COGNITIVE SCIENCE 14, 483–509 (1990)

Learning Plan Schemata From Observation: Explanation-Based Learning for Plan Recognition

RAYMOND J. MOONEY
University of Texas

This article discusses how explanation-based learning of plan schemata from observation can improve performance of plan recognition. The GENESIS program is presented as an implemented system for narrative text understanding that learns schemata and improves its performance. Learned schemata allow GENESIS to use schema-based understanding techniques when interpreting events and thereby avoid the expensive search associated with plan-based understanding. Learned schemata also function as new concepts that can be used to cluster examples and index events in memory. In addition, experiments are reviewed which demonstrate that human subjects, like GENESIS, can learn a schema by observing, explaining, and generalizing a single specific instance presented in a narrative.

1. INTRODUCTION

Abstract knowledge of typical plans, generally called *plan schemata* or *scripts*, have been shown to play an important role in cognitive tasks ranging from text understanding (Schank & Abelson, 1977) and memory (Bower, Black, & Turner, 1979) to planning and problem solving (Chi, Feltovich, & Glaser, 1981). However, the issue of how plan schemata are learned has not received much attention. To the extent that the learning issue has been addressed, it has generally been assumed that plan schemata are learned by induction across numerous experiences (Lebowitz, 1980; Rumelhart, 1980; Schank & Abelson, 1977).

I would like to thank Jerry DeJong, my thesis advisor, for initiating and guiding this research and Woo-Kyoung Ahn and Bill Brewer for allowing me to collaborate in the design and execution of psychological experiments motivated by the work on GENESIS. Comments by the anonymous reviewers are appreciated and significantly improved the final version of the article. This research was primarily conducted while the author was at the Coordinated Science Laboratory of the University of Illinois where it was supported by the Office of Naval Research under grant N-00014-86-K-0309 and by a University of Illinois Cognitive Science/AI fellowship. The Department of Computer Sciences at the University of Texas at Austin provided important support during the writing of the final article.

Correspondence and requests for reprints should be sent to Raymond J. Mooney, Dept. of Computer Sciences, Taylor Hall 2.124, University of Texas, Austin, TX 78712.

This article concerns the acquisition of plan schemata from specific observed instances by means of *explanation-based learning* (EBL; DeJong & Mooney, 1986; Mitchell, Keller, & Kedar-Cabelli, 1986). EBL is capable of learning a general plan schema from a single observed instance by building and generalizing an explanation for how the observed plan achieves its goal. The ability of EBL to use existing knowledge to acquire a schema from a single instance distinguishes it from *similarity-based learning* methods which induce concepts from numerous examples and counter-examples (Medin, Wattenmaker, & Michalski, 1987; Quinlan, 1986; Rumelhart, Hinton, & Williams, 1986).

In particular, this article describes how EBL of plan schemata from observation can improve the performance of *plan recognition,* the task of explaining the observed actions of others. In contrast, other work in EBL has focused on improving performance of classification or problem solving. The process of constructing explanations for observed behavior is an important cognitive task which, compared to classification and problem solving, has not received the attention it deserves (Schank, 1986). In particular, very little research has addressed the problem of learning for plan recognition. This article presents the GENESIS program as an example of an EBL system that improves its abilities to explain observed behavior. GENESIS is a narrative text comprehension system that improves its performance by learning plan schemata from specific observed plans. The second section describes the problem of learning for plan recognition and presents an overview of GENESIS as well as an example of it learning a plan schema that improves its performance. The third section presents some details on the design and operation of the GENESIS system. The fourth section reviews psychological experiments demonstrating that human subjects, like GENESIS, are capable of using their existing knowledge to learn plan schemata by observing, explaining, and generalizing a single specific instance. The fifth section highlights the unique aspects of the current work compared to other research in the area. The sixth section notes several problems for future research and the final section presents some conclusions.

2. EBL FOR PLAN RECOGNITION

Methods for plan recognition can be classified into two basic approaches based on whether or not the system is assumed to have an explicit representation of the plan to be recognized. A *schema-based* (*script-based*) system attempts to access directly and efficiently a relatively specific knowledge structure, a plan schema, that accounts for the observed actions. Such a system cannot understand an observed plan if it does not already have a schema for it. Appropriate schemata are selected in a bottom-up fashion

based on cues in the observed input and then used in top-down manner to connect actions causally, fill in missing actions, resolve anaphoric references, and so on. Examples of narrative understanding systems that use schema-based plan recognition are SAM (Cullingford, 1978) and FRUMP (DeJong, 1982). Recent work on plan recognition by Kautz and Allen (1986) also assumes the presence of a complete set of plan schemata.

A *plan-based* mechanism, on the other hand, can be used to understand novel situations for which the system does not have an explicit plan schema. In this more bottom-up approach, plan recognition involves searching for a set of missing actions that causally connect to observed actions to form a plan which achieves a character's known or inferred goal. This process can be very search intensive since it effectively requires constructing a plan in order to recognize it; however, it has the advantage of being able to recognize a plan without an explicit schema. PAM (Wilensky, 1978) is an example of a system that performs plan-based recognition.

Since a robust plan recognition system must be able to deal with both mundane and novel situations, more recent narrative understanding systems have utilized both approaches (Dyer, 1983; Wilensky, 1983). GENESIS also uses both types of plan recognition mechanisms. It first tries to find a single schema that will directly explain the characters' actions. If this fails, it tries to connect individual actions causally in a plan-based manner.

Since a plan recognition system constructs explanations for observed actions, it is particularly well-suited for employing EBL to acquire plan schemata from observation. Since schema-based recognition already requires an existing schema in order to recognize a plan, schema acquisition is only possible for novel plans recognized using a plan-based approach. By using explanation-based generalization, schemata can be acquired from the explanations constructed by a plan-based system. Since schema-based recognition generally requires less search than plan-based recognition, learned schemata can be used to improve the performance of plan recognition.

Consider a system like GENESIS that first tries to access efficiently a known schema to explain observed actions, and resorts to plan-based search only if it cannot find an appropriate schema. In this situation, a schema acquired from an initial plan-based understanding of one instance of a novel plan will allow the system to use schema-based processing to understand subsequent instances of the plan. This is analogous to a problem-solving system which learns macro-operators (Fikes, Hart, & Nilsson, 1972) or chunks (Laird, Rosenbloom, & Newell, 1986). By avoiding the search associated with plan-based recognition, subsequent instances are understood more efficiently. Both plan-based and schema-based understanding are necessary; however, an intelligent system should learn and improve its performance by increasing its reliance upon the latter as it gains experience in a particular domain.

2.1 Overview of GENESIS

A diagram illustrating the architecture of the complete GENESIS system is shown in Figure 1. In the diagram, circles represent declarative data structures and rectangles represent procedural subsystems.

The *parser* translates Engish text into predicate calculus using an adaptation of McDYPAR (a version of the parser used in the BORIS system; Dyer, 1983). The *lexicon* stores word definitions and disambiguation demons used by the parser. The parsed input is passed to the *understander,* which performs plan-based and schema-based understanding using knowledge from the *schema library.* This process causally connects the inputs and fills in missing information. The embellished representation constructed for a narrative is called the *causal model.* An explanation for a particular goal achievement is the subset of the causal model supporting the given goal state.

When building a causal model for a piece of text, the understander first tries to find a plan schema that directly explains the characters' actions. If this fails, it tries to connect individual actions causally in a plan-based manner. However, the search it performs during plan-based understanding is very limited in order to prevent getting lost in a combinatorially explosive search of all possible explanations. Consequently, GENESIS cannot produce explanations for many narratives although it theoretically could explain them given an exhaustive search algorithm and unlimited time and space. Nevertheless, the system's limited ability to do plan-based understanding allows it to construct explanations for many novel plans.

GENESIS also has a number of components for demonstrating its understanding. The *question-answer* (Q/A) takes questions from the user after

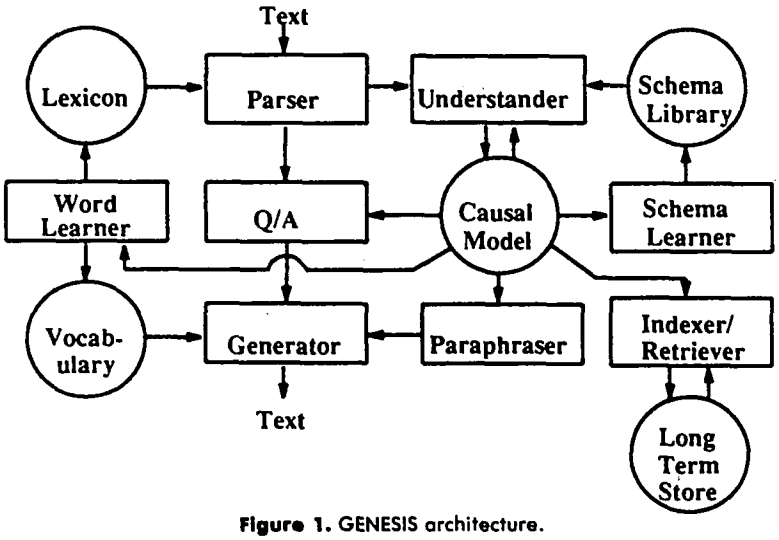

Figure 1. GENESIS architecture.

they have been parsed and employs a number of heuristics for retrieving answers from the causal model. Since the focus of the system is on the construction of generalizable explanations, this subsystem is primarily built to answer questions about why a character performed a certain action or why a particular state exists. The system also has a *paraphraser* which uses the most comprehensive schemata instantiated in the causal model to produce a paraphrase of the text. The paraphraser can also produce English descriptions of abstract schemata from the schema library. The output of both the Q/A and paraphraser are translated from predicate calculus to natural language by the *generator,* which is an adaptation of McMUMBLE (Schank & Riesbeck, 1981), a version of the generator used by the TALE-SPIN system (Meehan, 1976). The *vocabulary* contains information about how to translate instances of each predicate into natural language expressions.

The *schema learner* analyzes the causal model built by the understander in order to learn new plan schemata. It monitors the causal model and detects when a character has acheived an important goal in a novel manner. The explanations for how such goals are achieved are used to learn new schemata. A domain-indendent, explanation-based generalization system called EGGS (Mooney, 1990; Mooney & Bennett, 1986) is used to generalize these explanations and package them into plan schemata. Learned schemata are stored in the schema library and indexed so they can aid in the understanding of future narratives.

GENESIS also has the ability to learn word meanings for schema-related terms. The *word learner* can interpret unknown words as labels for learned schemata or as labels for the roles (slots) of learned schemata (Mooney, 1987). This subsystem uses techniques similar to those used by the FOUL-UP system (Granger, 1977). Finally, GENESIS is able to use learned schemata to index and retrieve specific events. The *indexer/retriever* stores events in the *long-term store* indexed under the plan schemata which they invoked or created. This allows for a dynamic memory that learns new ways of indexing events (Schank, 1982).

2.2 Sample Performance

A standard procedure is used to test the GENESIS system. This procedure illustrates both the schema-learning process itself as well as the ability of learned schemata to improve system performance. First, the system is given a *test narrative* that presents a sparse description of an instance of a schema. This description is missing a number of actions crucial to the overall plan. Consequently, the narrative is not detailed enough for the system to construct a causally complete explanation for characters' actions without a schema to supply missing actions and inferences. The system is therefore unable to answer questions that require making default inferences about

what must have taken place, and it is unable to produce an adequate paraphrase of the narrative. Next, the system is given a *learning narrative* that describes in detail a complete instance of the schema and contains the crucial actions and other information that were lacking in the test narrative. Using its existing knowledge and limited plan-based recognition abilities, GENESIS is able to construct a causally complete explanation for this narrative and can therefore answer questions about why actors performed certain actions as well as produce an adequate paraphrase of the text. Since the narrative presents a novel plan that achieves an important goal, the system generalizes its explanation of how the goal was achieved into a new plan schema. Finally, GENESIS is given the test narrative again. This time, it can use the schema it has just learned to fill in missing information, and as a result, it can now answer questions it was previously unable to answer, as well as produce a better paraphrase of the text.

The remainder of this section presents an example in which GENESIS learns a schema for a police officer impersonating a prostitute in order to arrest solicitors.[1] Before processing the narratives in this example, the system has schemata for people soliciting sexual favors and for police officers arresting lawbreakers; however, it does not have a schema for "solicitation entrapment." In order to illustrate the system's current ignorance of this schema and how this lack of knowledge affects its ability to understand certain narratives, the system is first given a test narrative. Below is the actual I/O trace produced by the system:

> Input: Alice was at a corner wearing tight blue jeans. Stan told Alice if she had sex with him then he would give her $75. Stan went to jail.
>
> Ready for questions:
>
> >Paraphrase.
> Stan solicited Alice's sexual favors for $75. Stan was put in a jail.
>
> >Why did Stan tell Alice if she had sex with him then he would give her money?
> Because Stan believed that Alice was a prostitute and because Stan needed to have sex.
>
> >Who arrested Stan?
> Answer unknown.

Since the arresting action is not explicitly mentioned, the limited plan-based procedures cannot connect Stan's incarceration to his solicitation. Consequently, it is unable to answer certain questions or produce an adequate paraphrase of the text. Next the system is given the following more detailed learning narrative.[2]

---

[1] Other schemata GENESIS has learned and used include: kidnapping for ransom, arson for insurance, and murder for inheritance. Mooney (1990) gives details on these other examples.

[2] Leading question marks indicate variables. Variables in the description of a schema are assumed to be existentially quantified.

Input: Jane is a policewoman. She dressed in a short red skirt and went to a corner. Bob approached the corner and told her if she had sex with him then he would give her $50. Jane arrested Bob for soliciting. Bob is Mary's husband and he told her that Jane entrapped him.

Thematic goal achieved: Jane is happy that Bob is under arrest for solicting Jane's sexual favors for the $50.
Explanation suitable for generalization. Pruning...Generalizing...Packaging...
Creating New Schema: (DressSolicitArrest ?b9 ?c2 ?s2 ?s4 ?l4 ?a9 ?c3)

?b9 is not equal to ?a9. ?b9 is a person. ?c2 is an apparel. ?b9 has ?c2. ?b9 puts on ?c2. ?s2 is a location. ?l4 is a corner. ?b9 is at ?s2. ?b9 goes from ?s2 to ?l4. ?a9 is a person. ?s4 is a location. ?a9 is at ?s4. ?s9 goes from ?s4 to ?l4. ?c3 is money. ?a9 needs to have sex. ?c2 is sexy. ?a9 solicits ?b9's sexual favors for ?c3. ?b9 is a police officer. ?b9 arrests ?a9 for soliciting ?b9's sexual favors for ?c3.

Having: ?b9 put on ?c2. suggests DressSolicitArrest
Having: ?a9 solicited ?b9's sexual favors for ?c3. suggest DressSolicitArrest
Having: ?b9 arrested ?a9 for soliciting ?b9's sexual favors for ?c3, suggest DressSolicitArrest

Unknown word '?b9 entrap ?a9' refers to DressSolicitArrest

Ready for questions:

>Paraphrase.
Jane had a skirt. Jane put on a skirt. Jane was at some place. Jane went to a corner. Bob was at some place. Bob went to the corner. Bob needed to have sex. The skirt was sexy. Bob solicited Jane's sexual favors for $50. Jane was a police officer. Jane arrested Bob for soliciting Jane's sexual favors for the $50.

>Why did Bob go to the corner?
So Bob could solicit Jane's sexual favors for the $50.

>Why did Bob solicit Jane's sexual favors?
Because Bob believed that Jane was a prostitute and because Bob needed to have sex.

>Why did Jane arrest Bob?
Because Jane believed that Bob solicited Jane's sexual favors for the $50 and because Jane was a police officer.

This time, because the arrest is explicitly mentioned, the system can connect it to the solicitation and thereby construct an adequate explanation for the story. The system generalizes its explanation for how Jane achieved her goal of arresting someone into a DressSolicitArrest schema and indexes it so that future solicitations *suggest* this new schema. The system also associates the unknown word "entrap" with the new schema. Although entrapment is actually a more general concept, this association is a good initial conjecture.

When the test narrative is processed again, Stan's solicitation suggests DressSolicitArrest and since this schema explains other facts in narrative

(e.g., why Stan went to jail), the test narrative is understood as an instance of the new schema.

Input: Alice was at a corner wearing tight blue jeans. Stan told Alice if she had sex with him then he would give her $75. Stan went to jail.

Thematic goal achieved: Alice is happy that Stan is under arrest for soliciting Alice's sexual favors for the $75.

Ready for questions:

>Paraphrase.

Alice had a pair of jeans. Alice put on the pair of jeans. Alice was at some place. Alice went to a corner. Stan was at some place. Stan went to the corner. Stan needed to have sex. The pair of jeans were sexy. Stan solicited Alice's sexual favors for $75. Alice was a police officer. Alice arrested Stan for soliciting Alice's sexual favors for the $75.

>Who arrested Stan?
Alice arrested Stan for soliciting Alice's sexual favors for the $75.

>Why did Alice arrest Stan?
Because Alice was a police officer and because Alice believed that Stan solicited Alice's sexual favors for the $75.

3. DETAILS OF THE GENESIS SYSTEM

This section provides further details on the design and operation of the GENESIS system. In particular, it describes the knowledge representation, plan recognition, and learning aspects of the system. Mooney (1990) provides further information, including detailed descriptions of the algorithms used by the system.

3.1 Knowledge Representation in GENESIS

All domain-specific knowledge in GENESIS is represented declaratively in the schema library. This knowledge is divided into information about *objects, attributes, states,* and *actions* which are further organized into taxonomic hierarchies. The hierarchies under each of these classes of knowledge support abstraction inferences such as: Isa(?x, Gun)→Isa(?x, Weapon), Mother(?x, ?y)→Parent(?x, ?y), and Poison(?x, ?y)→Murder(?x, ?y). Attributes refer to facts about an object that are not affected by actions while states refer to facts about objects that can be altered. Attributes and states can both have Horn-clause rules associated with them. For example, the example given above makes use of the following rule:

$$At(?x, ?l) \wedge At(?y, ?l) \wedge Isa(?l, Corner) \wedge Attire(?y, ?c) \wedge Seductive(?c) \rightarrow Believe(?x, Occupation(?y, prostitute))$$

This rule simply states that if ?x and ?y are at a corner and ?y is wearing seductive clothing then ?x will believe that ?y is a prostitute.

TABLE 1
Supports and Effects for Arrest(?a,?b,?d(?b))

Constraints	Preconditions	Motivations	Effects
Isa(?a,Character) Isa(?b,Character) ¬Equal(?a,?b) Illegal(?d(?b))	At(?b,?l) At(?a,?l)	Believe(?a,?d(?b)) Occupation(?a,police officer)	Arrested(?b,?d(?b),?a)

TABLE 2
Supports and Effects for Solicit(?a,?b,?o)

Constraints	Preconditions	Motivations	Effects
Isa(?a,Person) Isa(?b,Person) Isa(?o,Money)	CommPath(?a,?b)	NeedSex(?a) Believe(?a, Occupation(?b,prostitute))	Believe(?b, Solicit(?a,?b,?o))

Most of GENESIS' knowledge is in the form of action definitions. Actions are represented using an enhancement of STRIPS operators with precondition, add, and delete lists (Fikes et al., 1972). The first modification is that preconditions are divided into *constraints, preconditions,* and *motivations.* Constraints are required attributes or classes of the arguments of an action and therefore cannot be achieved by other actions. Preconditions are physical states of the world that enable an action and can be achieved by other actions. Motivations are mental states of the actor, such as goals, goal priorities, and beliefs, that motivate him to perform a volitional action. Together, constraints, preconditions, and motivations are called the *supports* of an action. *Effects* of an action are states resulting from the execution of an action. Deletions are simply negated effects. Actions also inherit supports and effects from more abstract actions in the hierarchy. Tables 1 and 2 give the supports and effects for the following actions used in the example: "?a arrests ?b for performing act ?d" and "?a solicits ?b's sexual favors for an amount ?o."[3] An Arrest requires that the officer and the offender be at the same place and is motivated by the officer believing that the offender performed an illegal act. A Solicit requires that the actor have a communication path to the recipient and is motivated by the actor needing to have sex and believing that the recipient is a prostitute.

Knowledge about an action may include information about its *expansion,* that is, its decomposition into more primitive actions. In this case an action is recursively defined as a plan schema that has an expansion in terms of other actions. The expansion of an action specifies the set of *subactions* that

[3] The tables use two predicates, Character and Person, which are subtly different. A Person is a human being while a Character is any agent capable of executing a volitional action and includes Persons as well as collective agents such as Companies.

comprise the plan schema and may also include other information such as temporal and causal ordering constraints on the subactions. For example, the expansion of the Arrest schema might include the following ordered list of subactions: Mtrans(?a, ?b, Miranda rights) (the officer reads the accused his rights), Handcuff(?a, ?b) (the officer handcuffs the accused), and Ptrans(?a, ?b, ?l1, ?l2) (the officer takes the accused to the police station).

3.2 Plan Recognition in GENESIS

As previously mentioned, GENESIS uses both schema-based and plan-based understanding mechanisms when constructing explanations for characters' actions. This section describes each of these processes, using the "solicitation entrapment" narratives as examples.

3.2.1 Schema Selection. If a schema-based mechanism is to be able to process a broad range of texts, it must have access to a large number of schemata. Therefore, in order to avoid repeated searching through the entire database of schemata, it must have an efficient method for selecting the particular schemata applicable to the current input. This is a difficult problem, frequently referred to as *schema selection* or *frame selection,* and has been the subject of several research efforts (Charniak, 1982; DeJong, 1982; Norvig, 1983).

In GENESIS, the process of selecting a schema has two stages: *suggestion* and *determination.* An action *A* in a narrative *suggests* a schema *B* if *A* is a subaction of *B* and the objects filling the case roles of *A* satisfy the constraints on the corresponding roles of *B.* For example, when processing the solicitation test narrative after acquiring the DressSolicitArrest schema, Stan's action of soliciting Alice suggests the new schema since Solicit is one of the schema's subactions.

Suggested schemata monitor the inputs in order to find confirming evidence in the form of additional inputs that match states or actions in its expansion. When all of the subactions of a suggested schema have been observed or inferred as filling a gap in a causal chain of events, it is *determined,* and its complete expansion is added to the causal model.[4] A single action may suggest several schemata, the determination process chooses which of these schemata actually to use in its interpretion of the text. In the sample test narrative, when the new schema is suggested, its Dress and Ptrans subactions are immediately inferred as explaining how Alice got to the corner wearing the jeans. When the input: "Stan went to jail" is processed, the system infers, as a necessary precondition, that he is under arrest. This fact matches the effect of the Arrest action in the suggested DressSolicitArrest schema thereby causing the Arrest action to be inferred as filling a gap be-

[4] The term *determination* is taken from Norvig (1983) who uses it to refer to a similar process in FAUSTUS.

tween Stan's solicitation and his incarceration. Since all of its subactions have been observed or inferred, the new schema is determined and its expansion is added to the causal model.

3.2.2 Plan Verification. While processing a narrative, if GENESIS encounters an input action that does not match part of an existing or suggested schema, nor suggest any schemata, it attempts to explain the action in a plan-based fashion by causally connecting it to previous actions. Using a process called *plan verification,* the system attempts to show that a set of known actions achieves a character's goal. In order to avoid combinatorially explosive search, it does not attempt to propose missing actions that might complete an explanation.

Specifically, plan verification uses backward chaining to try to infer the supports of observed actions from given facts and the effects of previous actions. By inferring the preconditions of actions from the effects of previous actions, causally connected sequences are identified. In the sample learning narrative, the preconditions and motivations of Bob's solicitation are inferred from the effects of previous actions. Specifically, the precondition requiring that there be a communication path between the actor and the recipient is inferred from the fact that they are both at the same corner, an effect of their previously mentioned movements. The motivation stating that Bob believes that Jane is a prostitute is inferred from the fact that she is standing on a corner dressed in seductive clothing, an effect of her previously mentioned Dress and Ptrans actions. The preconditions and motivations of Jane's Arrest action are connected to previous actions in a similar manner.

Complete plan recognition involves finding a connected set of actions which achieve a character's ultimate goal. Since narratives rarely explicitly mention characters' ultimate goals, the system must have a way of inferring them. *Thematic goals* (Schank & Abelson, 1977) are defined as goals arising from basic wants and needs and therefore requiring no further explanation. GENESIS has a set of inference rules for determining when a thematic goal has been achieved. For example, the following rule defines a police officer's thematic goal of wanting to arrest lawbreakers:

Arrested(?b,?d,?a)∧Occupation(?a, police officer)→
ThemeGoalMet(?a, Arrested(?b, ?d, ?a))

GENESIS checks input actions and determined schemata to see if they can be explained as achieving a thematic goal for some character. In the example, the above rule is used to determine that Bob's arrest has achieved a thematic goal for Jane. Figure 2 shows the final causal model constructed for this example.

Notice that, unlike a PAM-type (Wilensky, 1978) of system, plan verification does not conduct an exponential search for missing actions, which would create a causal chain of actions achieving a thematic goal. Rather, it

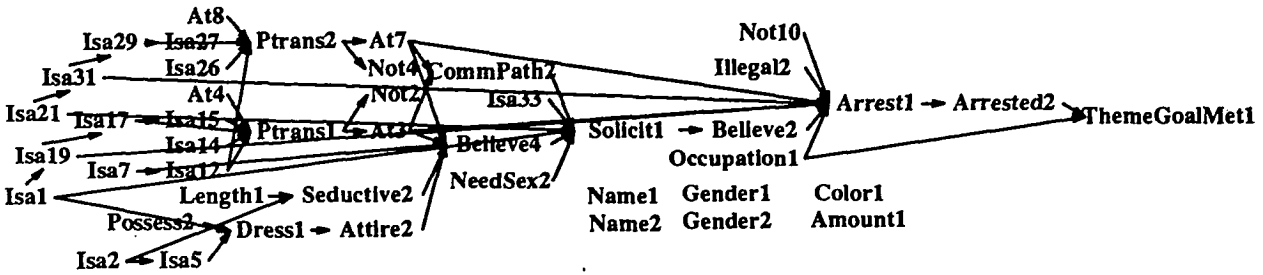

Name1	Jane's name is Jane.	Name2	Bob's name is Bob.
Gender1	Jane is female.	Gender2	Bob is male.
Color1	The skirt is red.	Amount1	The $50 is 50 dollars of money.
Isa5	The skirt is an apparel.	Isa2	The skirt is a skirt.
Isa15	Jane is a physical object.	Isa17	Jane is an animate object.
Isa19	Jane is a character.	Isa1	Jane is a person.
Isa27	Bob is a physical object.	Isa29	Bob is an animate object.
Isa31	Bob is a character.	Isa21	Bob is a person.

ThemeGoalMet1	Jane is happy that Bob is under arrest for soliciting Jane's sexual favors for the $50.
Occupation1	Jane is a police officer.
Arrested2	Bob is under arrest for soliciting Jane's sexual favors for the $50.
Arrest1	Jane arrested Bob for soliciting Jane's sexual favors for the $50.
Believe2	Jane believes that Bob solicited Jane's sexual favors for the $50.

Solicit1	Bob solicited Jane's sexual favors for the $50.
NeedSex2	Bob needs to have sex.
Believe4	Bob believes that Jane is a prostitute.
Attire2	Jane is wearing the skirt.
Dress1	Jane put on the skirt.
Possess2	Jane has the skirt.
Seductive2	The skirt is sexy.
Length1	The skirt is short.
At3	Jane is at the corner.
Ptrans1	Jane went to the corner.
Isa12	The corner is a location.
Isa7	The corner is a corner.
Isa14	Jane is a location.
At4	Jane was at some place.
Not2	Jane is not at some place.
At7	Bob is at the corner.
Ptrans2	Bob went to the corner.
Isa26	Bob is a location.
At8	Bob was at some place.
Not4	Bob is not at some place.
Isa33	The $50 is money.
CommPath2	Bob has a communication path to Jane.
Illegal2	It is illegal for Bob to solicit Jane's sexual favors for the $50.
Not10	Jane is not equal to Bob.

Figure 2. Causal model for the solicitation learning narrative.

simply tries to verify that existing actions causally connect to achieve such a goal. Therefore, the system is incapable of constructing explanations for narratives that have missing actions and do not suggest known schemata. For example, consider the processing of the solicitation test narrative *before* the system learns the schema. Since there is no explicit mention of an arrest, the system is incapable of connecting Bob's solicitation to his incarceration. There are lots of potential actions enabled by his solicitation: Alice could accept the offer, they could haggle over the price, Alice could be insulted and slap Bob, and so on. There are also many explanations for why Stan was sent to jail. He could have been in possession of drugs, he could have murdered someone, and so forth. A plan-based system with a large knowledge base of actions might search a long time before stumbling upon the right way to connect things. Since any real system is constrained by limitations on time and space, there will always be observations which it cannot explain even if it has all the relevant knowledge.

Of course, plan verification is not necessarily the best possible way to limit search during plan-based understanding. Plan verification probably limits search a little too much since people can often understand novel plans which require filling in one or two missing actions. However, searching for missing actions must be tightly constrained since it quickly becomes intractable.

3.3 Schema Learning in GENESIS

If the understander constructs an explanation for how a goal is achieved, GENESIS may be able to generalize this explanation into a new schema that can aid the understanding of future narratives. First, the system decides whether or not an explantion is worth generalizing. Second, a pruning procedure eliminates overly specific parts of the explanation. Third, explanation-based generalization is performed. Fourth, the generalized explanation is packaged into a new schema and indexed so that its subactions suggest it. The following subsections elaborate each of these steps of the learning process.

3.3.1 Deciding When to Learn. If every explanation GENESIS constructed were generalized into a schema, the system would eventually become overloaded with rarely used schemata. Most actions would suggest a large number of schemata, and selecting among these would require an excessive amount of processing time. This problem is an analogous to the performance deterioration noticed in problem-solving systems that learn too many useless macro-operators (Minton, 1985). In an attempt to avoid this problem, GENESIS has criteria for learning schemata selectively. Below is the set of criteria GENESIS uses to determine whether an explanation is worth generalizing.

1. It should be an explanation of how a thematic goal was achieved.
2. The highest level explanation for the goal achievement should not simply be an instantiation of a known schema.
3. All actions in the highest level explanation should rely upon the character whose thematic goal was achieved.

The first criterion is crucial for insuring that the acquired schema will be a useful one. A plan for achieving a state that satisfies normal human wants and desires is likely to make a schema that will arise again and again. Therefore, an explanation for a plan is considered to be worth generalizing into a schema only if it achieves a thematic goal. The explanation in Figure 2 meets this criteria since Alice achieves her thematic goal of arresting criminals.

The second criterion is the obvious one of not already possessing a schema for the combination of actions needed to achieve the goal. This simply involves checking the explanation for the goal achievement to make sure it consists of a chain of several schemata rather than a single schema. If the system already had a schema for this case, it would have used it in processing the narrative and the goal achievement would be explained by a single instance of this schema. The explanation in Figure 2 meets this criterion since it consists of a combination of several actions including Dress, Solicit, and Arrest.

The third and final criterion insures that the schema learned from the explanation is a volitional plan which agents can use to achieve their own thematic goals. Let the term, *main character,* refer to the character whose thematic goal was achieved. If the learned schema is to represent a plan that can be executed by the main character in order to achieve his thematic goal, actions in the explanation that are not volitionally performed by the main character should at least be motivated or initiated by actions that he performs. The solicitation example satisfies this criterion even though the Solicit action is performed by another agent, since it is motivated by the fact that the main character (Alice) is dressed seductively and standing on a corner. If the third criterion were not used, the system could learn schemata containing serendipitous actions over which the main character has no control. For example, assume GENESIS had the knowledge to explain the following narrative but did not have the third learning criterion.

John's rich uncle was killed in an earthquake. John inherited a million dollars.

Such a system would acquire an EarthquakeInherit schema (or possibly a NaturalDisasterInherit schema) from its explanation of this narrative. Such a nonvolitional schema would not be very useful for understanding later narratives since a natural disaster cannot normally be explained as part of a character's plan to inherit riches.

3.3.2 Pruning the Explanation. If an explanation meets all three learning criteria, GENESIS proceeds to generalize it into a new schema. Certain features of the example are immediately generalized away since they do not contribute to the explanation of how the goal was achieved. For example, the disconnected facts shown in Figure 2 concerning the names of the various characters and the color of the skirt are clearly irrelevant to the general schema. In order to increase the generality and applicability of the resulting schema even further, GENESIS has a pruning procedure that removes unnecessarily specific branches from the explanation prior to generalization. If such pruning were not performed, the system would frequently learn schemata too restrictive to be useful. For example, if the explanation in Figure 2 were not pruned, the resulting schema would only cover cases in which an officer used a short skirt to impersonate a prostitute.

The pruning procedure removes actions that only support the thematic goal achievement through more abstract actions in the taxonomic hierarchy. It also removes inference rules at the leaves of the explanation since such rules are only used to show how the supports of the overall plan were met in the specific example. In the solicitation example, the pruning process removes the rule stating that short skirts are seductive, thereby allowing the constraint of seductive attire to be achieved differently in future examples.

3.3.3 Generalizing and Packaging the Explanation. Once an explanation is pruned, it is generalized using standard explanation-based generalization (DeJong & Mooney, 1986; Mitchell et al., 1986). This process appropriately variablizes the explanation while maintaining any constraints needed for its validity. In particular, GENESIS uses the generalizer in the EGGS domain-independent EBL system (Mooney, 1990; Mooney & Bennett, 1986).[3] This generalizer computes the minimal variable bindings needed to allow the general rules and actions used in the construction of the explanation to connect in the way they did in the example. For the solicitation example, the final generalized explanation resulting from this process is shown in Figure 3.

The final step in acquiring a schema is packaging the generalized explanation into a form suitable for the schema library. First, the leaves of the generalized explanation are divided into constraints, preconditions, and motivations based on the manner in which they support actions in the generalized explanation. The subactions and internal structure of the generalized explanation are stored as the expansion of the new schema. A description of the DressSolicitArrest schema is shown in Table 3 and an English summary is shown in the trace in Section 2.2.

[3] The initial prototype implementation of GENESIS presented by Mooney and DeJong (1985) used an ad hoc generalizer which employed a representational language based on conceptual dependency forms.

TABLE 3
Learned Definition for DressSolicitArrest(?b9,?c2,?s2,?s4,?l4,?a9,?c3)

Subactions			
Dress(?b9,?c2)	Ptrans(?b9,?b9,?s2,?l4)	Prtans(?a9,?a9,?s4,?l4)	Solicit(?a9,?b9,?c3)
Arrest(?b9,?a9,Solicit(?a9,?b9,?c3))			
Constraints	**Preconditions**	**Motivations**	**Effects**
Isa(?c3,Money)	At(?a9,?s4)	Occupation(?b9, police officer)	Attire(?b9,?c2)
Isa(?a9,Person)	At(?b9,?s2)		¬At(?b9,?s2)
Isa(?b9,Person)	Possess(?b9,?c2)		¬At(?a9,?s4)
¬Equal(?a9,?b9)	NeedSex(?a9)		At(?b9,?l4)
Isa(?s2,Location)			At(?a9,?l4)
Isa(?s4,Location)			Believe(?b9, Solicit(?a9,?b9,?c3))
Isa(?l4,Corner)			Arrested(?a9, Solicit(?a9,?b9,?c3),?b9)
Isa(?c2,Apparel)			
Seductive(?c2)			

The final aspect of packaging a new schema is indexing the schema so that the system is "reminded" of it whenever it might be helpful in processing a future narrative. Indexing is performed by having the schema's volitional subactions suggest it. This is appropriate because the new schema may now provide an explanation for why someone is executing the subaction. For example, whenever the system encounters someone soliciting a prostitute, the new DressSolicitArrest schema is suggested. If this schema accounts for other observed information, it may eventually be determined, and add to the interpretation of the event. This process allows the system to understand the solicitation test narrative as described in Section 3.2.1.

4. PSYCHOLOGICAL EVIDENCE

Although a number of experiments have demonstrated people's ability to learn concepts or schemata from two examples using analogy (Gick & Holyoak, 1983) or from many examples using similarity-based induction (Medin et al. 1987), until very recently, there were apparently no experiments that demonstrated people's ability to learn a concept or schema by explaining and generalizing a single example. Nevertheless, as reviewed by Murphy and Medin (1985), there is a substantial amount of psychological research that reveals the important effect subjects' background knowledge and theories of the world have on the process of concept acquisition. This section reviews a series of recent experiments which indicate that people, like GENESIS, can use their existing knowledge to learn a plan schema by observing, explaining, and generalizing a single specific instance presented in a narrative. Ahn, Mooney, Brewer, and DeJong (1987) present a complete review of these experiments and Ahn (1987) details the experimental methodology.

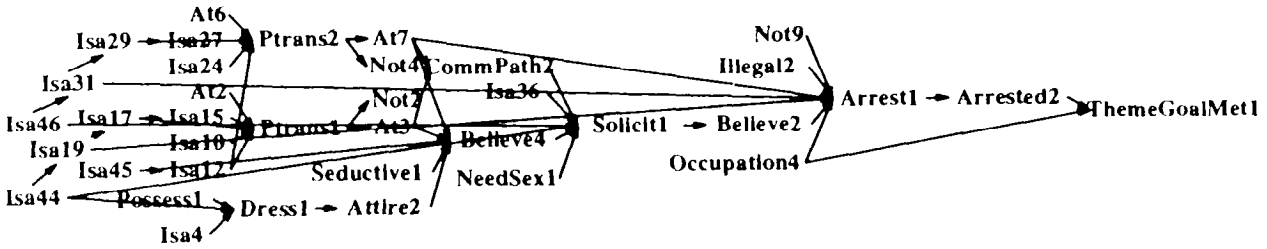

ThemeGoalMet1	?b9 is happy that ?a9 is under arrest for soliciting ?b9's sexual favors for ?c3.
Occupation4	?b9 is a police officer.
Arrested2	?a9 is under arrest for soliciting ?b9's sexual favors for ?c3.
Arrest1	?b9 arrested ?a9 for soliciting ?b9's sexual favors for ?c3.
Believe2	?b9 believes that ?a9 solicited ?b9's sexual favors for ?c3.
Solicit1	?a9 solicited ?b9's sexual favors for ?c3.
NeedSex1	?a9 needs to have sex.
Believe4	?a9 believes that ?b9 is a prostitute.
Attire2	?b9 is wearing ?c2.
Dress1	?b9 put on ?c2.
Isa4	?c2 is an apparel.
Isa44	?b9 is a person.
Possess1	?b9 has ?c2.
Seductive1	?c2 is sexy.

At3	?b9 is at ?l4.
Ptrans1	?b9 went from ?s2 to ?l4.
Isa12	?l4 is a location.
Isa45	?l4 is a corner.
Isa10	?s2 is a location.
Isa15	?b9 is a physical object.
Isa17	?b9 is an animate object.
Isa19	?b9 is a character.
At2	?b9 was at ?s2.
Not2	?b9 is not at ?s2.
At7	?a9 is at ?l4.
Ptrans2	?a9 went from ?s4 to ?l4.
Isa24	?s4 is a location.
Isa27	?a9 is a physical object.
Isa29	?a9 is an animate object.
Isa31	?a9 is a character.
Isa46	?a9 is a person.
At6	?a9 was at ?s4.
Not4	?a9 is not at ?s4.
Isa36	?c3 is money.
CommPath2	?a9 has a communication path to ?b9.
Illegal2	It is illegal for ?a9 to solicit ?b9's sexual favors for ?c3.
Not9	?b9 is not equal to ?a9.

Figure 3. Generalized explanation for the solicitation learning narrative.

The overall design of the experiments involved subjects reading a single narrative describing a specific instance of a novel plan and performing a variety of tests constructed to determine whether or not they had acquired an abstract schema from this single example. Three passages were constructed to present situations for which the subjects presumably did not already have a schema but which they could understand using plan-based understanding. For example, one passage involves a cooperative buying scheme used in some countries. In Korea the system is called a "kyeah," and in India it is called a "chit fund." The experimental narrative describing a single instance of this plan follows:

> Tom, Sue, Jane, and Joe were all friends and each wanted to make a large purchase as soon as possible. Tom wanted a VCR, Sue wanted a microwave, Joe wanted a car stereo, and Jane wanted a compact disk player. However, they each only had $50 left at the end of each month paying after their expenses. Tom, Sue, Jane, and Joe all got together to solve the problem. They made four slips of paper with the numbers 1,2,3, and 4 written on them. They put them in a hat and each drew out one slip. Jane got the slip with the 4 written on it, and said, "Oh darn, I have to wait to get my CD player." Joe got the slip with the 1 written on it and said, "Great, I can get my car stereo right away!" Sue got the number 2, and Tom got number 3. In January, they each contributed the $50 they had left. Joe took the whole $200 and bought a Pioneer car stereo at Service Merchandise. In February, they each contributed their $50 again. This time, Sue used the $200 to buy a Sharp 600 watt 1.5 cubic foot microwave at Service Merchandise. In March, all four again contributed $50. Tom took the money and bought a Sanyo Beta VCR with wired remote at Service Merchandise. In April, Jane got the $200 and bought a Technics CD player at Service Merchandise.

In addition to a group given specific narratives (the instance group), some experiments also used a control group given abstract descriptions of the schemata underlying each of the example narratives (abstract group). The description of the kyeah schema given to the abstract group follows:

> Suppose there are a number of people (let the number be n) each of whom wants to make a large purchase but does not have enough cash on hand. They can cooperate to solve this problem by each donating an equal small amount of money to a common fund on a regular basis. (Let the amount donated by each member be m.) They meet at regular intervals to collect everyone's money. Each time money is collected, one member of the group is given all the money collected ($n \times m$) and then, with that money, he or she can purchase what he or she wants. In order to be fair, the order in which people are given the money is determined randomly. The first person in the random ordering is therefore able to purchase their desired item immediately instead of having to wait until they save the needed amount of money. Although the last person

> does not get to buy their item early, this individual is no worse off then they would have been if they waited until they saved the money by themselves.

Since subjects in the abstract group had been directly told the content of the schema, they were presumed to have learned the schema. Consequently, if the instance group performed as well as the abstract group on a task requiring knowledge of the general schema, then it is reasonable to assume that the subjects in the instance group had also acquired the schema.

In one experiment, subjects in the instance group were asked to "write, in abstract terms, a description of the general technique illustrated in the narrative." In general, subjects were able to produce good descriptions of a plan schema after observing only one instance. The following is the description of the kyeah schema written by one subject:

> Suppose in a group of people, each person would like to buy something expensive, but over a period of time, each person cannot earn enough to buy what he would like. By using random selection, each person could be assigned a number, when the group had saved enough money *together* to purchase an item, the person with the first number would get his item. This would continue for the rest of the group until everyone had gotten what he wished.

In another experiment, subjects in the instance group were told to "write another story in which characters use the general method illustrated in the story but that is otherwise as different as possible" while subjects in the abstract group were told to "write a story in which particular individuals use the technique described in the passage in a specific case." In general, both groups produced equally good examples of the schema and the subjects in the instance group produced examples quite different from the original instance. A third experiment involved having subjects in both groups answer a set of questions constructed to test their understanding of the abstract schema. The two groups were able to answer these questions equally well. The instance group correctly answered 85% of the questions compared with 81% for the abstract group. In general, the equivalent performance of the two groups on tasks designed to test knowledge of the general schema, indicate that subjects in the two groups had learned the schema equally well.

These experiments support the claim that, like GENESIS, people can learn a plan schema by explaining and generalizing a single narrative. However, also like GENESIS, people can only perform EBL when they can explain all the actions composing the schema. Ahn and Brewer (1988) present additional experiments demonstrating subjects' inability to learn plan schemata from a single example when they cannot construct a complete explanation. For example, subjects could not learn a schema for an American Indian potlatch ceremony from one example because they do not have the appropriate cultural knowledge to understand the motivations for all of the actions.

5. RELATION TO OTHER WORK

5.1 Relation to Work in Learning

The small amount of previous work in learning for understanding and plan recognition has employed similarity-based methods. IPP (Lebowitz, 1980) was a text-understanding system which used incremental conceptual clustering methods to learn specializations of existing plan schemata. GENESIS, on the other hand, uses explanation-based methods to learn novel plans that are combinations of existing schemata. However, like IPP, the schemata GENESIS learns affect its understanding of subsequent narratives.

As previously mentioned, most work in EBL has been concerned with classification and problem solving as opposed to understanding. One other EBL system that improves its ability to explain observed behavior is ARMS (Segre, 1988), which learns robot assembly plans from observation. However, unlike GENESIS, ARMS assumes that the system can observe all executed actions and is incapable of using schemata to infer missing actions. A story-understanding system must handle missing information since narratives rarely mention all of the actions that took place.

Most EBL systems construct explanations explicitly for the sole purpose of learning and do not use their explanations for other purposes. In GENESIS, explanations are constructed primarily for the task of understanding. The same explanations used by the learning process also support question answering, paraphrasing, and other understanding tasks. Explanation is seen as a fundamental process that serves many functions, only one of which is learning.

The standard view of EBL is improving efficiency by operationalizing existing concepts (Mitchell et al., 1986). Consequently, EBL has been criticized for its inability to learn facts that are not logically entailed by its existing knowledge, or in other words, for its inability to learn at the *knowledge level* (Dietterich, 1986). Although GENESIS is probably best viewed as operationalizing knowledge for plan recognition, it can also be viewed as knowledge-level learning of specialized concepts like kidnapping, murder for inheritance, and so on. Under the operationalization view, the existing concept being operationalized when learning such plans is something like WEALTH-ACQUISITION-SEQUENCE (Mitchell et al., 1986). This is a very general concept that includes such diverse plans as working for a living, investing in the stock market, stealing, and so on. However, GENESIS' learned-plan schemata are specializations of this concept supported by a particular explanation structure. In general, explanation-based generalization learns specializations of the target concept, that is, it learns only sufficient conditions for being a member of the target concept rather than necessary and sufficient conditions. Since GENESIS has the ability to associate inductively new words like "kidnap" with its specialized plan schemata and to categorize subsequent events as instances of these schemata, it can be

viewed as performing knowledge-level concept learning. Flann and Dietterich (1989) have referred to this type of learning as *theory-based concept specialization.*

The fact that GENESIS learns such specialized concepts from observation means that it can also be viewed as an explanation-based form of *conceptual clustering* (Michalski & Stepp, 1983). The system is given unclassified examples in the form of narratives and categorizes them into conceptual clusters such as kidnapping, murder for inheritance, and so forth. Examples with the same explanation structure are grouped together in the same cluster. Since learned schemata are used to index events in memory, examples of a cluster are also stored in memory. The clusters are "conceptual" because the features in the explanation structure function as a concept definition. This approach to conceptual clustering is very different from the standard similarity-based approach (Fisher, 1987; Michalski & Stepp, 1983).

Generally, EBL is only considered to be useful for improving efficiency. Since understanding is a subjective task (Carbonell, 1979), in addition to improving efficiency, learned schemata can affect how GENESIS interprets an event. For a given narrative, there are generally several possible explanations connecting the individual facts together. Each explanation represents a different interpretation of the text. Most plan-based understanding systems use heuristic search and settle on the first explanation they find. Since GENESIS prefers to understand a narrative using an existing schema, learning a schema can affect interpretation. Like people, the system is influenced by its experience and tends to prefer explanations that were successful in the past. For example, without an explicit schema for "solicitation entrapment," it is unlikely that a system would interpret the sample test narrative the way GENESIS does after learning DressSolicitArrest. A nonlearning system would probably find other explanations first, like ones involving the presence of unmentioned police officers in the vicinity of the crime.

5.2 Relation to Work in Text Understanding

Several recent projects in text understanding have focussed on the use of marker passing and intersection search to find connections between inputs and perform plan recognition (Charniak, 1986; Norvig, 1987). This approach allows the search performed in plan-based understanding to take advantage of the parallelism inherent in marker passing. Although it can greatly improve efficiency, marker passing does not change the basic intractability of plan-based search. Consequently, it does not eliminate the need for schema acquisition.

Marker passing is a simple computational process which does not keep track of variable bindings nor examine the consistency and explanatory power of the paths it constructs. Consequently, it tends to find a large number of irrelevant paths which need to be weeded out during a serial

path-evaluation phase. Charniak (1986) states that for the small knowledge bases already tested, only about 1 in 10 paths found by marker passing constitute reasonable explanations. This factor would undoubtedly increase with the size of the knowledge-base and the length of the paths allowed. Therefore, if the gaps between inputs are large, a marker-passing system capable of finding the long paths representing meaningful explanations is also likely to find a large number of irrelevant paths which have to be evaluated and discarded. Consequently, even marker passing is incapable of efficiently understanding certain events in the absence of a schema.

However, marker passing is not incompatible with schema learning. In fact, a marker-passing system could be used to improve GENESIS's ability to explain and learn from novel events. The schemata learned by generalizing such events would alter the knowledge base by shortening the paths between facts in the schema. This would improve the efficiency of understanding future instances of the schema and increase the chance that the schema is used to explain an event.

6. PROBLEMS FOR FUTURE RESEARCH

The effect of learning on performance has not been adequately studied in GENESIS since the current system does not perform enough plan-based search to permit a fair comparison of learning and nonlearning versions. Schema acquisition improves performance because plan verification is inherently incapable of explaining certain events in the absence of a schema. The claim is that although exhaustive search can find explanations when plan verification cannot, its run-time performance would be much worse than a schema-based approach. To demonstrate this claim conclusively, experiments are needed comparing learning and nonlearning versions of a plan recognition system that conducts an exhaustive search. This would constitute a direct empirical analysis of the trade-off between search and knowledge for the task of plan recognition.

On a related topic, research in learning for problem solving has shown, that, unless the utility of learned information is monitored, learning can actually degrade rather than improve overall performance (Minton, 1988). Although GENESIS has several criteria for insuring the utility of learned schemata, additional techniques may be needed to insure that learning actually improves overall understanding performance.

Another area for future research concerns schemata with unexplainable aspects. Unlike plan schemata that can be acquired from one instance using EBL, many scripts, such as a wedding ceremony or a birthday party, contain conventional as well as causally necessary aspects. Recent psychological experiments suggest that people use a combination of explanation-based and similarity-based methods to acquire such concepts (Ahn & Brewer,

1988). Integrated learning methods (Lebowitz, 1986) are needed that can acquire concepts with both explainable and unexplainable aspects. Mooney and Ourston (1989) present an initial approach to this problem.

7. CONCLUSION

This article has shown how EBL can acquire plan schemata from a single observed instance, as well as how learned schemata can improve the performance of plan recognition. The GENESIS system was presented as an example of a narrative understanding system that uses EBL to improve its performance. Learned schemata allow GENESIS to use schema-based understanding techniques when interpreting events, thereby avoiding the expensive search associated with plan-based understanding. In addition to improving efficiency, schema acquisition effects how the system interprets events. The system prefers to reuse explanations from prior experience. Learned schemata also function as new concepts that can be used to cluster examples and index events in memory.

REFERENCES

Ahn, W. (1987). *Schema acquisition from a single example.* Unpublished master's thesis, University of Illinois, Department of Psychology, Urbana.

Ahn, W., & Brewer, W.F. (1988). Similarity-based and explanation-based learning of explanatory and nonexplanatory information. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* (pp. 524–530). Hillsdale, NJ: Erlbaum.

Ahn, W., Mooney, R.J., Brewer, W.F., & DeJong, G.F. (1987). Schema acquisition from one example: Psychological evidence for explanation-based learning. *Proceedings of the Ninth Annual Conference of the Cognitive Science Society* (pp. 50–57). Hillsdale, NJ: Erlbaum.

Bower, G.H., Black, J.B., & Turner, T.J. (1979). Scripts in memory for text. *Cognitive Psychology, 11,* 177–220.

Carbonell, J.G. (1979). *Subjective understanding: Computer models of belief systems* (Tech. Rep. No. 150). New Haven, CT: Yale University, Department of Computer Science.

Charniak, E. (1982). Context recognition in language comprehension. In W.G. Lehnert, & M.H. Ringle, (Eds.), *Strategies for natural language processing.* Hillsdale, NJ: Erlbaum.

Charniak, E. (1986). A neat theory of marker passing. *Proceedings of the National Conference on Artificial Intelligence* (pp. 584–588). San Mateo, CA: Morgan Kaufman.

Chi, M.T., Feltovich, P.J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5,* 121–152.

Cullingford, R.E. (1978). *Script application: Computer understanding of newspaper stories* (Tech. Rep. No. 116). New Haven, CT: Yale University, Department of Computer Science.

DeJong, G.F. (1982). An overview of the FRUMP system. In W.G. Lehnert & M.H. Ringle (Eds.), *Strategies for natural language processing.* Hillsdale, NJ: Erlbaum.

DeJong, G.F., & Mooney, R.J. (1986). Explanation-based learning: An alternative view. *Machine Learning, 1,* 145–176.

Dietterich, T.G. (1986). Learning at the knowledge level. *Machine Learning, 1,* 287–316.

Dyer, M.J. (1983). *In-depth understanding*. Cambridge, MA: MIT Press.

Fikes, R.E., Hart, P.E., & Nilsson, N.J. (1972). Learning and executing generalized robot plans. *Artificial Intelligence, 3,* 251-288.

Fisher, D.H. (1987). Knowledge acquisition via incremental conceptual clustering. *Machine Learning, 2,* 139-172.

Flann, N.S., & Dietterich, T.G. (1989). A study of explanation-based methods for inductive learning. *Machine Learning, 4,* 187-226.

Gick, M.L., & Holyoak, K.L. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15,* 1-38.

Granger, R.H. (1977). FOUL-UP: A program that figures out meanings of words from context. *Proceedings of the Fifth Internatinal Joint Conference on Artificial Intelligence* (pp. 172-178). San Mateo, CA: Morgan Kaufman.

Kautz, H.A., & Allen, J.F. (1986). Generalized plan recognition. *Proceedings of the National Conference on Artificial Intelligence* (pp. 32-37). Los Altos, CA: Morgan Kaufman.

Laird, J., Rosenbloom, P., & Newell, A. (1986). Chunking in Soar: The anatomy of a general learning mechanism. *Machine Learning, 1,* 11-46.

Lebowitz, M. (1980). *Generalization and memory in an integrated understanding system* (Tech. Rep. No. 186). New Haven, CT: Yale University, Department of Computer Science.

Lebowitz, M. (1986). Integrated learning: Controlling explanation. *Cognitive Science, 10,* 219-240.

Medin, D.L., Wattenmaker, W.D., & Michalski, R.S. (1987). Constaints and preferences in inductive learning: An experimental study of human and machine performance. *Cognitive Science, 11,* 299-239.

Meehan, J. (1976), The metanovel: Writing stories by computer (Tech. Rep. No. 74), New Haven, CT: Yale University, Department of Computer Science.

Michalski, R.S., & Stepp, R.E. (1983). Learning from observation: Conceptual clustering. In R.S. Michalski, J.G. Carbonell, & T.M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach.* Palo Alto, CA: Tioga.

Minton, S.N. (1985). Selectively generalizing plans for problem solving. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* (pp. 596-599). San Mateo, CA: Morgan Kaufman.

Minton, S. (1988). Quantitative results concerning the utility of explantion-based learning. *Proceedings of the National Conference on Artificial Intelligence* (pp. 564-569). San Mateo, CA: Morgan Kaufman.

Mitchell, T.M., Keller, R., & Kedar-Cabelli, S. (1986). Explanation-based generalization: A unifying view. *Machine Learning, 1,* 47-80.

Mooney, R.J. (1987). Integrated learning of words and their underlying concepts. *Proceedings of the Ninth Annual Conference of the Cognitive Science Society* (pp. 974-978). Hillsdale, NJ: Erlbaum.

Mooney, R.J. (1990). *A general explanation-based learning mechanism and its application to narrative understanding.* San Mateo, CA: Morgan Kaufman.

Mooney, R.J., & Bennett, S.W. (1986). A domain independent explanation-based generalizer. *Proceedings of the National Conference on Artificial Intelligence* (pp. 551-555). San Mateo, CA: Morgan Kaufman.

Mooney, R.J., & DeJong, G.F. (1985). Learning schemata for natural language processing. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* (pp. 681-687). San Mateo, CA: Morgan Kaufman.

Mooney, R.J., & Ourston, D. (1989). Induction over the unexplained: Integrated learning of concepts with both explainable and conventional aspects. *Proceedings of the Sixth International Workshop on Machine Learning* (pp. 5-7). San Mateo, CA: Morgan Kaufman.

Murphy, G.L., & Medin, D.L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92,* 289-316.

Norvig, P. (1983). Frame activated inferences in a story understanding program. *Proceedings of the Eighth International Joint Conference on Artificial Intelligence* (pp. 624–626). San Mateo, CA: Morgan Kaufman.

Norvig, P. (1987). *Unified theory of inference for text understanding*. Unpublished doctoral dissertation, University of California, Department of Computer Science, Berkeley, CA.

Quinlan, J.R. (1986). Induction of decision trees. *Machine Learning, 1*, 81–106.

Rumelhart, D. (1980). Schemata: The building blocks of cognition. In R.J. Spiro, B.C. Bruce, & W.F. Brewer (Eds.), *Theoretical issues in reading comprehension*. Hillsdale, NJ: Erlbaum.

Rumelhart, D.E., Hinton, G.E., & Williams, J.R. (1986). Learning internal representations by error propagation. In D.E. Rumelhart, & J.L. McClelland (Eds.), *Parallel distributed processing* (Vol. 1). Cambridge, MA: MIT Press.

Schank, R.C. (1982). *Dynamic memory*. New York: Cambridge University Press.

Schank, R.C. (1986). *Explanation patterns: Understanding mechanically and creatively*. Hillsdale, NJ: Erlbaum.

Schank, R.C., & Abelson, R.P. (1977). *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*, Hillsdale, NJ: Erlbaum.

Schank, R.C., & Riesbeck, C. (1981). *Inside computer understanding*. Hillsdale, NJ: Erlbaum.

Segre, A.M. (1988). *Machine learning of robot assembly plans*. Boston, MA: Kluwer.

Wilensky, R.W. (1978). *Understanding goal-based stories* (Tech. Rep. No. 140). New Haven: CT: Yale University, Department of Computer Science.

Wilensky, R.W. (1983). *Planning and understanding: A computational approach to human reasoning*. Reading, MA: Addison-Wesley.