

Chapter in the book
Machine Learning: An Artificial Intelligence Approach, vol. 3,
Y.Kodratoff and R.Michalski (eds), Morgan Kaufmann, 1990.

Apprenticeship Learning in Imperfect Domain Theories

Gheorghe Tecuci
Research Institute for
Computers and Informatics
71316, Bd. Miciurin 8-10
Bucharest 1, Romania

Yves Kodratoff*
George Mason University
AI Center
4400 Univ. Drive
Fairfax VA 22030-4444

Abstract

This chapter presents DISCIPLÉ, a multi-strategy integrated learning system illustrating a theory and a methodology for learning expert knowledge in the context of an imperfect domain theory. DISCIPLÉ integrates a learning system and an empty expert system, both using the same knowledge base. It is initially provided with an imperfect (nonhomogeneous) domain theory and learns problem solving rules from the problem solving steps received from its expert user, during interactive problem solving sessions. In this way, DISCIPLÉ evolves from a helpful assistant in problem solving to a genuine expert. The problem solving method of DISCIPLÉ combines problem reduction, problem solving by constraints, and problem solving by analogy. The learning method of DISCIPLÉ depends of its knowledge about the problem solving step (the example) from which it learns. In the context of a complete theory about the example, DISCIPLÉ uses explanation-based learning to improve its performance. In the context of a weak theory about the example, it synergistically combines explanation-based learning, learning by analogy, empirical learning, and learning by questioning the user, developing its competence. In the context of an incomplete theory about the example, DISCIPLÉ learns by combining the above mentioned methods, improving both its competence and performance.

* On leave from CNRS, Univ. Paris-Sud, LRI, Bat. 490, F-91405 Orsay France

1. INTRODUCTION

The present success of Artificial Intelligence is mostly due to the knowledge-based systems which proved to be useful almost anywhere. As the name suggests, the power of a knowledge-based system comes from its knowledge. However, building a knowledge base for such a system is a very complex, time-consuming, and error-prone process. Moreover, the resulting system lacks or has only poor abilities to update its knowledge or to acquire new knowledge.

One promising solution to this "knowledge acquisition bottleneck" is represented by the Learning Apprentice Systems (LAS). A LAS is an interactive knowledge-based consultant which is provided with an initial domain theory and is able to assimilate new problem-solving knowledge by observing and analyzing the problem solving steps contributed by its users, through their normal use of the system (Mitchell, Mahadevan & Steinberg, 1985).

Representative examples of this approach are the systems LEAP (Mitchell, Mahadevan & Steinberg, 1985) and GENESIS (DeJong & Mooney, 1986). The domain of expertise of LEAP is the VLSI design and that of GENESIS is story understanding. A common feature of LEAP and GENESIS is that they are based on a strong (complete) domain theory which allows them to learn a general rule or schemata from a single example, by reducing learning to deductive reasoning.

Nevertheless, such beautifully tailored domains are seldom available. A *typical real world domain theory is nonhomogeneous* in that it provides complete descriptions of some parts of the domain, and only incomplete or even poor (weak) descriptions of other parts of the domain. A learning episode, however, uses only one part of the domain theory, and this part may have the features of a complete, incomplete or weak theory, even if, globally, the theory is nonhomogeneous. Therefore, a learning system should be able to learn a general rule or concept not only when disposing of a complete theory about an example, but also when disposing of an incomplete or even weak theory about it. An illustration of such a learning system is DISCIPLINE. *DISCIPLINE is a multi-strategy integrated learning system*. It has the same general purpose as a learning apprentice system, but uses a multi-strategy approach to learning, instead of being based on deductive reasoning.

DISCIPLINE is a tool for building practical Expert Systems. It integrates an empty expert system and a learning system, both using the same knowledge base. To build an Expert System with DISCIPLINE, one has to first introduce into DISCIPLINE's knowledge base, elementary knowledge about an application domain, knowledge constituting a nonhomogeneous theory of the domain. Next, DISCIPLINE may be used to interactively solve problems, according to the following scenario:

The user gives DISCIPLINE the problem to solve and the expert subsystem starts solving this problem by showing the user each problem solving step (which we shall call partial solution). The user may agree or reject it. In the latter case, or when DISCIPLINE is unable to propose any partial solution, the user is compelled to give his own solution. Once this solution is given, a learning process will take place. DISCIPLINE will try to learn a general rule so that, when faced with problems similar with the current one (which it has been unable to solve), it will become able to propose a solution similar to the solution given by the user to the current problem. In this way, DISCIPLINE progressively evolves from a useful assistant in problem solving to a genuine expert.

2. DISCIPLINE AS AN EXPERT SYSTEM

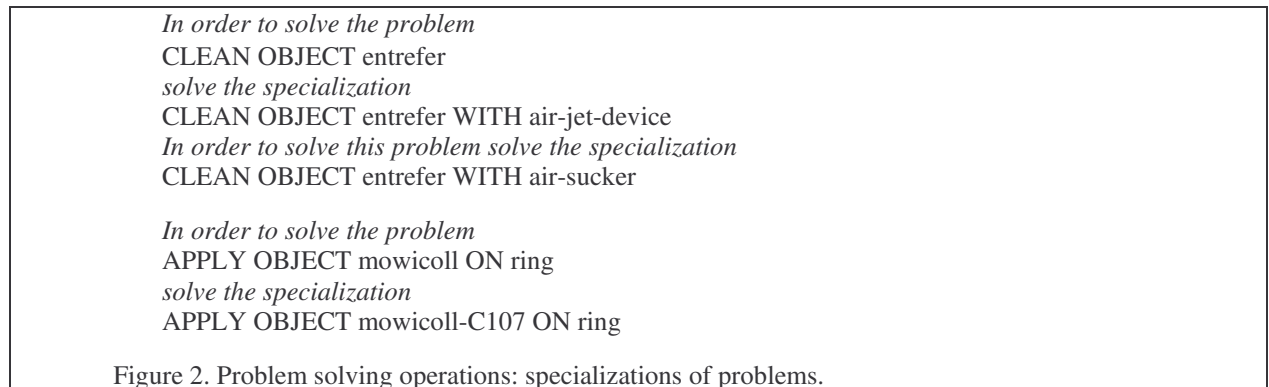
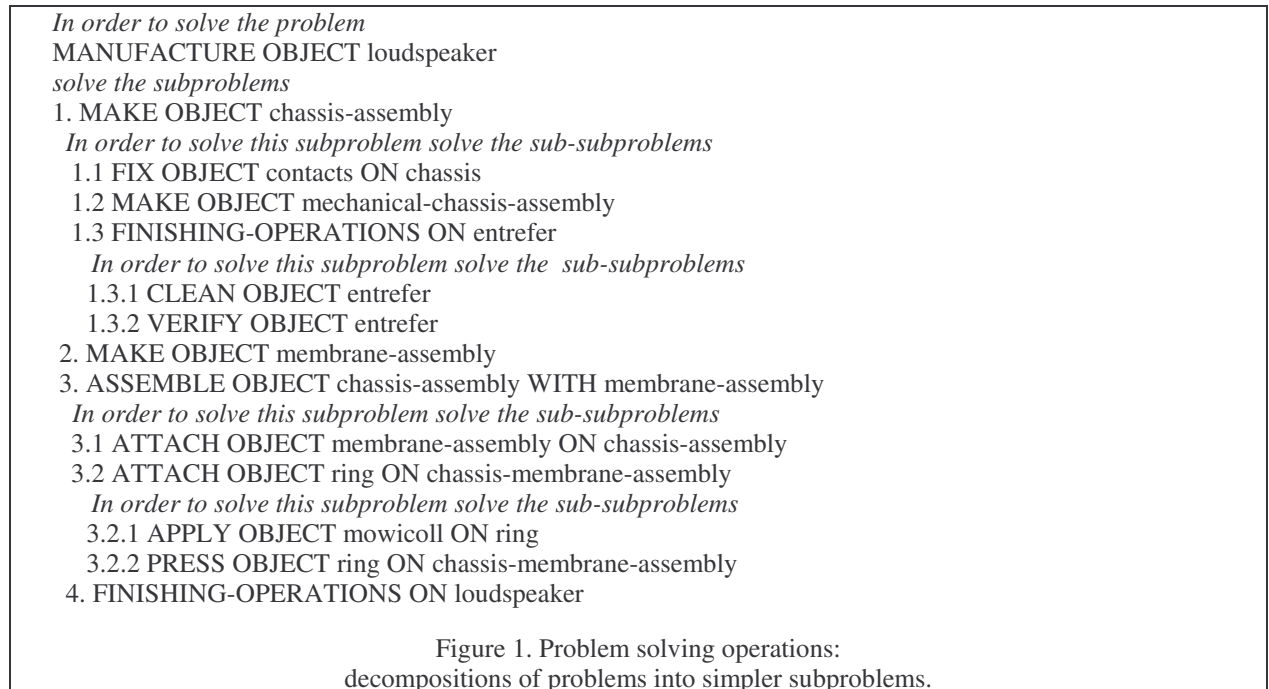
In DISCIPLINE we have adopted a problem reduction approach to problem solving. That is, a problem is solved by successively reducing it to simpler subproblems. This process continues until the initial problem is reduced to a set of elementary problems, that is, problems with known solutions. Moreover, the problem to solve may be initially imprecisely formulated, becoming better and better formulated as the problem solving process advances. To this purpose, DISCIPLINE formulates, propagates, and evaluates constraints ((Tecuci, 1988), (Tecuci, Kodratoff, Bodnaru & Brunet, 1987))

Problem Reduction is a general method, suitable for solving a large variety of problems. In the following, however, we shall consider only problems of designing action plans for achieving partially specified goals. These problems are similar to those solved by PLANX10 (Sridharan & Bresina, 1982), NONLIN (Tate, 1977), and others. An example of such a problem is the following one:

- given the incomplete specifications of a loudspeaker;
- design the actions needed to manufacture the loudspeaker.

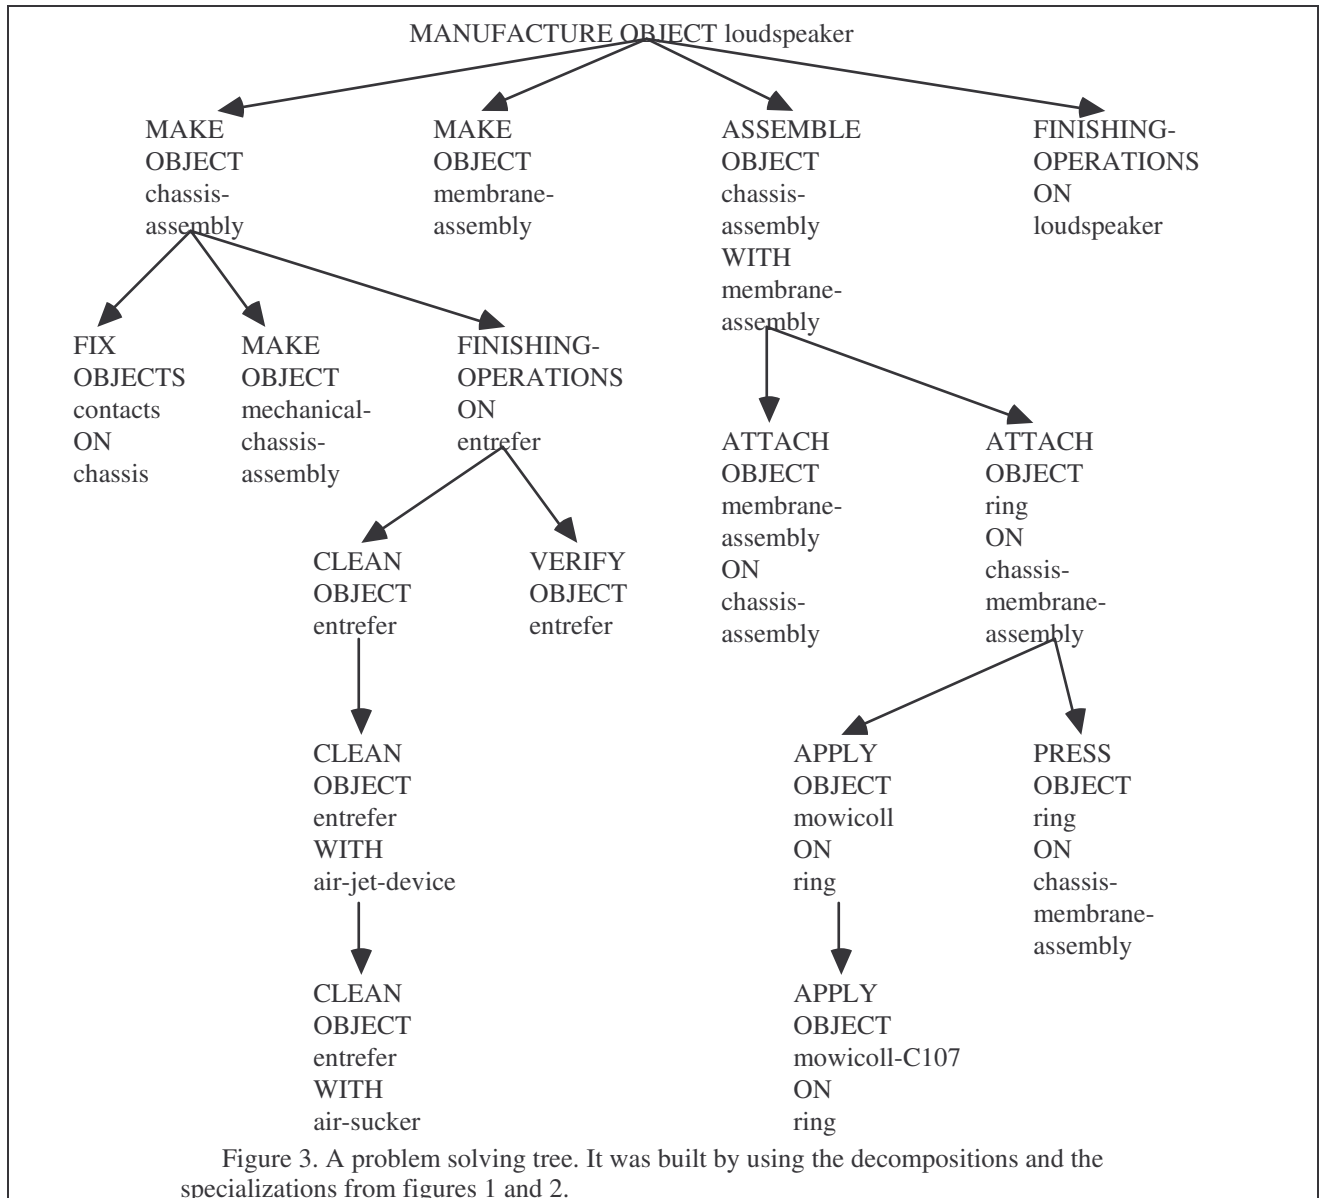
DISCIPLINE may start with the following top-level operation, which can be seen as the current goal:
 MANUFACTURE OBJECT loudspeaker

It will try to solve this problem by successive decompositions and specializations, as illustrated in figure 1 and in figure 2.



DISCIPLE will combine such decompositions and specializations, building a problem solving tree like the one in figure 3.

This process continues until all the leaves of the tree are elementary actions, that is, actions which may be executed by the entity manufacturing the loudspeaker.



This is a standard AND tree, the solution to the problem from the top of this tree consisting of the leaves of the tree. That is, to manufacture the loudspeaker, one has to perform the following sequence of operations:

FIX OBJECTS contacts ON chassis
 MAKE OBJECT mechanical-chassis-assembly
 CLEAN OBJECT entrefer WITH air-sucker
 VERIFY OBJECT entrefer
 MAKE OBJECT membrane-assembly
 ATTACH OBJECT membrane-assembly ON chassis-assembly
 APPLY OBJECT mowicoll-C107 ON ring
 PRESS OBJECT ring ON chassis-membrane-assembly
 FINISHING-OPERATIONS ON loudspeaker

Let us notice that the decompositions and the specializations model in fact the main operations used in design, where one usually starts with a very general specification of an object and successively imposes different constraints on the specification and reduces object design to sub-parts design.

3. THE LEARNING PROBLEM

The decompositions and the specializations from figure 3 were the result of the application of general reduction rules or were directly indicated by the user. From each solution received from the user, DISCIPLE is trying to learn a general problem solving rule. Therefore, the learning problem of DISCIPLE may be formulated as follows:

Given:

Domain Theory

The domain theory contains:

- a specification of the types of objects in the world and their properties and relations;
- a set of inference rules for inferring properties and relations from other properties and relations;
- a set of action models describing the actions that may be performed in the domain. An action model specifies the preconditions of the action (i.e. the states of the world in which the action may be executed), the effects of the action (i.e. the states that result after the execution of the action), as well as the objects that may play certain roles in the action (the agent executing the action, the object on which the action is performed, the instrument used, etc.).

Problem Solving Episode

It consists of

- P, a problem to solve, and
- S, a (partial) solution to P.

Determine:

A General Problem Solving Rule.

According to this rule, problems similar to P will receive solutions similar to S.

Figure 4. The learning problem of DISCIPLE.

For instance,

Given:

The theory of loudspeaker manufacturing;

The problem of attaching two parts of the loudspeaker (the 'ring' and the 'chassis-membrane-assembly') and the decomposition of this problem into two simpler subproblems expressing the gluing of the two parts with mowicoll:

Example 1:

Solve the problem

ATTACH OBJECT ring ON chassis-membrane-assembly

by solving the subproblems

APPLY OBJECT mowicoll ON ring

PRESS OBJECT ring ON chassis-membrane-assembly

Figure 5. A decomposition indicated by the user.

Determine:

A *general decomposition rule* indicating the conditions under which one may reduce an 'attachment' problem to a process of gluing:

<p>IF</p> <p>(x TYPE solid) & (y TYPE solid) & (x PARTIALLY-FITS y) & (z ISA adhesive) & (z TYPE fluid) & (z GLUES x) & (z GLUES y)</p> <p>THEN</p> <p>General Rule 1: <i>solve the problem</i> ATTACH OBJECT x ON y <i>by solving the subproblems</i> APPLY OBJECT z ON x PRESS OBJECT x ON y</p>

Figure 6. The general decomposition rule learned from Example 1.

If 'x' and 'y' are two solid objects that partially fits each other, and there is a fluid adhesive 'z' that glues both 'x' and 'y', then one may attach 'x' on 'y' by first applying 'z' on 'x' and then by pressing 'x' on 'y'.

As one may notice, the structure of **General Rule 1** is identical with the structure of **Example 1**. Therefore, rule learning is reduced to learning the features that the objects 'x', 'y', and 'z' should have so that the attachment of 'x' and 'y' may be reduced to a process of gluing them with 'z'. Otherwise stated, one should learn the concepts represented by these objects.

The method of learning this rule depends on the system's theory (knowledge) about **Example 1**. We distinguish between three types of theories: **complete**, **weak**, and **incomplete**.

A **complete theory** about Example 1 consists of the complete descriptions of the objects and actions from this problem solving episode.

In such a case, DISCIPLE uses an explanation-based learning method, being able to learn at once a general rule from Example 1 alone.

A **weak theory** about Example 1 consists only of incomplete descriptions of the objects. It differs qualitatively from a complete theory in that it does not contain action descriptions.

In this case, DISCIPLE uses an interactive learning method that synergistically combines explanation-based learning, learning by analogy, empirical learning, and learning by questioning the user.

The intermediate case, between a complete theory and a weak theory, is the **incomplete theory**. It contains incomplete descriptions of the objects and the actions from Example 1.

In the case of an incomplete theory about Example 1, DISCIPLE learns a general rule by combining the method corresponding to the weak theory with the one corresponding to the complete theory.

A side effect of rule learning in the context of a weak or incomplete theory is that of developing the domain theory.

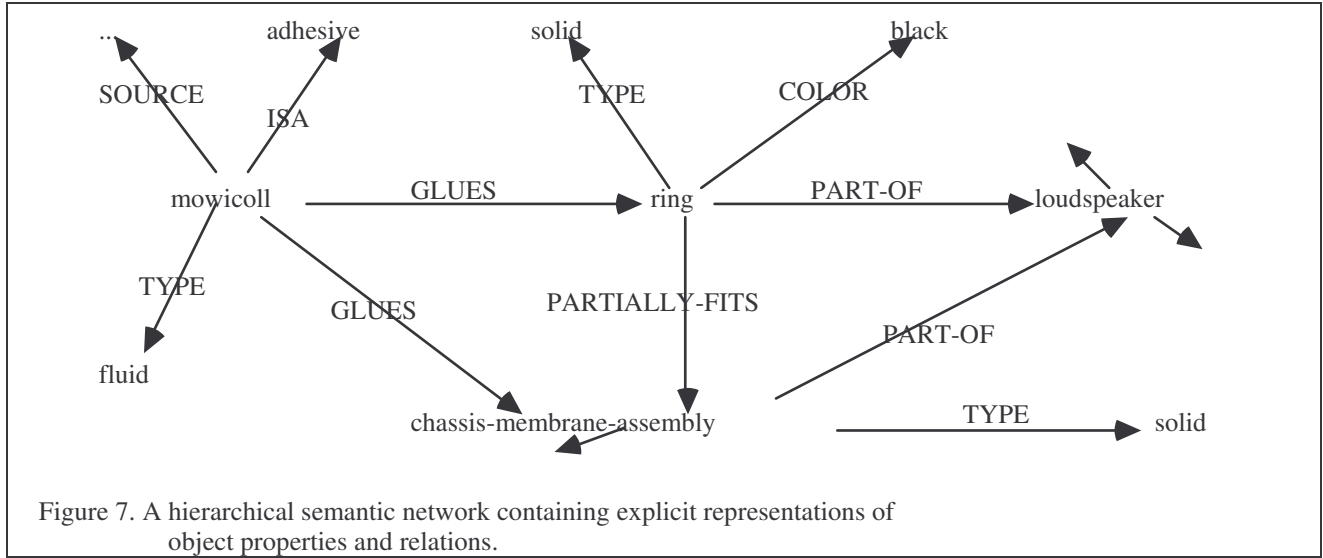
In the following sections we shall present these three learning methods of DISCIPLE.

4. LEARNING IN A COMPLETE THEORY DOMAIN


4.1 A sample of a complete theory

In the case of DISCIPLE, a complete theory of a domain consists of complete descriptions of all the objects and actions of the domain. In particular, a complete theory about the problem solving episode in figure 5, contains the complete descriptions of the objects 'ring', 'chassis-membrane-assembly', and 'mowicoll', as well as the complete descriptions (models) of the actions 'ATTACH', 'APPLY', and 'PRESS'.

The objects are described by specifying all the relevant factual properties and relations. Some of these may be explicitly specified, as indicated in figure 7.



Other properties and relations may be implicitly specified by using inference rules for deducing them from other properties and relations, as indicated in fig. 8.



The action models describe the actions that may be performed in the domain. A complete action model specifies all the necessary preconditions of the action (i.e. all the states of the world in which the action may be executed), all its effects (i.e. the states that result after the execution of the action), as well as all the objects that may play certain roles in the action (the agent executing the action, the object on which the action is performed, the instrument used, etc.).

Figure 9 presents the models of the actions from the problem solving episode in figure 5. For instance, the action 'APPLY' may be performed if and only if 'x' is a solid object and 'z' is a fluid object which is adherent on 'x'. As an effect of performing this action, 'z' will be applied on 'x'. One may notice that the necessary features of the objects are specified in the action's preconditions.

Action	Preconditions	Effects
ATTACH OBJECT x ON y	(x TYPE solid) & (y TYPE solid)	(x ATTACHED-ON y)
APPLY OBJECT z ON x	(z TYPE fluid) & (z ADHERENT-ON x) & (x TYPE solid)	(z APPLIED-ON x)
PRESS OBJECT x ON y	(z APPLIED-ON x) & (x PARTIALLY-FITS y) & (y TYPE solid)	(z BETWEEN x y)

Figure 9 Action models.

4.2 General presentation of the learning method

In the case of a complete theory about Example 1, the learning method of DISCIPLE follows the explanation-based learning paradigm developed by (DeJong & Mooney, 1986), (Fikes, Hart & Nilsson, 1972), (Mitchell, Keller & Kedar-Cabelli, 1986) and others:

1. Prove that the solution indicated by the user is indeed a solution of the problem to solve. This proof isolates the relevant features of the objects in Example 1, that is, those features which will be present in the condition of General Rule 1.
2. Generalize the proof tree as much as possible so that the proof still holds. This is done as in (Mooney & Bennet, 1986) by replacing each instance of action model or inference rule with its general pattern and by unifying these patterns. By generalizing the proof tree, one generalizes the problem, its solution, and the relevant features.
3. Formulate the learned rule from the generalized proof by extracting the generalized problem, its generalized solution, and the generalized relevant features, which constitute the applicability condition of the rule.

In the following sections we shall briefly illustrate this method with the aid of **Example 1** (figure 5).

4.3 Proving the example

To prove **Example 1** means to show that the sequence of the actions

APPLY OBJECT mowicoll ON ring

PRESS OBJECT ring ON chassis-membrane-assembly

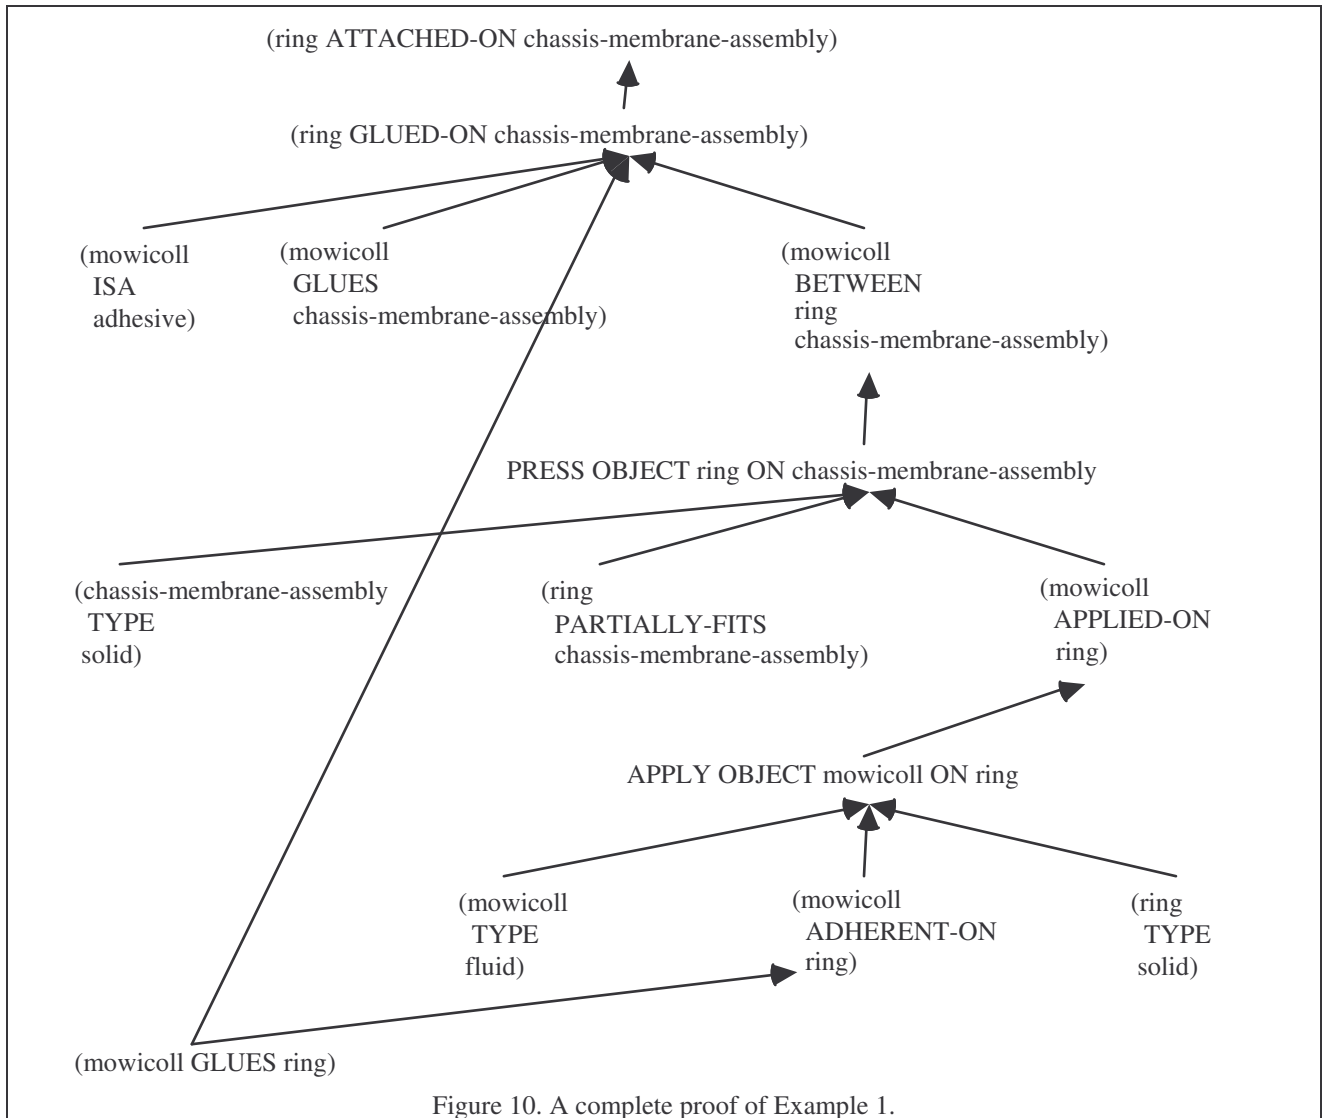
achieves the goal of the action

ATTACH OBJECT ring ON chassis-membrane-assembly

that is, achieves the goal

(ring ATTACHED-ON chassis-membrane-assembly).

The proof is indicated in figure 10. It was obtained by using the object descriptions in figure 7, the inference rules in figure 8, and the action models in figure 9.



The leaves of this tree are those features of 'ring', 'chassis-membrane-assembly', and 'mowicoll' which allowed one to reduce the problem of attaching the 'ring' on the 'chassis-membrane-assembly', to the process of gluing them with 'mowicoll'. Thus, by proving the example, one isolates the relevant features of it:

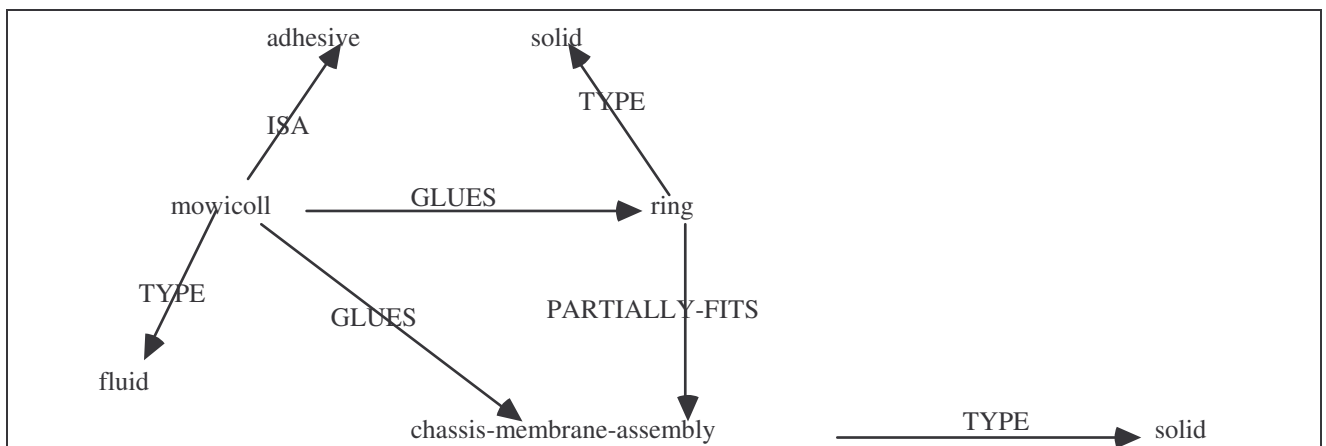


Figure 11. The relevant features of Example 1.

The 'color' of the 'ring' or the 'source' of the 'mowicoll' were not useful in proving the validity of the example. Therefore, these features are not important for this example.

4.4 Generalization of the proof

The next step consists in the generalization of the proof, as much as possible, so that the proof still holds.

Since the proof in figure 10 was obtained by using instances of inference rules and action models, one may generalize the proof by generalizing these instances.

One way to do this is to first replace each instantiated inference rule or action model with its general pattern, and then to unify these patterns (Mooney & Bennet, 1986):

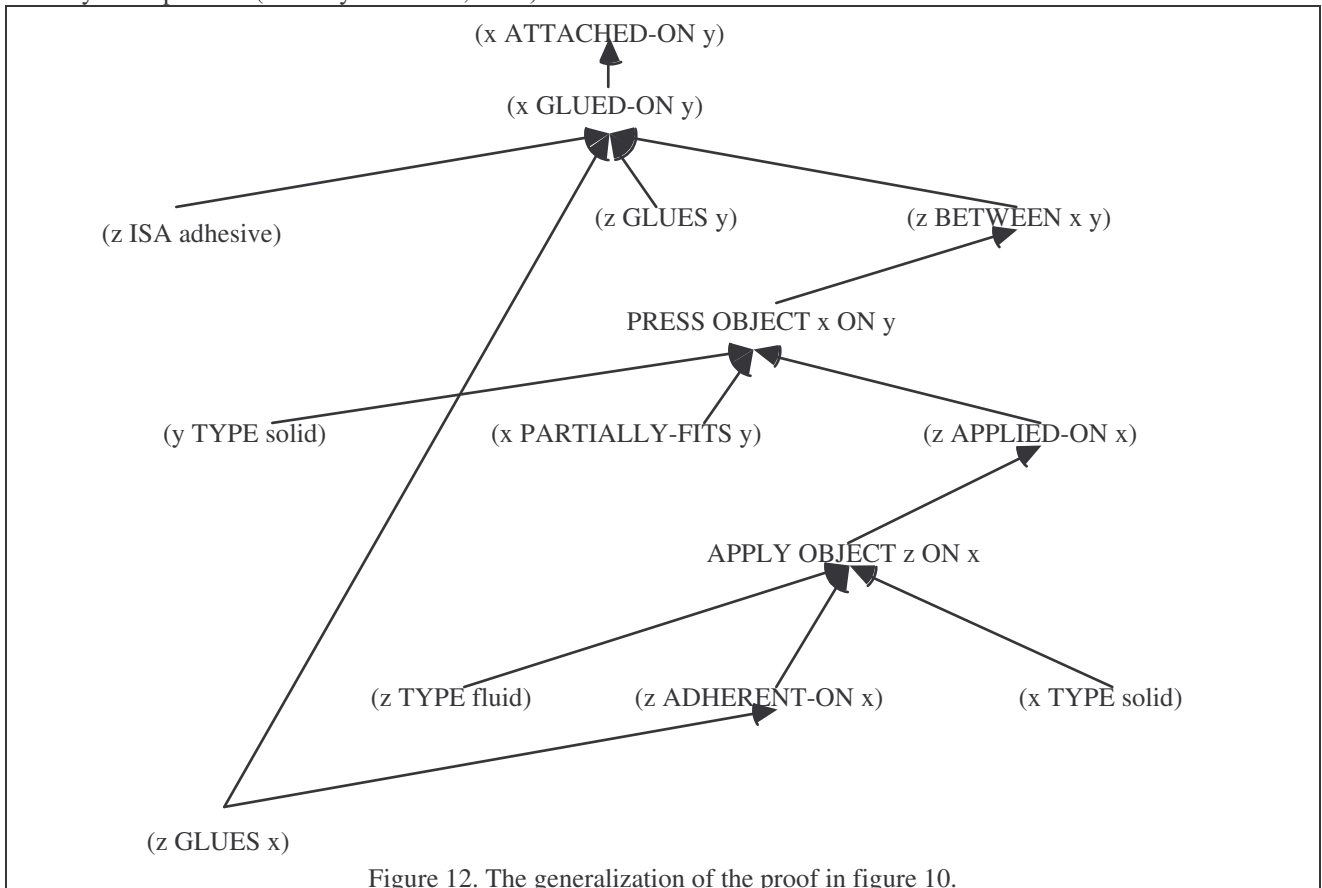


Figure 12. The generalization of the proof in figure 10.

The leaves of this generalized tree represent a justified generalization of the relevant features in figure 11:

(x TYPE solid) & (y TYPE solid) & (x PARTIALLY-FITS y) &
(z ISA adhesive) & (z TYPE fluid) & (z GLUES x) & (z GLUES y)

They also represent a general precondition for which the sequence of the actions 'APPLY OBJECT z ON x', 'PRESS OBJECT x ON y' achieves the goal of the action 'ATTACH OBJECT x ON y'.

That is, one has learned the general decomposition rule in figure 6.

5. LEARNING IN A WEAK THEORY DOMAIN

5.1 A sample of a weak theory

A weak theory about the problem solving episode in figure 5 (Example 1) consists of the incomplete descriptions of the objects from this episode. It does not contain any action model.

A sample of such a theory is represented in figure 13.

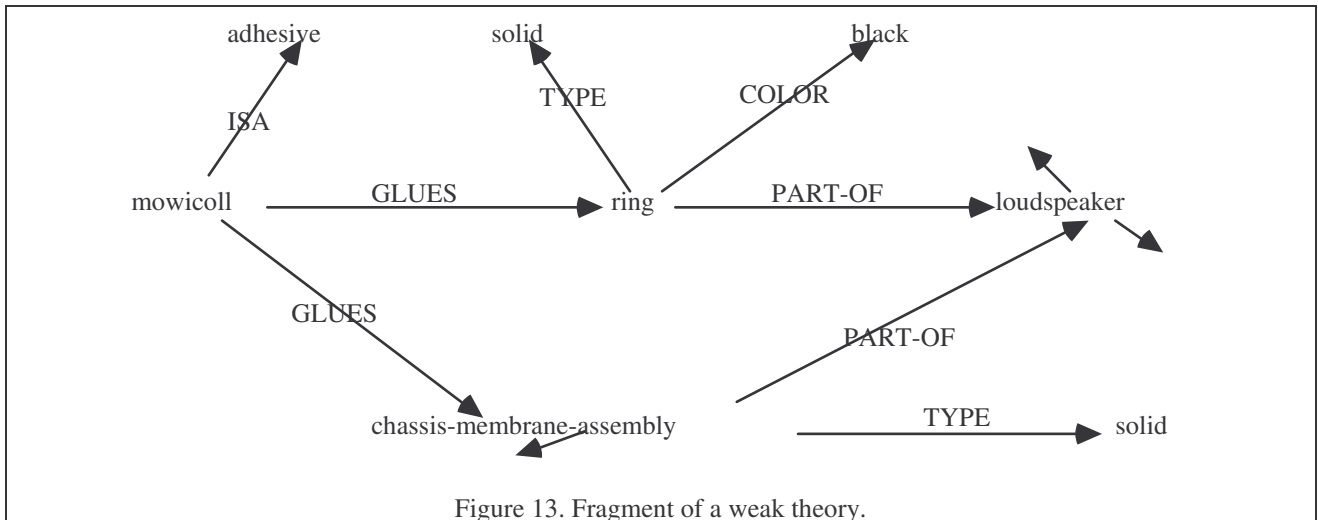


Figure 13. Fragment of a weak theory.

Considering such a theory is justified because it is very difficult for an expert to describe the actions in terms of their preconditions and effects. On the other hand, it is much easier for him to describe the objects and to give examples of decompositions and specializations.

Therefore, instead of forcing the expert to completely formalize his knowledge, we decided to accept the theory which was easily provided by him and to learn the rest of the necessary knowledge.

5.2 General presentation of the learning method

In the context of a weak theory, DISCIPLÉ will try to balance the lack of knowledge by using an integrated learning method whose power comes from the synergism of different learning paradigms: explanation-based learning, learning by analogy, empirical learning, and learning by questioning the user.

Rule learning takes place in several stages which are illustrated in figure 14.

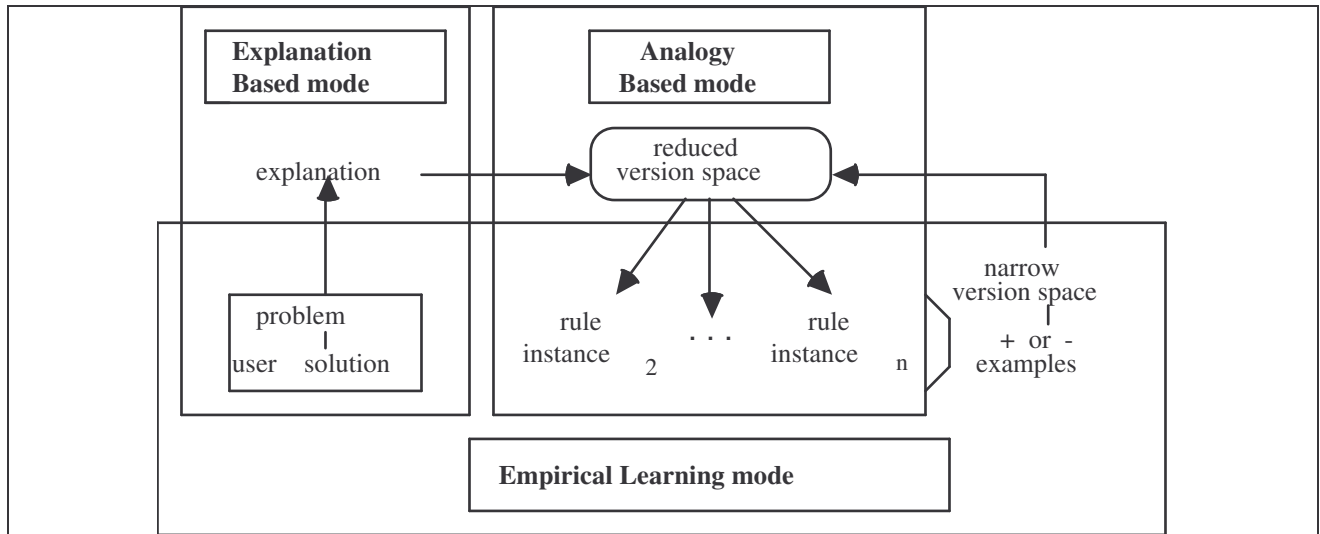


Figure 14 The learning method in the context of a weak theory.

First DISCIPLE looks for a shallow explanation of user's solution. Then it uses this explanation to formulate a reduced version space for the rule to be learned. Each rule in this space covers only instances which are analogous with the user's example. DISCIPLE carefully generates analogous instances to be characterized as positive examples or as negative examples by the user. These are used to further narrow the version space until it contains only the rule illustrated by the user's solution.

More formally, the learning method is the following one:

Explanation-Based Mode

1. Find an explanation of the user's solution (Example 1) and call it Explanation 1

Analogy-Based Mode

2. Over-generalize Example 1, by simply turning all the objects into variables, and call it General Rule 1.
3. Take Explanation 1 as a Lower Bound for the applicability condition of General Rule 1.
4. Over-generalize Explanation 1 to the most general expression that may still be accepted by the user as an explanation of General Rule 1.
5. Take the over-generalized explanation as an Upper Bound for the applicability condition of General Rule 1. The Upper Bound, the Lower Bound, and the General Rule 1 define a reduced version space for the rule to be learned.
6. Look in the knowledge base for tuples of objects which satisfy the Upper Bound but do not satisfy the Lower Bound.

If there are such objects then call Explanation-*i* the properties of these objects which were used to prove that they satisfy the Upper Bound and go to step 7.

If there are no such objects then show the Upper Bound, the Lower Bound, and the General Rule 1 to the user as an uncertain rule and stop.

7. Use the objects found in step 6 to generate an instance of General Rule 1. Call it Instance-*i*. This instance is analogous with Example 1.

8. Propose Instance-*i* to the user and ask him to characterize it as a valid or as an invalid reduction. If Instance-*i* is rejected by the user then go to step 9. Otherwise go to step 14.

Explanation-Based Mode

9. Take Instance-*i* as a near miss (negative example) of the rule to be learned.
10. Find an explanation of why Instance-*i* was rejected by the user and call it Failure-Explanation-*i*.

Empirical Learning Mode

11. Specialize the Upper Bound as little as possible, so that not to cover Failure-Explanation-*i*. If the new Upper Bound is identical with the Lower Bound then take it as a necessary and sufficient condition of General Rule 1, show them to the user and stop, else go to step 12.

12. Specialize (if necessary) the Lower Bound as little as possible, so that not to cover Failure-Explanation-*i*.

13. Go to step 6.

14. Take Instance-*i* as a new positive example of the rule to be learned and Explanation-*i* as a true explanation of Instance-*i*.

15. Look for a maximally specific common generalization of the Lower Bound and Explanation-i. Two cases may occur:
- if this generalization is not identical with the Upper Bound, then take it as the new Lower Bound and go to step 6;
 - if this generalization is identical with the Upper Bound, then take it as a necessary and sufficient condition of General Rule 1, show them to the user and stop.

In the following sections we shall illustrate and justify this learning method by using again **Example 1** from figure 5.

5.3 Explanation-based mode

In its first learning step, DISCIPLE enters the Explanation Based Mode and tries to find an explanation (within its weak domain theory) of the validity of the solution in figure 5 (Kodratoff & Tecuci, 1989).

We shall first define what we mean by an explanation in a weak theory and then we shall indicate a heuristic method to find such explanations.

5.3.1 Explanations in a weak theory domain

Let 'P' be the problem to solve and 'S' a solution to this problem. As has been shown in section 4, an explanation of the problem solving episode '**solve P by S**' is a proof that 'S' solves 'P'.

In the case of a complete theory about this problem solving episode, the learning system is able to find itself such a proof. In the case of a weak theory, however, the system is no longer able to find such a proof because it lacks the models of the actions from 'P' and 'S'. In such a case, the explanation may be regarded as being the premise of a single inference whose conclusion is '**S solves P**'.

For instance, in the context of a weak theory, a complete explanation of the problem solving episode in figure 5 would be the network from figure 11. Indeed, the fact that the 'ring', the 'chassis-membrane-assembly', and the 'mowicoll' have the features in figure 11 'explains' (in a weak theory) why the process of gluing the 'ring' and the 'chassis-membrane-assembly' with 'mowicoll' solves the problem of attaching them together.

5.3.2 A heuristic to find explanations

The explanation of Example 1 consists of the leaves of the proof tree in figure 11. Since such a tree cannot be built in a weak theory, DISCIPLE uses heuristics to propose plausible partial explanations to be validated by the user who may himself indicate other pieces of explanations. One heuristic is to *look for an explanation expressible in terms of the relations between the objects from the example, ignoring object features*. Therefore, to find an explanation of Example 1, DISCIPLE will look in its knowledge base for the links and for the paths (i.e. sequences of links) connecting 'ring', 'chassis-membrane-assembly', and 'mowicoll', and will propose the found connections as pieces of explanations of the Example 1. It is the user's task to validate them as true explanations:

Do the following justify your solution:
 mowicoll GLUES ring ? Yes
 mowicoll GLUES chassis-membrane-assembly ? Yes
 ring PART-OF loudspeaker &
 chassis-membrane-assembly PART-OF loudspeaker ? No

All the pieces of explanations marked by a user's yes form the explanation of the example rule:




Figure 15. The explanation of Example 1.

Notice that this explanation is incomplete. This is partially a consequence of using heuristics, and partially a consequence of the incompleteness of the domain theory (which may not contain all the relevant object properties and relations). Nevertheless, it shows some important features of the objects, features justifying the user's solution.

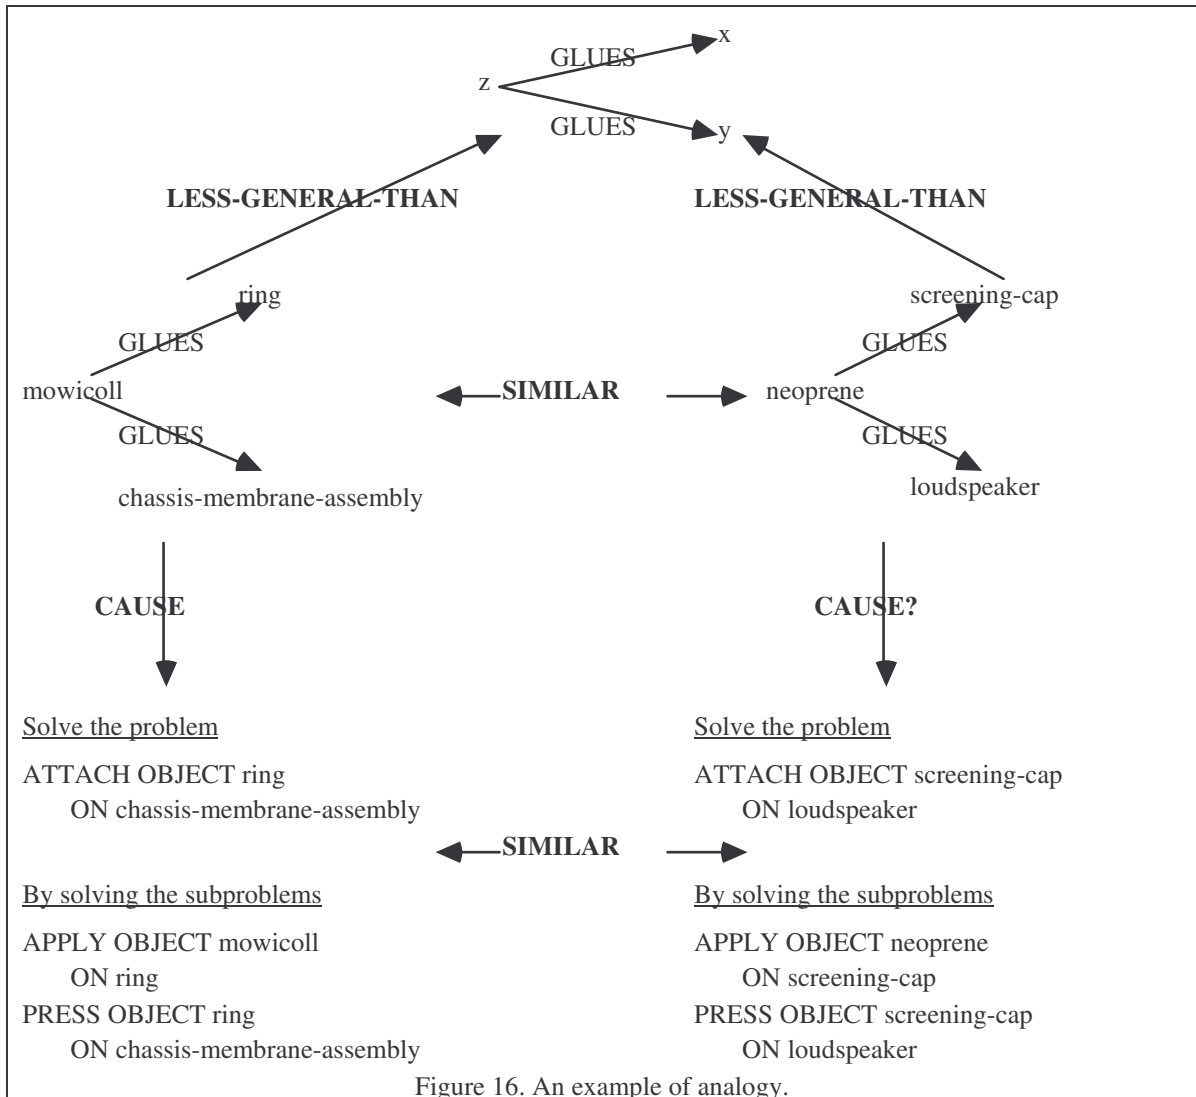
This explanation will be used in the next learning mode (the analogy-based mode) which will be described in the following section. There we shall also give a justification of the above presented heuristic.

5.4 Analogy-Based Mode

The central intuition supporting the learning by analogy paradigm is that if two entities are similar in some respects then they could be similar in other respects as well. An important result of the learning by analogy research ((Bareiss & Porter 1987), (Burstein, 1986), (Carbonell, 1983; 1986), (Chouraqui, 1982), (Gentner, 1983), (Kedar-Cabelli, 1985), (Russel, 1987), (Winston, 1986)) is that the analogy involves mapping some underlying causal network of relations between analogous situations. The idea is that similar causes are expected to have similar effects.

In DISCIPLE, the explanation of a problem solving operation may be regarded as a cause for performing the operation. Therefore, two similar explanations are supposed to 'cause' similar problem solving episodes. Moreover, the explanations are considered to be similar if they are both less general than an over-generalized explanation which is taken as the analogy criterion.

Figure 16 contains an example of such an analogy. The fact that the 'mowicoll' glues both the 'ring' and the 'chassis-membrane-assembly' 'CAUSED' the reduction of the problem of attaching the 'ring' to the 'chassis-membrane-assembly' to a process of gluing them with 'mowicoll'. Because the 'neoprene' glues both the 'screening-cap' and the 'loudspeaker' we may expect (reasoning by analogy) to be able to reduce the problem of attaching the 'screening-cap' and the 'loudspeaker' to a process of gluing them with 'neoprene'.



According to the structure-mapping theory of Gentner (1983), analogy usually involves mapping higher order relations (as the 'CAUSE' relation, in our case). Looking for an explanation in terms of relations between objects, DISCIPLINE ensures that the 'CAUSE' relation, which it imports by analogy, is a higher order relation.

5.4.3 Determining a reduced version space for the rule to be learned

The purpose of the previous sections was to justify the following procedure for determining a reduced version space for the rule to be learned.

First of all DISCIPLINE over-generalizes Example 1 by turning all the objects into variables, thus obtaining:

General Rule 1:
solve the problem
 ATTACH OBJECT x ON y
by solving the subproblems
 APPLY OBJECT z ON x
 PRESS OBJECT x ON y

Next Explanation 1 is rewritten as a lower bound of the applicability condition of General Rule 1 (S bound in fig. 17). Notice that it is indeed a lower bound because it reduces General Rule 1 to Example 1, which is known to be true.

Further, DISCIPLINE determines an analogy criterion which will allow it to generate instances analogous to Example 1.

The analogy criterion is a generalization of Explanation 1. In the case of our example, it was obtained by simply transforming the constants of Explanation 1 into variables, or, if we consider the form of Explanation 1 in figure 17, by dropping the 'ISA' predicates.

In general, *the analogy criterion should be the most general generalization of Explanation 1 that may still be accepted by the user as an explanation of General Rule 1.*

The analogy criterion is taken by DISCIPLINE as an upper bound for the applicability condition of General Rule 1 (G bound in fig.17).

Thus, the analogy criterion, Explanation 1, and General Rule 1 define a reduced version space (Mitchell, 1978) for the rule to be learned:

IF
G:upper bound (analogy criterion)
 (z GLUES x) & (z GLUES y)

S:lower bound (Explanation 1)
 (x ISA ring) & (y ISA chassis-membrane-assembly) & (z ISA mowicoll) &
 (z GLUES x) & (z GLUES y)
 THEN
General Rule 1:
solve the problem
 ATTACH OBJECT x ON y
by solving the subproblems
 APPLY OBJECT z ON x
 PRESS OBJECT x ON y

Figure 17. A reduced version space for the rule to be learned.


Each rule in this space has an applicability condition that is less general than the analogy criterion and more general than Explanation 1. Therefore, it covers only instances that are analogous with Example 1.

5.4.4 Generation of instances

To search the rule in the space from figure 17, DISCIPLINE needs positive and negative instances of it. *These instances may be provided by future problem solving episodes or may be generated by the system itself.*

To generate an instance, DISCIPLINE looks in the knowledge base for objects satisfying the analogy criterion.

The objects 'screening-cap', 'loudspeaker', and 'neoprene' are such objects. DISCIPLINE calls **Explanation-i** the properties of these objects that were used to prove that they satisfy the analogy criterion:



It uses the found objects to generate an instance of General Rule 1 (see figure 17) and asks the user to validate it:

May I solve the problem
 ATTACH OBJECT screening-cap ON loudspeaker
by solving the subproblems
 APPLY OBJECT neoprene ON screening-cap
 PRESS OBJECT screening-cap ON loudspeaker ? Yes

Figure 18. An instance generated by analogy with Example 1.

5.5 Empirical Learning mode

The instances generated in the analogy mode are accepted or rejected by the user, being thus characterized as positive examples or as negative examples of the rule to be learned. These instances are used to search the rule in the version space from figure 17.

5.5.1 The use of the positive examples

Each positive example shows a true explanation. All these explanations are generalized (Kodratoff & Ganascia, 1986) and the obtained generalization is used as a new lower bound of the condition version space.

Let us suppose, for instance, that the user accepts the decomposition in figure 18. Then, Explanation-i, computed in section 5.4.4, is a true explanation which may also be rewritten as a lower bound for the applicability condition of General Rule 1:

Explanation i:

(x ISA screening-cap) & (y ISA loudspeaker) & (z ISA neoprene) &
 (z GLUES x) & (z GLUES y)

Therefore, DISCIPLE computes a maximally specific common generalization of the lower bound in figure 17 and of Explanation-i and takes it as a new lower bound of the condition to be learned:

<p>IF</p> <p>G:upper bound (z GLUES x) & (z GLUES y)</p> <p>S:lower bound (x TYPE solid) & (y TYPE solid) & (z ISA adhesive) & (z GLUES x) & (z GLUES y)</p> <p>THEN</p> <p>General Rule 1: <i>solve the problem</i> ATTACH OBJECT x ON y <i>by solving the subproblems</i> APPLY OBJECT z ON x PRESS OBJECT x ON y</p>
--

Notice that DISCIPLINE generalized '(z ISA mowicoll)' and '(z ISA neoprene)' to '(z ISA adhesive)', by applying the well-known rule of climbing the generalization hierarchies (Michalski, 1983). But it generalized '(x ISA ring)' and '(x ISA screening-cap)' to '(x TYPE solid)' because there is no common generalization of 'ring' and 'screening-cap', and the only relevant property common to 'ring' and 'screening-cap' is that they are both 'solid'. Another common property of 'ring' and 'screening-cap' is that they are both PART-OF 'loudspeaker'. DISCIPLINE considers that this property is not relevant because it was not accepted as explanation of **Example 1** (see section 5.3.2).

Notice also that the new lower bound is always more specific than the upper bound because both the previous lower bound and Explanation 1 are less general than the upper bound. However, the generalization of the lower bound was made in the context of an incomplete knowledge. Therefore it could be an over-generalization, to be latter particularized, when new knowledge becomes available.

5.5.2 The use of the negative examples

Each negative example shows the incompleteness of Explanation 1 and of its over-generalization (the analogy criterion). The explanation of why the instance is a negative example points to the features which were not present in Explanation 1. These new features are used to particularize both bounds of the version space.

Let us consider the objects 'screening-cap', 'loudspeaker' and 'scotch-tape' (an adhesive tape). They also satisfy the analogy criterion (the upper bound of the condition version space) but the corresponding instance is rejected by the user:

<p><i>May I solve the problem</i> ATTACH OBJECT screening-cap ON loudspeaker <i>by solving the subproblems</i> APPLY OBJECT scotch-tape ON screening-cap PRESS OBJECT screening-cap ON loudspeaker ? <u>No</u></p> <p>Figure 19. A negative example of the rule to be learned.</p>
--

In this case, DISCIPLINE looks for an explanation of the failure because this explanation points to the important object features which were not contained in Explanation 1.

The explanation is that 'scotch-tape' is not fluid (therefore, it might not be applied on a curved surface):

Failure Explanation: NOT (scotch-tape TYPE fluid)

That is, the concept represented by 'z' must fluid. Therefore, DISCIPLINE will specialize both bounds of the version space by adding the '(z TYPE fluid)':

<p>IF</p> <p>G:upper bound (z GLUES x) & (z GLUES y) & (z TYPE fluid)</p> <p>S:lower bound (x TYPE solid) & (y TYPE solid) & (z ISA adhesive) &</p>
--

<p style="text-align: center;">(z GLUES x) & (z GLUES y) & (z TYPE fluid)</p> <p>THEN</p> <p>General Rule 1: <i>solve the problem</i> ATTACH OBJECT x ON y <i>by solving the subproblems</i> APPLY OBJECT z ON x PRESS OBJECT x ON y</p>

In another situation, failing to glue two objects whose surfaces do not fit each other, DISCIPLE discovers the condition that the objects should partially fit:

<p>IF</p> <p>G:upper bound (z GLUES x) & (z GLUES y) & (z TYPE fluid) & (x PARTIALLY-FITS y)</p> <p>S:lower bound (x TYPE solid) & (y TYPE solid) & (z ISA adhesive) & (z GLUES x) & (z GLUES y) & (z TYPE fluid) & (x PARTIALLY-FITS y)</p> <p>THEN</p> <p>General Rule 1: <i>solve the problem</i> ATTACH OBJECT x ON y <i>by solving the subproblems</i> APPLY OBJECT z ON x PRESS OBJECT x ON y</p>
--

The learning process decreases the distance between the two bounds of the version space. This process should, in principle, continue until the lower bound becomes identical with the upper one.

In our case, other negative examples will show that '(x TYPE solid) & (y TYPE solid) & (z ISA adhesive)' are necessary features of the objects 'x', 'y', and 'z'. Thus one learns the rule in figure 6.

However, since the domain theory is weak, we should expect that this will not always happen. Therefore, we will be forced to preserve two conditions (the upper bound and the lower bound), instead of a single applicability condition. We propose to define such a case as being typical of an **uncertain explanation** (in which uncertainty is not expressed by numerical means).

5.5.3 Active experimentation

In the Analogy-Based Mode DISCIPLE may generate many instances of the rule to be learned. However, they are not equally useful for searching the version space. Therefore, in the Empirical Learning Mode, DISCIPLE will determine the features of the most useful instances, asking for the generation of such instances. Its strategy is to generalize the lower bound of the version space by generalizing the referred objects (i.e. 'mowicoll', 'ring', and 'chassis-membrane-assembly'). It will therefore try to climb the generalization hierarchy of these objects in such a way as to preserve consistency with the necessary condition.

During this generalization process, several situations may occur:

- there are different ways to generalize;
- the generalization may cover objects that are not guaranteed to produce positive examples of the rule.

When faced with such problems, DISCIPLE will ask the user "clever" questions (as, for instance, in (Sammut & Banerji, 1986)) whose answers allow it to take the right decision. This process is illustrated in (Tecuci, 1988).

5.6 Developing the domain theory

As has been shown in section 5.3.2, DISCIPLE looks for explanations in its knowledge base. Because the domain theory is weak, we may expect that it will not always contain the right pieces of explanations. In such situations the pieces of the explanation must be provided by the user.

Let us consider, for instance, that the explanation of the failure in figure 19 was provided by the user. In this case the domain theory will be enriched by storing this explanation: 'NOT (scotch-tape TYPE fluid)'.

More significantly, as a consequence of updating the Lower Bound of the version space, the following relations between the objects that previously generated positive examples of the rule (and are therefore supposed to satisfy the Lower Bound) are added to the domain theory: '(mowicoll TYPE fluid) & (neoprene TYPE fluid)'.

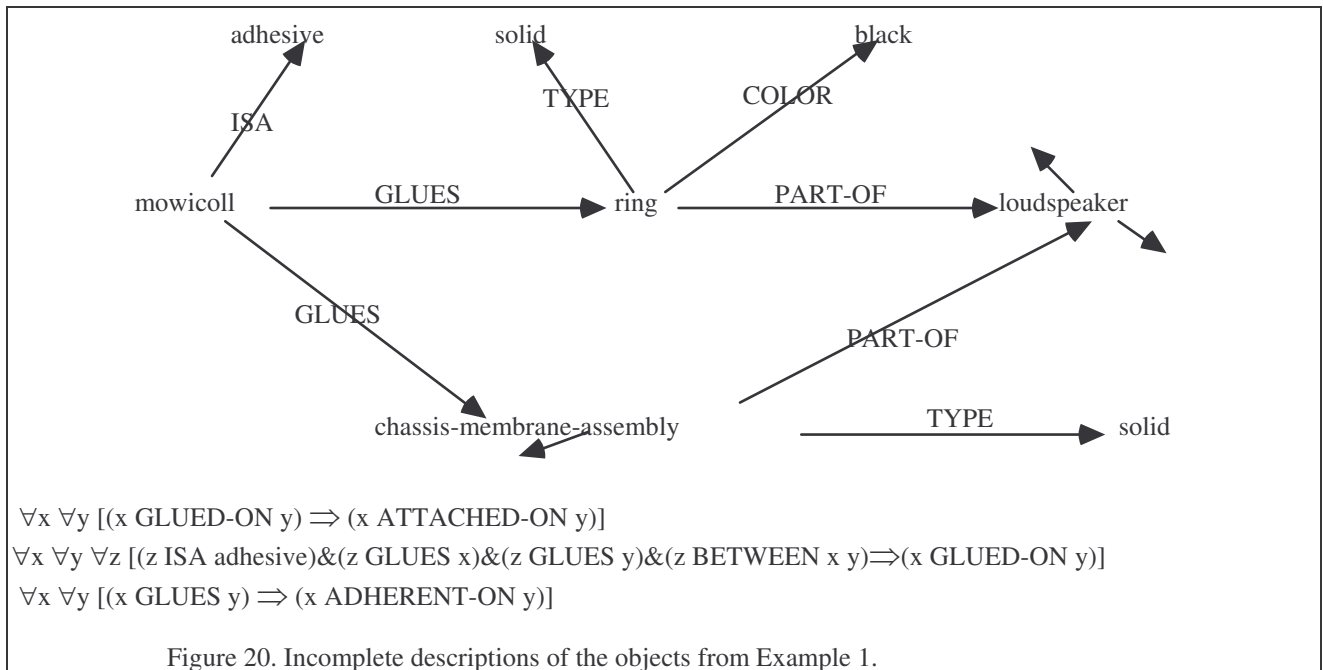
6. LEARNING IN AN INCOMPLETE THEORY DOMAIN

6.1 A sample of an incomplete theory

In the case of DISCIPLE, an incomplete theory of a domain may lack some object descriptions, inference rules, or action models. Also, it may contain incomplete descriptions of these.

An incomplete description of an object lacks certain properties or relations with other objects, an incomplete action model lacks some precondition predicates or some effect predicates, and an incomplete inference rule lacks some left hand side or right hand side predicates.

A sample of an incomplete theory about Example 1 (figure 5) is given in the figures 20 and 21.



Action	Preconditions	Effects
ATTACH OBJECT x ON y	(x TYPE solid) & (y TYPE solid)	(x ATTACHED-ON y)
APPLY OBJECT z ON x	(z ADHERENT-ON x) & (x TYPE solid)	(z APPLIED-ON x)

Figure 21. Incomplete models of two actions from Example 1.

As one may notice, the explicit properties and relations of the objects 'ring', 'chassis-membrane-assembly' and 'mowicoll' are the ones considered in the case of the weak theory (see figure 13).

Let us also notice that this incomplete theory lacks entirely the model of the action 'PRESS'. It also contains an incomplete model of the action 'APPLY', model lacking the precondition predicate '(z TYPE fluid)'.

6.2 General presentation of the learning method

In this case, the learning method combines the two learning methods presented previously.

First, the system will construct an incomplete proof of Example 1 and will generalize it, as in a complete theory. In this way, it will determine an **over-generalized explanation** of Example 1. Then, the system will use the over-generalized explanation as an **analogy criterion** to perform experiments and to synthesize the general rule, as in a weak theory:

1. Prove that the solution indicated by the user is indeed a solution of the problem to solve. Because the domain theory is incomplete, the system may ask the user focused questions, in order to fill the possible gaps in the proof. The leaves of the proof tree represents an incomplete explanation of Example 1.

2. If the user's solution contains new actions, then use the proof found in step 1 in order to define initial version spaces for the models of these actions. As a side effect of rule learning, DISCIPLÉ will learn the models of these new actions.

3. Over-generalize the proof tree found in step 1, as in a complete theory. If an action model is incompletely learned then use the upper bound of its preconditions and effects. The leaves of the over-generalized proof tree represent an over-generalized explanation of Example 1, being taken by DISCIPLÉ as an analogy criterion.

4. Formulate a reduced version space for the rule to be learned, as in a weak theory, by using the explanation found in step 1 and the over-generalized explanation found in step 3.

5. Search the rule in the version space defined in step 4 by performing experiments, as in a weak theory. Use the over-generalized proof determined in step 3 in order to find the explanations of the failures.

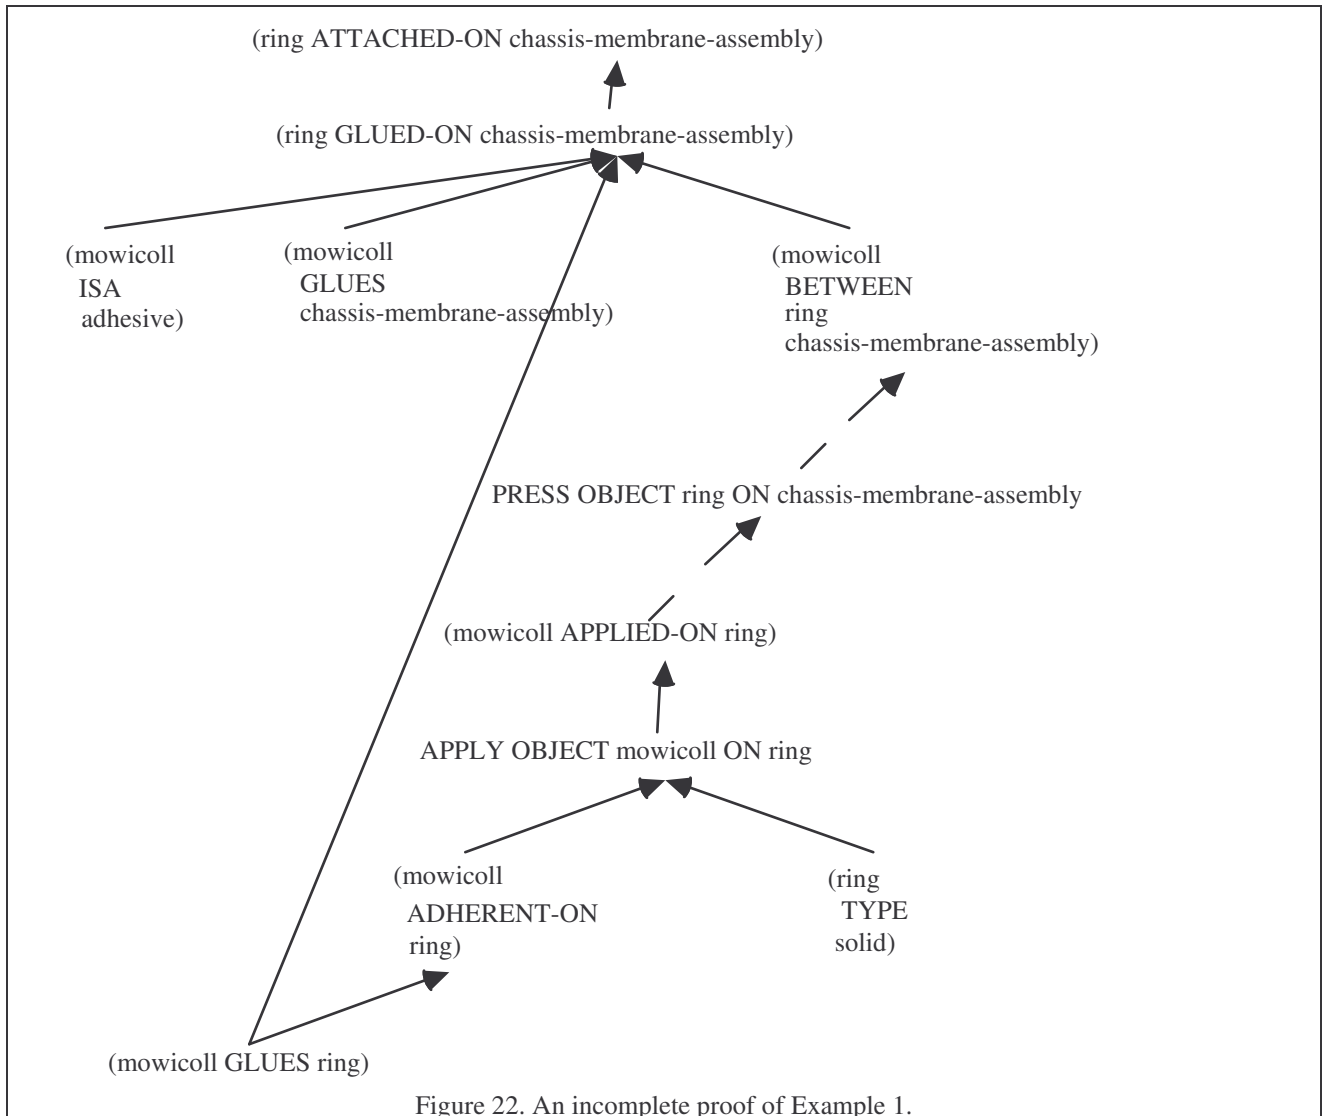
In the next section we shall illustrate this learning method.

6.3. Incomplete proving of the example

Even when the objects, the inference rules, and the actions are incompletely specified, one may be able to construct a proof tree, which lacks some parts of the complete proof tree (see the chapter of Wilkins in this book and (Wilkins, 1988)).

When the system lacks inference rules or action models, it will try to sketch the proof tree both top-down and bottom-up, and will ask the user focused questions, in order to connect the different parts of the proof.

Using the incomplete theory about Example 1, presented in the previous section, the system may build the following proof of Example 1:



The dotted lines from the above proof tree do not result from the domain theory but are hypotheses made by the system and confirmed by the user. For instance, the system makes the hypothesis that

(mowicoll BETWEEN ring chassis-membrane-assembly)

is an effect of the action

PRESS OBJECT ring ON chassis-membrane-assembly

from the fact that all the other left hand side literals of the inference rule

$$\forall x \forall y \forall z [(z \text{ ISA adhesive}) \& (z \text{ GLUES } x) \& (z \text{ GLUES } y) \& (z \text{ BETWEEN } x \ y) \Rightarrow (x \text{ GLUED-ON } y)]$$

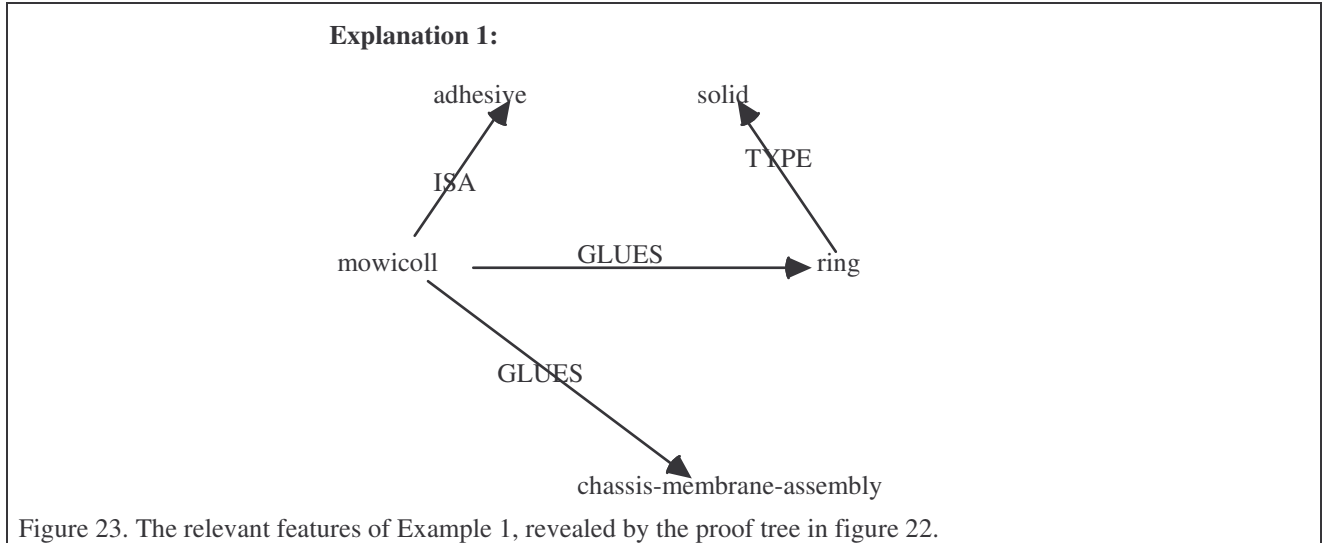
are true in the current situation, that is

[(mowicoll ISA adhesive) & (mowicoll GLUES ring) &

(mowicoll GLUES chassis-membrane-assembly)] = TRUE

and the literal '(mowicoll BETWEEN ring chassis-membrane-assembly)' is not known to be true.

Comparing the proof tree in figure 22 with the complete one in figure 11, one may easily notice that it lacks some leaves. Nevertheless, the leaves which are present represent some important features of the objects from Example 1, features which, in the case of a weak theory, would correspond to the following explanation of Example 1:



6.4 Defining version spaces for the unknown actions

The incomplete proof allows one to define initial version spaces for the models of the unknown actions used in the proof. For instance, one may define the following version space for the action 'PRESS':

Action	Preconditions	Effects
PRESS OBJECT x ON y	<p><i>upper bound:</i> (z APPLIED-ON x)</p> <p><i>lower bound:</i> (z APPLIED-ON x) & (x ISA ring) & (y ISA chassis-membrane-assembly) & (z ISA mowicoll)</p>	<p><i>upper bound:</i> (z BETWEEN x y)</p> <p><i>lower bound:</i> (z BETWEEN x y) & (x ISA ring) & (y ISA chassis-membrane-assembly) & (z ISA mowicoll)</p>

The *lower bounds* for the preconditions and effects are taken directly from the proof tree.

The *upper bound* of the effects is the generalization of the lower bound (mowicoll BETWEEN ring chassis-membrane-assembly) taken from the premise of the inference rule

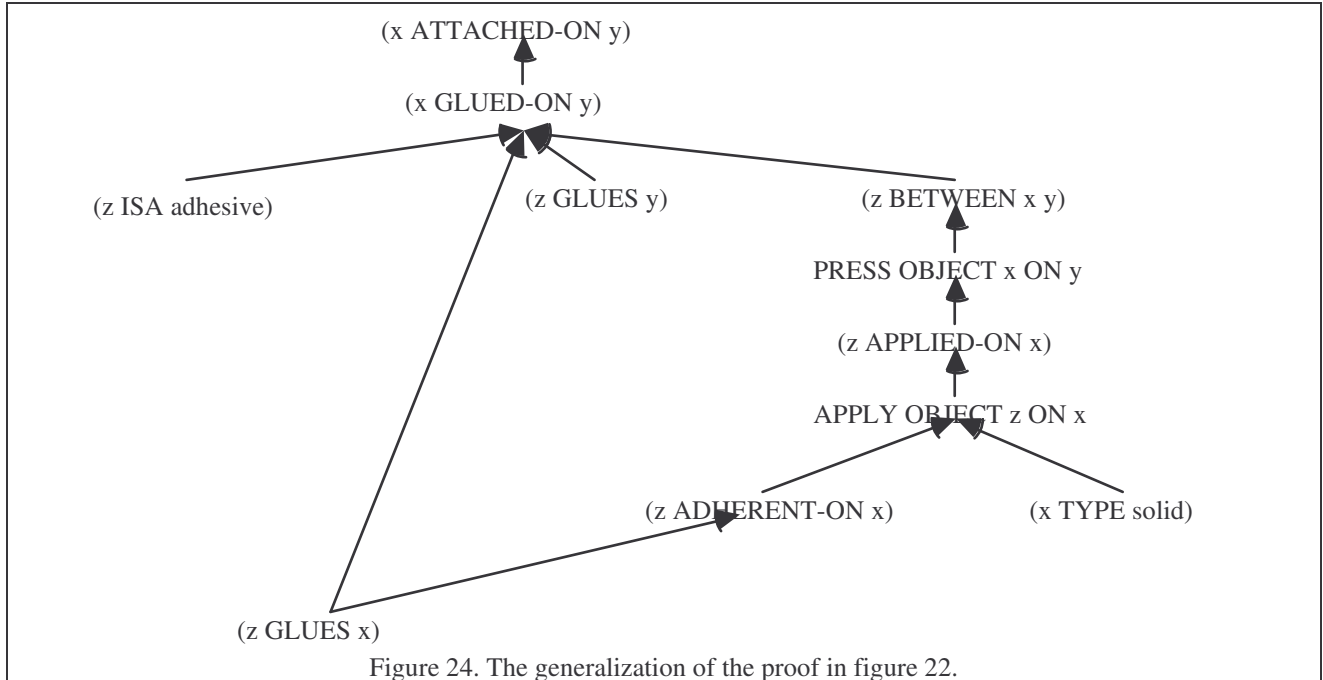
$$\forall x \forall y \forall z [(z \text{ ISA adhesive}) \& (z \text{ GLUES } x) \& (z \text{ GLUES } y) \& (z \text{ BETWEEN } x \ y) \Rightarrow (x \text{ GLUED-ON } y)]$$

The *upper bound* of the preconditions is the generalization of the lower bound, taken from the effects of the model of the action 'APPLY OBJECT z ON x '.

During the learning of the decomposition rule in figure 6, the system will also refine the model of the action 'PRESS'.

6.5 Generalization of the incomplete proof

Once the proof in figure 22 is built, the system will generalize it, as in a complete theory:



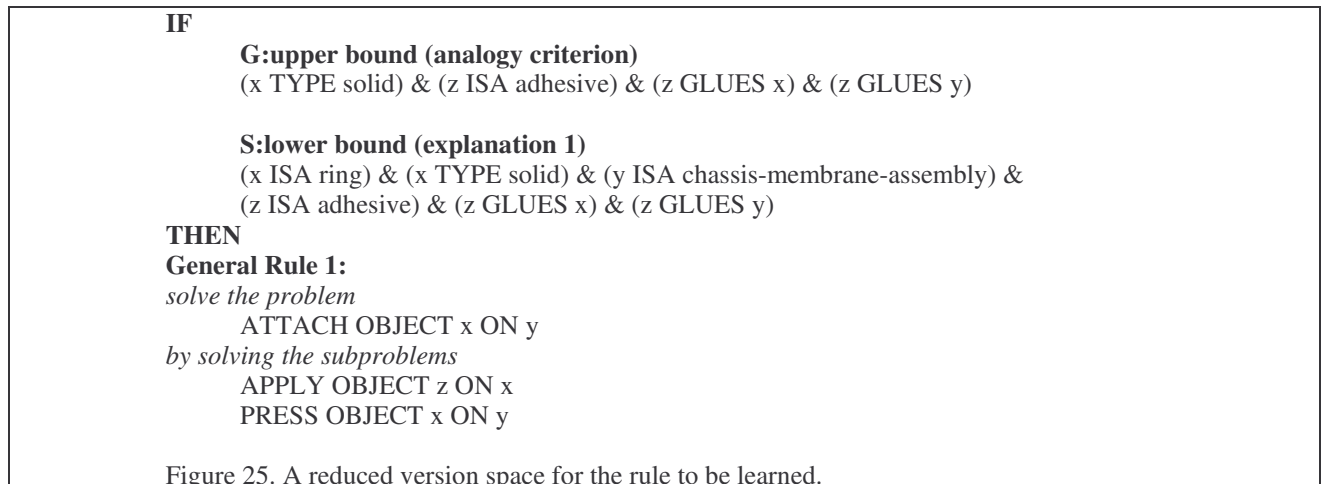
Let us notice that, for generalizing the proof, the system used the upper bounds of the preconditions and effects of the action 'PRESS'.

6.6 Determining a reduced version space for the rule to be learned

As in the case of a weak theory, the Explanation 1 in figure 23 may be rewritten as a Lower Bound for the applicability condition of General Rule 1 (figure 25).

Also, the leaves of the generalized proof tree in figure 24 provide an **over-generalized explanation** of Example 1. This over-generalized explanation corresponds to the **analogy criterion** from a weak theory and is taken by DISCIPLE as an Upper Bound for the applicability condition of General Rule 1 (see figure 25).

Therefore, as in a weak theory, the system is able to formulate the following version space for the rule to be learned:



6.7 Searching the rule in the version space

As soon as the version space from figure 25 has been determined, rule learning will continue as in a weak theory. This time, however, the generalized proof tree in figure 24 provides a focus for the process of finding the explanations of the failures. To illustrate this, let us consider again the failure in figure 19.

In this case, the system generates the instance of the generalized proof in figure 24, corresponding to this problem solving episode (by replacing 'x', 'y', and 'z' with 'screening-cap', 'loudspeaker', and 'scotch-tape', respectively).

The fact that the user rejected the solution proposed by the system proves that the leaves of the instantiated tree do not imply the top of the tree (the leaf predicates are true but the top predicate is not).

This means that some action models or inference rules are faulty (incomplete, in our case). To detect them, the system follows the proof tree from bottom up, asking the user to validate each inference step. If the user says that the effect of an action or the consequent of an inference rule is not true, then the corresponding action model (inference rule) may be the incomplete one.

Therefore, in an incomplete theory, finding the explanations of the failures reduces to finding the knowledge which is lacking from the knowledge pieces. In this case, the generalized proof in figure 24 plays the role of a justification structure for the rule to be learned, as in (Smith, Winston, Mitchell & Buchanan, 1985).

7. EXPERIMENTS WITH DISCIPLE

We have implemented a version of DISCIPLE in COMMON LISP (Steele, 1984) and we have used it to learn rules in several domains as, for instance, manufacturing, common sense planning, chemistry, and architecture ((Kodratoff & Tecuci, 1987), (Tecuci & Kodratoff, 1990)).

With a very poor theory of chemistry¹ DISCIPLE was able to learn, starting from the example of the chemical reaction $(\text{NaOH} + \text{HCl} \rightarrow \text{H}_2\text{O} + \text{NaCl})$, that, in general, $(\text{Base} + \text{Acid} \rightarrow \text{Water} + \text{Salt})$.

More precisely, starting from the example

the problem

COMBINE SUBSTANCE1 NaOH SUBSTANCE2 HCl

has the following solution

RESULT SUBSTANCE1 H2O SUBSTANCE2 NaCl

DISCIPLE learned the following rule:

IF

G:upper bound

(b COMPOSED-OF x1) & (b COMPOSED-OF x2) &
 (a COMPOSED-OF x3) & (a COMPOSED-OF x4) &
 (w ISA water) & (w COMPOSED-OF x1) (COMPOSED-OF x3) &
 (s ISA salt) & (s COMPOSED-OF x2) (s COMPOSED-OF x4) &
 (s (COMPOSED-OF ANION-OF) a) & (s (COMPOSED-OF CATION-OF) b))

S:lower bound

(b ISA base) & (b COMPOSED-OF x1) & (b COMPOSED-OF x2) &

¹ This application was suggested by D.Sleeman.

(a ISA acid) & (a COMPOSED-OF x3) & (a COMPOSED-OF x4) &
 (w ISA H2O) & (w COMPOSED-OF x1) & (COMPOSED-OF x3) &
 (s ISA salt) & (s COMPOSED-OF x2) & (s COMPOSED-OF x4) &
 (s (COMPOSED-OF ANION-OF) a) & (s (COMPOSED-OF CATION-OF) b)) &
 (x1 ISA OH) & (x2 ISA METAL) & (x3 ISA H) & (x4 ISA METALLOID)

THEN

General Rule

the problem

COMBINE SUBSTANCE1 b SUBSTANCE2 a

has the following solution

RESULT SUBSTANCE1 w SUBSTANCE2 s

The lower bound of this rule says that all the positive examples share the structure (Base + Acid --> Water + Salt), together with the appropriate components: for instance, x2 is the metal of the base b (because "b COMPOSED-OF x2") that will go to the salt s (because "s COMPOSED-OF x2").

The upper bound of the rule says that none of the negative examples met shared the structure (Something + Something --> Water + Salt) together with the appropriate components, i.e. there might be compounds leading to water and salt, other than bases and acids.

The application from architecture² consists in designing a building.

In this application, DISCIPLE may learn rules to refine specifications of objects. For instance, starting from an example of a door separating a hall from a sleeping-room, for which the expert established that it should be opened towards the sleeping-room, DISCIPLE learned a general rule for establishing the direction of opening of the doors: "A door separating a house-piece from a room should be opened towards the room".

From these experiments we have learned the followings:

- firstly, it is fairly easy to define a small initial domain theory for DISCIPLE;
- secondly, although DISCIPLE has knowledge to generate hundreds of examples it actually generates a few of them only in order to learn a rule.

CONCLUSIONS

Trying to cope with the complexity of the real world applications, we have made the hypothesis that DISCIPLE's domain theory is nonhomogeneous, describing completely some parts of the domain, but only incompletely or even poorly, the other parts. The use of DISCIPLE is tuned to problem solving situations in which some variabilization is meaningful. For instance, a set of zeroth-order rules solving a problem will not yield interesting results under DISCIPLE.

We have shown how the system is able to learn the same general rule from the same example, by using a method corresponding to its theory about the example.

In the context of a complete theory, DISCIPLE uses explanation-based learning. It is thus able to learn a justified rule from a single example, and may also reject incorrect examples.

The learning method in the context of a weak theory integrates different learning paradigms: explanation based learning, learning by analogy, empirical learning, and learning by questioning the user. Among the most relevant features of this learning method one could mention: the notion of "explanation" in a weak theory and a heuristic method to find such explanations, the use of analogy to define a reduced version space for the rule to be learned, the use of both the explanations of the successes and the explanations of the failures to search the rule in its version space, the formulation of "clever" questions, in order to extract useful knowledge from the expert, the possibility of hiding the learned rules to the expert, a great confidence in the human expert.

In the context of an incomplete theory, DISCIPLE learns by combining the method corresponding to the complete theory with the method corresponding to the weak theory. This method borrows features from both the learning method in a complete theory (may reject incorrect examples, learns justified rules) and from the learning method in a weak theory (use of analogy, clever questions to the user, etc.).

It is interesting to notice that, although in each of the presented cases the system learned the same general rule, the effect of this rule on the future behavior of the system depends of the domain theory: in a *complete theory*, the learned

² Suggested by F.Guena and K.Zreik.

rule improves only the *performance* of the system, in a *weak theory* it develops the *competence* of the system, and, in an *incomplete theory*, it may develop both the *performance* and the *competence*.

Let us also notice that, by the integration of these three learning methods, DISCIPLE proposes a solution to the so called "falling off the knowledge cliff" problem of the current systems. This problem is that a system performs well within the scope of the knowledge provided to it, but any slight move outside its narrow competence causes the *performance to deteriorate rapidly* (Michalski, 1986). On the contrary, in DISCIPLE, the move from one part of the application domain, characterized by a complete theory, to another part, characterized by an incomplete theory or by a weak theory, causes only a *slight deterioration of the performance*, this effect being obtained by a corresponding replacement of the learning method used.

There are also several **weaknesses** of DISCIPLE, on which will shall direct our future research.

For instance, the expressions DISCIPLE deals with are made of predicates, constants and variables, but no actual function evaluation is going to take place.

A semantic limitation comes from the fact that the *generality of the learned rule is limited by the generality of the over-generalized explanation* (the analogy criterion) which may not be in the most general form. However, the rule may be further generalized, in response to a problem solving situation in which the rule does not apply and the user says that it should apply. In this case, the condition of the rule may be generalized to cover the new situation as well.

Also, the method of finding an explanation in a weak theory is not powerful enough. Other sources of knowledge are needed, as well as meta-rules for finding far off explanations. One possible extension of the current method is suggested by the way CLINT (De Raedt & Bruynooghe, 1989) changes its description language in order to be able to learn a concept. DISCIPLE might also use an ordered series of explanation-schemas E1, E2, ... , En, Firstly it will be looking for an explanation of the form E1 (for instance, a path of length 1 between two concepts). If no such explanation is found, then it will be looking for an explanation of the form E2 (for instance, a path of length 2 between two concepts), and so on.

While DISCIPLE uses control knowledge in the form of meta-rules (Tecuci, 1988), such knowledge is not learned, having to be provided by the user. Therefore, if two experts provide different solutions to the same problem, DISCIPLE simply generates two different rules. The learning mechanisms of DISCIPLE should be used to propose explanations of this difference and find meta-explanations that can become meta-preconditions on the use of the rules.

An important future direction of research consists in developing the learning methods of DISCIPLE in order to be able to deal with other types of imperfections in the domain theory ((Mitchell, Keller & Kedar-Cabelli, 1986), (Rajamoney & DeJong, 1987)). We shall consider, for instance, imperfections resulting from the fact that certain pieces of knowledge (objects, inference rules, action models) contain minor errors in their definitions in that parts of these definitions may be either more general or less general than they should be.

A weakness of all the learning apprentice systems is that they need an initial domain theory and provide no means of defining it. Although DISCIPLE facilitates this task by accepting a nonhomogeneous domain theory, the task remains very difficult. A solution here is that proposed by BLIP (Morik, 1989) which is an interactive learning system mainly concerned with the construction of a domain theory, as a first phase of building a knowledge-based system. Therefore, a very promising research direction seems that of building a system incorporating the capabilities of BLIP and DISCIPLE. Such a system could be an effective tool for building knowledge-based systems:

- in the first stage the system and the user will build together an initial theory of the application domain. This theory, containing elementary knowledge about the domain (basic concepts and inference rules), will be neither complete nor entirely correct;

- in the second stage, the system and the user solve problems together and, during this cooperative problem solving, the system will learn general problem solving rules. Since this learning takes place in the context of an imperfect domain theory, the learned rules will accumulate exceptions. These exceptions correspond in fact to lacking concepts in the domain theory. When too many exceptions are accumulated, the domain theory has to be refined. Therefore, one reenters the first stage and reformulates the domain theory so that to better characterize the current knowledge of the system.

There are also several **lessons** we have learned from the design of DISCIPLE.

One is *to cope with the complexity of real-world applications, one should use any available learning technique*. Indeed, the different learning paradigms have many complementary prerequisites and effects. Therefore they may be synergistically combined.

Another lesson is that *full formalization of imperfect theories is short-time harmful*. Indeed, forcing the expert to completely formalize a domain theory (which may even not have such a complete theory) may result in a degradation of the knowledge provided by him/her.

Lastly, we have discovered that *over-generalization is not only harmless, but also useful and necessary, when interacting with a user*, allowing the identification of features usually neglected by the expert.

Acknowledgments

This work has been sponsored by PRC-GRECO 'Intelligence Artificielle' and the Romanian CNST. The chapter has been written while one of the authors was at LRI, on leave from his institute. His expenses have been taken in charge through an agreement between the French CNRS and the Romanian Academy of Sciences. We wish to express our gratitude to all these institutions. We also wish to thank Mr. Zani Bodnaru, for his indispensable contribution as domain expert.

References

- Bareiss E. & Porter B., PROTOS: An Exemplar-Based Learning Apprentice, in Langley P.(ed) Proc. 4th Int. Workshop on Machine Learning, Irvine, 1987.
- Burstein M. H. Concept Formation by Analogical Reasoning and Debugging, in Michalski R., Carbonell J. & Mitchell T. (eds) *Machine Learning: An Artificial Intelligence Approach*, Vol. 2, Morgan Kaufmann 1986, pp.351-370.
- Carbonell J., Derivational Analogy: A Theory of Reconstructive Problem Solving and Expertise Acquisition, in Michalski R., Carbonell J. & Mitchell T. (eds) *Machine Learning: An Artificial Intelligence Approach*, Vol. 2, Morgan Kaufmann 1986, pp.371-392.
- Carbonell J. & Gil Y., Learning by Experimentation, in Langley P.(ed) Proc. 4th Int. Workshop on Machine Learning, Irvine, 1987.
- Chouraqui E., Construction of a Model for Reasoning by Analogy, in Proceedings of the European Conference on Artificial Intelligence, Orsay, France, 1982.
- DeJong G. & Mooney R., Explanation-Based Learning: An Alternative View, in Michalski R., Carbonell J. & Mitchell T. (eds) *Machine Learning: An Artificial Intelligence Approach*, Morgan Kaufmann 1983, pp. 145-176.
- De Raedt L. & Bruynooghe M., Towards Friendly Concept-Learners, In Proceedings of IJCAI-89, Detroit, Morgan Kaufmann, 1989.
- Fikes R.E., Hart P.E. & Nilsson N.J., Learning and Executing Generalized Robot Plans, *Artificial Intelligence*, 3, pp. 251-288, 1972.
- Gentner D, Structure-Mapping: a Theoretical Framework for Analogy, *Cognitive Science*, 7, pp.155-170.
- Kedar-Cabelli S., Purpose-Directed Analogy, In Proceedings of the Cognitive Science Society, pp.150-159, Irvine, California, 1985.
- Kodratoff Y. & Ganascia J-G., Improving the Generalization Step in Learning, in Michalski R., Carbonell J. & Mitchell T. (eds) *Machine Learning: An Artificial Intelligence Approach*, Vol. 2, Morgan Kaufmann 1986, pp. 215-244.
- Kodratoff Y. & Tecuci G., Techniques of Design and DISCIPLINE Learning Apprentice, *International Journal of Expert Systems: Research and Applications*, vol.1, no.1, pp. 39-66, 1987.
- Kodratoff Y., *Introduction to Machine Learning*, Pitman, 1988.
- Kodratoff Y. & Tecuci G., The Central Role of Explanations in DISCIPLINE, in Morik K. (ed), *Knowledge Representation and Organization in Machine Learning*, Springer-Verlag, Berlin 1989.
- Michalski R.S., A Theory and a Methodology of Inductive Learning, *Artificial Intelligence* 20(1983), pp 111-161.
- Michalski R.S., Inductive Learning as Rule-guided Transformation of Symbolic Descriptions: a Theory and Implementation, in *Automatic Program Construction Techniques*, A. W. Biermann, G. Guiho & Y. Kodratoff (eds), Mac-Millan Publishing Company, 1984, pp. 517-552.
- Michalski R.S., Understanding the Nature of Learning: Issues and Research Directions, in Michalski R., Carbonell J. & Mitchell T. (eds) *Machine Learning: An Artificial Intelligence Approach*, Vol. 2, Morgan Kaufmann 1986, pp. 3-25.
- Mitchell T.M., Version Spaces: An Approach to Concept Learning, Doctoral dissertation, Stanford University, 1978.
- Mitchell T., Mahadevan S. & Steinberg L., LEAP: a Learning Apprentice System for VLSI Design, Proc. IJCAI-85, Los Angeles 1985, 573-580.
- Mitchell T.M., Keller R.M. & Kedar-Cabelli S.T., Explanation-Based Generalization: A Unifying View, *Machine Learning* 1, pp. 47-80, 1986.
- Mooney R. & Bennet S., A Domain Independent Explanation Based Generalizer, in Proceedings AAAI-86, Philadelphia, 1986, pp.551-555.
- Morik K., Sloppy modeling, in Morik K. (ed), *Knowledge Representation and Organization in Machine Learning*, Springer Verlag, Berlin 1989.
- Pazzani M.J., Integrating Explanation-based and Empirical Learning Methods in OCCAM, in Sleeman D. (ed), Proceedings of the Third European Working Session on Learning, Glasgow, 1988.
- Quillian M.R., Semantic Memory, in *Semantic Information Processing*, Minsky M., editor, Cambridge, Mass: MIT Press, 1968, pp. 227-270.

- Rajamoney S. & DeJong G, The Classification, Detection and Handling of Imperfect Theory Problems, Proc. IJCAI-87, Milan, pp. 205-207.
- Russel S.J., Analogy and Single-Instance Generalization, in Langley P. (ed) Proc. 4th Int. Workshop on Machine Learning, Irvine, 1987.
- Sammut C. & Banerji R.B., Learning concepts by asking questions, in *Machine Learning: An Artificial Intelligence Approach*, Volume 2, Michalski R.S., Carbonell J.G., Mitchell T.M. (eds), Morgan-Kaufmann 1986, pp. 167-191.
- Smith R., Winston H., Mitchell T. & Buchanan B., Representation and Use of Explicit Justifications for Knowledge Base Refinement, Proc. IJCAI-85, Los Angeles.
- Sridharan N. & Bresina J., A Mechanism for the Management of Partial and Indefinite Descriptions, Technical Report CBM-TR-134, Rutgers Univ., 1983.
- Steele G., COMMON LISP: The Language, Digital Press, 1984.
- Tate A., Generating Project Networks, Proc. IJCAI-77, Massachusetts, pp. 888-893.
- Tecuci G., Kodratoff Y., Bodnaru Z. & Brunet T., DISCIPLINE: An expert and learning system, Expert Systems 87, Brighton, December, 14-17, in D. S. Moralee (ed): *Research and Development in Expert Systems IV*, Cambridge University Press, 1987.
- Tecuci G., DISCIPLINE: A Theory, Methodology, and System for Learning Expert Knowledge, Ph.D. Thesis, University of Paris-Sud, 1988.
- Tecuci G. & Kodratoff Y., How to learn it with DISCIPLINE, LRI Research Report, Orsay, 1990.
- Wilkins D.C., Knowledge Base Refinement Using Apprenticeship Learning Techniques, Proceedings AAAI-88.
- Winston P.H., Learning by Augmenting Rules and Accumulating Censors, in Michalski R.S., Carbonell J.G., Mitchell T.M. (eds), *Machine Learning: An Artificial Intelligence Approach*, Volume 2, Morgan-Kaufmann 1986, pp. 45-61.