LEARNING SCHEMATA FOR NATURAL LANGUAGE PROCESSING

Raymond Mooney
Gerald DeJong

Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana,IL 61801

ABSTRACT

This paper describes a natural language system which improves its own performance through learning. The system processes short English narratives and is able to acquire, from a single narrative, a new schema for a stereotypical set of actions. During the understanding process, the system attempts to construct explanations for characters' actions in terms of the goals their actions were meant to achieve. When the system observes that a character has achieved an interesting goal in a novel way, it generalizes the set of actions they used to achieve this goal into a new schema. The generalization process is a knowledge-based analysis of the causal structure of the narrative which removes unnecessary details while maintaining the validity of the causal explanation. The resulting generalized set of actions is then stored as a new schema and used by the system to correctly process narratives which were previously beyond its capabilities.

I INTRODUCTION

A natural language system requires extensive knowledge about the world Clearly, if a computer system is to summarize, translate, or answer questions about a text, it must have knowledge about the concepts expressed in the text. Imagine trying to process a narrative describing a bank robbery without knowledge of money and why people want it. This is a conceptual rather than linguistic requirement, and it means that at the heart of a natural language processor there must be a problem solver to infer missing but important concepts, to insure that the narrative phrases are causally related in an appropriate way, and perhaps to guide the linguistic processing [DeJong82l

Schema-based problem solvers [Charniak77, Minsky75, Schank77] have proven themselves more workable for natural language processing applications than their heuristic search counterparts. In order to process a wide range of text, a schema-based natural language processor must possess many schemata, perhaps hundreds of thousands. This presents both practical and theoretical problems. Somehow these schemata must find their way into the system. They cannot all be built in by hand; there are simply too many. Furthermore, hand coding does not allow for dynamic augmentation of world knowledge. This is an important facet of language processing. For example, the word "Skyjacking," is now an accepted newspaper term but was unheard of twenty years ago. Readers have learned it as a by-product of their normal newspaper reading and natural language processing systems must be able to do the same.

We have taken the first steps in this direction at the University of Illinois. A natural language processing system called GENESIS (for GENeralizing Explanations of Stories Into Schemata) has been designed and implemented which acquires new schemata in the normal course of processing narratives.[1] After acquiring new schemata, the system is able to correctly process narratives that were previously beyond its capabilities.

We call the learning process used by GENESIS *explanatory schema acquisition* [DeJong8 3). It is a form of *explanation-based learning* [DeJong85l which can be briefly defined as learning a new problem solving method by analyzing the causal structure of a problem solution. The system is fully implemented and an example sequence demonstrating the system's learning is given later in the paper. A longer version of this paper appears as [Mooney85l

D GENERAL SYSTEM ORGANIZATION

The general organization of the GENESIS narrative processing system is shown in figure 1. First, English input is processed by a parser into a *conceptual representation* (CRep), a case-frame representation which uses some *conceptual dependency primitives* [Schank75] as well as predicates for complex schemata. Currently, we are using an adaptation of Dyer's McDYPAR [Dyer83] for this purpose; however, since the focus of our research is learning, we make no claims about parsing and alternative approaches could be used for this task (eg. [Marcus80, Waltz84]).

The basic task of the *understander* is to construct a causally complete representation called the *model* A model for a narrative has explicit representations for all the inputs as well as the many inferences that must be made to causally connect them together. There are four types of causal links for connecting assertions in the model of a narrative. These are:

precondition: A link between a state and an action it enables.
effect. A link between an action and a resulting state.

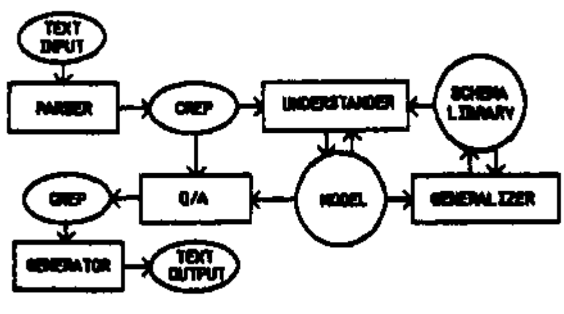

Figure 1: General System Organization

This work wa* supported in pin by the Air Force Office of Scientific Research under gram F4962O-82-K-00O9 ind in pan by the National Science Foundation under graniNSF-IST-17889.

[1] The use of the term "itory" in natural language preening ha* been controvertial (Brewer82] Therefore, In this paper we have adopted the term "narrative" to refer to connected text which may tack a plot or other defining aspect of a "Story."

motivation: A link between a volitional action and the beliefs and goals of the actor which motivated him to perform the action.

inference: A link between a state and another state which it implies.

To avoid confusion, such "causal" links between assertions in the model will be called *support* links, since a precondition of an action *supports* the performance of that action but does not *cause* it The closely related term *data dependency link* [Doyle78] is not used since it is normally reserved for the support of inferences, not for the support of both inferences and actions. Inferring causal connections necessarily employs a large amount of background knowledge which is stored in the *schema library.* The techniques and representations used in this process are similar to those used in past work in narrative understanding [Charniak77, Cullingford78, DeJong82, Dyer83, Wilensky83] and are discussed in sections IV and V.

In order to demonstrate the abilities of the understander, a simple question answering system is used to inspect the model. Since our interests lie in guiding the generalization process through the use of causal relationships, this subsystem is primarily used for accessing the reasons *why* an actor performed a certain action or *why* a particular state exists. This information is easily retrieved by inspecting the support links between the various states and actions in the model. Of course, there are many issues involved in retrieving the most appropriate answers to questions (see lLehnert78]) which we do not fully confront in this subsystem. A simple natural language generator for translating replies into English is also included as part or the system.

if an actor in a narrative achieves an important goal through a novel combination of actions, the *explanation* for how the goal was achieved is generalized into a new schema. The combination of actions which supports the achieved goal state is generalized as far as possible without breaking any of the connecting support links or violating the well-formedness of individual actions and states. This generalized structure is then stored as a new schema in the library where it is used to facilitate the processing of future narratives. This generalization process is discussed in section VI and is the key to the learning technique of explanatory schema acquisition.

III AN EXAMPLE

Currently, GENESIS has acquired two new schemata. In one example, the system learns a schema for someone kidnapping an individual and holding them for ransom. In the other, it learns a schema for someone burning his own building to collect the insurance. Here we will show the performance of the system on the kidnapping example. Before processing the following narratives, GENESIS contains information in its schema library about bargaining, capturing and confining individuals, threatening, and many other concepts; however, it does not have a schema for kidnapping for ransom. First it receives the following "tesT narrative:

INPUT: Ted is the husband of Alice, He won $ 100000 in the lottery. Bob imprisoned Alice in his basement Bob got $75000 and released Alice,

Processing Story.
Finished processing.

No thematic goals achieved by a novel combination of actions: no generalization.

Ready for questions:

? (Who gave Bob the money)
Answer unknown.

? (Why did Bob lock Alice in his basement)
Cannot find sufficient reason for his action.

? (Why did Bob release Alice)
Cannot find sufficient reason for his action.

Notice that in this narrative it is not mentioned how Bob got the money or why Bob imprisoned Alice and then released her. Since the system does not have a schema for kidnapping, it cannot infer the missing information and construct a causally complete explanation of the narrative. Next, it is given the following narrative.

INPUT: Fred is the father of Mary and is a millionaire. John approached Mary. She was wearing blue jeans. John pointed a gun at her and told her he wanted her to get into his car. He drove her to his hotel and locked her in his room. John called Fred and told him John was holding Mary captive. John told Fred if Fred gave him $ 250000 at Trenos then John would release Mary. Fred gave him the money and John released Mary.

Processing Story-
Finished processing.

John achieved the thematic goal: John has $ 250000.
Generalizing- Assembling new schema: CAPTURE-BARGAIN.

Ready for questions:

9 (Why did John imprison Mary in his room)
So John and Fred could make a bargain in which John released Mary and Fred gave John 250000 dollars at Trenos restaurant

? (Why did John make the bargain with Fred)
Because John wanted to have 25000 dollars more than he wanted to hold Mary captive.

9 (Why did Fred make the bargain with John)
Because Fred wanted Mary to be free more than he wanted to have 250000 dollars.

? (Why did Fred want Mary free)
Because Fred was Mary's father.

Unlike the first narrative, this one is detailed enough to allow GENESIS to causally connect the individual actions. The resulting causal structure is then generalized into a new schema of kidnapping for ransom (which GENESIS calls CAPTURE BARGAIN). Next, the system is given the first narrative again, and using the schema it has just acquired, it is able to infer the missing information and causally connect the actions. Consequently, it is able to answer the questions which previously it could not answer.

INPUT: Ted is the husband of Alice. He won $ 100000 in the lottery. Bob imprisoned Alice in his basement Bob got $75000 and released Alice.

Processing Story-
Finished processing.

No thematic goals achieved by a novel combination of actions no generalization.

Ready for questions:

? (Who gave Bob the money)
Ted gave Bob 75000 dollars.

? (Why did Bob lock Alice in his basement)
So Bob and Ted could make a bargain in which Bob released Alice and Ted gave Bob 75000 dollars.

? (Why did Bob release Alice)
Because Bob and Ted made a bargain in which Bob released Alice and Ted gave Bob 75000 dollars.

IV KNOWLEDGE REPRESENTATION

GENESIS* knowledge is represented in a library of schemata: packets of general information about stereotypical objects, situations, and actions. All schemau in the library are arranged in a hierarchical inheritance net under the three major classes of ACTION, STATE, and OBJECT (the highest level class is simply called SCHEMA). Each schema has a set of *roles* associated with it which can be filled by other schemata to create an insunce of the schema. The type of information associated with a schema depends on whether it is an ACTION. STATE, or OBJECT so each of these will be discussed in turn.

ACTION schemau represent dynamic events which change the sute of the world. The following pieces of information are atuched to ACTION schemau. In addition to the information atuched directly to a particular schema, each ACTION inherits the information atuched to ACTIONS above it in the abstraction hierarchy.

Role Constraints:	Each role is marked with the type of schema which can legally fill it.
Defaults	Default fillers can be specified for each role.
Preconditions:	States which must be true in order for the action to uke place.
Motivations:	Sutes (BELIEFS and GOALS) which explain why an actor would perform this action.
Effects:	Sutes which are true after the action is performed.
Terminations:	Sutes which are no longer true after the action is performed. (These are similar to the delete-lists in STRIPs but sutes are temporally marked as no longer holding instead of being deleted from the model.)
Expansion Schemau:	A set of lower-level sutes and actions which actually make up this action along with the support relationships between them (similar to the body of a script).
Suggested Schemata:	Larger composite actions which this action may be a part of.
Determining Conditions:	A set of lower-level actions and sutes which if all present indicate the occurrence of this action.

STATES, on the other hand, represent relatively sutic situations in the world, such as an individual being someone's father or being in possession of some object The following pieces of information are atuched to STATE schemata. In addition to the information atuched directly to a particular schema, each STATE inherits the information atuched to STATES above it in the abstraction hierarchy.

Role Constraints:	Each role is marked with the type of schema which can legally fill it
Defaults:	Default fillers can be specified for each role.
Inferences:	Other sutes which are reasonable inferences to make from this sute.
Achieving Actions:	Actions which can be used to achieve this sute.

OBJECTS represent types of things in the world The information atuched to OBJECT schemau varies from class to class. Cornmon examples for physical objects would be defaults for size, shape, and other physical attributes.

V THE UNDERSTANDING PROCESS

Since applying explanatory schema acquisition depends on having a causal chain of actions to generalize, the "undemanding" ability of GENESIS is concentrated on constructing this chain by inferring missing information and causally connecting inputs together. We do not attempt to deal with other important issues which have recently occupied researchers in narrative undemanding such as *plot units* [Lehnert82l *thematic abstractions units* [Dyer83] *story points* [Wilensky83l affect [DyerS3l and *integrated parsing* [DeJong82, Dyer83).

In accomplishing the usk of constructing causal connections, GENESIS, like FAUSTUS [Norvig83, WilenskyS3l uses a combination of top-down and bottom-up processing techniques. If a set of inputs in a narrative matches a schema which the system already has, then it uses top-down processing to fill in the *expansion* of this schema with the particular inputs of this narrative, much like a script driven program such as SAM [Cullingford78] or FRUMP (DeJong82). However, if an action in the narrative is not explained by a known schema, it attempts to conned it 10 other actions and sutes in the narrative by searching for existing sutes which fulfill the preconditions for this action, or by hypothesizing intermediate actions which causally connect it to existing sutes or actions. In this way, it also operates in a more bottom-up fashion like plan-based programs such as PAM [Wilensky83).

A. Schema Activation and Determination

If a schema-based system is to be able to process a range of possible inputs, it must have access to a large number of schemata. Therefore, in order to avoid repeated searching through the entire daubase of schemata, it must also have an efficient method for selecting the particular schemau which are applicable to the current input Several researchers have addressed this difficult problem [Charniak78, DeJong82, Norvig83) and below is a brief description of the approach GENESIS uses.

When GENESIS processes an input, it adds it to the model and *activates* all the schemau in the list of *suggested schemata* atuched to the schema class of the input. Active schemau then monitor subsequent inputs and check if they match parts of its expansion and can therefore be considered part of this active schema. When all the *determining conditions* of an active schema are met, it is *determined* or considered to have occurred in the narrative and is added to model along with the schemau and support relationships given in its expansion.[2] If a determining condition is an action, then it is also considered to have occurred if all of its effects are in the model.

B. Bottom-up Construction of Support Relationships

When a new schema insunce is added to the model (either as the result of an input or an inference on the part of the system), the system first tries to explain it as part of a known schema. However, if the new insunce does not suggest any higher-level schemau nor match part of any already active schemata, then GENESIS tries to causally connect it to other actions and sutes in the model using planning information.

The first step in integrating a new schema insunce into the model is to add any *primary* inferences or effects. The effects and inferences atuched to a schema are divided into *primary* and *secondary* categories. Primary ones are used in a forward inferencing fashion while secondary ones are used in a backwards inferencmg fashion and only added to the model if they are required by the explanation.

If the new insunce is an action, then its preconditions must be *reconciled* with the model. This means that it first searches the model for each precondition and if it finds it, it adds an appropriately labeled support link from it to the new action. If it does not find a precondition, it next attempts to infer it by searching for secondary effects or inferences which match this precondition. If this also fails, it hypothesizes the existence of an action which can be used to achieve this precondition (using the *achieving actions* atuched to this sute) and attempts to reconcile its preconditions with the model.

GENESIS also attempts to find *motivations* for volitional

[2] The term determined *is* borrowed from FAUSTUS (Norvig83) which also uses a multi-ftep schema selection process.

actions. It does this by checking if the action achieves a state which is a goal for the actor, or if it the actor has the specific goals and beliefs marked as possibly motivating this action. Goals which arise from known *themes* [Schank77] can be automatically inferred if they will motivate a character's action. Such goals will be called *thematic goals* and represent the highest level goals which motivate a person. Goals of possessing money, satisfying hunger, and preserving one's health are examples of thematic goals.

Other plan-based understanding systems such as PAM [Wilensky83) used planning information to predict future courses of actions a character might take. However, searching through a space of possible future actions is combinatorially explosive. Consequently, such an approach is intractable if a system's knowledge of actions is large, which it obviously must be if it is to be able to understand a wide range of narratives. For example, if a PAM-like system were used to process the detailed kidnapping narrative, it would conduct an exhaustive search for an explanation of why John captured Mary before continuing to process the rest of the text. If the system had a large knowledge base of actions, it would be a long time before it stumbled upon the idea of using the action of releasing Mary as part of a bargain with another person.

Since GENESIS does not conduct a complete search for an explanation, it is incapable of "understanding" narratives which have large gaps and do not suggest known schemata. When the first kidnapping narrative is processed without a schema for "kidnapping for ransom," very little of the missing structure can be inferred and the only support links the system can construct are shown in figure 2. As a result, it is unable to answer the questions shown earlier. However, with the same initial knowledge, it is able to understand the second narrative because the gaps and missing information are not too severe. Consequently, the understander is

POSSESS1	Ted has $100,000.
WIN1	Ted wins $100,000 in the lottery.
POSSESS3	Bob has $75,000.
ATRANS1	Someone gives Bob $75,000.
FREE2	Alice is free.
SET-FREE1	Bob releases Alice.
HELD-CAPTIVE1	Bob is holding Alice captive.
CAPTURE1	Bob captures Alice.
FREE1	Alice is free.

Figure 2: Support Network for Narrative #1 Before Learning

able to construct the support network shown in figure 3. It should be noted that the support networks shown in this paper (called *highest-level support networks)* contain only the highest level schemata which were determined to be in the narrative. Most of ihe representation at the level of the inputs and their connecting inferences is contained in the expansions of the CAPTURE and BARGAIN schemata which were activated bottom-up from the inputs.

As indicated earlier, this structure is then generalized into a new schema. When the first narrative is processed again, Bob's action of imprisoning Alice in his basement determines a CAPTURE schema, and this in turn suggests the new "kidnap" schema. The new schema is then used in a top-down fashion to fill in missing information. It is finally determined when both of the effects of the BARGAIN: Alice becoming free again and Ted receiving money, are added to the model. The final support network (the expansion of the new schema for this narrative) is shown in figure 4. This causal structure allows the system to answer the questions it could not answer before learning the schema.

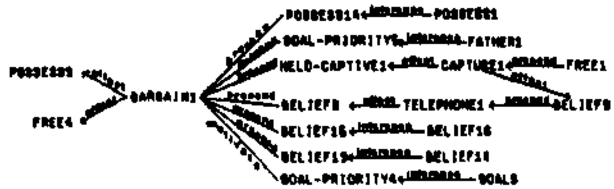

POSSESS9	John has $250,000.
FREE4	Mary is free.
BARGAIN1	John makes a bargain with Fred in which John releases Mary and Fred gives $250,000 to John.
POSSESS14	Fred has $250,000.
POSSESS1	Fred has millions of dollars.
GOAL-PRIORITY5	Fred wants Mary free more than he wants to have $250,000.
FATHER1	Fred is Mary's father.
HELD-CAPTIVE1	John is holding Mary captive.
CAPTURE1	John captures Mary.
FREE1	Mary is free.
BELIEF8	Fred believes John is holding Mary captive.
TELEPHONE1	John calls Fred and tells him that he is holding Mary captive.
BELIEF9	John believes he is holding Mary captive.
BELIEF15	John believes Fred has $250,000.
BELIEF16	John believes Fred has millions of dollars.
BELIEF13	John believes Fred wants Mary to be free more than he wants to have $250,000.
BELIEF14	John believes Fred is Mary's father.
GOAL-PRIORITY4	John wants to have $250,000 more than he wants to hold Mary captive.
GOAL9	John wants to have $250,000.

Figure 3. Highest-level Support Network for Narrative #2

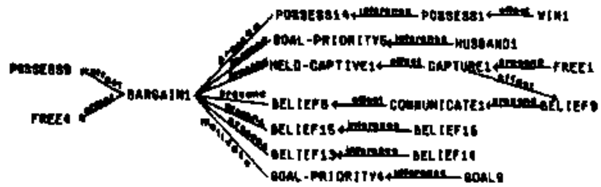

POSSESS9	Bob has $75,000.
FREE4	Alice is free.
BARGAIN1	Bob makes a bargain with Ted in which Bob releases Alice and Ted gives $75,000 to Bob.
POSSESS14	Ted has $75,000.
POSSESS1	Ted has $100,000.
WIN1	Ted wins $100,000 in the lottery.
GOAL-PRIORITY5	Ted wants Alice free more than he wants to have $75,000.
HUSBAND1	Ted is Alice's husband.
HELD-CAPTIVE1	Bob is holding Alice captive.
CAPTURE1	Bob captures Alice.
FREE1	Alice is free.
BELIEF8	Ted believes Bob is holding Alice captive.
COMMUNICATE1	Bob contacts Ted and tells him that he is holding Alice captive.
BELIEF9	Bob believes he is holding Alice captive.
BELIEF15	Bob believes Ted has $75,000.
BELIEF16	Bob believes Ted has $100,000.
BELIEF13	Bob believes Ted wants Alice to be free more than he wants to have $75,000.
BELIEF14	Bob believes Ted is Alice's husband.
GOAL-PRIORITY4	Bob wants to have $75,000 more than he wants to hold Alice captive.
GOAL9	Bob wants to have $75,000.

Figure 4: Support Network for Narrative #1 After Learning

VI THE GENERALIZATION PROCESS

Once a causally complete explanation has been constructed to learn a useful new schema. If so, it generalizes the causal structure in the model into a new schema and stores it in the schema library where it can be used in the processing of future narratives.

A. When to Generalize

If every combination of actions the system encountered was generalized into a new schema, the system would soon become overloaded with rarely used schemata. Most actions would activate a large number of schemata and selecting among these would require an excessive amount of processing time. In order to avoid this problem, certain conditions must be met before a combination is generalized.

Tim, the combination of actions should achieve a goal for one of the characters in the narrative. In the process of motivating actions, the understander checks if an action achieves a goal for a character, so finding achieved eoals is a simple matter of inspecting the model. Second, this goal should be a common one which is likely to be encountered again. A goal is considered to be common enough if it is a thematic goal, for example satisfying hunger or acquiring money. Since in the sample narrative John achieves the thematic goal of possessing money, it satisfies both of these conditions.

The final condition for generalization is the obvious one of not already possessing a schema for the combination of actions which achieves the thematic goal. This simply involves checking the highest-level support for the achievement of the goal and making sure it contains a combination of actions. If the system already had a schema for this case, it would have used it in processing the narrative and the goal would be supported by an instance of this schema instead of a combination of actions. Additional conditions for generalization were discussed in [DeJong83i however, these are currently not implemented.

B. Generalizing the Support Network

After deciding to generalize, the system extracts the explanation for the goal state, isolating the actions and states which actually contribute to its achievement. This simply involves extracting the highest-level support for the achieved state. In the example, the achieved state is John possessing $250,000 and the support for this state is shown in figure 3. This step eliminates extraneous information in the narrative which does not contribute to the achievement of the goal, such as the fact in the example that Mary was wearing blue jeans.

Once the support network is extracted, there are several steps involved in constructing a generalized version of this causal structure. The overall approach is to initially generalize as far as possible and then re-introduce only the constraints necessary to maintain the causal connections between schemata and the well-formedness of individual schemata. Initially, the class of each schema instance is generalized to SCHEMA (the highest level in the hierarchy) and each role filler is replaced by a unique new parameter. Constraints are then imposed on this over-generalized structure to make it a well-formed causal network. These constraints progressively refine the class of each instance and constrain certain role fillers to be equal.

First, the goal which the support network achieves is constrained to be a thematic goal. This is accomplished by constraining it to match the pattern for the thematic goal which was achieved in the original narrative. In the kidnapping example, this constrains the goal state to be the kidnapper acquiring money.

Next the *inierschema* constraints are imposed. These involve maintaining the validity of each connecting support link in the network. We will use the kidnapping narrative to illustrate this process by showing how the FATHER relationship in its support network is only constrained by the explanation to be a POSITIVE-IPT (for positive-interpersonal-theme, a superclass ofPrARENT, SPOUSE, etc.). Every time a support link is added during understanding, it is annotated with the pattern from the schema library used to construct it and the class in the schema hierarchy where it was inherited from. In the example, when GENESIS infers that Fred wants Mary free more than he wants to have $250000 as a secondary inference from the fact that he is her father, it annotates with the corresponding inference pattern from the schema library and the fact that this inference was inherited from the schema POSITIVE-IPT. When the interschema constraints for this link are imposed, the ransom payer's GOAL-PRIORITY is constrained to match the system's inference pattern, and the instance which was a FATHER state in the original narrative is constrained to be a POSITIVE-IPT. Thus, the system only imposes the required relationship between the individuals filling the roles of kidnap victim and ransom payer. The fact that there was a specific father-daughter relationship in this particular narrative is recognized as incidental and not crucial in maintaining the validity of the explanation.

Next, the *intraschema* constraints are imposed. These concern maintaining the well-formedness of each individual schema instance. This is accomplished by imposing on the filler of each role the appropriate *role constrain!* from the schema library. For example, since the role constraints specify that the SUBJECT of a POSSESS schema be a PERSON, the SUBJECT role filler of each POSSESS schema in the support network is constrained to be a PERSON.

The final step in constructing a generalized support network is to merge parameterized instances which have been constrained to be equal. The resulting instances form a set of conceptual roles for the overall schema. In the example, this collects together all the individual occurrences of the kidnapper, the victim, the ransom payer, and the ransom money and creates a unique OBJECT for each one.

The result of this generalization process is a general causal structure which achieves a common goal. The generalized support network generated for the kidnapping example is shown in figure 5.

POSSESS17	Person4 has Money1.
FREE10	Person5 is free.
BARGAIN2	Person4 makes a bargain with Person8 in which Person4 releases Person5 and Person8 gives Money3 to Person4.
POSSESS18	Person8 has Money2.
GOAL-PRIORITY7	Person8 wants Person5 free more than he wants to have Money3.
POSITIVE-IPT2	There is a positive interpersonal relationship between Person8 and Person5.
HELD-CAPTIVE6	Person4 is holding Person5 captive.
CAPTURE2	Person4 captures Person5.
FREE7	Person5 is free.
BELIEF17	Person8 believes Person4 is holding Person5 captive.
COMMUNICATE1	Person4 contacts Person8 and tells him that he is holding Person5 captive.
BELIEF18	Person4 believes he is holding Person5 captive.
BELIEF19	Person4 believes Person8 has Money3.
BELIEF21	Person4 believes Person8 wants Person5 to be free more than he wants to have Money3.
BELIEF22	Person4 believes there is a positive interpersonal relationship between Person8 and Person5.
GOAL-PRIORITY9	Person4 wants to have Money3 more than he wants to hold Person5 captive.
GOAL10	Person4 wants to have Money3.

Figure 5: Generalized Support Network for Narrative #2

C Packaging Into a New Schema

The final step in acquiring a new schema is separating the generalized support network into preconditions, effects, expansion schemata, etc., which can be added to the schema library. Following is an outline of how this information is extracted

Roles (with constraints):	The subject of the achieved thematic goal becomes the actor of the new schema. New roles are created for each remaining person and object in the generalized support network. In the example, roles are created for the kidnapper, victim, ransom payer, and ransom money.
Preconditions:	STATES which are leaves of the generalized support network but not *motivations* of actions by the main actor. In the example, the ransom payer possessing the ransom money is a precondition.
Motivations:	STATES which are leaves of the generalized support network and *motivations* of actions by the main actor. In the example, the kidnapper wanting to have money is a motivation.

Effects:	Effects of all actions within the generalized support network which are not terminated by other internal actions. In the example, the kidnapper possessing the ransom money is an effect
Terminations:	STATES which are terminated by an action within the generalized support network but not produced by another internal action. In the example, the ransom payer no longer possessing the ransom money is a termination.
Determining Conditions:	The set of all ACTIONS within the generalized support network. In the example, the bargain between the ransom payer and the kidnapper is a determining condition.
Expansion Schemata:	All remaining schemata in the generalized support network along with their connecting support relationships.
Suggested Schemata:	The schema for each action within the generalized support network is marked as suggesting the new composite action. In the example, the schema CAPTURE now suggests the new schema.

VII RELATION TO OTHER WORK

The process just described uses a database of background knowledge to generalize the causal structure or *explanation* of a single example. This approach differs dramatically from most approaches to learning (e.g. [Michalski83, Mitchell78, Winston70]) in which generalization is accomplished by extracting features which are shared by a number of examples.

GENESIS* generalization process is most similar to the method used by STRIPS to generalize planning sequences into new MACROPs [Fikes72] However, unlike STRIPS, GENESIS generalizes actions and states as well as objects and locations, and generalizes the order of independent actions (since it uses a dependency network instead of a linear ordering of steps).

The general technique used by GENESIS, explanatory schema acquisition, is also being applied to learning theorem proving strategies [0'Rorke84l robot assembly tasks lSegre85l and concepts in physics problem solving [Shavlik85]. Explanatory schema acquisition is closely related to a growing body of recent work in *explanation-based or analytic* learning [Minton 84, Mitchell83, Silver83, Winston82] which is characterized by learning from a single example through the analysis of its causal structure.

However, there are also important differences between explanatory schema acquisition and some of the other work referenced above. While GENESIS learns from the problem solving behavior of other agents, STRIPS and Mitchell's LEX learn from their own problem solving. Although learning from external behavior makes a system less autonomous, it allows a system to learn plans which are beyond its own ability to generate. In addition, LEX only learns heuristics for applying operators it already possesses and not new combinations of operators which achieve important goals (MACROPs or schemata). Although Winston's system, like GENESIS, learns from short narratives, it learns if-then rules and not schemata. In addition, Winston's system does not need to infer causal connections during "understanding" since all causal connections are given explicitly in the input text

VIII CONCLUSION

Unlike most learning systems, explanatory schema acquisition does not depend on correlational evidence. Thus, it is capable of one trial learning. Also, it avoids the problem of searching through a large space of features for ones which are relevant to a new concept. Only features which contribute to the explanation of an achieved goal are considered for inclusion in the description of a concept. The approach is heavily knowledge-based; a great deal of background knowledge must be present for learning to take place. Finally, the system does not increase its representation power with this kind of learning. The learning results in greatly improved efficiency of processing by avoiding combinatonally explosive searches.

In the future we plan to address the issue of schema refinement. Clearly, the system ought to have the capability of refining existing schemata if the svstem is presented with an example which violates its expectations. We also hope to explore language learning. It should be possible to acquire the English names for these new problem solving schemata from context. This direction of parallel and interacting language and concept development should complement existing work on inducing grammars [Berwick82] and learning to attach new names to known concepts lSelfndgeM].

ACKNOWLEDGEMENTS

We would like to thank Paul ORorke and Alberto Segre for helpful comments on various versions of this paper.

REFERENCES

[Berwick82] R. Berwick. "Locality Principles and the Acquisition of Syntactic Knowledge," Ph.D. Thesis, MIT Department of Electrical Engineering and Computer Science, Cambridge, MA, 1982.

[Brewer82] W. F. Brewer, "Plan Understanding, Narrative Comprehension, and Story Schemas," *Proceedings of the National Conference on Artificial Intelligence,* Pittsburgh, PA, August 1982, 262-264.

[Charniak77] E Charniak, "MS MALAPROP, A Language Comprehension System," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence,* Cambridge, MA, 1977.

[Charniak78] E Charniak, "With Spoon in Hand this Must be the Eating Frame," *Theoretical Issues in Natural Language Processing* 2, Urbana, IL, 1978.

[Cullingford78] R. E Cullingford, "Script Application: Computer Understanding of Newspaper Stories," Technical Report 116, Department of Computer Science, Yale University, New Haven, CT, January 1978.

[DeJong82] G. F. DeJong, "An Overview of the FRUMP System," in *Strategies for Natural Language Processing,* W. Lehnert and M. Ringle (ed.), Lawrence Erlbaum and Associates, Hillsdale, NJ , 1982.

[DeJong83] G. DeJong, "Acquiring Schemata through Understanding and Generalizing Plans," *Proceedings of the Eighth International Joint Conference on Artificial Intelligence,* Karlsruhe, West Germany, 1983.

[DeJong85] G. DeJong, R. Mooney, P. O'Rorke, S. Rajamoney, A. M. Segre and J. Shavlik, "A Review of Explanation-Based Learning," Technical Report in preparation. Coordinated Science Laboratory, Urbana. IL, 1985.

[Doyle78] J. Doyle, "Truth Maintenance Systems for Problem Solving." Technical Report Technical Report-419, MIT Artificial Intelligence Ub, Cambridge, MA, 1978.

[Dyer83] M. Dyer, In Depth Understanding, MIT Press, Cambridge, MA, 1983.

[Fikes72] R. E. Fikes, P. E Hart and N. J. Nilsson, "Learning and Executing Generalized Robot Plans," Artificial Intelligence 3, (1972), 251-288.

[Lehnert78] W. Lehnert, The Process of Question Answering, Lawrence Erlbaum and Associates, Hillsdale, NJ, 1978.

[Lehnert82] W. G. Lehnert, "Plot Units: A Narrative Summarization Strategy," in Strategies for Natural Language Processing, M. H. Ringle (ed.), Lawrence Erlbaum and Associates, Hillsdale, NJ, 1982,375-414.

[Marcus80] M. P. Marcus, A Theory of Syntactic Recognition for Natural Language, MIT Press, Cambridge, MA, 1980.

[Michalski83] R. S. Michalski, "A Theory and Methodology of Inductive Learning," in Machine Learning: An Artificial Intelligence Approach, R. S. Michalski, J. G. Carbonell, T. M. Mitchell (ed.), Tioga Publishing Company, Palo Alto, CA, 1983, 83-134.

[Minsky75] M. L Minsky, "A Framework for Representing Knowledge," in The Psychology of Computer Vision, P. Winston (ed.), McGraw-Hill New York, NY, 1975, 211-277.

[Minton84] & Minton, "Constraint-Based Generalization: Learning Game-Playing Plans from Single Examples," Proceedings of the National Conference on Artificial Intelligence, Austin, TX, August 1984,251-254.

[Mitchell78] T. M. Mitchell, "Version Spaces: An Approach to Concept Learning," STAN-CS-78-711, Stanford University, Palo Alto, CA, 1978.

[Mitchell83] T. Mitchell. "Learning and Problem Solving," Proceedings of the Eighth International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany, 1983, 1139-1151.

[Mooney85] R. Mooney, "Generalizing Explanations of Narratives into Schemata," Technical Report T-147, Coordinated Science Laboratory, University of Illinois, Urbana, IL, May 1985.

[\orvig83] P. Norvig, "Frame Activated Inferences in a Story Understanding Program," Proceedings of the Eighth International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany, August 1983,624-626.

[(VRorke84] P. CVRorke, "Generalization for Explanation-based Schema Acquisition," Proceedings of the National Conference On Artificial Intelligence, Austin, TX, August 1984.

[Schank75] R. C. Schank, Conceptual Information Processing, North-Holland/American Elsevier, Amsterdam, 1975.

[Schank77] R. Schank and R. Abelson, Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures, Lawrence Erlbaum and Associates, Hillsdale, XJ, 1977.

[Segre85] A. M. Segre and G. F. DeJong, "Explanation Based Manipulator Learning: Acquisition of Planning Ability Through Observation," Proceedings of the IEEE International Conference on Robotics and Automation, Su Louis, MO, March 1985.

[Selfridge8l] M. Selfndge, "A Computer Model of Child Language Acquisition," Proceedings of the Seventh International Joint Conference on Artificial Intelligence, Vancouver, B.C. , Canada, August 1981,446-451.

[Shavlik85] J. Shavlik, "Learning about Momentum Conservation," Proceedings of the Ninth International Joint Conference on Artificial Intelligence, Los Angeles, CA, 1985.

[Silver83] B. Silver, "Learning Equation Solving Methods from Worked Examples," Proceedings of the I9S3 I nternaiwnal Machine Learning Workshop, Urbana, IL, June 1983, 99-104.

[Waltz84] D. L Waltz and J. B. Pollack, "Phenomenologically Plausible Parsing," Proceedings of the National Conference on Artificial Intelligence, Austin, TX, August, 1984.

[Wilensky83] R. W. Wilensky, Planning and Understanding: A Computational Approach to Human Reasoning, Advanced Book Program, Addison-Wesley, Reading, MA. 1981

[Winston70] P. H. Winston, "Learning Structural Descriptions from Examples," AI Technical Report-231, MIT Artificial Intelligence Lab, Cambridge, MA, 1970.

[Winston82] P. H. Winston, "Learning New Principles From Precedents and Exercises," Artificial Intelligence 79,(1982X321-350.