

# Incremental Learning of Explanation Patterns and their Indices

Ashwin Ram

College of Computing  
Georgia Institute of Technology  
Atlanta, Georgia 30332-0280  
(404) 853-9372  
E-mail: ashwin@cc.gatech.edu

## Abstract

This paper describes how a reasoner can improve its understanding of an incompletely understood domain through the application of what it already knows to novel problems in that domain. Recent work in AI has dealt with the issue of using past explanations stored in the reasoner's memory to understand novel situations. However, this process assumes that past explanations are well understood and provide good "lessons" to be used for future situations. This assumption is usually false when one is learning about a novel domain, since situations encountered previously in this domain might not have been understood completely. Instead, it is reasonable to assume that the reasoner would have gaps in its knowledge base. By reasoning about a new situation, the reasoner should be able to fill in these gaps as new information came in, reorganize its explanations in memory, and gradually evolve a better understanding of its domain.

We present a story understanding program that retrieves past explanations from situations already in memory, and uses them to build explanations to understand novel stories about terrorism. In doing so, the system refines its understanding of the domain by filling in gaps in these explanations, by elaborating the explanations, or by learning new indices for the explanations. This is a type of *incremental learning* since the system improves its explanatory knowledge of the domain in an incremental fashion rather than by learning new XPs as a whole.

## 1 Case-based learning

Case-based reasoning and learning programs deal with the issue of using past experiences or *cases* to understand, plan for, or learn from novel situations [Kolodner, 1988; Hammond, 1989]. This happens according to the following process: (a) Use problem description to get reminded of old case. (b) Retrieve the

results (lessons, explanations, plans) of processing the old case and give them to the understander, planner or problem-solver. (c) Adapt the results from the old case to the specifics of the new situation. (d) Apply the adapted results to the new situation.

The intent behind case-based reasoning is to avoid the effort involved in re-deriving these lessons, explanations or plans by simply reusing the results from previous cases. However, this process assumes that past cases are well understood and provide good "lessons" to be used for future situations, since it is these very cases that determine the performance of the system in new situations. This assumption is usually false when one is learning about a novel domain, since cases encountered previously in this domain might not have been understood completely. Instead, it would be reasonable to assume that the reasoner would have *gaps* in the knowledge represented by these cases.

Even if past cases are not well understood, they can still be used to guide processing in new situations. However, in addition to using the past case to understand the *new* situation, a reasoner can also learn more about the *old* case itself, and thus improve its understanding of the domain. This is an important problem that has not been addressed in case-based reasoning research, and one that is suited to a machine learning approach in which learning occurs incrementally as these gaps are filled in through experience.

This paper describes a case-based story understanding system that retrieves past explanations from situations already in memory, and uses them to build explanations to understand novel situations encountered in newspaper stories about terrorism. The system learns in an incremental manner, by filling in the gaps in the retrieved explanation that is being used as a precedent in understanding the new situation. What is done with the newly learned information depends on the kind of "knowledge gap" the system is trying to fill. The new piece of knowledge could result in a new explanation in memory; it could be used to fill in a gap in an existing explanation; it could be used to elaborate an existing explanation if that explanation was not detailed enough to deal with the new situation; or it could be used to reorganize or re-index knowledge in memory to allow the reasoner to use what it already knows in novel situations to which that piece of knowledge had

not been applied before. Each type of learning leaves the system a little closer to a complete understanding of its domain. Each type of learning could also result in a new set of gaps as the system realizes what else it needs to learn about, which in turn drives the system towards further learning.

Much of real-world learning is an incremental process of this type. A reasoner learns by modifying what it already knows using little pieces of new information that it comes across during its experiences. This paper presents a theory of incremental learning for case-based story understanding.

## 2 Explanation patterns

Before we can discuss the learning process, we must describe what needs to be learned. This depends on the purpose to which the learned knowledge will be put. Consider the problem of building motivational explanations for the purpose of understanding stories. An understander could construct such explanations by using rules connecting typical goals and plans of people (e.g., [Wilensky, 1978]). However, this would be very inefficient in complicated situations, where motivational causal chains could be several steps long. To get around this problem, a case-based understander uses pre-stored explanations for stereotypical situations. These explanations represent standard patterns that are observed in these situations, and hence are called *explanation patterns* [Schank, 1986]. When the understander sees a situation for which it has a canned explanation pattern (XP), it tries to apply the XP to avoid detailed analysis of the situation from scratch. Thus an XP is like an abstract case; it represents a generalization based on the understander's experiences that can be used as a paradigmatic case for similar situations in the future.

For example, a "blackmail" situation may be represented by the following XP (`xp-blackmail`):<sup>1</sup>

- (1) The blackmailee has a goal G1.
- (2) The blackmailer has a goal G2, and the blackmailee does not have the goal G2 (since otherwise he or she would satisfy the goal without needing to be threatened).
- (3) The blackmailee has a goal G3, which he or she values above goal G1.
- (4) The blackmailer threatens to violate G3 unless the blackmailee performs an action A that satisfies G2, even though the action would have a negative effect of violating G1.

## 3 Learning explanation patterns

How are stereotypical XPs formed in memory? The work in explanation-based learning focusses on the problem of learning through the generalization of causal structures underlying novel situations [DeJong and Mooney, 1986; Mitchell *et al.*, 1986]. However,

<sup>1</sup>Details of XP representations may be found in [Ram, 1989].

it is difficult to determine the correct level of generalization. Furthermore, many stories do not provide enough information to prove that the explanation is correct. The understander must often content itself with two or more competing hypotheses, or otherwise jump to a conclusion. This means that in a real world situation, an explanation-based learning system may still need to deal with the problem of incomplete or incorrect domain knowledge.

Thus the system's memory of past experiences will not always contain "correct" cases or "correct" explanations, but rather one or more hypotheses about what the correct explanation might have been.<sup>2</sup> These hypotheses often have questions attached to them, representing what is still not understood or verified about those hypotheses. As the understander reads new stories, it is reminded of past cases, and of old explanations that it has tried. In attempting to apply these explanations to the new situation, its understanding of the old case gradually gets refined. New indices are learned as the understander learns more about the range of applicability of the XP. The XP is re-indexed in memory and is more likely to be recalled only in relevant situations.

Thus XP learning is an incremental process of theory formation, involving both case-based reasoning and explanation-based learning processes.

## 4 The AQUA program

AQUA is a story understanding program which learns about terrorism by reading newspaper stories about terrorist incidents in the Middle East [Ram, 1987; Schank and Ram, 1988; Ram, 1989]. AQUA reads stories about suicide bombing and attempts to understand them by constructing causal and motivational explanations for the events in the stories.

AQUA's case memory is based on XPs that have been used to explain past situations. AQUA improves its explanatory knowledge of the domain through a process of re-indexing and incremental modification of these XPs. For example, suppose AQUA has just read the following suicide bombing story (New York Times, April 14, 1985):

### **Boy Says Lebanese Recruited Him as Car Bomber.**

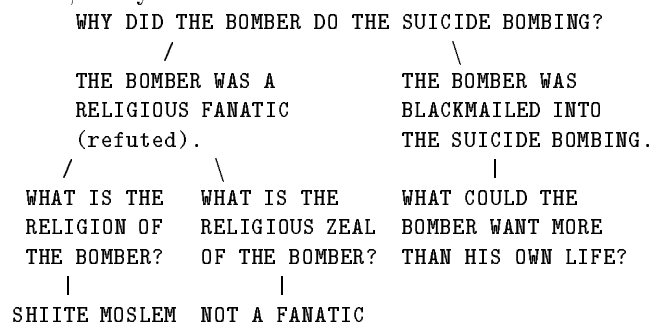
JERUSALEM, April 13 — A 16-year-old Lebanese was captured by Israeli troops hours before he was supposed to get into an explosive-laden car and go on a suicide bombing mission to blow up the Israeli Army headquarters in Lebanon. ...

What seems most striking about [Mohammed] Burro's account is that although he is a Shiite Moslem, he comes from a secular

<sup>2</sup>Actually, a single story or episode can provide more than one "case," each case being a particular interpretation or dealing with a particular aspect of the story. For an explanation program, each anomaly in a story, along with the corresponding set of explanatory hypotheses, can be used as a case.

family background. He spent his free time not in prayer, he said, but riding his motorcycle and playing pinball. According to his account, he was not a fanatic who wanted to kill himself in the cause of Islam or anti-Zionism, but was recruited for the suicide mission through another means: blackmail.

After reading this story, AQUA builds the following *hypothesis tree* in memory, representing an anomaly (*Why would the bomber perform an action that resulted in his own death?*), alternative hypotheses constructed by applying known XPs to the anomalous situation (*religious fanatic* and *blackmail*), questions that would verify these hypotheses, and answers to these questions, if any:<sup>3</sup>



The final explanation built for this story involves a novel application of a stereotypical XP, *xp-blackmail*. Even though the system already knows about blackmail, it learns a new variant of this XP (*xp-blackmail-suicide-bombing*), based on the particular manner in which *xp-blackmail* was adapted to fit the story. AQUA also learns indices to the new XP. Both kinds of learning are important in a case-based reasoning system. Let us start with the latter.

## 5 Learning indices for explanation patterns

Regardless of whether a new XP is learned from scratch or by applying an existing XP to a new situation, the XP needs to be indexed in memory appropriately so that it can be used in future situations in which it is likely to be useful. Ideally, an XP should be indexed in memory such that it is retrieved only in those situations in which it is applicable. But this is impossible in practice. For example, consider the applicability conditions for “blackmail.” In general, blackmail is a possibility whenever “someone does something he doesn’t want to do because not doing it results in something worse for him.” But trying to show this in general is very hard. Thus, in addition to general applicability conditions, an understander must learn specific, sometimes superficial, features that suggest possibly relevant XPs even though they may not completely determine the applicability of the XP to the situation. For example, a classic blackmail situation is one where a rich businessman who is cheating

<sup>3</sup>The understanding process by which AQUA builds this hypothesis tree is irrelevant for the purposes of this paper. Details may be found in [Ram, 1989].

on his wife is blackmailed for money using the threat of exposure. If one read about a rich businessman who suddenly began to withdraw large sums of money from his bank account, one would expect to think of the possibility of blackmail. However, one does not normally think of blackmail when one reads a story about suicide bombing, although theoretically it is a possible explanation.

The point is that XPs are associated with stereotypical situations and people in memory. An understander needs to learn the stereotypical categories that serve as useful indices for motivational explanations. This is a type of inductive category formation [Dietrich and Michalski, 1981]; however, the generalization process is constrained so that the features selected for generalization are those that are causally relevant to the explanations being indexed [Flann and Dietterich, 1989].

AQUA indexes motivational XPs in memory using typical contexts in which the XPs might be encountered (*situation indices*), as well as character stereotypes representing typical categories of people to whom the XPs might be applicable (*stereotype indices*) [Ram, 1989]. In the above example, AQUA learns a new context for blackmail (suicide bombing), as well as a new character stereotype representing the type of person who one might expect to see involved in a “blackmailed into suicide bombing” explanation. Let us discuss how AQUA learns these indices.

### 5.1 Learning situation indices

AQUA learns new contexts (e.g., “suicide bombing”) for stereotypical XPs (e.g., “blackmail”) which are then used as situation indices for these XPs in the future. The main issue here is how far the context should be generalized before it is used as an index. In the above example, should the new situation index for blackmail be *suicide-bombing*, *suicide*, *bombing*, *destroy*, or indeed any MOP (action) with a negative side effect for the actor? The issue here isn’t one of correctness but of utility. As discussed earlier, *xp-blackmail* is a possibility whenever the actor does something he would ordinarily not do because of a negative side effect. However, XP theory tries to replace generalized reasoning of this form with specific reasoning about stereotypical situations. The latter is more efficient even though it is less general.

After reading the above story, for example, one would expect to think of blackmail when one reads another story about a suicide bombing attack. However, one would probably not think of blackmail on reading any story about *suicide*, say, a teenager killing himself after failing his high school examinations, even though theoretically it is a possible explanation. Furthermore, it would not be useful to index the new XP under *bombing* in general (as opposed to *suicide-bombing* in particular), since the particular goal violation of the *p-life* goal is central to this explanation.

Thus in the above example, AQUA uses *suicide-bombing* as the situation index for the new variant of *xp-blackmail* (figure 1). After reading several stories

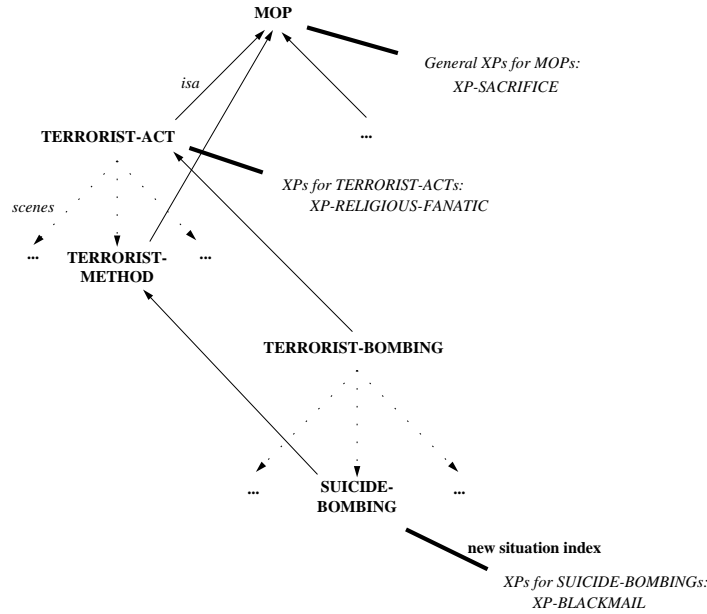


Figure 1: Learning situation indices for XPs. Upward lines represent *isa* links, and downward dotted lines represent *scenes* of MOPs. Heavy lines represent situation indices, which point from MOPs to XPs. Here, AQUA has just built a situation index from suicide-bombing to a copy of xp-blackmail.

about blackmail, AQUA would know about different stereotypical situations in which to use the blackmail explanation, rather than a generalized logical description of every situation in which blackmail is a possible explanation. In other words, AQUA would have indexed a copy of `xp-blackmail` under all the MOPs for which it has seen `xp-blackmail` used as an explanation. Whenever these MOPs are encountered, AQUA would retrieve the new blackmail XP (if the other indices are also present).<sup>4</sup> The reason that a copy of the original XP is used is that the XP, once copied, will need to be modified for that particular situation, as discussed below.

## 5.2 Learning stereotype indices

The main constraint on a theory of stereotype learning is that the kinds of stereotypes learned must be useful in retrieving explanations. In other words, they must provide the kinds of discrimination that are needed for indexing XPs in memory. Since volitional explanations are concerned with goals, goal orderings, plans and beliefs of characters, the learning algorithm must produce typical collections of goals, goal-orderings, plans and beliefs, along with predictive features for these elements. Such a collection is called a *character stereo-*

<sup>4</sup>AQUA can still understand other blackmail situations that it has not learned about as yet, as it did the story in the above example. Thus not having a situation index for an XP does not necessarily mean that the XP cannot be applied to the situation, but rather that this XP is not one that would ordinarily come to mind in that situation.

*type*.

Character stereotypes serve as motivational categories of characters and are an important index for XPs in memory. In the above example, AQUA learns a new stereotype (`stereotype.79`) representing a typical Lebanese teenager who might be blackmailed into suicide bombing, which is used to index the blackmail XP. The stereotype is built from the novel blackmail explanation by generalizing the features of the character involved in that explanation:

Answering question: WHY DID THE BOY DO THE SUICIDE BOMBING?  
with: THE BOY WAS BLACKMAILED INTO DOING THE SUICIDE BOMBING.

Novel explanation for A SUICIDE BOMBING!

Building new stereotype STEREOTYPE.79:

Typical goals:  
P-LIFE (in)  
A-DESTROY (OBJECT) (out)  
AVOIDANCE-GOAL (STATE) (question)  
Typical goal-orderings:  
AVOIDANCE-GOAL (STATE) over P-LIFE (question)

Typical beliefs:  
RELIGIOUS-ZEAL = NOT A FANATIC (in)  
Typical features:  
AGE = TEENAGE AGE (hypothesized)  
RELIGION = SHIITE MOSLEM (hypothesized)  
GENDER = MALE (hypothesized)  
NATIONALITY = LEBANESE (hypothesized)

Indexing XP-BLACKMAIL-SUICIDE-BOMBING  
Stereotype index = STEREOType.79  
Situation index = SUICIDE-BOMBING

The label **in (out)** marks features that are known to be true (false) of this stereotype [Doyle, 1979]. These features are definitional of the stereotype. The label **question** marks features that are **in** but incomplete. In this case, (**AVOIDANCE-GOAL (STATE)**) refers to an unknown goal that needs to be filled in when the information comes in. This is represented as a **goal** with an unknown **goal-object**. Finally, the label **hypothesized** marks features that were true in this story but were not causally relevant to the explanation. These features are retained for the purposes of recognition and learning. Since AQUA does not assume that its explanations are complete, there is the possibility of learning more about this explanation in the future that would help to determine whether these features have explanatory significance. This has not yet been implemented in AQUA.

The stereotype is used to index the new explanation in memory (figure 2). After reading this story, AQUA uses the new stereotype to retrieve the blackmail explanation when it reads other stories about Lebanese teenagers going on suicide bombing missions.

This stereotype is built through generalization under causal constraints from the hypotheses that were considered, including the ones that were ultimately refuted. The causal constraints are derived both from the successful explanation (blackmail) as well as from unsuccessful hypotheses, if any (here, religious fanaticism).

### 5.2.1 Learning from successful explanations

Clearly, much of **stereotype.79** comes from the motivational aspects of the blackmail explanation. AQUA retains those goals, goal orderings and beliefs of the character in the story that are causally implicated in the blackmail explanation. Since blackmail relies on a goal ordering between two goals, one of which is sacrificed for the other, the stereotype must specify that the character has a goal that he or she values above **p-life**. The stereotype also specifies that the character would normally not have the goal of performing terrorist missions, since this is part of the blackmail explanation. In the above story, AQUA infers the following goals and goal-orderings for the actor (corresponding to (1), (2) and (3) of **xp-blackmail**, page 2):

Building new stereotype (STEREOType.79):

Inferring GOALS  
from XP-BLACKMAIL (successful):  
THE ACTOR WANTED TO PRESERVE HIS OWN LIFE.  
THE ACTOR DID NOT WANT TO PERFORM THE  
TERRORIST MISSION.  
THE ACTOR WANTED TO AVOID SOMETHING.

Inferring GOAL-ORDERINGS  
from XP-BLACKMAIL (successful):  
THE GOAL OF THE ACTOR TO AVOID SOMETHING  
WAS MORE IMPORTANT THAN THE GOAL OF  
THE ACTOR TO PRESERVE HIS OWN LIFE.

These goals and goal-orderings are added to the stereotype being built. At this point, **stereotype.79** has the following features:

Typical goals:  
P-LIFE (in)  
A-DESTROY (OBJECT) (out)  
AVOIDANCE-GOAL (STATE) (question)  
Typical goal-orderings:  
AVOIDANCE-GOAL (STATE) over P-LIFE (question)

### 5.2.2 Learning from failed explanations

Many explanation-based learning programs learn only from positive examples (e.g., [Mooney and DeJong, 1985; Segre, 1987]). However, it is also possible to apply this technique to learn from negative examples (e.g., [Mostow and Bhatnagar, 1987; Gupta, 1987]). AQUA uses refuted hypotheses to infer features that should *not* be present in the newly built stereotype. These are features which, if present, would have led to the hypothesis being confirmed.

For example, in the blackmail story, AQUA knows that the person being blackmailed is not a religious fanatic, since the religious fanatic explanation, which depended on this fact, has been refuted. The kind of person likely to be blackmailed into suicide bombing is, therefore, not a religious fanatic.<sup>5</sup> This feature is recorded in the newly built stereotype.

Building new stereotype (STEREOType.79):

XP-RELIGIOUS-FANATIC failed because:  
THE BOY DID NOT BELIEVE FANATICALLY IN THE  
SHIITE MOSLEM RELIGION.

Inferring BELIEFS  
from XP-RELIGIOUS-FANATIC (failed):  
THE ACTOR DID NOT BELIEVE FANATICALLY IN  
A RELIGION.

This results in a new belief being added to **stereotype.79**:

<sup>5</sup>As before, this is a stereotypical inference and not a logically correct one. A religious fanatic could indeed be blackmailed into suicide bombing; however, on reading a story about a religious fanatic going on a suicide bombing mission, blackmail would not normally come to mind. This means that **xp-blackmail-suicide-bombing** should not be indexed under **religious-fanatic**, at least on the basis of this example.

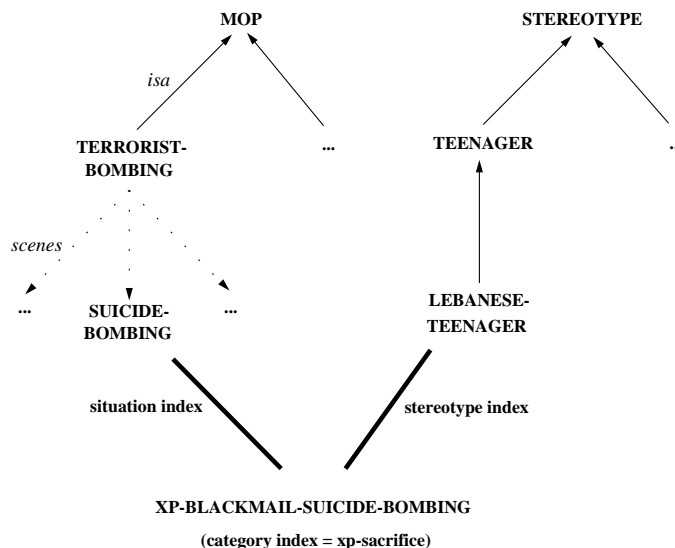


Figure 2: Learning stereotype indices for XPs. Upward lines represent *isa* links, and downward dotted lines represent *scenes* of MOPs. Heavy lines represent indices to XPs. Here, AQUA has just built a stereotype index from `stereotype.79`, representing a `lebanese-teenager`, to `xp-blackmail-suicide-bombing`.

Typical beliefs:

RELIGIOUS-ZEAL = NOT A FANATIC (in)

The reason that learning from the failed explanation works in this example is that the blackmail explanation specifies that the person being blackmailed would normally not have the goal to perform that action. This rules out other explanations which would result in this goal. Our theory does not deal with the issue of multiple successful explanations; more research needs to be done in this area.

## 6 Modifying existing explanation patterns

### 6.1 Associating new questions with XPs

Suppose AQUA reads the blackmail story with only the religious fanatic XP for suicide bombing in memory. When reading this story, AQUA is handed an explanation for the suicide bombing: the story explicitly mentions that the bomber was blackmailed. In a sense, then, the story has been understood since an explanation for the bombing has been found. However, one could not really say that AQUA had understood the story if it didn't ask the question, *What could the boy want more than his own life?* Unless this question is raised while reading the story, one would have to say that AQUA had missed the point of the story.

Such questions correspond to gaps in the explanation structures that are built during the understanding process (figure 3). These questions are associated with the XP, and may be answered later in the story or when the XP is applied to a future story. When they are answered, the understander can elaborate and modify the XP, thus achieving a better understanding

of the causality represented by the XP.

### 6.2 Incremental refinement of XPs by answering questions

In addition to raising new questions, of course, an understander must answer the questions that is already has in order to improve its knowledge of the domain. AQUA uses its questions to focus the understanding process, and learns when these questions are answered.

For example, consider the following story:

JERUSALEM — A young girl drove an explosive-laden car into a group of Israeli guards in Lebanon. The suicide attack killed three guards and wounded two others. ...

The driver was identified as a 16-year-old Lebanese girl. ... Before the attack, she said that a terrorist organization had threatened to harm her family unless she carried out the bombing mission for them. She said that she was prepared to die in order to protect her family.

When this story is read, AQUA retrieves the new `xp-blackmail-suicide-bombing` and applies it to the story. The question that is pending along with this explanation is also instantiated. When the question is answered, it is replaced by a new node representing the `protect-family` goal, and becomes part of `xp-blackmail-suicide-bombing`. Since no explanations are known for the newly added node, this in turn becomes a new question about the elaborated XP (not shown in the figure). The question is seeking a reason for the unusual `goal-ordering` of the actor, in which `protect-family` is given a higher priority than `p-life`.

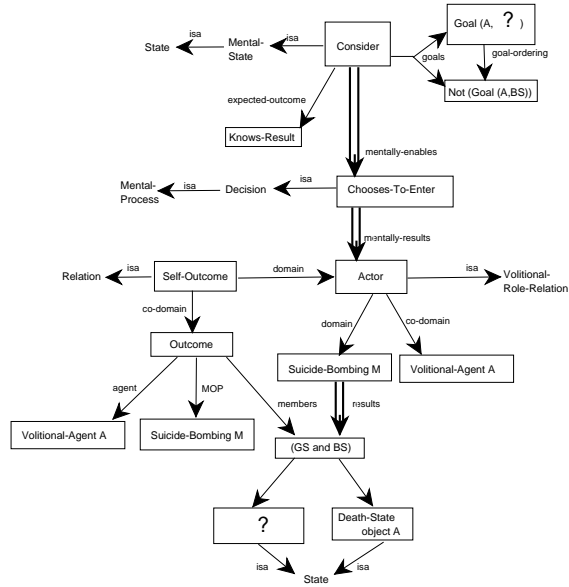


Figure 3: Associating new questions with XPs. The XP represents a situation in which an agent A volitionally performs (chooses-to-enter) an action whose outcome is known (knows-result) to be the death-state of A, as well as an unknown state that A wants more than he wants to avoid his death-state (the goal-ordering). The unknown goal represents the new question, *What could the actor want more than his own life?* This is depicted as an empty box, representing a gap in the program's knowledge. The XP is elaborated by filling in this gap when this question is answered.

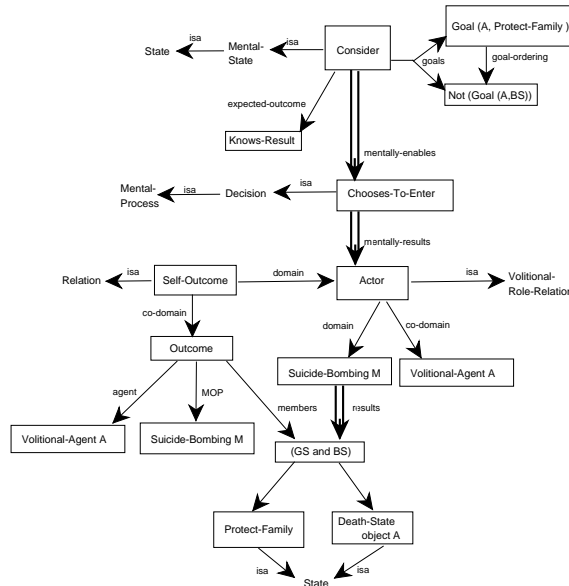


Figure 4: Elaborating an XP through incremental learning. The changed portion is depicted as a newly filled-in box, representing the answering of the question that was indexed at that point (compare with figure 3).

When the elaborated XP is applied to a new suicide bombing story, the new node will now be one of the premises of the hypothesis, causing AQUA to ask whether the actor was trying to protect his family. This reflects a deeper understanding of this particular scenario and is shown in figure 4. The new question will also be instantiated, causing AQUA to look for an explanation for the unusual **goal-ordering**. Should new questions be raised and then answered during future stories, AQUA will again be able to elaborate this XP in a similar manner. Thus AQUA evolves a better understanding of the “blackmailed into suicide bombing” scenario through a process of question asking and answering.

## 7 Conclusions

Explanation patterns are used for constructing explanations for anomalous situations by applying stereotypical packages of causality from similar situations encountered earlier. Thus XPs are abstract cases that are used as paradigmatic examples of stereotypical situations.

This paper presents a theory of XP learning through the incremental modification of existing XPs, using explanation-based learning techniques to constrain the modification process. The modifications involve the adaptation and elaboration of XPs, as well as the learning of indices for XPs. Both types of knowledge are essential in any case-based reasoning system. The theory is implemented in the AQUA program, which learns about terrorism by reading newspaper stories about unusual terrorist incidents in the Middle East.

## References

- [DeJong and Mooney, 1986] G. F. DeJong and R. J. Mooney. Explanation-Based Learning: An Alternative View. *Machine Learning*, 1(2):145–176, 1986.
- [Diettrich and Michalski, 1981] T. G. Diettrich and R. S. Michalski. Inductive Learning of Structural Descriptions: Evaluation Criteria and Comparative Review of Selected Methodologies. *Artificial Intelligence*, 16:257–294, 1981.
- [Doyle, 1979] J. Doyle. A Truth Maintenance System. *Artificial Intelligence*, 12:231–272, 1979.
- [Flann and Diettrich, 1989] N. S. Flann and T. G. Diettrich. A Study of Explanation-Based Methods for Inductive Learning. *Machine Learning*, 4:187–226, November 1989.
- [Gupta, 1987] A. Gupta. Explanation-Based Failure Recovery. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 606–610, Seattle, WA, July 1987.
- [Hammond, 1989] K. J. Hammond, editor. *Proceedings: Case-Based Reasoning Workshop*. Morgan Kaufmann, Inc., Pensacola Beach, FL, May 1989.
- [Kolodner, 1988] J. L. Kolodner, editor. *Proceedings of a Workshop on Case-Based Reasoning*. Morgan Kaufmann, Inc., Clearwater Beach, FL, May 1988.
- [Mitchell *et al.*, 1986] T. M. Mitchell, R. Keller, and S. Kedar-Cabelli. Explanation-Based Generalization: A Unifying View. *Machine Learning*, 1(1):47–80, 1986.
- [Mooney and DeJong, 1985] R. J. Mooney and G. F. DeJong. Learning Schemata for Natural Language Processing. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 681–687, Los Angeles, CA, August 1985.
- [Mostow and Bhatnagar, 1987] J. Mostow and N. Bhatnagar. FAILSAFE – A Floor Planner that uses EBG to Learn from its Failures. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 249–255, Milan, Italy, August 1987.
- [Ram, 1987] A. Ram. AQUA: Asking Questions and Understanding Answers. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, pages 312–316, Seattle, WA, July 1987. Morgan Kaufman Publishers, Inc., Los Altos, CA.
- [Ram, 1989] A. Ram. *Question-driven understanding: An integrated theory of story understanding, memory and learning*. Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, May 1989. Research Report #710.
- [Schank and Ram, 1988] R. C. Schank and A. Ram. Question-driven Parsing: A New Approach to Natural Language Understanding. *Journal of Japanese Society for Artificial Intelligence*, 3(3):260–270, May 1988.
- [Schank, 1986] R. C. Schank. *Explanation Patterns: Understanding Mechanically and Creatively*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1986.
- [Segre, 1987] A. M. Segre. *Explanation-Based Learning of Generalized Robot Assembly Tasks*. Ph.D. thesis, University of Illinois at Urbana-Champaign, Urbana, IL, January 1987. Technical Report UILU-ENG-87-2208.
- [Wilensky, 1978] R. Wilensky. *Understanding Goal-Based Stories*. Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, 1978.