Deep and Qualitative Knowledge for Expert Systems*, MIT Press, Cambridge, MA, 1989.

Carey, S., *Conceptual Change in Childhood*, MIT Press, Cambridge, MA, 1985.

Chestnik, B., Kononenko, I. and Bratko, I., ASSISTANT 86: A Knowledge Elicitation Tool for Sophisticated Users, *Proceedings of the 2nd European Working Session on Learning*, I. Bratko and N. Lavrac (ed.), Bled, Yugoslavia, May 1987.

Collins, A. and Michalski, R.S., "The Logic of Plausible Reasoning: A Core Theory," *Cognitive Science*, Vol. 13, 1989.

Collins, A. and Smith, E.E., *Readings in Cognitive Science: A Perspective from Psychology and Artificial Intelligence*, Morgan Kaufmann, 1988.

Estes, W.K., "Memory Storage and retrieval Processes in Category Learning," *Journal of Experimental Psychology: General*, Vol. 115, pp.155-175, 1986.

Flannagan, M.J., Fried, L.S. and Holyoak, K.J., "Distributional Expectations and the Induction of Category Structure," *Journal of Experimental Psychology: Learning, Memory and Cognition*, Vol. 12, pp. 241-56, 1986.

Gentner, D., and Landers, R., "Analogical Reminding: A Good Match is Hard to Find," *Paper prepared for the Panel on Commonsense Reasoning at the International Conference on Systems, Man and Cybernetics*, Tucson, Arizona, 1985.

Gross, P.K., "Incremental Multiple Concept Learning Using Experiments," *Proceedings of the Fifth International Conference on Machine Learning*, Ann Arbor, June 12-14, 1988.

Hoff, W., Michalski, R.S., and Stepp, R.E., "INDUCE.3: A Program for Learning Structural Descriptions from Examples," *Report ISG 83-4, UIUCDCS-F-83-904*, Dept. of Computer Science, University of Illinois, Urbana, 1983.

Hofstadter, D.R., "Analogies and Roles in Human and Machine Thinking," chapter in *Metamagical Themas: Questing for the Essence of Mind and Pattern*, Basic Books, Inc., 1985.

Holyoak, K.J., "The Pragmatics of Analogical Transfer," *The Psychology of Learning and Motivation*, 19, 1985.

Hong, J., Mozetic, I., and Michalski, R.S., "AQ15: Incremental Learning of Attribute-Based Descriptions from Examples, The Method and User's Guide," *Report ISG 86-5, UIUCDCS-F-86-949*, Dept. of Computer Science, University of Illinois, Urbana, 1986.

Iba, W., Woogulis, J. and Langley, P., "Trading Simplicity and Coverage in Incremental Concept Learning," *Proceedings of the Fifth Intern. Conference on Machine Learning*, Ann Arbor, June 12-14, 1988.

Keil, F. "Concepts, Word Meanings, and Cognitive Development," Harvard University Press, Cambridge, MA, 1989.

Kempler-Nelson, D.G., "The Effect of Intention on What Concepts are Acquired," *Journal of Verbal Learning and Verbal Behavior*, Vol. 23, pp. 734-59, 1984.

Lakoff, G., *Women, Fire and Dangerous Things: What Categories Tell us About the Nature of Thought*, Chicago University Press, 1987.

Lenat, D.B., Hayes-Roth, F. and Klahr, P., "Cognitive Economy in Artificial Intelligence Systems," *IJCAI*, Vol. 6, Tokyo, Japan, 1979.

McCloskey, M. and Glucksberg, S., "Natural Categories: Well-defined or Fuzzy Sets?" *Memory and Cognition*, Vol. 6, pp. 462-472, 1978.

Medin, D.L. and Smith, E.E., "Concepts and Concept Formation," in *Annual Review of Psychology*, M.R. Rosenzweig and L.W. Porter (Eds.), Vol. 35, pp. 113-118, 1984.

Medin, D., "Concepts and Conceptual Structure," *American Psychologist*, Vol. 44, pp. 1469-1481, 1989.

Mehler, G., Bentrup, J., and Riedesel J., "INDUCE.4: A Program for Incrementally Learning Structural Descriptions from Examples," *Reports of Intelligent Systems Group*, Vol. 86, Department of Computer Science, University of Illinois, Urbana, 1986.

Michalski, R.S., "AQVAL/1--Computer Implementation of a Variable-Valued Logic System VL_1 and Examples of its Application to Pattern Recognition," *Proceedings of the First International Joint Conference on Pattern Recognition*, pp. 3-17, Washington, DC, October 30 - November 1, 1973.

Michalski, R.S., "Variable-Valued Logic: System VL_1 ," *Proceedings of the 1974 International Symposium on Multiple-Valued Logic*, West Virginia University, Morgantown, pp. 323-346, May 29-31, 1974.

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

Michalski, R.S., "Designing Extended Entry Decision Tables and Optimal Decision Trees Using Decision Diagrams," *Report No. 898*, Department of Computer Science, University of Illinois, Urbana, March 1978b.

Michalski, R.S., "Theory and Methodology of Machine Learning." In R.S. Michalski, J.G. Carbonell, T.M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach*, Morgan Kaufmann, Palo Alto, Tioga, 1983.

Michalski, R.S., "Concept Learning," *Encyclopedia of Artificial Intelligence*, John Wiley & Sons, 1986.

Michalski, R.S., "Two-tiered Concept Representation, Analogical Matching and Conceptual Cohesiveness," *Invited paper for the Workshop on Similarity and Analogy*, Allerton House, University of Illinois, June 12-14, 1986. [An extended and improved version is in *Analogy and Similarity*, S. Vosniadou and A. Ortony (eds.), Cambridge University Press, 1989]

Michalski, R.S., "How to Learning Imprecise Concepts: A Method for Employing a Two-tiered Knowledge Representation in Learning," *Proceedings of the Fourth International Workshop on Machine Learning*, University of California at Irvine, June 22-25, 1987.

Michalski, R.S., "Dynamic Recognition: An Outline of a Theory on How to Recognize Concepts without Matching Rules," *Reports of Machine Learning and Inference Laboratory*, Center for Artificial Intelligence, George Mason University, June 1989.

Michalski, R.S. "Toward A Unified Theory of Learning: An Outline of Basic Ideas," *Proceedings of the First World Conference on the Fundamentals of Artificial Intelligence*, Paris, France, July 1-5, 1991.

Michalski, R.S. and Chilausky, R.L., "Learning by Being Told and Learning from Examples: An Experimental Comparison of the Two Methods of Knowledge Acquisition in the Context of Developing an Expert System for Soybean Disease Diagnosis," *International Journal of Policy Analysis and Information Systems*, Vol. 4, No. 2, pp. 125-16, 1980.

Michalski, R.S. and Larson, J., "AQVAL/1 (AQ7) User's Guide and Program Description," *Report No. 731*, Dept. of Computer Science, University of Illinois, Urbana, 1975.

Michalski, R.S. and Larson, J.B., "Selection of Most Representative Training Examples and Incremental Generation of VL₁ Hypotheses: the Underlying Methodology and the Description of Programs ESEL and AQ11," *Reports of the Department of Computer Science*, Vol. 867, University of Illinois, Urbana, May 1978.

Michalski, R.S. and McCormick, B.H., "Interval Generalization of Switching Theory," *Reports of the Department of Computer Science*, Vol. 442, University of Illinois, Urbana, 1971.

Michalski, R.S. and Negri, P., "An Experiment on Inductive Learning in Chess End Games," *Machine Representation of Knowledge, Machine Intelligence*, Vol. 8, pp. 175-192, E. W. Elcock and D. Michie (eds.), Ellis Horwood Ltd., New York, 1977.

Michalski, R.S. and Stepp, R.E., "Learning from Observations: Conceptual Clustering," in *Machine Learning - An Artificial Intelligence Approach*, R.S. Michalski, J.G. Carbonell, T.M. Mitchell (eds.), Palo Alto, Tioga, 1983.

Michalski, R.S. and Winston, P.H., "Variable Precision Logic." *Artificial Intelligence Journal*, 29, 1986.

Michalski, R.S., Mozetic, I., Hong, J. and Lavrac, N., "The AQ15 Inductive Learning System: An Overview and Experiments," *Proceedings of the American Association for Artificial Intelligence Conference (AAAI)*, 1986.

Mumdani, E.H. and Gaines, B.R. (eds.), *Fuzzy Reasoning and its Applications*, Academic Press, 1981.

Murphy, G.L. and Medin, D.L., "The Role of Theories in Conceptual Coherence," *Psychological Review*, Vol. 92, 1985.

Nosofsky, R.M., "Exemplar-based Accounts of Relations between Classification, Recognition and Typicality," *Journal of Experimental Psychology: Learning, Memory and Cognition*, 1988.

Popper, K.R., *Objective Knowledge: An Evolutionary Approach*, Oxford at the Clarendon Press, 1979.

Quinlan, R.J., "Learning Efficient Classification Procedures and their Application to Chess End Games," in *Machine Learning: An Artificial Intelligence Approach*, R.S. Michalski, J.G. Carbonell, T.M. Mitchell (eds.), Palo Alto, Tioga, 1983.

Quinlan, R.J., "Probabilistic Decision Trees," in *Machine Learning: An Artificial Intelligence Approach III*, Y. Kodratoff and R.S. Michalski (eds.), Morgan Kaufmann Publishers, 1990.

Reinke, R.E. and Michalski, R.S., "Incremental Learning of Decision Rules: A Method and Experimental Results." in J.E. Hayes, D. Michie, J. Richards (eds.), *Machine Intelligence*, Oxford University Press, 1988.

Rips, L.J. "Similarity, Typicality, and Categorization," in *Analogy and Similarity*, S. Vosniadou and A. Ortony (eds.), Cambridge University Press, 1989.

Rosch, E. and Mervis, C.B., "Family Resemblances: Studies in the Structure of Categories," *Cognitive Psychology*, Vol. 7, 1975.

Schank, R.C., Collins, G.C. and Hunter, L.E., "Transcending Induction Category Formation in Learning," *The Behavioral and Brain Sciences*, 1986.

Smith, E.E., "Concepts and Induction," in *Foundation of Cognitive Science*, Michael I. Posner (ed.), A Bradford Book, The MIT Press, Cambridge, MA, pp. 501-526, 1989.

Smith, E.E. and Medin, D.L., *Categories and Concepts*, Harvard University Press, Cambridge, MA, 1981.

Sowa, J.F., *Conceptual Structures: Information Processing in Mind and Machine*. Addison-Wesley Publishing Company, 1984.

Sturt, E., "Computerized Construction in Fortran of a Discriminant Function for Categorical Data", *Applied Statistics*, Vol. 30, pp. 213-22, 1981.

Watanabe, S., "Knowing and Guessing, a Formal and Quantitative Study," Wiley Publishing Co., 1969.

Winston. P.H., "Learning Structural Descriptions from Examples," *The Psychology of Computer Vision*, Winston P.H. (Ed.), McGraw Hill, New York, Ch. 5, 1975.