Title
Integrating explanation-based and empirical learning methods in OCCAM

Author
Pazzani, Michael J.

Publication Date
1988-10-18

Peer reviewed

Integrating Explanation-based and Empirical Learning Methods in OCCAM

Michael J. Pazzani

Department of Information & Computer Science
University of California, Irvine, CA 92717

Technical Report 88–23

October 18, 1988

To appear in *Proceedings of the Third European Working Session on Learning*, 1988.

Integrating Explanation-based and Empirical Learning Methods in OCCAM

Michael J. Pazzani
Department of Information and Computer Science
University of California
Irvine, CA 92717 USA
pazzani@ics.uci.edu

Abstract

This paper discusses an approach to integrating empirical and explanation-based learning techniques. The paper focuses on OCCAM, a program that has the capability to acquire via empirical means the knowledge needed for analytical learning. Two examples of this capability are discussed:

- The ability to use empirical techniques to acquire a domain theory for explanation-based learning.
- The ability to use empirical learning techniques to find common patterns for causal relationships. These patterns encode a theory of causality (i.e., a set of general principles for recognizing causal relationships). Once acquired, a theory of causality can facilitate later learning by focusing on hypotheses which are consistent with the theory.

1. Introduction

I present a theory of learning to predict and explain the outcome of events. This theory is implemented in a computer program called OCCAM. OCCAM is able to learn to predict and explain possible outcomes in a variety of domains from simple physical causality (breaking glass and inflating balloons) to more complex events (kidnapping and economic sanctions). OCCAM acquires a hierarchy of explanatory schemata similar to Schank's MOPS (Schank, 1982, Lebowitz, 1980, Kolodner, 1984). The hierarchy of schemata serves as a discrimination net for making predictions and finding explanations.

OCCAM can utilize three sources of information when acquiring new schemata. First, OCCAM uses *inter-example relationships*, regularities among a number of examples that reveal the conditions under which a cause produces an effect. The empirical learning component of OCCAM makes use of this source of information.

IF an action $A1$ on an object O occurs before a subsequent action $A2$ precedes a state change $S2$ of O
THEN $A1$ results in a state $S1$ which enables $A2$ to result in $S2$

Figure 1: A causal pattern used by the theory-driven learning module of OCCAM.

Second, OCCAM makes use of *intra-example relationships*, temporal and spatial relationships between a cause and an effect which constrain the search for a hypothesis. The intra-example relationships in OCCAM consist of a set of patterns which suggest an explanation (in terms of a causal mechanism) for an effect. One such causal pattern is shown in Figure 1. This pattern matches the situation in which a person shakes a bottle of soda ($A1$) and then opens the bottle ($A2$), and the soda sprays out. In this situation, the pattern suggest that shaking the bottle results in some intermediate state which enables the opening of the bottle to result in the soda spraying out. The pattern indicate what actions and components of the actions might play a part in the causal relationship. For example, since the pattern does not mention the actor of any action, it encodes the assumption that the actor (as well as the actor's hair color) is irrelevant. However, since the pattern mentions the object of the action (in this example, the bottle of soda), the type of object can affect the outcome. The hypothesis that *soda will shoot out of a bottle of carbonated soda if a blond person shakes it before opening* might be consistent with a number of training examples (and postulated by an empirical learning program). However, by requiring that the hypothesis be consistent with the data and the theory of causality as represented by the pattern in Figure 1, the theory-driven learning component of OCCAM will postulate the hypothesis that *soda will shoot out of a bottle of carbonated soda if the bottle is shaken before opening.* The theory-driven learning (TDL) component of OCCAM is discussed more fully in (Pazzani, 1987). In section 4, I discuss an extension to OCCAM that acquires via empirical learning the causal patterns used by TDL.

Finally, OCCAM can also use *domain knowledge*, prior knowledge which predicts and explains regularities in events. OCCAM utilizes prior knowledge in its explanation-based learning (or SBL) module.

OCCAM always applies the most knowledge-intensive learning strategy which is applicable to a problem. First, EBL is attempted. If there is not enough knowledge to produce an explanation, theory-driven learning is attempted. If the example does not match a known causal pattern, the empirical methods are attempted. There are two rationales for integrating the learning techniques in this manner. First, the more knowledge-intensive strategies have stronger justifications. The information encoded in schemata can be accessed to make predictions about future events. Therefore, a schema should only contain features that an understander has a justification for believing will appear in future events. EBL demonstrates deductively that a set of features are sufficient to produce the predicted outcome. TDL has a general theory of what configurations of events might be causally related. Correlations that are not consistent with the theory of causality can justifiably be treated as coincidences and ignored. The justification that empirical techniques use for including a feature in a schema is that the feature has always appeared in previous events. Although there has been considerable progress recently in determining the accuracy of empirical techniques (Valiant, 1984, Etzioni, 1988), Hume's problem has not been solved: there is no logical justification for predicting a future occurrence from past observations (Hume, 1739).

The second rationale for preferring knowledge-intensive learning strategies over data-intensive strategies is that the design of OCCAM has been inspired by findings in cognitive and developmental psychology. These findings indicate that people exhibit this same preference. See (Pazzani, Dyer & Flowers, 1986) for a summary of these findings.

In the remainder of this paper, I first discuss some frameworks for integrated learning. Next, I discuss the acquisition of a domain theory for explanation-based learning in OCCAM. Finally, I show how a theory of causality can be acquired by empirical means.

2. Integrated approaches to learning

Purely empirical or purely analytical (e.g., explanation-based) approaches both fail as general theories of learning. Explanation-based learning methods (e.g., (Silver, 1986, Mitchell, Kedar-Cabelli & Keller, 1986, DeJong, 1986, Minton, Carbonell, Etzioni, Knoblock & Kuokka, 1987)), rely on prior knowledge and cannot fully account for how this prior knowledge is acquired. Empirical learning methods (e.g., (Mitchell, 1982, Michalski, 1977, Holland, Holyoak, Nisbett, and Thagard., 1986)) cannot account for a number of psychological studies (e.g., (Murphy and Medin, 1985, Nisbett & Ross, 1978)) which attribute differences in learning rates or learning accuracy to differences in prior knowledge. Even if one is not interested in psychological modeling, a major shortcoming of empirical learning methods is that they require a large number of training examples so that coincidental regularities are unlikely. Integrated approaches to learning, which combine empirical and explanation-based learning methods have the potential of overcoming the inadequacies of either method applied individually.

OCCAM incrementally forms a concept hierarchy that explains and organizes previous experiences. There are two tasks which must be accomplished:

- Aggregation-- experiences are grouped into clusters of related events (Fisher & Langley, 1985).
- Generalization-- A general description is created for a cluster of events.

Various combinations of explanation-based and empirical methods prefer different sources of information for each of these tasks. For example, one integrated strategy is to first use empirical means to form a generalization, and then use EBL to verify that the generalization is consistent with existing knowledge. Those parts of the generalization which are not supported by the existing knowledge are to be discarded. The flowchart in Figure 2 illustrates this approach. This strategy is employed by the UNIMEM program (Lebowitz, 1986a, Lebowitz, 1986b).

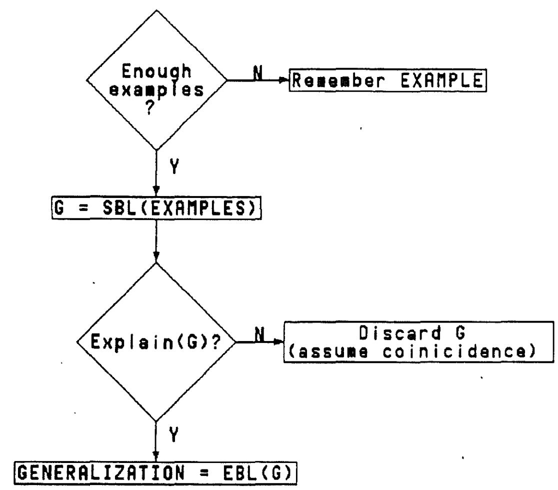

Figure 2: Preferring empirical learning to explanation-based learning.

This strategy uses a syntactic clustering algorithm for the aggregation of events into useful clusters. The clustering algorithm may group together events which have different structures if the events share many surface similarities. A syntactic clustering algorithm may ignore important, meaningful regularities in the explanation structure. For example, consider the following three economic sanctions incidents:

Economic-Sanction-1

In 1983, Australia refused to sell uranium to France, unless France ceased nuclear testing in the South Pacific. France paid a higher price to buy uranium from South Africa and continued nuclear testing.

Economic-Sanction-2

In 1980, the US refused to sell grain to the Soviet Union unless the Soviet Union withdrew troops from Afghanistan. The Soviet Union paid a higher price to buy grain from Argentina and did not withdraw from Afghanistan.

Economic-Sanction-3

In 1961, the Soviet Union refused to sell grain to Albania if Albania did not rescind economic ties with China. Albania continued the ties with China, and China gave Albania a discount on wheat imported from Canada.

Into what clusters should these three events be divided? If all three were included in one cluster an important regularity concerning the price of the commodity would be ignored since the price varies in the three events. If the events are divided into two clusters by syntactic means, then Economic-Sanction-2 and Economic-Sanction-3 would most probably be grouped together since these have the most features in common. For example, in both events, the commodity in dispute is grain, the <u>source</u> of the threat is a superpower, and the <u>target</u> of the threat is a communist country. However, the explanation of these two events is entirely different. In Economic-Sanction-2, Argentina sold the grain to make a large profit (Brown, 1985). In Economic-Sanction-3, China sold the grain at a loss to gain political influence (Freedman, 1970). In the other two case, the motive of country which sold the commodity was to profit monetarily rather than politically from the sale (Hufbauer & Schott, 1985, Brown, 1985). If events are clustered according to the explanation for their outcome, then Economic-Sanction-2 should be grouped with Economic-Sanction-1 since they share the same explanation structure. A syntactic clustering algorithm followed by EBL of empirical generalization can prevent the EBL algorithm from detecting, explaining and generalizing meaningful regularities.

The problem arises because a syntactic clustering algorithm groups instances according the number of features they have in common. There is no guarantee that these clusters will share the same explanation structure. In general, the clusters found by a syntactic clustering algorithm will not share the same explanation structure when "relevant" similarities are uncorrelated with the "coincidental" similarities. A system such as OCCAM which prefers EBL to empirical methods will not have this problem provided it has enough domain knowledge to produce an explanation.

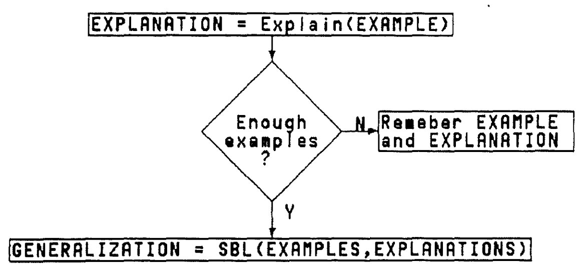

Figure 3: Empirically combining explanation structures.

Another strategy to combine existing knowledge and correlational information is to use existing knowledge to produce an explanation. However, rather than using EBL for the generalization, empirical learning methods can find regularities between explanations. The flowchart in Figure 3 illustrates this approach. Such a strategy has been implemented in the WYL program (Flann & Dietterich, 1986) which performs inductive generalization on explanation structures. A similar strategy is a component of Purpose-Directed Analogy (Kedar-Cabelli, 1985) and part of DISCIPLE (Kodratoff & Tecuci, 1987). Since this strategy uses an empirical generalization technique, irrelevant coincidental information is likely to appear in the generalized explanations. For example, the WYL system learns a concept of a *trap* in checkers. WYL contains an explicit bias that indicates that no regularities are coincidental. If it is presented with a set of training examples with a coincidental regularity in which a red man always traps a white man, it cannot determine the correct color relationship between the men in a *trap*.

2.1. Empirically acquiring a domain theory for EBL

The final alternative I consider is the one implemented in OCCAM. In this approach, EBL is applied if applicable and empirical learning is used as a last resort. This approach is illustrated in Figure 4.

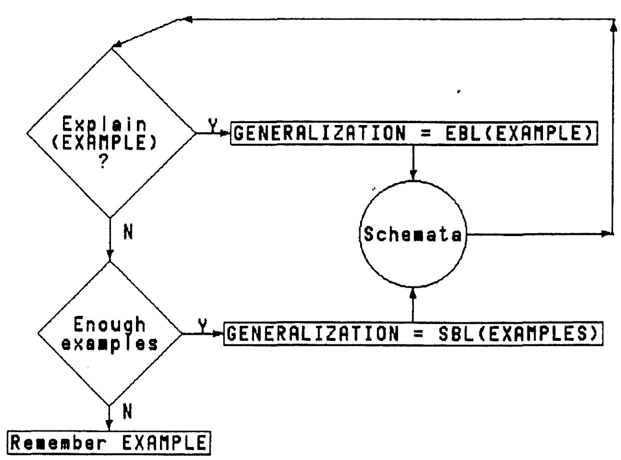

Figure 4: Preferring explanation-based to empirical learning. Schemata are formed by empirical or explanation-based learning and serve as background knowledge for explanation-based learning.

The primary benefit of this approach is that the knowledge necessary to perform EBL can be acquired by the learning system. In OCCAM, EBL is preferred, but if there is not enough knowledge to produce an explanation, empirical techniques can acquire this knowledge. For example, by empirical techniques OCCAM learns that parents have a goal of preserving the health of their children. This social knowledge provided an explanation for a parent paying the ransom in kidnapping, which enables OCCAM to create a kidnapping schema by explanation-based learning techniques (see Section 3.4). Another example of OCCAM acquiring knowledge for analytical learning is described in Section 4. It is important to realize that OCCAM is much more than a switch which decides when to run each type of learning program. The separate parts of OCCAM cooperate by utilizing the same memory for learning and explanation. Michalski calls this type of system "closed-loop" learning (Michalski, 1987).

3. OCCAM: Acquiring a domain theory for EBL

The primary advantage of the approach to learning implemented in OCCAM is that the learner gets better at learning. At first, the learner relies on empirical and theory-driven learning. As a consequence, learning is slow and requires a large number of examples to rule out incorrect hypotheses which are consistent with a small number of initial observations. However, the knowledge which the learner acquires through these data-intensive mechanisms enables later knowledge-intensive learning. In this section, I discuss the knowledge representation in OCCAM and I demonstrate how OCCAM learns information to produce an explanation of how the kidnapper's goal is achieved in a kidnapping episode.

3.1. Knowledge representation in OCCAM

Conceptual Dependency (CD) (Schank, 1977) structures are used to represent goals, plans, states and simple actions. A complex situation, such as an economic sanction incident, that consists of several events are represented by a number of CD structures. The relationships between these structure are represented by causal or intentional links such as **enables** or **achieves** (Dyer, 1983).

A schema is comprised of several components. These components are:

A generalized event-- A schema includes a template for recognizing instances of the schema. This template is a Conceptual Dependency pattern which is an abstract description of a class of events.

A causal chain-- Typically, a schema can be decomposed into a small set of actions or scenes. For example, a schema that represents coercion can be broken down into a number of simpler components that represent a threat, a demand and an outcome. An abstract description of these actions and their interaction is an integral part of the schema. The relationship between these events are specified by causal, temporal and intentional links. The generalized event of a schema indicates *what* the effect of a certain action is. The causal chain indicates *how* the action brings about the effect.

Indices--
A schema may organize more specialized schemata and specific instances. These are indexed by features which elaborate on the generalized event of the schema. An index consists of a feature name (e.g., **actor**) and a feature value (e.g., **human**).

Support--
A schema also includes information about the source and the confidence in the schema. This information can include the justification for the schema (if the schema was created by generalizing an explanation) or the number of successful and unsuccessful times a schema has made a prediction (if the schema was created by empirical methods).

OCCAM starts with an initial hierarchy of schemata which represent the Conceptual Dependency actions, goals and states. As OCCAM learns the hierarchy is extended by creating specializations of the existing schema. The more specialized schemata are treated identically to the initial schemata.

Occasionally, I will refer to schemata as "rules". Schemata that OCCAM acquires that have a causal chain consisting of exactly two events connected by one intentional link, can be treated as backward chaining rules for making inferences. For example, OCCAM learns a rule that states that when a glass object is dropped, then the object shatters. This is represented as a specialization of the **propel** schema. The causal chain for this schema is shown in Figure 13.

One additional facet of OCCAM's representation system may need further explaining. The feature values of the generalized event of a schema can contain *role tokens* which are used to indicate that the feature value must be identical to another feature value. For example, OCCAM acquires a specialization of **coerce** which represents a common pattern of economic sanction incidents: one country that exports a commodity threatens a country with a strong economy that imports the commodity by refusing to sell them the commodity. A response to this threat is to purchase the commodity at a higher price from another country. Figure 5 shows the generalized event of this schema. This schema contains several role tokens. For example, the object of the threat must be the same as the object exported by the actor and imported by the target.

```
COERCE actor POLITY exports =OBJECT
       object COMMODITY
       target POLITY economic-health STRONG
                     imports =OBJECT
       threat ACT type SELL
                  actor =ACTOR
                  object =OBJECT
                  to =TARGET
                  mode NEG
       response ACT type SELL
                    actor POLITY business-relationship =TARGET
                                 exports =OBJECT
                    object =OBJECT
                    price MONEY value >MARKET
                    to =TARGET
       result STATE type POSSESS
                    actor =TARGET
                    value YES
                    object =OBJECT
```

Figure 5: An economic sanctions schema produced by OCCAM.

This economic sanction schema also serves as an example of how a schema can be used for prediction and for explanation. The generalized event indicates the result of this class of situations. The causal chain indicates why this result occurs (see Figure 6).

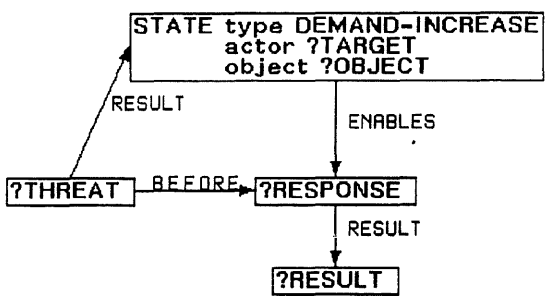

Figure 6: The causal chain of an economic sanction schema. This schema indicates that a threat results in an increased demand for a product enables another country to sell the product at an inflated price.

3.1.1. Learning and explanation processes

OCCAM has two different processes that it can use to explain events. The first process for producing explanations is a recognize and instantiate strategy. This strategy explains a new event by recognizing the event as an instance of a known schema and instantiating a causal chain associated with the schema. Recognition is a search process through a memory of hierarchically organized schemata. Search starts at the most general schema in memory and follows a depth-first search to find the most specific applicable schema (see (Lebowitz, 1987)). An example of OCCAM's hierarchically memory is illustrated in Figure 15.

The second process for producing explanations is a constructive strategy which builds new explanations by chaining together several schemata. The construction of explanations is guided by abstract knowledge of the class of causal explanations represented as causal patterns. For example, the explanation for why a state change occurs can be that the state change is the result of an action; the explanation for why an actor performs an action after a certain event occurs can be that the action will fulfill a goal for the actor which was motivated by the event.

There are two different ways in which learning can improve the explanation process. First, learning can increase the number of situations that one can explain by a recognition process. Explanation-based learning performs exactly this task by exploiting the interactions among existing knowledge and creating a schema which recognizes the class of situations in which these interactions will occur. Second, learning can increase the number of situations that one can explain by a chaining process. Empirical techniques can detect regularities in data and hypothesize that these regularities will hold in future cases. New schemata created by empirical techniques increase the ability of a learner to construct new explanations via chaining as well as the ability to recognize new explanations. An important point of this research is that when new knowledge is acquired by empirical methods, new explanations can be constructed via chaining. When new explanations are constructed, they can be generalized by explanation-based techniques so that the situations in which these explanations apply can be recognized. Empirical learning techniques provide the necessary background knowledge for explanation-based learning. EBL alone is not sufficient for increasing the number of situations that one can explain by constructing explanations (Dietterich, 1986).

3.1.2. The process of explanation-based learning

The process of explanation-based learning in OCCAM consists of the following steps:

1. Match the situation against the set of causal patterns which can suggest an abstract explanation.
2. Verify and refine the abstract explanation with specific domain knowledge. The abstract explanation can also be denied by existing knowledge.
3. A general description of the situations in which the explanation will apply is constructed by retaining only those features of the example which were needed to produce the explanation. This general description serves as the generalized event of the new schema.
4. The generalized explanation is saved as the causal chain of the schema. This can be instantiated to explain the prediction made by the generalized event.

EBL can fail if the example does not match any causal pattern or if there is not sufficient domain knowledge to verify the abstract explanation.

3.1.3. The process of theory-driven learning

In OCCAM, theory-driven learning is attempted only when EBL fails. The process of theory-driven learning consists of the following steps:

1. Match the situation against the set of causal patterns which can suggest an abstract explanation. The patterns are ordered by simplicity[1]. If more than one pattern matches, the simplest is used.
2. Instantiate the consequent of the causal pattern by replacing the variables in the pattern by the most specific conjunctive generalization of the components of each event which has the same outcome. If some events match the pattern, but have a different outcome, the components are further specialized to account for the different outcome.
3. The instantiated antecedent serves as the generalized event of a new schema. The schema is indexed in memory by the generalized event.
4. The instantiated consequent is connected to the instantiated antecedent and serves as the causal chain for the new schema.

In TDL, a causal pattern which matches a training example proposes a hypothesis. A proposed hypothesis will be tested against new

[1]This is why I named the system OCCAM. The simplest causal pattern produces the simplest hypothesis.

data and either accepted or rejected depending upon the accuracy of the hypothesis. Theory-driven learning can be viewed as a form of explanation-based learning in which the domain knowledge (i.e., the set of causal patterns) is known to be overly general. Since the domain theory can propose hypotheses which are not true, these hypotheses are evaluated against further examples. A schema constructed by TDL contains a counter that is incremented during memory search when a successful prediction is made, and another counter that is incremented when an incorrect prediction is made. When the ratio of these counters is lower than a certain value[2], then the schema is eliminated. Typically, when a schema is eliminated, either a more specialized version is created by the same causal pattern, or a more complex causal pattern creates an alternative explanation.

3.1.4. The process of empirical learning

Empirical learning is attempted only in EBL and TDL cannot be applied. The empirical learning component creates clusters of events which share the largest number of features and creates a generalized event for these clusters by finding the most specific conjunctive generalization which accounts for the examples. Schemata formed in this manner are subject to revision when further examples are seen. See (Pazzani, 1988a) for a more complete description of this component.

[2]This is a parameter in OCCAM. The current value of the parameter is 0.95.

3.2. Learning about preservation goals

OCCAM acquires several rules which are needed to explain why the ransom is paid by the kidnapper. OCCAM must acquire one rule which indicates that a certain class of persons have a goal of preserving the health of another class of persons. This rule will explain why the target pays the ransom, since paying the ransom in kidnapping is a means of preserving the health of the hostage.

OCCAM is presented with the CD representation of the following example:

- Lynn is playing on the swing and she falls off and scuffs her knee. Her mother, Chris, gets a band-aid and puts it on her knee. Her neighbor, Tiffany, gets on the swing and rides it.

The CD representation for this event includes a listing of the attributes of each person involved. The attributes of Chris include her family relationships (mother of Lynn), age (adult), hair color (brown) and height (tall). The attributes of Tiffany include age (child), hair color (blonde) and height (short).

```
IF an event E motivates a goal G for P
and H observes the event E
and performs an action A
which achieves G for P,
THEN E motivates the goal G for H
```

Figure 7: A causal pattern for inferring the motivation for an action

OCCAM contains a causal pattern to deal with this type of situation (see Figure 7). In this example, OCCAM postulates that a difference between Chris and Tiffany is responsible for the different goal and, therefore, the different response. There are a number of differences between Chris and Tiffany (e.g., age, height, family relationship, hair color). Without any knowledge to favor one feature over another, OCCAM selects one feature at random: height. OCCAM constructs a rule which indicates that when a tall person observes an action which motivates a goal of preserving the health of another person, then the tall person will also have a goal of preserving the health of that person.

OCCAM contains a mechanism to evaluate its knowledge (see (Pazzani, 1987) for more details). First, OCCAM must detect that an inference rule is making an incorrect prediction. If the rule was formed with empirical techniques, each incorrect prediction reduces confidence in the rule until it is finally eliminated. Then, a new rule is created which is consistent with the data. After OCCAM creates an incorrect inference rule from the previous example, it is presented with the following example:

- Lynn is playing on the monkey bars and she falls off and scuffs her elbow. Tiffany's mother, Loreli, who is eating an ice cream near the monkey bars does not help.

The attributes of Loreli include her family relationships (mother of Tiffany), age (adult), hair color (blond) and height (tall). Since there was very little support for the inference rule which predicts that Loreli will help because she is tall, the rule is abandoned.

In the current example, after abandoning the hypothesis that tall people have the goal of preserving the health of others, OCCAM must come up with a new hypothesis. Unfortunately, the new hypothesis is not much better than the first. One difference between those persons who helped and those who didn't is that Chris has brown hair while Loreli and Tiffany both have blond hair. Once again, this illustrates the hazards of guessing: one is likely to guess wrong.

Finally, OCCAM is presented with another example which allows it to formulate a correct hypothesis:

- Karen falls off her bike and bruises her lip. Her sister, Lynn, gets an ice cube to put on Karen's lip.

In this example, since the person who helped (Lynn) has blond hair, the inference rule which indicates that blonds will not help must be discarded. OCCAM finds another difference between those people who helped and those who did not. In all the cases that a person helped, they were related to the person who was injured. OCCAM creates a new rule which indicates that members of the same family have a goal of preserving the health of other family members.

3.3. Learning that "give" requires "have"

OCCAM acquires another rule which indicates that possessing an object is an enabling condition for giving someone an object. This rule explains how the target is able to pay the ransom in kidnapping.

By pure similarity-based learning, OCCAM acquires a `delta-agency` (Schank, 1977) schema. `delta-agency` is a particular plan for achieving a goal by asking another agent to perform an action which achieves the goal. From the following examples, OCCAM acquires a specialized version of `delta-agency`:

- `pizza-1`: Karen has a goal of possessing a slice of pizza. Her plan is to ask her father, Mike for a slice of pizza. Her goal succeeds.
- `zoo-1`: Karen wants to go to the zoo. She asks Mike to take her to the zoo. Her goal succeeds.
- `play-doh-1`: Lynn wants some Play Doh. She asks Mike to give

her some, and her goal succeeds.

- **apple-1**: Karen wants an apple. She asks her mother Chris for one and Chris gives her one.
- **oil-1**: Chris wants some peanut oil. She asks her husband, Mike, who gets her some peanut oil from the store.

The SBL component of OCCAM clusters these events together and creates a schema that indicates that when a person asks a relative to do something, then the relative will. What happens, for example, when a person asks a relative for an object which the relative doesn't have? For example, what should OCCAM do when it encounters **apple-2**?.

- **apple-2**: Karen wants an apple. She asks her mother Chris for one and Chris tells her that she doesn't have an apple.

The CD representation of **apple-2** indicates that the goal is blocked by Chris not having the apple. OCCAM has a causal pattern to deal with the situation in which a goal is blocked. This causal pattern is:

```
If a goal to perform an action
is blocked by a state
THEN the opposite state is an
     enabling condition for the action
```

Instead of simply decreasing support for **delta-agency**, OCCAM is able to come up with a hypothesis for why **delta-agency** fails. The causal pattern suggests a possible explanation: in order to give an object to someone, you must possess the object. Note that this explanation is dependent on the mother informing the child what state is blocking her goal. If the mother just didn't give the child an apple, then OCCAM would not be able to come up with an explanation and would decrease support for **delta-agency**. The inductive leap that OCCAM makes is that the opposite of state[3] which is blocking the goal is an enabling condition of performing the action which achieves the goal. Figure 8 illustrates the rule that OCCAM learns in this situation.

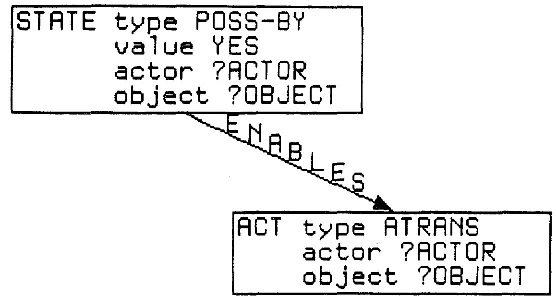

Figure 8: A rule acquired by OCCAM when an exception to delta-agency is encountered: in order to give an object to someone, you must possess the object.

This example also emphasizes an important design consideration of OCCAM. Wherever possible, OCCAM uses the most knowledge-intensive mechanism possible. In learning, it prefers explanation-based learning, followed by learning with a general theory of causality. Similarity-based learning is used only as a last resort. Similarly, when an expectation fails, OCCAM first tries to explain why it failed. Only if an explanation cannot be found, does OCCAM resort to statistical means of evaluating its knowledge.

3.4. Explanation-based learning: An example

Once OCCAM has acquired these rules, it is in a position to learn a kidnapping schema. OCCAM is presented with the CD representation

[3] i.e., the state which is blocking the goal in this example is that Chris does not have an apple. The opposite of this state, that Chris have an apple is required for Chris to give Karen the apple.

of Kidnapping-1:

Kidnapping-1

John, a 10-year-old child was abducted on his way to church on Sunday morning. His father, Richard, received a phone call that evening. The kidnapper threatened that John would be killed unless Richard paid a $100,000 ransom. Monday at noon, Richard left the money in a locker at the train station.

The first problem for OCCAM is to explain how the kidnapper's goal of obtaining money was achieved. The following explanation chain is constructed:

1. The goal was achieved when the target of the threat paid the ransom.
2. The target possessing the money (i.e., being rich) is a state which enables the payment.
3. The threat to kill the child motivates a goal to preserve the health of the child (because the father is related to the hostage).
4. The demand motivates a goal for the target to retain his wealth.
5. The action of paying the ransom is a realization of a plan to prefer the achievement of the more important goal: preserving the child's life.

Knowledge acquired by empirical means is used by OCCAM to infer the second and third links in this explanation chain. Once OCCAM explains the target's action, it can now generalize the plan it observed for obtaining money. The example is generalized by removing those features from the original example which were not needed to produce the explanation. For example, the target is required to be rich, so that he can afford the ransom, and the target is required to be related to the hostage so that he is willing to pay the ransom. On the other hand, the age of the hostage or the hostage's destination do not play a part in the explanation and are not included in the generalization. The generalization which OCCAM constructs is illustrated in Figure 9.

```
COERCE the-threat ACT type KILL
                      actor =THE-ACTOR
                      object =THREAT-OBJ
       the-actor HUMAN
       threat-obj HUMAN
       the-target HUMAN relation IPT type FAMILY-REL
                                     of =THREAT-OBJ
                        income-class RICH
       ... ...
```

Figure 9: Part of the kidnapping schema formed from generalizing Kidnapping-1. Note that the target is required to have an interpersonal relationship with the hostage (i.e., the `threat-obj`) so that he is willing to pay the ransom and that the target is required to be rich so that hew can afford to pay the ransom.

It is instructive to find out what happens if OCCAM is presented with complex examples such as kidnapping before it has acquired a background theory. If there is not sufficient knowledge to construct an explanation, OCCAM resorts to pure similarity-based learning and looks for regularities among examples. The initial attempts at SBL will produce schemata that are overly specific. Incorrectly classified examples will cause the initial schemata to be deleted. If at any time, OCCAM has acquired the sufficient background knowledge, EBL will be used to create a schema to organize the events. If EBL fails, SBL will created a more specialized schema.

If OCCAM has an incorrect theory (e.g., if it is presented with a kidnapping example when it has a rule which states that tall people want to preserve the health of others), then it will learn an incorrect schema with explanation-based learning. This occurs because the explanation is incorrect. When further examples force OCCAM to revise the underlying theory, the schemata formed by explanation-based learning with an incorrect theory will also be revised (Pazzani, 1988b).

3.4.1. Specializations of kidnapping

Once OCCAM has formed a kidnapping schema, it is ready to learn about some specializations of kidnapping. Since kidnapping is a complex event, there are many goals involved. For example, in addition to the central goal of the kidnapper (to obtain money) and the target (to ensure the safety of the hostage), the kidnapper also wants to avoid going to jail, the hostage wants to remain alive, the police want to arrest the kidnapper, etc. The specializations of kidnapping will focus on the features of the various agents which determine the outcome of these subordinate goals.

For example, OCCAM forms a specialization of kidnapping when it is presented with the following episode (Alix, 1978):

Kidnapping-2

In May 1933, Mary McElroy, twenty-five-year-old daughter of the city manager of Kansas City, Missouri was abducted. The abductors demanded $60,000 for her safe return. They accepted a $30,000 ransom and released the hostage unharmed from a farm in Kansas where she had been held for twenty-nine hours. The kidnappers were arrested by the FBI. The testimony of the victim was largely responsible for their conviction. The kidnappers received a sentence of life in jail.

In this episode, the kidnappers' goal of preserving their freedom was thwarted when they received the punishment of life in jail. To create a specialized kidnapping schema, OCCAM must identify the circumstances which led to this goal failure. A possible explanation is suggested by the following causal pattern: *if a preservation goal is thwarted after an action which is needed to perform a plan which achieves a goal, then the action results in a state which enables the preservation goal to fail.* This causal pattern suggests that abducting the hostage results in a state which enables the conviction of the kidnappers. OCCAM's domain knowledge is needed to complete the explanation. The complete explanation indicates that abducting the hostage results in the hostage seeing the kidnapper which enables the hostage to testify against the kidnapper. OCCAM generalizes this explanation and uncovers an inherent flaw in kidnapping: the hostage sees the kidnapper when he is abducted and can testify against the kidnapper. A new schema is created and indexed in memory under the kidnapping schema. The explanation is saved as the causal chain for the specialized kidnapping schema (see Figure 10).

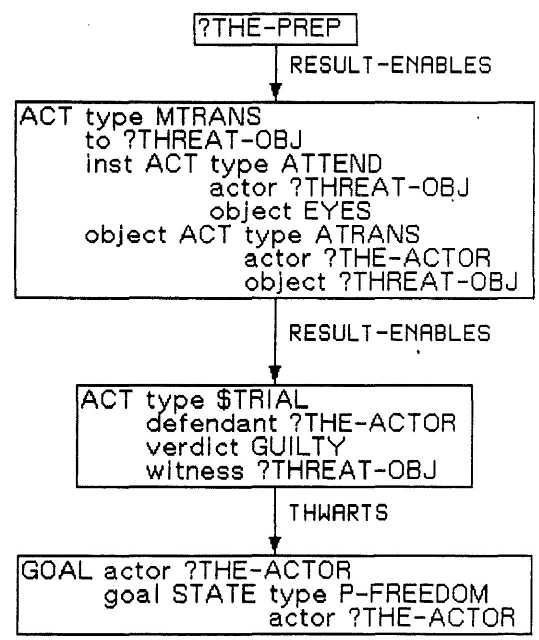

Figure 10: Part of the generalized explanation stored with a specialization of kidnapping. Since the hostage sees the kidnapper during the abduction, the hostage can testify against the kidnapper.

Another kidnapping episode results in a different specialization of kidnapping that avoids the problem of the previous incident (Moorehead, 1980):

Kidnapping-3

On June 2, 1920, Blakely Coughlin, the thirteen-month-old son of a wealthy Pennsylvania family vanished from his bedroom. A ladder was found abandoned near the window to the nursery. Several nights later, a ransom letter arrived and instructed Mr. Coughlin to throw $12,000 from a moving train when he saw a white flag being waved.

When this kidnapping episode is added to memory, a causal pattern suggests an explanation for the selection of the hostage: *if a preparation is performed on an object, look for other schemata which have a goal failure. Postulate the preparation avoids the goal failure.* OCCAM searches memory and finds the specialization of kidnapping in which the kidnapper is convicted by the testimony of the hostage. Since the hostage does not testify in this case, the causal pattern suggests that this particular victim was chosen to avoid the goal failure. OCCAM next tries to determine if the hostage in this episode would be able to testify against the victim. However, the explanation which worked in the previous case will not work in this case because the hostage is an infant. Therefore, OCCAM constructs an explanation which indicates that the kidnapper selected this particular hostage as a plan to avoid the failure of the kidnapper's goal to preserve his freedom. The generalized event which OCCAM constructs for this situation is illustrated in Figure 11.

The new kidnapping schema is indexed under the kidnapping schema by the age of the hostage and the preparation (i.e., abducting the hostage) since these are the only features needed to construct the explanation. Alternative examples might focus on other reasons that the hostage might not be able to testify by interfering with other locations in the causal chain. For example, by killing the hostage the kidnapper can prevent the hostage from testifying as well as preventing the hostage from assisting the police by providing information which might lead to the kidnapper's capture.

```
COERCE threat-obj HUMAN age INFANT
       the-prep ACT type ATRANS
                    actor =THE-ACTOR
                    to =THE-ACTOR
                    object =THREAT-OBJ
```

Figure 11: A specialized version of kidnapping which avoids a potential problem with kidnapping by selecting an infant as the hostage.

4. Empirical acquisition of causal patterns

In previous versions of OCCAM there was a fixed set of causal patterns which never changed as the program learns. When the program starts, it has its complete theory of causality. While there is evidence that very young infants are able to perceive causal relationships (Leslie & Keeble, 1987), there is no question that older children are better at attributing causality than younger children (Piaget, 1930, Bullock, 1979). It might be better if OCCAM started with a few very simple causal patterns, and learned the more complex ones. Certainly, I would not want to claim that the more complex causal patterns such as the one in Figure 7 are innate.

A general theory of causality can be acquired empirically by noticing common patterns of causal relationships. For example, OCCAM can start with a very simple theory of causality with only one causal pattern:

IF an action A
precedes a state change S
THEN A results in S

This causal pattern is necessary to warrant the inference of a causal relationship from a temporal relationship. This simple causal pattern doesn't contain any constraints between causes and effect. Although, it will make some correct causal inferences, it also allows a number of mistakes. For example, if the cat meows shortly before the doorbell rings, the inference that the cat caused the doorbell to ring will be made. Eventually, with further examples of the doorbell ringing without the cat meowing and the cat meowing without the doorbell ringing, this mistake will become apparent. Other proposed causal relationships will be confirmed after a large number of examples have been observed. Similarities can be detected between the examples in which the simple causal pattern successfully proposes causal relationships and the causal pattern could be specialized.

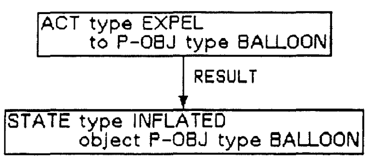

Figure 12: A confirmed causal relationship: A balloon is inflated when air is blown into it.

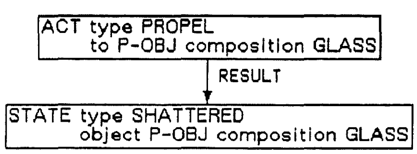

Figure 13: A glass object shatters when it is struck.

To illustrate how this scheme works consider the following example. A causal relationship is noticed about balloons: when air is blown into balloons, they get bigger. This relationship is shown in Figure 12. Other examples suggest another causal relationship: when a glass object is struck, it shatters. This relationship is illustrated in Figure 13.

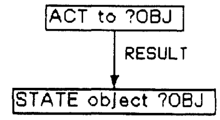

Figure 14: A causal pattern formed by detecting the common features of confirmed causal relationships.

By noticing the common features between these two generalizations, a new causal pattern is created which is illustrated in Figure 14. This generalization provides an additional constraint which facilitates the detection of future causal relationships by preferring those relationships in which the object that changes was acted upon by the action:

IF an action A on an object O
precedes a state change S of O
THEN A results in S

The ability to acquire new causal patterns to guide generalization enables OCCAM to adapt to new domains. These causal patterns suggests causal mechanisms for state changes. One would not want to claim that people are born with a theory of electric switches. However, as adults we are more likely to attribute a change in an electronic device to a pushing of a button or the flicking of a switch than to an other random action (such as a cat meowing). This is true even if the wires are hidden (as in a light switch) or the connection is not observable (as in the remote control for a television). This knowledge could be acquired by noticing similarities among the control of electrical devices and could be represented as a new causal pattern:

```
IF a switch is pressed A
immediately before a change of state
of an electronic device S
THEN A results in S
```

In some respects, OCCAM's use of causal patterns to come up with an explanation is similar to SWALE's explanation patterns (Schank, 1986, Leake & Owens, 1986). However, there are many more explanation patterns in SWALE than there are causal patterns in OCCAM. In addition, the explanation patterns in SWALE can be quite specific, referring to particular goals, actions and attributes. For example, SWALE contains the following explanation pattern:

```
Being a star performer
can result in stress.
Taking drugs can relieve stress.
Taking too much drugs
can result in death.
```

In contrast, OCCAM's causal patterns are much more general and refer to relationships between goals and actions. The following causal pattern would suggest an explanation for situations similar to the one handled by SWALE's explanation pattern:

```
An action that achieves a goal
can result in a side effect.
A plan can achieve the goal
motivated by the side effect.
Executing the plan can result
in the failure of another goal.
```

A difference between SWALE's use of explanation patterns and OCCAM's use of causal patterns is the interpretation process. The explanations produced by SWALE are too specific and must be tweaked (i.e. modified) to apply in future situations. In contrast, the explanations produced by OCCAM's causal patterns are too general and must be refined with additional knowledge. One way to reconcile these two approaches in the future might be to create a hierarchy of explanation patterns. OCCAM's causal patterns would serve as general nodes in the hierarchy and the more specific explanation patterns would be indexed in memory under the causal patterns.

Figure 16 shows the flow of information when causal patterns are acquired via empirical techniques. Background knowledge for explanation-based learning is acquired through similarity-based or theory-driven learning. Causal patterns are learned by similarity-based learning and used by theory-driven learning.

5. Conclusion

In this paper, I have argued that both empirical and explanation-based learning techniques are necessary components of a system that learns causal relationships. The strength of empirical techniques is to acquire background knowledge such as simple causal rules by noticing regularities in observed data. The strength of explanation-based techniques is to recognize new interactions among existing knowledge and determine the class of situations in which the interactions occur. In this paper, I provided two examples which demonstrate how the complementary nature of these learning processes is exploited by OCCAM. As a consequence of its architecture, as OCCAM learns, it acquires knowledge which facilities future learning.

Acknowledgments

This research was performed while the author was in the Artificial Intelligence Laboratory at the University of California, Los Angeles and was supported in part by the UCLA-RAND Artificial Intelligence Fellowship. I am grateful to Professor Michael Dyer for many

fruitful discussions during the development of OCCAM.

References

Alix, E. *Ransom kidnapping in the United States, 1874-1974: The creation of a capital crime*. Southern Illinios University Press, 1978.

Brown, M.L. *Economic sanctions: A cost-benefit approach.* Doctoral dissertation, Harvard University, 1985.

Bullock, M. *Aspects of the young child's theory of causality.* Doctoral dissertation, University of Pennsylvania, 1979.

DeJong, G. An approach to learning from observation. In Michalski, R., Carbonell, J., & Mitchell, T. (Ed.), *Machine Learning, Vol 2.* Los Altos, Ca.: Morgan Kaufmann Publishers, 1986.

Dietterich, T. Learning at the knowledge level. *Machine Learning*, 1986, *1*(3), 287-315.

Dyer, M. *In Depth Understanding.* MIT Press, 1983.

Etzioni, O. Hypothesis filtering: A practical Approach to reliable learning. In *Proceedings of the Fifth International Machine Learning Conference.* Ann Arbor, MI: , 1988.

Fisher, D. & Langley, P. Approaches to conceptual clustering. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence.* Los Angeles, CA: Morgan-Kaufmann, 1985.

Flann, N. & Dietterich, T. Selecting appropriate representations for learning from examples. In *Proceedings of the National Conference on Artificial Intelligence.* American Association for Artificial Intelligence, 1986.

Freedman, R. O. *Economic Warfare in the Communist Bloc: A study of Soviet economic pressure against Yugoslavia, Albania, and Communist China.* New York: Praeger, 1970.

Holland, J., Holyoak, K., Nisbett, R., and Thagard, P. *Induction.* MIT Press, 1986.

Hufbauer, G.C., & Schott, J.J. *Economic Sanctions Reconsidered: History and Current Policy.* Washington, D.C.: Institute For International Economics, 1985.

Hume, D. *A treatise of human nature.* Oxford: Claredon, 1739.

Kedar-Cabelli, S. Purpose-directed analogy. In *Proceedings of the Seventh Annual Conference of the Cognitive Science Society.* Irvine, CA: Lawrence Erlbaum, 1985.

Kodratoff, Y. & Tecuci, G. The central role of explanations in DISCIPLE. In *Proceedings of the International Workshop on Knowledge Representation and Organization in Machine Learning.* Geseke, Germany: , 1987.

Kolodner, J. *Retrieval and organizational strategies in conceptual memory: A computer model.* Hillsdale, NJ.: Lawrence Erlbaum Associates, 1984.

Leake, D. & Owens, C. Organizing memory for explanations. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society.* Lawrence Erlbaum, 1986.

Lebowitz, M. *Generalization and memory in an integrated understanding system* (Computer Science Research Report 186). Yale University, 1980.

Lebowitz, M. Integrated learning: Controlling explanation. *Cognitive Science*, 1986, Vol. *10*.

Lebowitz, M. Not the path to perdition: The utility of similarity-based learning. In *Proceedings of the National Conference on Artificial Intelligence.* American Association for Artificial Intelligence, 1986.

Lebowitz, M. Experiments with incremental concept formation: UNIMEM. *Machine Learning*, 1987, *2*(2), 103-138.

Leslie, A., & Keeble, S. Do six-month-old infants perceive causality? *Cognition*, 1987, *25*, 265-288.

Michalski, R. A system of programs for computer aided induction: A summary. In *Proceedings of the Fifth International*

Joint Conference on Artificial Intelligence. Cambridge, Mass.: Morgan-Kaufmann, 1977.

Michalski, R. The Inference-based theory of learning: Principles and approaches. In *Proceedings of the International Workshop on Knowledge Representation and Organization in Machine Learning.* Geseke, Germany: , 1987.

Minton, S., Carbonell, J., Etzioni, O., Knoblock C., & Kuokka D. Acquiring effective search control rules: Explanation-based learning in PRODIGY. In *Proceedings of the Fourth International Machine Learning Workshop.* Irvine, CA: , 1987.

Mitchell, T. Generalization as search. *Artificial Intelligence*, 1982, Vol. *18*(2).

Mitchell, T., Kedar-Cabelli, S. & Keller, R. Explanation-based learning: A unifying view. *Machine Learning*, 1986, Vol. *1*(1).

Moorehead, C. *Hostages to Fortune: A Study of Kidnapping in the World Today.* New York: Atheneum, 1980.

Murphy, G. and Medin, D. The Role of Theories in Conceptual Coherence. *Psychology Review*, 1985, *92*(3), 289-316.

Nisbett, Richard & Ross, Lee. *Human Inference: Strategies and Shortcomings of Social Judgments.* Engelwood Cliffs, NJ: Prentiss-Hall, Inc., 1978.

Pazzani, M. Inducing causal and social theories: A prerequisite for explanation-based learning. In *Proceedings of the Fourth International Machine Learning Workshop.* Irvine, CA: , 1987.

Pazzani, M. J. *Learning causal relationships by integrating analytical and empirical methods.* Doctoral dissertation, University of California, Los Angeles, 1988.

Pazzani, M. Integrated Learning with Incorrect and Incomplete Theories. In *Proceedings of the Fifth International Machine Learning Conference.* Ann Arbor, MI: , 1988.

Pazzani, M., Dyer, M. & Flowers, M. The role of prior causal theories in generalization. In *Proceedings of the National Conference on Artificial Intelligence.* American Association for Artificial Intelligence, 1986.

Piaget, J. *The child's conception of physical causality.* London: Kegan Paul, 1930.

Schank, R.C. & Abelson, R.P. *Scripts, plans, goals, and understanding.* Hillsdale, NJ.: Lawrence Erlbaum Associates, 1977.

Schank, R. *Dynamic Memory: A Theory of Reminding and Learning in Computers and People.* Cambridge University Press, 1982.

Schank, R.C. *Explanation patterns: understanding mechanically and creatively.* Hillsdale, NJ.: Lawrence Erlbaum Associates, 1986.

Silver, B. *Meta-level inference.* North Holland, 1986.

Valiant, L. A theory of the learnable. *Journal of the Association of Computing Machinery*, 1984, *27*(11), 1134-1142.

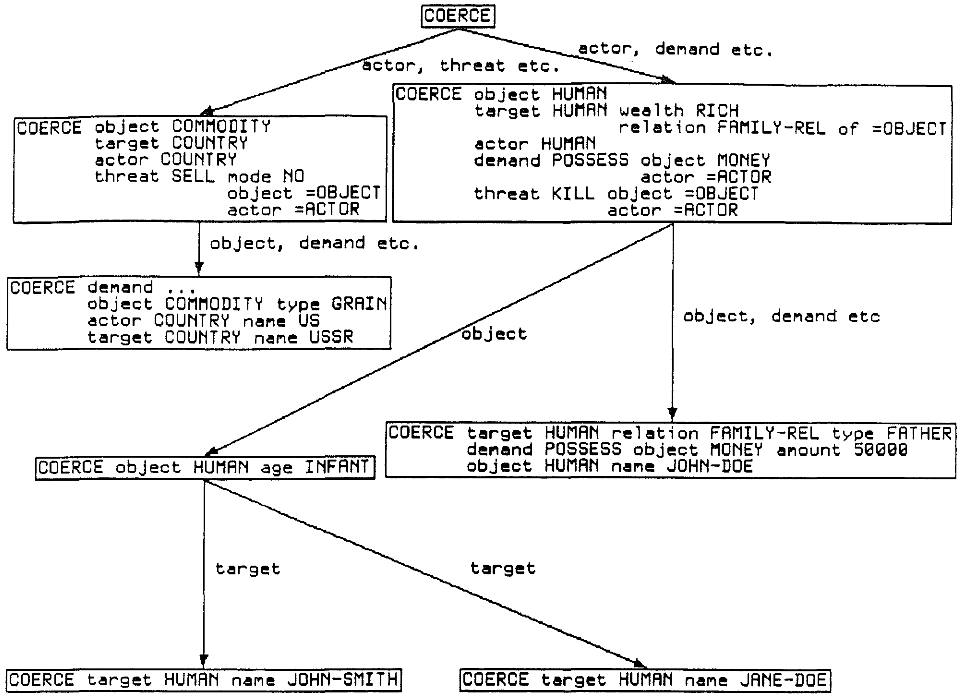

Figure 15: Hierarchical organization of schemata in memory.

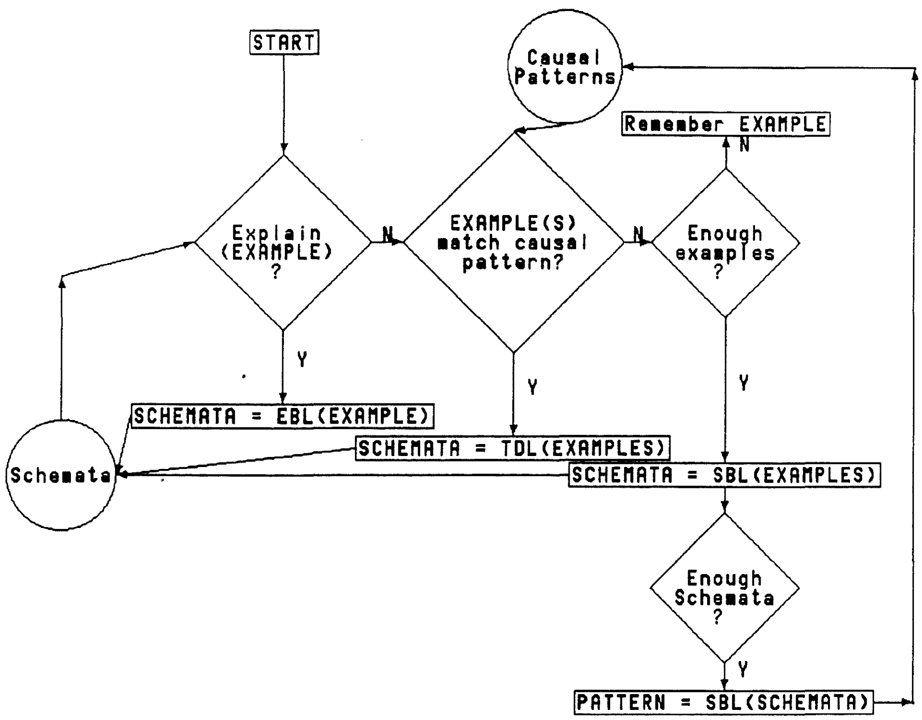

Figure 16: The flow of information in an extended version of OCCAM. Background knowledge for explanation-based learning is acquired through similarity-based or theory-driven learning. Causal patterns are learned by similarity-based learning and used by theory-driven learning.