Learning as Understanding the External World

Gheorghe D. Tecuci
Center for Artificial Intelligence, Department of Computer Science
George Mason University, 4400 University Drive, Fairfax, VA 22030-4444, USA
and Research Institute for Informatics, 71316, Bd. Mrs. Averescu 8-10, Bucharest 1, Romania

Abstract

In this paper we propose an abstract model of a learning system that represents a suitable framework for multistrategy task-adaptive learning, a type of learning that integrates dynamically a whole range of learning strategies, depending of the features of the current learning task. The system has an incomplete and partially incorrect representation of the external world, in the knowledge base KB, and receives new input information from the environment. The learning goal is to extend, update and/or improve the KB, so as to consistently integrate the information contained in the input. This means that, after learning from an input I, the KB should be such that a generalization of I is inferable from it. The learning method is based on the idea of "understanding" the input through an exploration of the KB, and an employment of different inference types.

Key words: multistrategy task-adaptive learning, understanding, knowledge base improvement, plausible reasoning,

1. Introduction

A major goal of artificial intelligence is that of creating systems capable of intelligent behavior. One defining feature of the intelligent behavior is its gradual improvement, or learning. From a general point of view, one may distinguish between two dimensions along which an intelligent system could improve itself. These are the competence of the system in solving some class of problems, and its performance in problem solving. A system is improving its competence if it learns to solve a broader class of problems, and to make fewer mistakes in problem solving. A system is improving its performance, if it learns to solve more efficiently (for instance, by using less time or space resources) the problems from its area of competence.

A largely accepted model of the learning systems is the one from Figure 1. The Knowledge Base (KB) is shared by the Learning Element (LE) and by the Performance Element (PE). The PE solves problems, by using the KB. The LE improves the KB, by using information supplied by the external world (the environment, a teacher, etc.), information from the PE attempts to solve problems, and information from the KB.

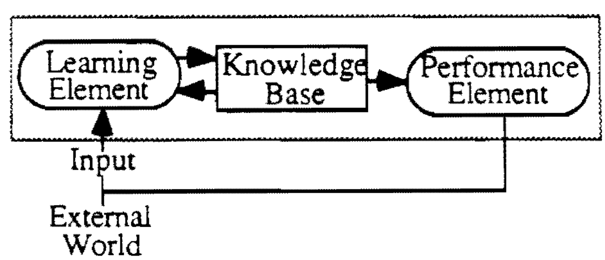

Figure 1: A model of learning systems.

Research in machine learning has elaborated and investigated in detail several single-strategy learning methods like, for instance, empirical induction, explanation-based learning, learning by abduction, learning by analogy, case-based learning, and others (Michalski, Carbonell and Mitchell, 1983,86; Kodratoff and Michalski, 1990; Shavlik and Dietterich, 1990). However, as the field of machine learning evolves and concentrates more and more on solving complex real-world learning problems, becomes more and more clear that the single-strategy learning methods provide solutions to overly-simplified problems. One kind of over-simplification consists of specific requirements imposed to the input information and to the content of the KB. For instance, empirical induction requires many input examples and a small amount of background knowledge. Explanation-based learning requires one input example and a complete background knowledge. Learning by analogy and case-based learning require background knowledge analogous with the input. Learning by abduction requires causal background knowledge related to the input. Another kind of over-simplification consists of the limited result of the single-strategy learning process. This is a hypothetical generalization of several input examples (in the case of empirical induction), or an operational generalization of an input example (in the case of explanation-based learning), or new knowledge about the input (in the case of learning by analogy or case-based learning), or new background knowledge (in the case of learning by abduction). From the above characterization, one may notice however the complementarity of the requirements and of the results of these strategies. This strongly suggests that real-world learning problems may be solved by systems that can apply different strategies in an integrated fashion, and explains the increasing interest in building such systems (e.g., Bergadano and Giordana, 1990; Danyluk, 1987; Flann and Dietterich, 1989; Genest, Matwin and Plante, 1990; Hirsh, 1989; Lebowitz, 1986; Minton and Carbonell, 1987; Mooney and Ourston, 1991; Morik, 1991; Pazzani, 1988; De Raedt and Bruynooghe, 1991; Shavlik and Towell, 1989; Tecuci and Kodratoff, 1990; Whitehall, 1990; Wilkins, 1990). Due to the integration of different strategies, these systems are able to solve more complex learning problems. However, these systems integrate only a small number of learning strategies (most frequently empirical induction, explanation-based learning, and a simple form of analogy), the integration is predetermined inside the learning method, and the result of learning is limited (usually a concept or a rule). Our goal is to elaborate a general multistrategy learning method that includes the capabilities of the known learning methods. This multistrategy learning method is neither a switcher between single-strategy learning methods, nor even an elaborated fixed integration of such strategies. On the contrary, we try to identify the basic reasoning mechanisms of the single strategy learning methods (at the level of single inference steps like deduction, analogy, abduction, generalization, specialization, abstraction, concretion, determination, mutual dependency, etc.) and to define a new learning method having all these reasoning capabilities. This learning method should be able to integrate dynamically the different types of elementary reasoning mechanisms, depending of the features of the current learning problem. Consequently, it should allow learning in situations in which previous (single-strategy or integrated) learning methods were insufficient. Moreover, if a particular learning task corresponds to a single-strategy method, then the behavior of the system should correspond to the application of such a method. Such an adaptive integration of different learning strategies (that seems also to be a characteristic of human learning) has been called multistrategy task-adaptive learning, or MTL for short (Michalski, 1990, 1991; Tecuci and Michalski, 1991a,b).

Building general MTL systems is a long term goal. In this paper, we propose an initial approach to this goal. First, in section 2, we present an architecture of a learning system at a level of abstraction suitable for MTL. The next section of the paper contains a general presentation of a proposed MTL learning problem and method. Then, in sections 4 and 5, we illustrate the proposed MTL method. We conclude the paper by analyzing the strength and the limitations of our approach to MTL, and by indicating what we consider to be the most promising directions of future research.

2. Toward a More Concrete Model of the Learning Systems

The main idea of our approach to building MTL systems is to define the architecture of a learning system at a level of abstraction that would allow a common view on the single-strategy learning methods, and would therefore facilitate their dynamic integration. The starting point will be the architecture in Figure 1, which will be made more concrete so as to allow the identification of the roles of different learning strategies.

First of all, let us notice that the KB may be regarded as a representation of the real world, that is, a world model. In general, it will contain representations of the real world objects, situations, laws, as well as representations of the learner's goals, capabilities and actions in the world. In order to solve a problem with such a system, the user has to represent the problem in the language of the world model. Then, the PE will look for a solution in its world model, and will show the solution to the user, who will interpret it in real world terms.

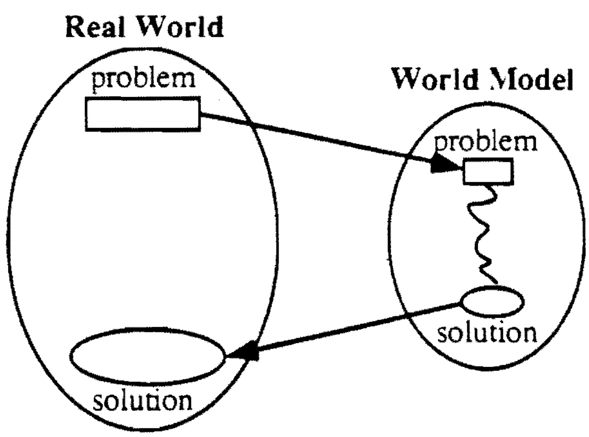

Figure 2: The use of the world model in solving real world problems.

From Figure 2 it is obvious that the behavior of the system depends of the model of the world (environment) in which it operates. The more adequately this model approximates the environment, the better is the system's behavior. Therefore, the goal of the LE is to continuously improve its world model so as to better represent the real world. It is interesting to notice that this general learning goal may be specialized to different specific goals, each corresponding to a specific single strategy learning method. For instance, empirical induction extends the world model with a generalization of specific facts. Learning by analogy extends the description of an entity with new features transferred from a similar entity. Learning by abduction adds new knowledge into the domain model, so as to explain a given piece of knowledge. Explanation-based learning improves the organization of the world model, by learning more efficient ways of solving problems, in the form of operational definitions of concepts.

We further consider that the main activity of the system is that of "understanding" the external world. Consequently, the abstract learning and problem solving scenario we are investigating is the following one:

The system has an incomplete and partially incorrect representation of the real world, and receives new input information about it.

The input information may consist of a fact that is true in the real world. It may consist of one or several positive (and/or negative) examples of a concept. It may even consist of one or several positive (and/or negative) examples of problem solving episodes, each episode specifying a problem and its solution.

The goal of the system is to extend, update and/or improve its KB (world model) so as to consistently integrate the information contained in the input. This means that, after learning from an input I, the KB should be such that a generalization of I is inferable from the KB.

The general learning strategy is based on understanding the input in terms of the current KB. This means that the LE will try to build a plausible justification tree that demonstrates that the input is a plausible consequence of the current KB. To build such a plausible proof tree, the system may need to use different types of inference, like deduction, analogy, abduction, empirical generalization, etc. Thus, it will hypothesize new pieces of knowledge (through analogy, abduction, empirical generalization and deduction, etc.). The hypothesized knowledge is supported by the fact that it is a plausible connection between a KB that represents parts of the real world, and a piece of knowledge (the input) that is known to be true in the real world. Moreover, aiming to learn as much as possible from any input it receives, the LE will try to generalize the understanding of the current input, as much as allowed by the knowledge used to build it in the first place. By this, it will generalize the hypothesized knowledge, so as the resulting KB will entail not only the received input, but also similar ones.

The result of this learning process is the increase of the problem solving abilities of the system (both in terms of competence and performance). Indeed, if the input was a new fact, then the PE will be able to predict that certain similar facts are true in the real world. Or, if the input was a set of examples (and/or counterexamples) of some concept, then the PE will be able to predict if a new instance is (or is not) an example of the learned concept. If the input was an example of a problem solving episode, represented by a problem P and its solution S, then the PE will be able to solve problems similar to P, by proposing solutions similar to S.

The next section contains a general presentation of the multistrategy task-adaptive learning problem and method.

3. General Presentation of the MTL Problem and Method

The learning problem we are considering is presented in Table 1.

Table 1: The multistrategy task-adaptive learning problem

Given

• *Incomplete and partially incorrect KB*

The KB may include a variety of knowledge types (facts, examples, implicative or causal relationships, hierarchies, etc.).

• *A (noise-free) input*

The input may consist of one or several examples, facts, or problem solving episodes.

Determine

• *An extended, updated, and/or improved KB*

The KB may be developed by adding new pieces of knowledge, modifying other pieces of knowledge, and/or reorganizing parts of the KB, so as to consistently integrate a generalization of the input.

As indicated in Table 1, the system may learn from one or from several examples. Figure 3 and Table 2 describe the learning method in the case of a single positive example.

Table 2: Multistrategy task-adaptive learning from one example

• *Understand the input*

1. Build a plausible justification tree T demonstrating that the input I is a plausible consequence of the knowledge from the KB. The inference steps in the tree T may be the result of different types of inference: deduction, analogy, abduction, inductive generalization, etc.
2. Introduce into the KB the knowledge pieces hypothesized (through analogy, abduction, inductive generalization and deduction, etc.) during the building of the plausible justification tree T.

• *Generalize the understanding*

3. Build an explanation structure ES, by replacing each inference step from T with the most general plausible generalization of it. The generalization of each inference step will depend of the type of inference, and of the knowledge used to derive it. It will correspond to the generalization of all the similar inference steps that the system would consider plausible.
4. Introduce into the KB the generalizations hypothesized during the building of the explanation structure ES.
5. Determine the most general unification of the general inference steps from ES. The obtained tree T_u is the most general plausible generalization of T. It shows how a generalization of the input is a plausible consequence of the extended KB of the system.
6. Extract different knowledge pieces from the tree T_u as, for instance, operational definitions of the concept illustrated by the input, abstract definitions of the concept illustrated by the input, or abstractions of the input.

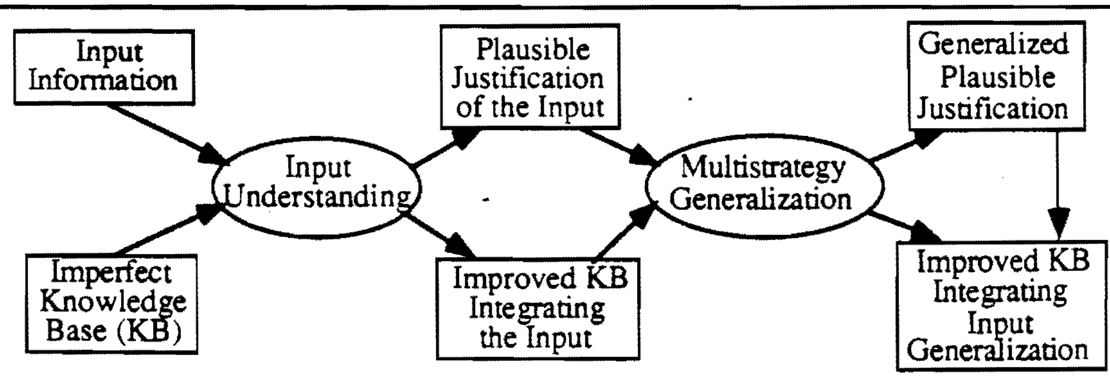

Figure 3: Learning from one positive example.

As will be shown in section 4.2.4, the learning method from Table 2 generalizes explanation-based learning, learning by abduction, and learning by analogy.

If the learning system receives additional examples illustrating a same concept, then learning continues as described in Figure 4 and in Table 3.

Table 3: Multistrategy task-adaptive learning from several examples

For each new example repeat steps 1 through 7

1. Let T_u be the generalized plausible justification tree of the previous positive examples that does not cover any justification of the previous negative examples. Determine the instance T_i of T_u corresponding to the current input I_i. If I_i is a positive example then go to step 2. Otherwise, if it is a negative example, go to step 5.

2. Analyze T_i. If all the inference steps in T_i are plausible then T_u already covers a plausible justification tree of I_i. This ends the processing of the current example. Otherwise, if some of the inferences from T_i are not valid, then make minimum modifications to T_i so as to become a plausible justification of the current positive example.

3. Introduce into the KB the knowledge pieces hypothesized (through analogy, abduction, inductive generalization and deduction, etc.) during the correction of the plausible justification tree Ti.

4. Generalizes T_u so as to cover the corrected justification tree T_i. This ends the processing of the current example.

5. Analyze T_i. If at least one inference steps in T_i is wrong, then T_u does not covers a plausible justification tree of I_i. This ends the processing of the current example. Otherwise, if all the inferences from T_i are plausible, then go to step 6.

6. T_i is a wrong justification tree that is covered by T_u. Indeed, T_i "proves" that the current input is a positive example, although it is known that it is a negative example. Hypothesize that R_{ki} is the wrong inference step in T_i, and correct the KB so as no longer to entail R_{ki}. The identification of R_{ki} is based on the following assumptions:

- there is only one wrong inference step in T_i;
- the inference step is wrong because it does not specify a necessary left hand side predicate;
- the hypothesis that R_{ki} is the wrong inference step requires minimum corrections to the KB.

7. Correct T_u so as no longer to cover T_i, and KB so as no longer to entail R_{ki}. This ends the processing of the current example.

8. All the input examples have been processed. Extract different knowledge pieces from T_u like, for instance, operational definitions of the concept illustrated by the input, abstract definitions of the concept illustrated by the input, or abstractions of different input examples.

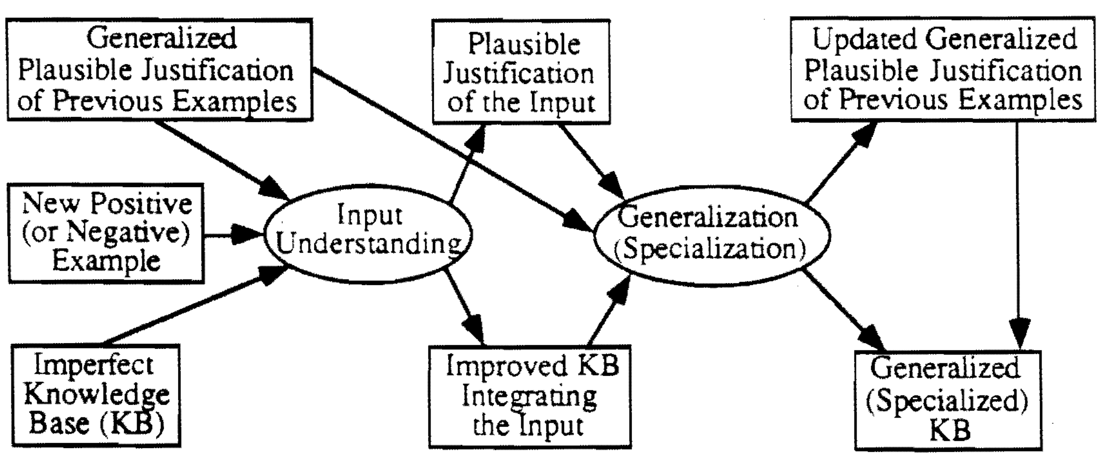

Figure 4: Learning from several positive and negative examples.

As will be shown in section 4.3.4, the learning method from Table 3 generalizes empirical induction, constructive induction, and mEBG.

The next sections illustrate the above learning methods, by using an intuitive example from the area of geography.

4. An Illustration of the Multistrategy Task-adaptive Learning

4.1 An exemplary problem

In order to illustrate the multistrategy task-adaptive learning method, we shall consider the case of a learning system in the area of geography. The purpose of the system is that of acquiring geographical data and rules in order to answer questions about geography. We suppose that the system disposes of an incomplete and partially incorrect KB, consisting of facts, examples, determinations (Davies and Russell, 1987), and deductive rules, as shown in Figure 5.

Facts:

terrain(Philippine, flat),
rainfall(Philippine, heavy),
water-in-soil(Philippine, high),
climate(Nepal, subtropical),
temperature(Nepal, warm),
terrain(Nepal, abrupt),
soil(Greece, red-soil),
soil(Greece, fertile-soil),
terrain(Greece, abrupt),
water-in-soil(Egypt, high)

Examples:

soil(Greece, red-soil)
: : > soil(Greece, fertile-soil)
terrain(Egypt, flat) & soil(Egypt, red-soil)
: : > soil(Egypt, fertile-soil)

Determination:

rainfall(x, y) >- water-in-soil(x, z)

Deductive rules:

∀ x, soil(x, loamy) → soil(x, fertile-soil)
∀ x, climate(x, subtropical)
→ temperature(x, warm)
∀ x, water-in-soil(x, high) &
temperature(x, warm) & soil(x, fertile-soil)
→ grows(x, rice)

Figure 5. A sample of an incomplete and partially incorrect KB.

In the next section will shall illustrate multistrategy task-adaptive learning from a single positive example. Then, we shall show how learning proceeds when the system receives new positive and negative examples.

4.2 Multistrategy task-adaptive learning from a positive example

Let us suppose that the system receives the following positive example of the relationship "grows(x, y)":

Positive Example 1:

rainfall(Vietnam, heavy) &
climate(Vietnam, subtropical) &
soil(Vietnam, red-soil) &
terrain(Vietnam, flat) &
location(Vietnam, SE-Asia)
: : > grows(Vietnam, rice)

Figure 6. A sample input for MTL.

4.2.1 Understanding the input

As shown in Figure 3 and Table 2, the system tries to "understand" the input by building a *justification tree* which demonstrates that the input is a plausible consequence of the KB. If the system cannot build such a tree, then the input represents entirely new knowledge that is stored explicitly into the KB.

The justification tree for the Example 1 in Figure 6 is shown in Figure 7. It demonstrates that the input is indeed an example of the relationship "grows(x, y)". One may notice that the individual inference steps are a result of different inference types. For example, two inference steps in Figure 7 are the results of deduction, as shown in Figures 8 and 9.

Another inference step is the result of the simple analogy illustrated in Figure 10. This form of analogy is based on the determination rule: "rainfall(x, y) >- water-in-soil(x, z)". According to this rule, the type of rainfall determines the quantity of water in soil. That is, if two entities x1 and x2 are characterized by the same type of rainfall y0, then it is very likely that they are also characterized by the same quantity z0 of water in soil. As can be seen in Figure 10, Philippine and Vietnam are similar from the point of view of "rainfall" (in both cases this is heavy). Therefore, one may

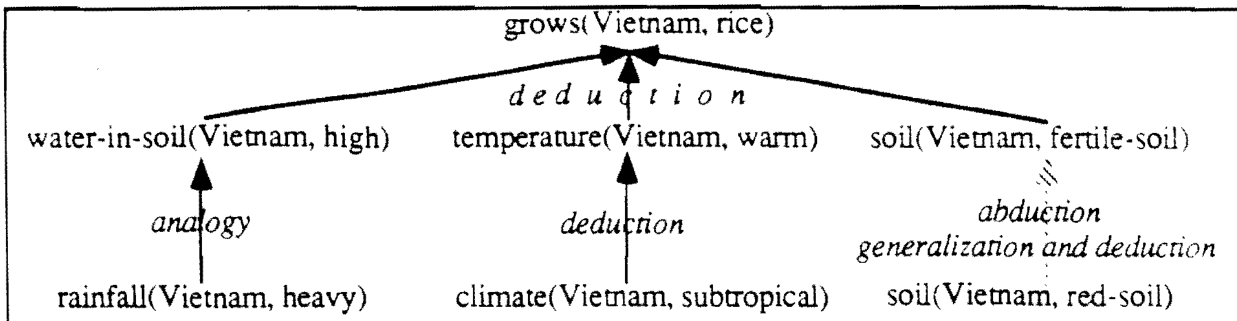

Figure 7. A plausible justification of Example 1.

hypothesize that these two countries are also similar from the point of view of "water-in-soil". Thus, the system concluded that "water-in-soil(Vietnam, high)" from the fact "water-in-soil(Philippine, high)". This form of analogy is described in (Davies and Russell, 1987). In general, the MTL method is intended to incorporate different forms of analogy, based on different kinds of similarities, such as similarities among causes, relations, and meta-relations (Carbonell, 1986; Gentner, 1983; Kedar-Cabelli, 1985; Kodratoff, 1990; Porter, Bareiss and Holte, 1990; Winston, 1986).

Deduction 1:

∀ x, water-in-soil(x, high) &
temperature(x, warm) &
soil(x, fertile-soil)
→ grows(x, rice)

water-in-soil(Vietnam, high) &
temperature(Vietnam, warm) &
soil(Vietnam, fertile-soil)

----

grows(Vietnam, rice)

Figure 8: Proving that rice grows in Vietnam.

Deduction 2:

∀ x, climate(x, subtropical)
→ temperature(x, warm)

climate(Vietnam, subtropical)

----

temperature(Vietnam, warm)

Figure 9: Proving that temperature of Vietnam is warm.

Another type of plausible reasoning is induction, which includes abduction, empirical generalization, and constructive induction (Michalski, 1990). Abduction consists in hypothesizing a fact or an inference step so as to build a plausible proof of the input. For instance, in order to demonstrate "grows(Vietnam, rice)", the system needed to prove that "soil(Vietnam, fertile-soil)". However, no deductive or analogical knowledge exists to infer this fact. Therefore, the system made the hypotheses that "soil(Vietnam, fertile-soil)" is a direct consequence of the input fact "soil(Vietnam, red-soil)", and abduced the inference step shown in Figure 11.

The inference step from Figure 11 resulted also from a combination of empirical generalization and deduction. Indeed, the system looked into the KB for examples in which the predicate "soil" is true. Then it inductively generalized them to a rule that was used (deductively) to produce the same inference step as abduction (see Figure 12).

Analogy:

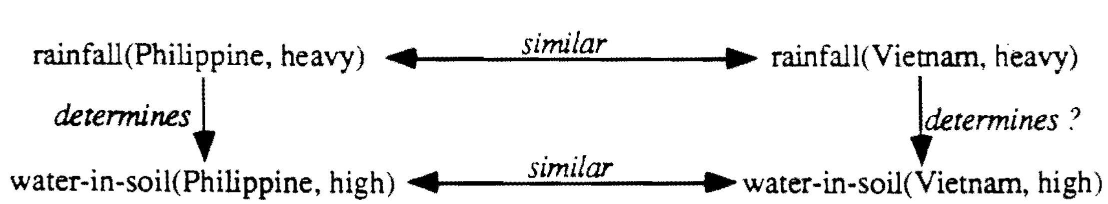

Figure 10: Hypothesizing Vietnam's water-in-soil through analogy.

This is because, when the system cannot justify a given predicate by deduction, it will try to justify it by several plausible reasoning methods that are supporting each other.

Abduction:
soil(Vietnam, red-soil) $\rightarrow$ soil(Vietnam, fertile-soil)
Figure 11: Abducing an inference step.

Examples:
soil(Greece, red-soil) : : > soil(Greece, fertile-soil)
terrain-type(Egypt, flat) & soil(Egypt, red-soil) : : > soil(Egypt, fertile-soil)
Inductive generalization:
$\forall$x, soil(x, red-soil) $\rightarrow$ soil(x, fertile-soil)
Deduced inference:
soil(Vietnam, red-soil) $\rightarrow$ soil(Vietnam, fertile-soil)
Figure 12: Making an inference through inductive generalization and deduction.

The system builds a plausible proof tree from top-down, by using the following strategy to control the different types of inference it can perform. First, it tries to justify a given predicate (for instance, "grows(Vietnam, rice)" in Figure 7) by deduction. If this attempt succeeds, then the justification of the given predicate is reduced to the justification of other predicates ("water-in-soil(Vietnam, high)", "temperature(Vietnam, warm)" and "soil(Vietnam, fertile-soil)"). However, if this attempt fails, then the system tries to justify the given predicate by using as many plausible reasoning methods as possible. It will try these methods in the following order: analogy, abduction, and inductive generalization followed by deduction. If one of them produces a plausible inference step, then the system tries to use the remaining ones in order to confirm or contradict it. If no contradiction is found, the inference step is accepted.

4.2.2 Generalization of the plausible justification tree

Once a justification tree was successfully created, the system analyzes the individual inference steps to determine if they could be locally generalized within the constraints of the KB that were used to make these steps. After the inference steps are generalized locally, the system unifies them globally, and builds a generalized justification tree. This technique is an extension of the one elaborated by (Mooney and Bennet, 1986). The extension concerns the way individual inference steps are (deductively or inductively) generalized, by using the knowledge from which they were derived. The idea is to replace each inference step S with the generalization of all the similar inference steps that could be derived from the knowledge that produced S (Tecuci and Michalski, 1991b).

The deductive steps are replaced by the deductive rules that generated them.

The analogy steps are generalized by considering the knowledge used to derive them. In our example, the inference step
rainfall(Vietnam, heavy)
$\rightarrow$ water-in-soil(Vietnam, high)
was obtained by analogy with "rainfall(Philippine, heavy)" and "water-in-soil(Philippine, high)", based on the determination
rainfall(x, y) >- water-in-soil(x, z).
Because the system would infer "water-in-soil(x, high)" for any x such that "rainfall(x, heavy)", the analogical inference is generalized to: rainfall(x, heavy) $\rightarrow$ water-in-soil(x, high)

The generalization of the inductive steps depends on the type of induction performed. In the case of the abduction step illustrated in Figure 11, there is no domain knowledge that could be used to generalize it. Therefore, if this step would have been obtained only through this type of abduction, it would not be generalized (abductions based on tracing backward causal relationships can be generalized (Tecuci and Michalski, 1991b)). However, this inference step was also a result of applying a rule generalized from examples (see Figure 12), and is therefore replaced with this rule. In general, when an inference step is a result of different types of inference, all the involved knowledge is used to generalize it.

The generalization of the inference steps from Figure 7 form the explanation structure shown in Figure 13. The most general unification of

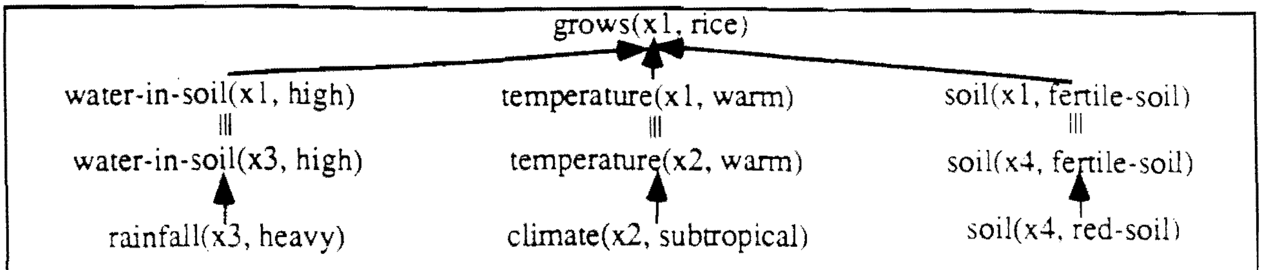

Figure 13: Explanation structure.

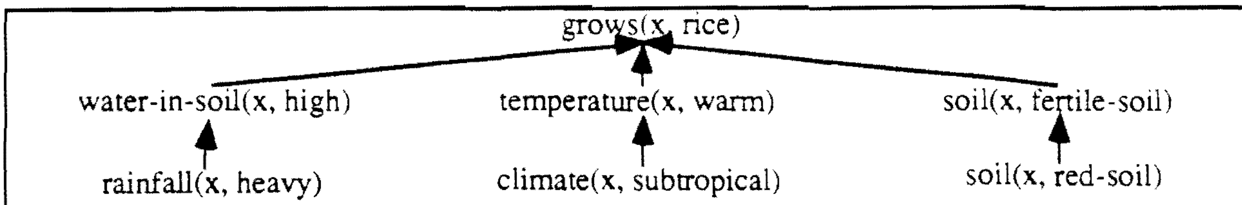

Figure 14: A generalization of the plausible justification tree from Figure 7.

the connection patterns in Figure 13 is shown in Figure 14, and represents a generalization of the plausible justification from Figure 7.

4.2.3 The learned knowledge

As indicated in Figure 3, the system may improve the KB by learning different types of knowledge. For instance, during the understanding of the input, the system learned a new fact (water-in-soil(Vietnam,high)) and a rule ($\forall$x,soil(x,red-soil) $\rightarrow$ soil(x,fertile-soil)). Also, it learned different pieces of knowledge from the generalized justification tree in Figure 14 (Michalski, 1990). For instance, it stored the leaves of this tree as an *operational definition* of the relationship "grows(x, rice)":

rainfall(x, heavy) &
climate(x, subtropical) & soil(x, red-soil)
: : > grows(x, rice)

It also stored the upper part of the generalized tree as the *abstract definition* of the relationship "grows(x, rice)":

water-in-soil(x, high) &
temperature(x, warm) & soil(x, fertile-soil)
: : > grows(x, rice)

Since this rule was already known, the new knowledge is just that it represents an abstract definition of the relationship "grows(x, rice)". The system has also learned an *abstraction* of Example 1:

water-in-soil(Vietnam, high) &
temperature(Vietnam, warm) &
soil(Vietnam, fertile-soil)
: : > grows(Vietnam, rice)

4.2.4 Basic cases

The presented MTL method is a generalization of explanation-based learning, learning by abduction, and learning by analogy. Indeed, it behaves as such a method whenever the relationship between the input and the KB satisfies the applicability conditions of it, and the general learning goal of MTL is specialized to the goal of the single-strategy method, as illustrated in the following sections.

4.2.4.1 Explanation-based learning

Let us suppose that, in addition to the rules in Figure 5, the KB also contains the following deductive rules:

$\forall$x, rainfall(x, heavy) $\rightarrow$ water-in-soil(x, high)
$\forall$x, soil(x, red-soil) $\rightarrow$ soil(x, fertile-soil)

In such a case, the justification trees in Figures 7 and 14 become logical proofs, and the result of the learning process is an operational definition of the relationship "grows(x, rice)". Thus, the MTL method reduces to explanation-based learning (DeJong and Mooney, 1986; Mitchell, Keller and Kedar-Cabelli, 1986).

4.2.4.2 Learning by abduction

Let us now suppose that the relationship between "rainfall" and "water-in-soil" is not a determination, but an implication

$\forall$x, rainfall(x,heavy) $\rightarrow$ water-in-soil(x,high)

and the KB does not contain examples of the predicate "soil". In this case, in order to build the justification tree of the input, the system

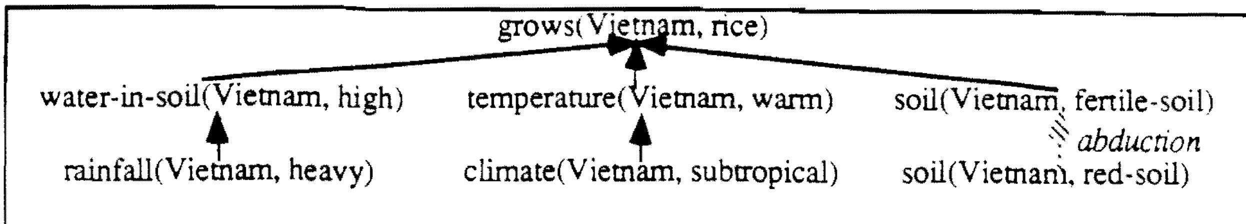

Figure 15: Using abduction to build a justification tree of the input.

needs only to create the explanatory hypothesis
soil(Vietnam, red-soil)
→ soil(Vietnam, fertile-soil)
as shown in Figure 15. Therefore, the result of learning is the created explanatory hypothesis, and the MTL method reduces to abductive learning.

4.2.4.3 Learning by analogy

Let us finally suppose that the only background knowledge that is related to Example 1 is:

Facts:
rainfall(Philippine, heavy),
water-in-soil(Philippine, high)
Determination:
rainfall(x, y) >- water-in-soil(x, z)

In this case, the system can only infer that "water-in-soil(Vietnam, high)", by analogy with "water-in-soil(Philippine, high)", as shown in Figure 10. Thus, the MTL method reduces to analogical learning.

4.3 Multistrategy task-adaptive learning from several examples

The basic idea of learning from several examples is to generalize and/or specialize the plausible justification tree in Figure 14, so as to cover explanations of the new positive examples, and not to cover explanations of the negative examples. In the same time, the MTL system will generalize and/or specialize some knowledge pieces from the KB.

4.3.1 Learning from new positive examples

Let us first suppose that the system receives a new positive example of the relationship "grows(x, y)" (see Figure 16).

First of all, the system tries to understand the current example, by using the understanding of the previous examples. To this purpose, it determines the instance of the general tree in Figure 14, corresponding to the example in Figure 16 (see Figure 17).

Then the system analyzes the inference steps from this proof tree. If all the inference steps are plausible, then the tree in Figure 17 is a plausible justification of the new positive example that is already covered by the general justification tree in Figure 14. This ends the processing of the current example.

Positive Example 2:
rainfall(Madagascar, heavy) &
soil(Madagascar, loamy) &
climate(Madagascar, subtropical) &
terrain(Madagascar, flat) &
in(Madagascar, Pacific-Ocean)
: : > grows(Madagascar, rice)

Figure 16: A new positive example of the relationship "grows(x, y)".

However, the tree in Figure 17 is not a correct justification of Example 2 because the leaf predicate "soil(Madagascar, red-soil)" is not true. Therefore, the system uses the deductive rule "∀x, soil(x,loamy) → soil(x,fertile-soil)", from Figure 5, and builds the plausible justification tree in Figure 18.

Next, the system builds the explanation structure in Figure 19 that has three general components to be unified: the part of the tree in Figures 14 that covers part of the tree in Figure 18, the part that is specific to the tree in Figure 14, and the generalization of the part of the tree in Figure 18 that is specific to it (this generalization being done according to the procedures described in section 4.2.2).

As the result of the unification of the connection patterns in Figure 19, one obtains the general justification tree in Figure 20 that covers the justification trees of both Example 1

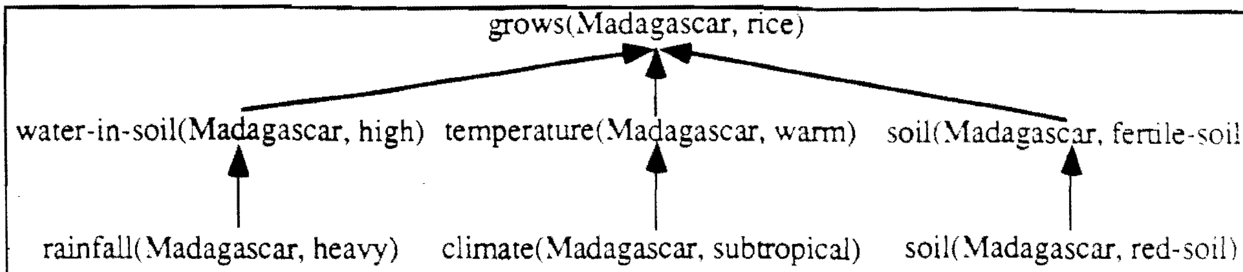

Figure 17: Instance of the general justification tree corresponding to Example 2.

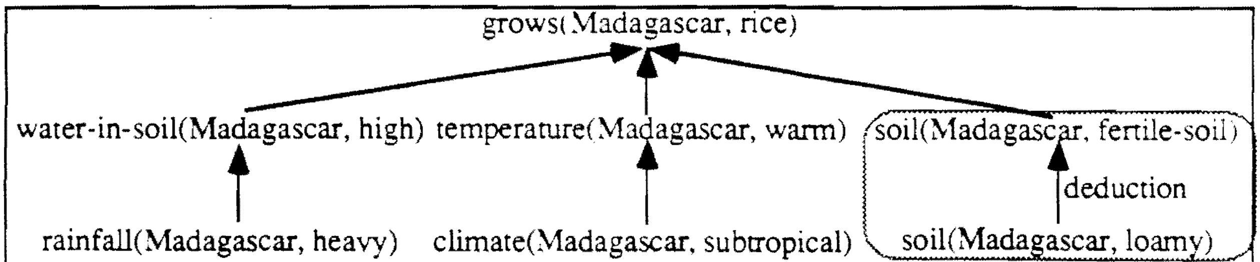

Figure 18: Plausible justification of Example 2 in Figure 16.

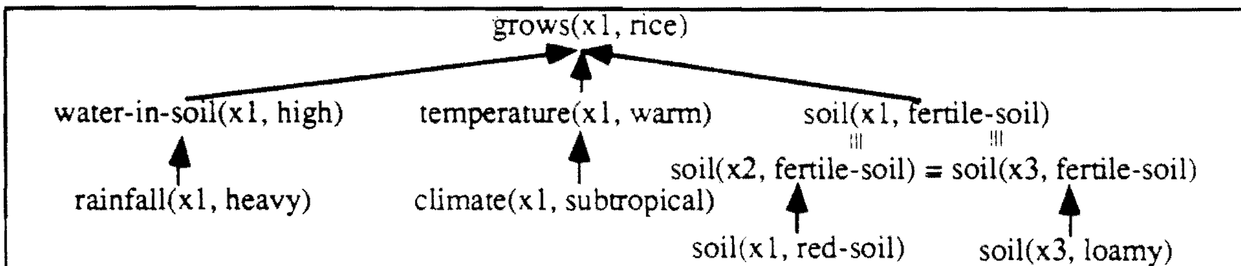

Figure 19: Explanation structure covering the plausible justification in Figure 18.

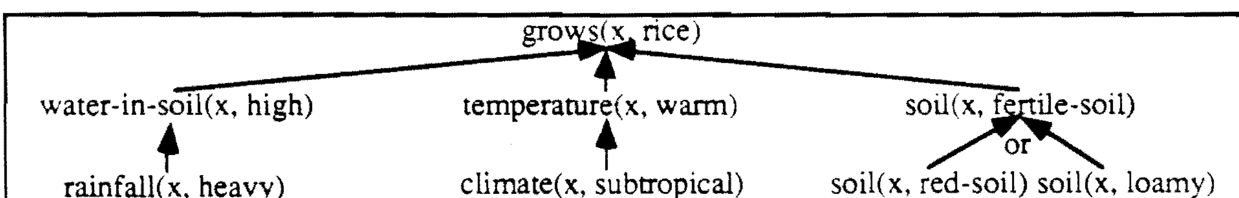

Figure 20: General justification tree covering the justifications of Example 1 and Example 2.

and Example 2. This is an AND/OR tree. Indeed, "grows(x, rice)" is an AND node and "soil(x, fertile-soil)" is an OR node.

4.3.2 Learning from negative examples

Let us now consider that the MTL system receives the negative example from Figure 21. Again the system builds the instance of the general justification tree in Figure 20, corresponding to this new example (see Figure 22). This tree would lead to the wrong conclusion that the current input is a positive example of the relationship "grows(x, y)".

Negative Example 3:

rainfall(Nepal, heavy) &
climate(Nepal, subtropical) &
soil(Nepal, loamy) &
Location(Nepal, Central-Asia) &
terrain(Nepal, abrupt)
: : > ¬ grows(Nepal, rice)

Figure 21: A negative example of grows(x, y).

Therefore, the tree must contain some wrong inferences. These wrong inferences have to be detected, and both the general justification tree and the KB should be specialized, so as no longer to contain them. This is an instance of the well known credit assignment problem (Minski, 1963). In the current version of MTL, the solution to this problem is based on the following simplifying assumptions:

- there is only one wrong inference step in the tree from Figure 22;
- the inference step is wrong because it does not specify a necessary left hand side predicate that has to be true in order for the right hand side predicate to be true.

If the system finds more than one inference step that could be wrong, it selects one of them by using, in order, the following criteria:

- select the weakest inferences (first abduction, then analogy, and lastly deduction);
- among the remaining ones select those for which the correction of the KB and of the general justification tree causes the minimum loss of coverage of previous examples;
- among the remaining ones, select those for which the corrections produce a minimum increase in the complexity of the modified knowledge pieces;
- choose arbitrarily one of the remaining hypotheses.

Let us illustrate the solution to the credit assignment problem for the current negative example. First, the MTL system will analyze all the previous examples, looking for a feature distinguishing between the positive examples and the negative ones. Such a feature is "terrain(x, flat)" that is true for the two positive examples seen, and false for the negative example.

Secondly, the system makes the hypothesis that "terrain(x, flat)" is an additional condition for inferring one of the non-leaf predicates from the proof tree in Figure 20: "water-in-soil(x, high)", "temperature(x, warm)", "soil(x, fertile-soil)", and "grows(x, rice)".

Thirdly, it tests each of these hypotheses against the facts from the KB. Let us consider, for instance, the predicate "temperature(x, warm)". The system will hypothesize that "terrain(x, flat)" is a necessary condition for inferring it only if the warm temperature is always associated with flat terrain, when the climate is subtropical (the other condition for warm temperature). However, because it is known that "temperature(Nepal, warm)", "terrain(Nepal, abrupt)", and "climate(Nepal, subtropical)" (see Figure 5) the system rejects this hypothesis. Similarly, the system rejects the hypothesis that "terrain(x, flat)" is a necessary condition for inferring "soil(x, fertile-soil)". Because "water-in-soil(x, high)" and "grows(x, rice)" are never associated with "terrain(x, non-flat)" these are the remaining hypotheses. "water-in-soil(x, high)" was the result of an analogical inference, and "grows(x, rice)" was the result of a deduction. Therefore, the system hypothesizes that the wrong inference step is

rainfall(x, heavy) $\rightarrow$ water-in-soil(x, high)

and corrects it to

rainfall(x, heavy) & terrain(x, flat) $\rightarrow$ water-in-soil(x, high)

Thus, the general justification tree in Figure 20 is specialized as indicated in Figure 23.

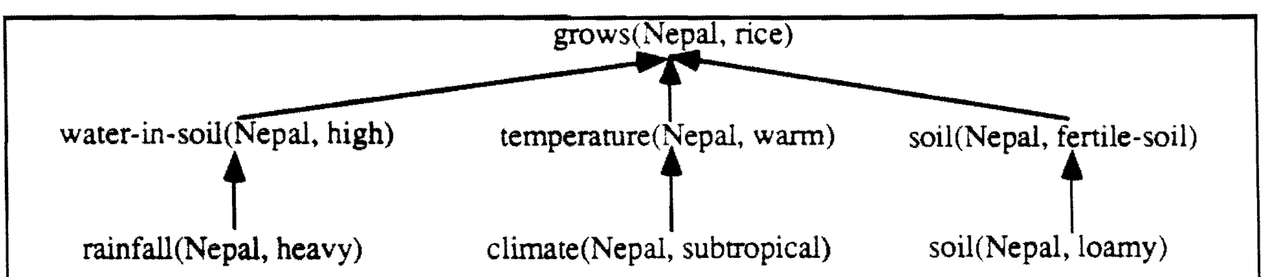

Figure 22: The instance of the justification tree in Figure 20 which is a wrong proof tree.

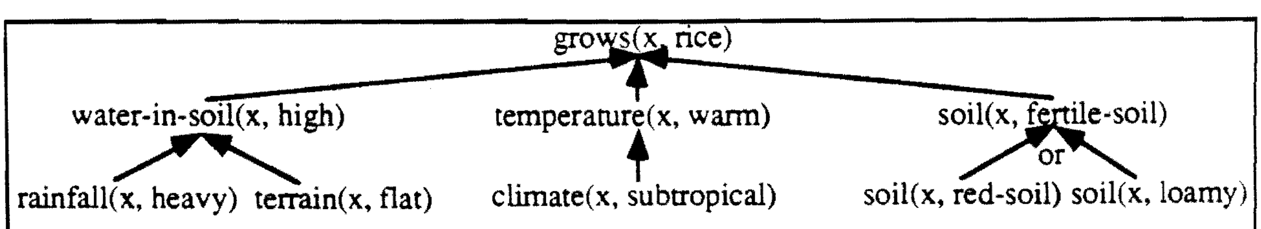

Figure 23: Specialization of the general justification tree.

After correcting the wrong inference step from the general justification tree, the system attempts to correct also the piece of knowledge from the KB that entailed this wrong inference step. The wrong inference step

rainfall(x, heavy) → water-in-soil(x, high)

was generated from the determination

rainfall(x, y) >- water-in-soil(x, z)

Consequently, the system applies the same corrective specialization to this determination that becomes:

rainfall(x, y) & terrain(x, flat)
>-water-in-soil(x, z)

4.3.3 The learned knowledge

As in the case of learning from one example, the system learned different kinds of knowledge that extend, correct and/or improve the KB. For instance, it improved the determination

rainfall(x, y) >- water-in-soil(x, z)

by adding a necessary left hand side predicate:

rainfall(x, y) & terrain(x, flat)
>- water-in-soil(x, z)

Also, it improved the operational definition of "grows(x, y)", by both specializing and generalizing it:

rainfall(x, heavy) & terrain(x, flat) &
climate(x, subtropical) &
(soil(x, red-soil) or soil(x, loamy))
: : > grows(x, rice)

This illustrates the way the imperfections of the knowledge pieces are interpreted by the presented method: an imperfect knowledge piece is one having parts that are too general and/or parts that are too specific. In (Tecuci, 1988) it is shown that this is a general way of approaching the imperfect theory problem (Mitchell, Keller and Kedar-Cabelli, 1986; Rajamoney and DeJong, 1987).

4.3.4 Basic cases

4.3.4.1 Empirical and constructive induction

Let us assume that the KB does not contain the determination and the deductive rules shown in Figure 5. In this case, each input is new, neither confirming nor contradicting the KB. Therefore, each example is interpreted as representing a single inference step, as shown in top part of Figure 24.

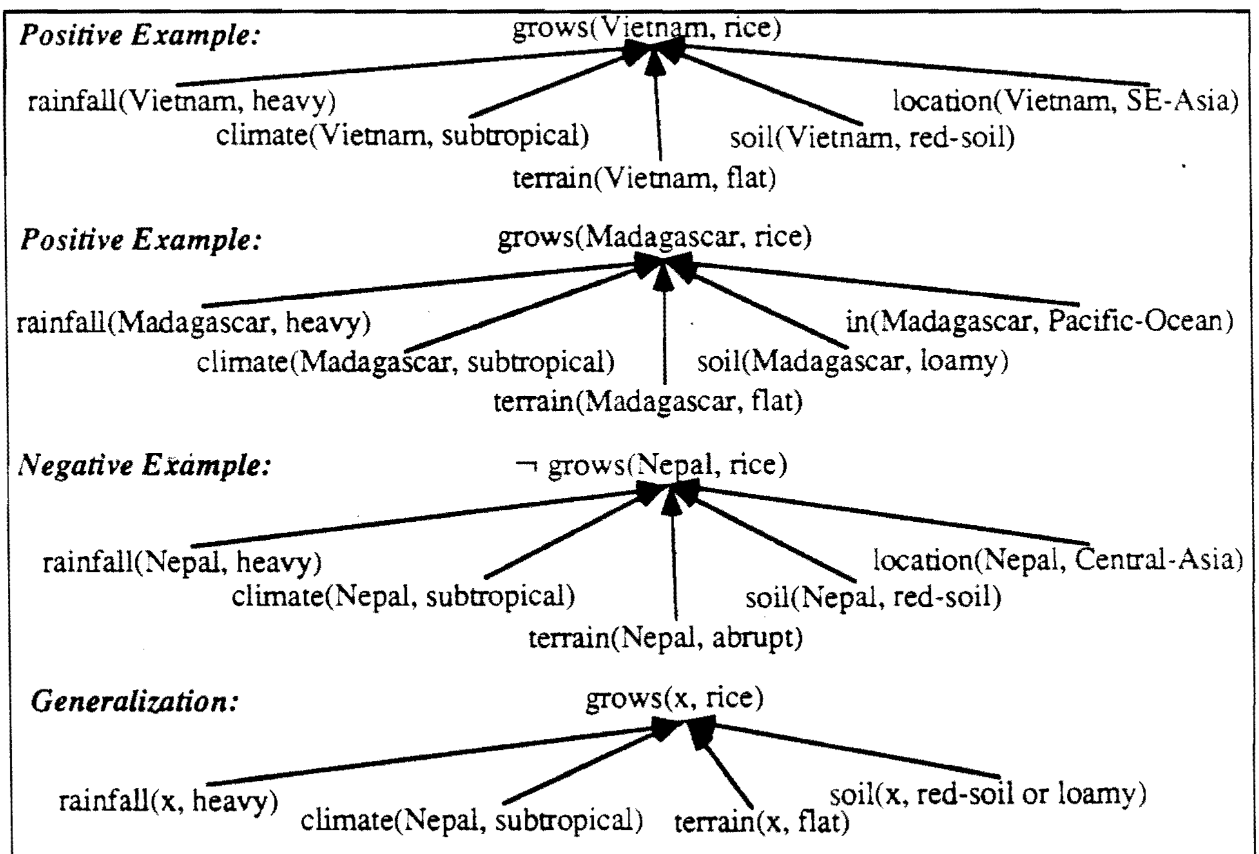

Figure 24. Empirical or constructive induction from examples.

The MTL method will compute a generalization of the trees corresponding to the positive examples, generalization that does not cover the trees corresponding to the negative examples (see the bottom of Figure 24). The result of learning is therefore a definition of the relationship "grows(x, rice)", that represents the common properties of the positive examples that are not properties of the negative examples. Thus, in this case, the MTL method behaves as empirical induction or as constructive induction (depending on the generalization process).

4.3.4.2 Multiple-example explanation-based generalization

If the input of the system consists only of positive examples, that are deductively entailed by the KB, then the presented MTL method behaves as the multiple example explanation-based generalization, or mEBG, that was developed, among others, by (Hirsh, 1989; Kedar-Cabelli, 1985; Pazzani, 1988).

5 Discussion and Conclusion

In this paper we have proposed an abstract model of a learning system that is improving its representation of the real world, so as to consistently integrate new input information (that represents true facts about the real world, or appropriate actions to be taken in the real world). This model allowed the elaboration of a general multistrategy task-adaptive learning method, that integrates dynamically a whole range of learning strategies, depending of the features of the current learning situation. The method is based on the idea of understanding the input through an exploration of the KB, and an employment of different inference types. One of the major advantages of this method is that it enables the system to learn in situations in which single-strategy learning methods, or even previous integrated learning methods were insufficient. Therefore, the proposed approach has a great potential to make machine learning programs applicable to a wider range of problems. Another important aspect of the method is that it behaves as a single-strategy method, whenever the applicability conditions for such a method are satisfied. In this respect, the proposed MTL method may be regarded as a generalization of the single-strategy methods.

Obviously, humans learn through a kind of multistrategy method. Although we do not claim that the presented method is a model of human learning, we would like to mention, however, that some of the aspects of the method, do model human learning. For instance, research in experimental psychology (Wisniewski and Medin, 1991) provides evidence that humans build explanations of new examples by using the explanations of the previous examples (as it is also done in the presented MTL method). Also, the method for controlling the application of different inference types (during the understanding of the input), although quite simple in the current version (and definitely a necessary topic of future research) is related to that employed by humans (Collins and Michalski, 1989). Indeed, Collins and Michalski argue that people solve problems by pursuing different "lines of reasoning". They estimate the "strength" of each line of reasoning, and make their conclusion on the basis of this evaluation. If the lines lead to the same conclusion, they have a strong belief in the result. If the lines lead to different conclusions, and the associated "strengths" are roughly similar, people restrain from making any decisive conclusion.

There are also many ways in which the presented MTL method may be improved. One is to integrate additional learning strategies (like, for instance, conceptual clustering and reinforcement learning), and to develop the included learning strategies (especially learning by analogy and learning by abduction, which are quite simple in the current version). This will also require further elaboration of the generalization techniques specific to each involved strategy. The method may also be extended with respect to the input. One extension is to learn also from other types of input (like general pieces of knowledge, or input already known). Another extension is to accept incomplete or even noisy input. If the input may contain errors, then the system has to decide if the cause of a contradiction between the input and the KB is the incorrectness of the input, or that of the KB, or both, and to take appropriate corrective actions. This, of course, complicates significantly the learning process. Another important research direction regards the extension and the application of the MTL method to the problem of knowledge acquisition from a human expert. In this case, the method would be extended with an important interactive component, that would allow the system to ask different questions to the human expert, in order to decide on the best learning actions to take (Tecuci, 1991a,b).

Acknowledgments

I have benefited from many useful discussions with Ryszard Michalski that helped me extend and better describe this work.
This research was done in the GMU Center for Artificial Intelligence. Research of the Center is supported in part by the National Science Foundation Grant No. IRI-9020266, in part by the Office of Naval Research grants No. N00014-88-K-0397, No. N00014-88-K-0226 and No. N00014-91-J-1351, and in part by the Defense Advanced Research Projects Agency grants No. N00014-87-K-0874 and No. N00014-91-J-1854, administered by the Office of Naval Research.

References

Bergadano F., and Giordana A., Guiding Induction with Domain Theories, in Y. Kodratoff and R.S. Michalski (eds), *Machine Learning: An Artificial Intelligence Approach*, vol. 3, Morgan Kaufmann, 1990.

Carbonell J., Derivational Analogy: A Theory of Reconstructive Problem Solving and Expertise Acquisition, in R.S. Michalski, J.G. Carbonell, and T.M. Mitchell (eds) *Machine Learning: An Artificial Intelligence Approach*, vol. 2, pp.371-392, Morgan Kaufmann 1986.

Collins, A., and Michalski, R.S., The Logic of Plausible Reasoning: A Core Theory, *Cognitive Science*, Vol. 13, No. 1, pp.1-49, 1989.

Danyluk, A.P., The Use of Explanations for Similarity-Based Learning, *Proceedings of IJCAI-87*, pp. 274-276, Milan, Italy, 1987.

Davies T.R. and Russell S.J., A Logical Approach to Reasoning by Analogy, *Proceedings of IJCAI-87*, pp. 264-270, Milan, Italy, 1987.

DeJong, G. and Mooney, R., Explanation-Based Learning: An Alternative View, *Machine Learning Journal*, vol. 2, 1986.

Flann, N., and Dietterich, T., A Study of Explanation-based Methods for Inductive Learning, *Machine Learning* 4, 1989.

Genest, J., Matwin, S., and Plante, B., Explanation-based Learning with Incomplete Theories: A Three-step Approach, in B. Porter and R. Mooney (eds), *Machine Learning: Proc. of the Eighth International Workshop*, Texas, Morgan Kaufmann, 1990.

Gentner D, Structure-Mapping: a Theoretical Framework for Analogy, *Cognitive Science*, 7, pp.155-170, 1983.

Hirsh, H., Incremental Version-space Merging: A General Framework for Concept Learning, *Doctoral dissertation*, Stanford University, 1989.

Kedar-Cabelli S., Purpose-Directed Analogy, in *Proceedings of the Cognitive Science Society*, pp.150-159, Irvine, California, 1985.

Kodratoff, Y., Using Abductive Recovery of Failed Proofs for Problem Solving by Analogy, in B. Porter and R. Mooney (eds), *Machine Learning: Proc. of the Eighth International Workshop*, Texas, Morgan Kaufmann, 1990.

Kodratoff, Y., and Michalski R.S., (eds), *Machine Learning: An Artificial Intelligence Approach*, vol. 3, Morgan-Kaufmann 1990.

Lebowitz, M., Integrated Learning: Controlling Explanation, *Cognitive Science*, vol. 10, no. 2, pp. 219-240, 1986.

Michalski, R. S., A Theory and Methodology of Inductive Learning, *Machine Learning: An Artificial Intelligence Approach*, in R.S. Michalski, J.G. Carbonell, and T.M. Mitchell (eds.), Tioga Publishing Co., 1983.

Michalski R.S., Carbonell J.G., Mitchell T.M. (eds), *Machine Learning: An Artificial Intelligence Approach*, vol.1, Tioga Publishing Co., 1983.

Michalski R.S., Carbonell J.G., Mitchell T.M. (eds), *Machine Learning: An Artificial Intelligence Approach*, vol. 2, Morgan-Kaufmann 1986.

Michalski, R.S., Toward a Unified Theory of Learning: Multistrategy Task-Adaptive Learning, in *Reports of Machine Learning and Inference Laboratory*, MLI 90-1, Center for Artificial Intelligence, George Mason University, 1990.

Michalski, R.S., Inferential Learning Theory as a Basis for Multistrategy Task-Adaptive Learning, in this volume, 1991.

Minski M., Steps Toward Artificial Intelligence, in E.A. Feigenbaum and J. Feldman (eds), *Computers and Thought*, New York, McGraw-Hill, 1963.

Minton, S., Carbonell, J.G., Strategies for Learning Search Control Rules: An Explanation-based Approach, *Proc. IJCAI-87*, Milan, Italy, 1987.

Mitchell T.M., Version Spaces: An Approach to Concept Learning, *Doctoral dissertation*, Stanford University, 1978.

Mitchell, T.M., Keller, T., and Kedar-Cabelli, S., Explanation-Based Generalization: A Unifying View, *Machine Learning Journal*, vol. 1, 1986.

Mooney R., Bennet S., A Domain Independent Explanation Based Generalizer, in *Proceedings of AAAI-86*, pp. 551-555, Philadelphia, 1986.

Mooney R. and Ourston D., A Multistrategy Approach to Theory Refinement, in this volume, 1991.

Morik K., Balanced Cooperative Modeling, in this volum, 1991.

Pazzani, M.J., Integrating Explanation-Based and Empirical Learning Methods in OCCAM, *Proceedings of EWSL-88*, pp. 147-166, Glasgow, Scotland, 1988.

Porter B., Bareiss R., and Holte R., Concept Learning and Heuristic Classification in Weak-Theory Domains, *Artificial Intelligence*, vol. 45, no. 1, 1990.

De Raedt L. and Bruynooghe M., CLINT: A Multistrategy Interactive Concept Learner and Theory Revision System, in this volume, 1991

Rajamoney S., and DeJong G., The Classification, Detection and Handling of Imperfect Theory Problems, *Proc. IJCAI-87*, pp. 205-207, Milan, 1987.

Shavlik J.W., & Towell G.G., An Approach to Combining Explanation-based and Neural Learning Algorithms, *Connection Science*, vol.1, no.3, 1989.

Tecuci G., DISCIPLE: A Theory, Methodology, and System for Learning Expert Knowledge, *Ph.D. Thesis*, University of Paris-South, 1988.

Tecuci G. and Kodratoff Y., Apprenticeship Learning in Imperfect Theory Domains, in Y. Kodratoff and R.S. Michalski (eds), *Machine Learning: An Artificial Intelligence Approach*, vol. 3, Morgan Kaufmann, 1990.

Tecuci G., A Multistrategy Learning Approach to Domain Modeling and Knowledge Acquisition, in Y. Kodratoff (ed), *Machine Learning - EWSL-91*, Springer-Verlag, 1991a.

Tecuci G., Steps Toward Automating Knowledge Acquisition for Expert Systems, in *Proceedings of the AAAI-91 Workshop on Knowledge Acquisition: From Science to Technology to Tools*, Anaheim, California, 1991b.

Tecuci G., and Michalski R.S., A Method for Multistrategy Task-adaptive Learning Based on Plausible Justifications, in L. Birnbaum, and G. Collins (eds), *Machine Learning: Proceedings of the Eighth International Workshop*, Chicago, June 1991, Morgan Kaufmann, 1991a.

Tecuci G., and Michalski R.S., Input Understanding as a Basis for Multistrategy Task-adaptive Learning, in *Proceedings of the International Symposium on Methodologies for Intelligent Systems*, Charlotte, North-Carolina, October 1991, *Lecture Notes in Artificial Intelligence*, Springer-Verlag, 1991b.

Whitehall, B.L., Knowledge-based Learning: Integration of Deductive and Inductive Leaning for Knowledge Base Completion, *PhD Thesis*, University of Illinois at Champaign-Urbana, 1990.

Wilkins, D.C., Knowledge Base Refinement as Improving an Incorrect and Incomplete Domain Theory, in Y. Kodratoff and R.S. Michalski (eds), *Machine Learning: An Artificial Intelligence Approach*, vol. 3, Morgan Kaufmann, 1990.

Winston P.H., Learning by Augmenting Rules and Accumulating Censors, in R.S. Michalski, J.G. Carbonell, and T.M. Mitchell(eds), *Machine Learning: An Artificial Intelligence Approach*, vol. 2, Morgan-Kaufmann 1986.

Wisniewski E.J., and Medin D.J., Is it a Pocket or a Purse ? Tightly Coupled Theory and Data Driven Learning, in L. Birnbaum and G. Collins (eds), *Machine Learning: Proceedings of the Eighth International Workshop (ML91)*, Morgan Kaufmann, 1991.