SEMANTIC MODELING FOR DEDUCTIVE QUESTION-ANSWERING

K. Biss , R. Chien, F. Stahl, and S. Weissman
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, Illinois 61801

Abstract

A description of techniques used for semantic modeling in a deductive question-answering system is given. The system maintains a dialog and is able to understand situations which can be expressed as a series of sequential time-frames. Specific relevant questions are asked by the system when it is unable to succeed in a given task. It can also provide reasons for its previous actions.

i. Introduction

This paper describes semantic modeling techniques developed for a deductive question-answering system. The problems treated include handling dynamic information, engaging the user in a meaningful dialog and handling a multiplicity of interpretations and assumptions at one time. The data base used to demonstrate these general approaches deals with driving situations. The driver's world was chosen as a data base because it possesses many of the features which make it well-suited as an experimental domain for artificial intelligence research in computer understanding. The processes involved are of sufficient complexity, dynamic in nature, and amenable to codification. The types of facts encountered and deductions performed in the driver's world are representative of those in the real world and would not normally be considered trivial. Since decisions are made based upon events in which the various objects are not stationary, a dynamic representation scheme is essential. In addition, the pertinent rules and regulations have been specified explicitly and relatively unambiguously in the form of laws.

The system described here has been limited to the solving of problems associated with deduction and man-machine interaction about the driver's world. It has been implemented in MICRO-PLANNER. Information is input to the system as MICRO-PLANNER assertions. At present a parser is being adapted to parse natural language input; it is not implemented and is therefore not discussed in the paper. It is based in part on Winograd's PROGRAMMAR.[6]

Research aimed at developing intelligent systems capable of communicating in natural language has been carried on for well over a decade. Simmons[3,4] has surveyed the earlier systems; of the more recent work, Winograd , Charniak[1], and Woods, et al.[7,8] have made significant contributions.

Winograd has shown how a model can be used in conjunction with procedures in a general way. His model has proved very effective in the blocks world. The blocks world does not, however, permit the AI researcher to experiment with a number of problems that can be found in domains such as the driver's world. In particular, in the blocks world every object has an exact known location; every object is completely under the control of the model; every object has distinguishing characteristics; usually only unambiguous information is filled in by the model; and the model does not ask the user for specific information when it is needed.

Charniak has relied on extensive use of demons, which are antecedent theorems, and has worked on a model for computer understanding of children's stories. With the demons he is able to fill in many assumptions that are not explicitly stated in the stories themselves. These are invoked whenever appropriate patterns occur in the predicate-like input.

Woods, et al. have demonstrated the practicality of natural language processing by computer in applying these techniques to a LUNAR SCIENCE INFORMATION SYSTEM.

II. Technical Problems Encountered

Researchers dealing with natural language communication with computers have constantly tried to extend the limits of competence of their systems over previous systems. Every system is limited to some extent in the discourse it can maintain with its user. Particular difficulty comes about when the system must either maintain a coherent dialog or cope with information which is not completely specified; or make subtle inferences on how conditions exist in the real world.

On Multiplicity of Interpretations and Assumptions

In an intelligent dialog it may be the case that a given input could generate many possible interpretations or models of the state of the world. As an example, consider the following statements which might be encountered in the driver's world.

> Four cars arrive at an intersection at the same time (1)

and

> Two cars arrive at an intersection at the same time. (2)

In (1) it is possible to generate many interpretations, some of which are shown in Figure 1. These are all valid Interpretations which are physically as well as legally possible. However, it is felt that in a great majority of cases, the following picture will be brought to mind: an intersection of unmarked two-lane highways with one car in the right hand lane in each part of the intersection (Figure 2).

In (2) above, excluding rotations, there are two valid interpretations (Figure 3). One does not appear to be more likely than the other.

It could be argued that when a statement is ambiguous or leads to more than one possible interpretation, that each interpretation should be considered. However, an intelligent approach would consider only the most plausible interpretations. So, while in (1)

This work is an outgrowth of research sponsored by the Joint Services Electronics Program (U.S. Army, U.S. Navy, and U.S. Air Force) under Contract No. DAAB-07-72-C-0259.

presently at Dept. of Electrical and Computer Engineering, University of Michigan, Ann Arbor, Michigan 48104 one interpretation is most plausible, in (2) both interpretations are equally possible.

Sot only is it necessary to choose between various interpretations, but it is also necessary to decide which are the reasonable assumptions that should be made in a given situation. Consider the case where a vehicle is approaching an intersection. In the absence of explicit information what should be assumed about the situation at the intersection? Are there traffic signals, a police offier, other vehicles, etc. that will affect any decisions concerning the car's progress through the intersection? What, if any, questions should be asked to clarify the situation?

In (2) the user may have to be asked which interpretation is intended. Many systems would report a standard insufficient information message, but a preferable response would be

> What are the relative directions of the vehicles?

This response is direct, to the point, and shows a true understanding of the meaning and ambiguities of the input.

It is not really important which assumptions or interpretations are actually made as long as the model that was used to arrive at these decisions corresponds to what we would consider as reasonable. The important point to be made is that any system should be capable of accepting a wide variety of reasonable models.

On Responding

The intelligence of any system is related to the naturalness of its responses. A system which responds with the correct answer inserted into a predefined format appears to be less intelligent than one in which responses are spontaneous and more to the point. Consider the following inputs:

> A car and a truck approach an intersection from different streets. The car has a yield sign. Which one can proceed first? (3)
>
> A car is approaching an intersection. Another car is approaching from a different street and has a yield sign. Which car has the right-of-way? (4)

Examples (3) and (4) describe exactly the same situation except that where there is a truck in (3) there is a car in (A). The answer to (3) is "the truck" while the answer to (4) is "the one without the yield sign". The answer, therefore, depends not only on the situation, but also on the way in which objects can be uniquely identified. Sometimes a question may have to be answered by using its relationship to other things as in (4) rather than by naming it explicitly.

On Deciding What to Ask

When trying to answer a question, the traffic laws usually have to be applied. There are two reasons for any given traffic law to be deemed inapplicable: one or more of the necessary conditions of the law contradict the known information or, alternatively, some of the necessary conditions of the law may be unspecified or unknown. Consider the following:

> A car is on a twolane highway. It is behind a truck which is going below the speed limit.
> Can the car pass the truck?

Any law which is not applicable cannot be considered. But what about the law which says you cannot pass on a curve, or the one which states that you cannot pass in a no-passing zone? It is not clear if these laws are appropriate. Both of these laws would fail because of lack of information. If it is possible to isolate which portion of the law fails because of lack of information, then it would be possible to formulate a question which would ask the user to supply a specified piece of information, such as "does this take place in *a* no-passing zone or on a curve?" This capability would allow the system to engage the user in an intelligent dialog.

In this paper we will describe a system based upon the driver's world and the approaches taken to handle a number of technical problems such as those already mentioned above. The techniques used to overcome these problems involve solutions to time handling, system made assumptions, asking relevant questions and setting up relative scenes.

III. Sample Dialog

The following is a sample dialog with the computer model for the driver's world. A more detailed discussion describing the performance will be given in a later section. Actual input is presently in the form of MICRO-PLANNER assertions rather than natural English, Pseudo-input is in lower case and output is in upper case.

> Three cars approach an intersection at the same time. The car in the middle has a yield sign. The car to the left of this car is going to go straight while the third car intends to make a left turn. Who has the right-of-way?
>
> *THE CAR WITHOUT THE YIELD SIGN AND WHICH DOES NOT INTEND TO TURN LEFT
>
> The car with the right-of-way proceeds. Who has the right-of-way now?
>
> *THE CAR WHICH INTENDS TO MAKE A LEFT TURN

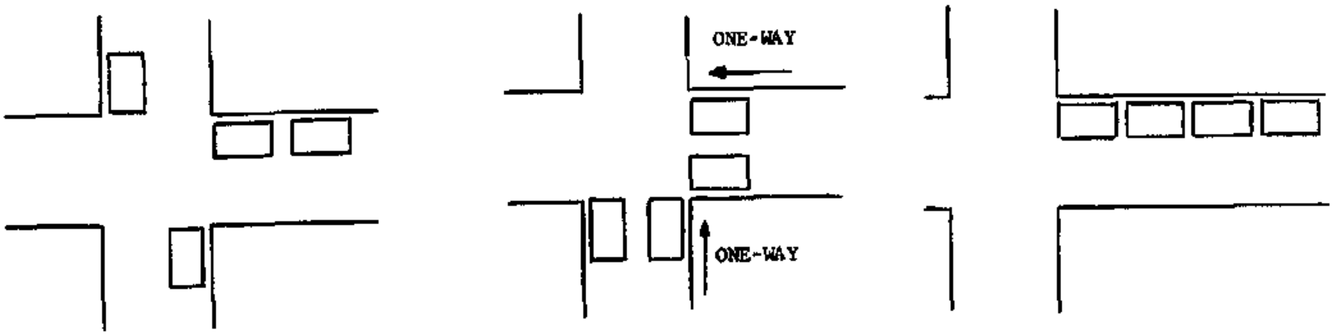

Figure 1. Possible interpretations of "four cars approach an intersection"

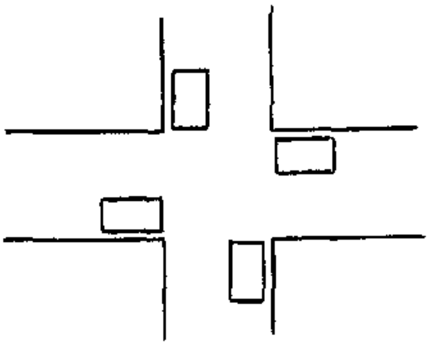

Figure 2. Assumed interpretation of "Four cars approach an intersection"

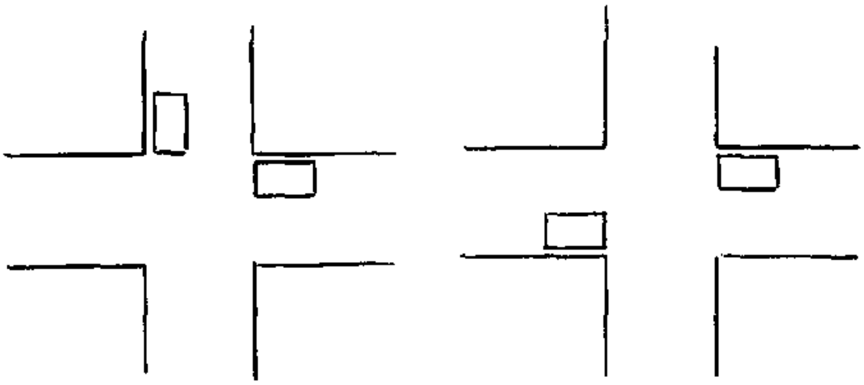

Figure 3. Two interpretations of "two cars approach an intersection"

Suppose, when the three cars approached the intersection the yield sign is replaced by a flashing yellow light. Which car may proceed?

*THE CAR AT THE FLASHING YELLOW LIGHT.

Why?

*THE OTHER TWO DIRECTIONS HAVE YIELD SIGNS.

If this car proceeds, who has the right-of-way?

*THE CAR WHICH INTENDS TO GO STRAIGHT.

If this car proceeds, who has the right-of-way?

*THE CAR WHICH INTENDS TO MAKE A LEFT TURN

A truck in front of a car is travelling below the speed limit. They are on a two-lane highway. Can the car pass the truck on the right?

*NO

Why?

*PASSING ON THE SHOULDER IS PROHIBITED.

Two cars are approaching an unmarked intersection from different directions. Who has the right-of-way?

*WHAT ARE THE RELATIVE DIRECTIONS OR POSITIONS OF THE VEHICLES?

One car is proceeding north. The other car is to its right.

*THE CAR PROCEEDING WEST.

Thank you.

*YOU'RE WELCOME.

IV. A Model for Computer Understanding

Structure

The model for the driver's world is implemented in MICRO-PLANNER. As the input is being accepted, antecedent theorems may be invoked in order to add information to the data base. After the data has been entered any question will be transformed into a goal or series of goals. In attempting to satisfy these goals, theorems which represent the laws and other facts concerning the driver's world are applied. In a few cases the original goals will not be satisfied due to the lack of sufficient information. In attempting to gain more complete information, routines called specialists are invoked in order to determine the most likely default conditions. The applicable specialist will take the input data and make appropriate assumptions based on the available information, and generate a series of time-frames which includes all of the possible interpretations. Another attempt is made to satisfy the original goals and the laws are reapplied. If the goals cannot be satisfied, then the user may be asked to supply additional information by use of the conditional failure mechanism. This added information is followed by repeated attempts to satisfy the goals.

After a question has been successfully answered, it is possible for the user to further interrogate the system, modify the situation, or add additional information, or completely change the direction of the dialogue. The previous interpretations will be updated to reflect the new situation.

A representation of the structure of the system is given in Figure 4. The arrows indicate which routines can be invoked by other routines. Since all information is expressed relative to time-frames, a description of this formalism will be discussed first.

Time-Frames

In this system a completely described state of the world is called an event. It consists of a list of participants (e.g. vehicles) and a list of assertions expressing relationships and attributes among these participants. Every continuous sequence of events is considered to be a series of still pictures called time-frames. A complete time-frame consists of the following: an event, a pointer to the time-frame immediately preceding the present one and a list of pointers to time-frames which the present time-frame precedes, which may be empty or have multiple entries.

There are primitives available to create and destroy time-frames. Whenever a time-frame is created, all of the statements from the previous frame are carried over. It is possible to add or delete assertions from just the present time-frame as well as from the present and all following time-frames. It is also possible to move to time-frames with respect to each other. It is felt that this provides a natural way to handle dynamic data and sequences of events.

In CONNIVER contexts are similar to the time-frames described. However, time-frames, which are embedded in MICRO-PLANNER, have permitted us to use

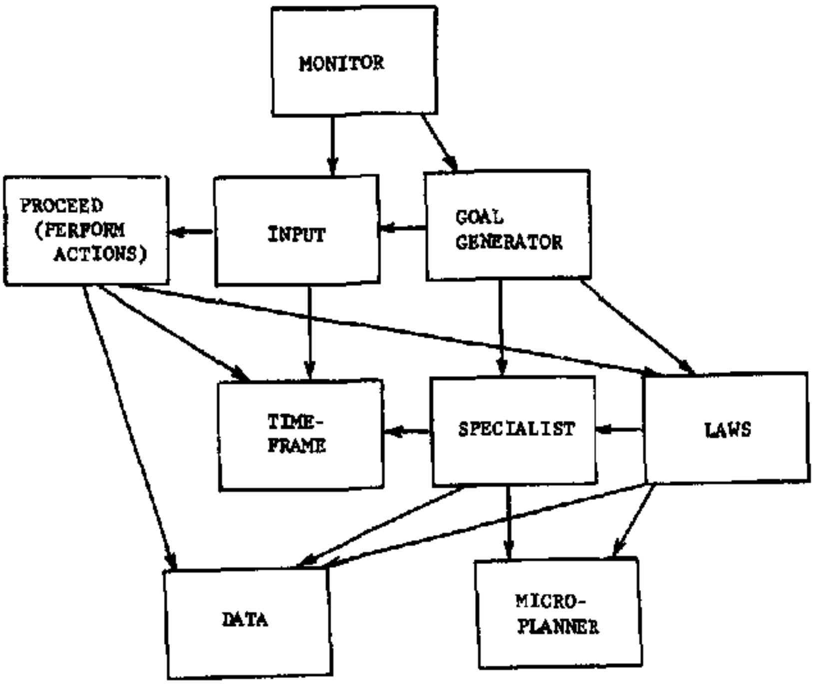

Figure 4. Organization of the System

the pertinent aspects of contexts and yet maintain the salient features of MICRO-PLANNER without excessive overhead as might be the case if implemented in some other manner.

Global Facts

Global facts are theorems which contain, among other things, the properties of objects such as vehicles as well as certain laws of nature in the driver's world such as

> If a car is about to enter an intersection then it is at the intersection.

These facts are expressed in terms of antecedent and consequent theorems.

The rules of law for the driver's world are stored as consequent theorems. An example of one of these laws is given in Figure 5. These lavs are called whenever a legal consequence, such as "who has the right-of-way?" has to be determined.

Some goals have filters associated with them that restrict the application to specific time-frames. In general, any theorem may be used to satisfy a goal, but all data will be restricted with respect to time-frames. (See examples in Figure 5.)

Structure of Assertions

The inputs to the system as well as system generated statementa are stored as a series of predicates and arguments. For example, the input "a car arrives at an intersection" will be represented as

(AI ARRIVE-AT CAR INTERSECTION)

All of the input assertions are stored as statements in the present time-frame.

The statement "two cars arrive at an intersection" would be represented as

(A2 ARRIVE-AT *GR0UP1 INTERSECTION)
(A3 IS *GR0UP1 *CAR1 *CAR2)

Whenever an assertion is made, any antecedent theorem containing global information may be invoked to alter or augment the assertion.

Some of the arguments of the predicate may themselves be names of predicates. For example, "Two cars arrive at an intersection at the same time" is represented as

(A4 ARRIVE-AT *GR0UP2 INTERSECTION)
(A5 AT A4 SAME-TIME)
(A6 IS *GR0UP2 *CAR3 *CARA)

Each of the vehicles or other objects in the system will have on its property list the special traits and modifiers which distinguish it from other objects.

Specialists

Each of the specialists are routines which contain detailed knowledge of some small aspect of the driver's world. The specialists use their knowledge to fill in information, to set up relative scenes and to keep track of scene modifiers. A specialist can be called upon at any time. Unlike Charniak's demons, they do not have to be pattern invoked.

As an example, consider the intersection specialist. This routine can examine scenes dealing with intersections and tries to place all vehicles in a standard location of a standard intersection. A standard intersection is made up of four "arms" each of which is a type of street (e.g. two-lane). The positions which make a standard intersection are shJQMtf in Fig. 6.

The station position is immediately adjacent to the intersection. When approaching, the position can only contain one vehicle, which is the next to enter the intersection. The approach position can contain any number of vehicles. Vehicles in the approach are heading toward the intersection but will not be the first to enter. The leave position will hold any number of vehicles which have left the intersection. The in position contains the one car which has entered the intersection. The sign position denotes the occurence of any traffic control devices. The specialist attempts to assign every vehicle a pair of values, (ARM POSITION), which uniquely describes its position in the time-frame.

The intersection specialist has only one input, the name of the time-frame which contains the assertions for the scene. When the specialist is called, a series of antecedent theorems are asserted. Each of the input assertions is then pseudo-asserted, i.e., any relevant theorem is used but the calling theorem is not asserted (it has been asserted previously).

The time-frame system described previously is used to create and save possible interpretations. As each new vehicle is encountered, the possible number interpretations is increased. As new information is evaluated, any interpretations which contradict this information is deleted.

For example, if it is known that two cars, A and *S,* are *at an intersection,* if the fact that A is to the left of B is input, then any interpretation which contradicts this, such as an interpretation which states that B is to the left of A, is deleted.

Whenever an antecedent theorem has been successfully applied, the assertion which has been used is said to have been processed. If an input assertion has predicates for arguments the predicates must be processed first. The net effect of this is that before a vehicle position or relative position can be modified, the vehicle has to be placed in the intersection. There is a constant check to eliminate equal interpretations (i.e., rotations).

The specialist returns a list of time-frames which are all possible interpretatlon(s). After this processing has been completed the most likely interpretation (or interpretations) is (are) selected.

Among the various specialists are those which deal with two-lane highways, four-lane highways, ramps, alleys, intersections, passing, parking, etc.

Giving Reasons

When some of the laws within the system are successfully applied, it is possible to assign a "reason" for success to some of the objects in the scene.

As an example consider LAW436 (Fig. 5). Every vehicle that is used when the law is satisfied can be placed in one of two classes. The vehicle which must yield has associated with it the "reason" intends to make a left turn while the vehicle which has the right-of-way is tagged with the "reason" does not intend to make a left turn.

In order for the question to be answered properly, it may be necessary to use the reasons to form an answer. If the objects do not have unique names, then the answer must be given in relative terms. In fact, the answer will be the reasons associated with that object.

ENGLISH:

A car which intends to make a left turn at an intersection must yield to all traffic approaching from the opposite direction.

MICRO-PLANNER:

```
(THCONSE LAW436 (X Y)($?A1 MUST-YIELD $?X $?Y)
(THPROG (($R S SCENE) A B C D A2 A3 A4 A5 A6 A7 A8 A9)
(THGOAL ($?A2 AT $?S INTERSECTION)(THDBF PTF))
(THGOAL ($?A3 INTEND $?X $?A4)(THDBF PTF))
(THGOAL ($?A4 MAKE $?X LEFT-TURN)(THDBF PTF))
(THNOT (THAND (THGOAL $?A8 INTEND $?Y $?A9)(THDBF PTF))
              (THGOAL $?A9 MAKE $?Y LEFT-TURN)
                        (THDBF PTF))))
(THGOAL ($?A5 AT $?Y $?A $?B)(THDBF PTF))
(THGOAL ($?A6 AT $?X $?C $?D)(THDBF PTF))
(THNOT (THOR (EQ $?B 'LEAVE) (EQ $?D 'LEAVE)))
(THGOAL ($?A7 OPPOSITE $?A $?C)(THTBF THTRUE))
(THPUTPROP $?X '(INTENDS TO MAKE LEFT TURN) 'REASON)
(THPUTPROP $?Y '(DOES NOT INTEND TO MAKE A LEFT TURN)
               'REASON)
(THSUCCEED)))
```

where PTF is the filter of all statements in the present time-frame.

Figure 5 A sample law

Consider the following:

> Two cars approach an intersection from opposite directions. One car intends to make a left turn. Which car has the right-of-way?

Because the vehicles do not have unique names, a relation must be used to supply the answer. The proper response is the car that does not intend to make a left turn. While the question:

> A car and a truck approach an intersection from opposite directions. The truck intends to make a left turn. Who must yield?

would be answered the car.

Generating Questions

In attempting to answer a question or satisfy a goal, it is possible to encounter failure in the traditional MICRO-PLANNER sense. This may lead to the case where it is impossible to pursue the deduction further. We have implemented in MICRO-PLANNER the notion of conditional failure to cope with a large number of cases of this type.

When a goal fails, there are two possible cases. In the first case, there is no positive information to support the goal but there exists information which would prevent the goal from succeeding no matter what additional information is supplied. In the second case, additional information could be used to prove the goal. Consider the following example: Suppose the data base contains the assertion

```
        (B ON A)                        ; DATA
```

and the theorems

```
(THCONSE X-ON-Z (X Y Z) ($?X ON $?Z)
     (THGOAL($?X ON $?Y)   )          ; GOAL1
     (THGOAL($?Y ON $?Z)   ))         ; GOAL2
(THCONSE NOT-X-ON-Y (X Y) (NOT $?X ON $?Y)
     (THGOAL ($?Y ON $?X) $T)         ; GOAL3
```

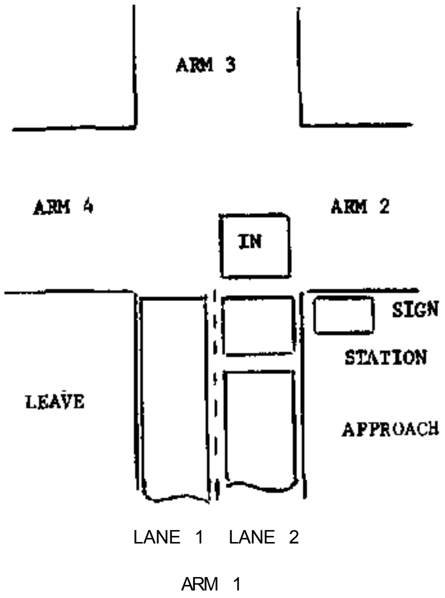

Figure 6. Standard by 2-lane intersection

and further suppose we wish to prove the goals

(THGOAL (A ON B) ST) (1)
(THGOAI. (B ON C) ST) (2)

These goals will ordinarily always fail.

Note, that a distinction can be made between these two failures. In CI) the goal could never succeed because we can prove

(THGOAL (NOT A ON B) $T)

but no such result can be proved for

(THGOAL (NOT B ON C) $T)

In order to take advantage of this distinction, consider the following procedure:

(1) Upon entering specially marked theorems, if a marked subgoal fails without having once succeeded and subsequently backed through, then record this goal on a list called GFAIL.

(2) If the major goal fails in the normal sense, then for each element of GFAIL construct a goal which is its converse.

(3) If any goal in (2) succeeds remove the corresponding goal from GFAIL.

(4) Any remaining elements on GFAIL represent conditional failures (those goals which could conceivably be determined to be true) which are used to construct relevant questions.

In the one example (2) with the goal

(THGOAL (B ON C) ST)

we proceed as follows:

Because (B ON C) is not an assertion, theorem X-ON-Z is invoked. GOALI succeeds with (B ON A). G0AL2, (A ON C) fails and is placed on GLIST. An attempt is made to find another instantiation for GOALI, (B ON S?Y). This fails but is not placed on GLIST because it has previously succeeded. X-ON-Z fails. GLIST now contains only (A ON C).

The converse, (NOT A ON C), is constructed, but this goal fails when using NOT-X-ON-Y. Consequently, the GLIST remains unchanged.

It is now possible to easily construct a relevant question the answer to which could provide the necessary information *to* satisfy the original goal, (B ON C). In this case the question is: IS (A ON C)? If the response is positive, (A ON C) is asserted and the original goal when reattempted will succeed.

It is possible to assign to each element of GFAIL a priority which will direct the order in which the questions are asked when GFAIL has more than one entry.

V. Performance of the Model

The performance of the model can best be judged by describing a sample dialog. The dialog given earlier will serve this purpose. Consider the first statement.

> Three cars approach an intersection at the same time. The car in the middle has a yield sign. The car to the left is going to go straight while the third car intends to make a left turn. Who has the right-of-way?

The actual input *to* the system would be the following MICRO-PLANNEH assertions.

```
(I1 ARRIVE-AT *GROUP1 INTERSECTION)
(I2 AT I1 SAME-TIME)
(I3 AT *CAR1 YIELD-SIGN)
(I4 LEFT-OF *CAR1 *CAR2)
(I5 INTEND *CAR2 GO-STRAIGHT)
(I6 RIGHT-OF *CAR1 *CAR3)
(I7 INTEND *CAR3 TURN-LEFT)
(Q1 HAS-RIGHT-OF-WAY X)
(I8 IS *GROUP1 *CAR1 *CAR2 *CAR3)
```

The notation *CAR1, *CAR2 and *CAR3 for the vehicle indicates the vehicle names are "system assigned" rather than user assigned (e.g. a car and a truck). In order to answer the question the yielding laws will be applied to determine which vehicles must yield. In this case all of the applicable laws will fail because the data is incomplete. The input statements which state that some vehicle is approaching the intersection will lead to the activation of one of the intersection specialists. Initially the model will be a two-lane by two-lane intersection. If the data is contradictory and cannot be accommodated in the model another specialist (e.g. intersection of highways) will be called.

When the two-lane by two-lane intersection specialist is called, the following scene is constructed.

```
(A1 AT *CAR1 ARM1 STATION)
(A2 AT *CAR2 ARM4 STATION)
(A3 AT *CAR3 ARM2 STATION)
(A4 AT ARM1 YIELD-SIGN)
(A5 *SCENE1 INTERSECTION)
```

and the time-frame is

```
(*TF1 (*CAR1 *CAR2 *CAR3)(A1 A2 A3 A4 A5)(GLOBAL)())
```

where GLOBAL is the global time-frame. The assumptions that are supplied by the specialist are that the intersection is composed of a pair of two-lane streets, one with a yield with one empty arm and one car in each of the others. Under these assumptions, the above is the only interpretation excluding rotations. Using the yield laws (e.g.. LAW436), the cars are compared pairwise, and it is determined that *carl must yield to *car2 and *car3 because of the yield sign, and *car3 must yield to *car2 because it intends to turn left. It is then determined that the only logical choice for the answer is *car2. Because this is a system supplied name, the reasons associated with *car2 will be used to answer the question instead of the system generated name. Thus the answer

THE CAR WITHOUT THE YIELD SIGN AND WHICH DOES
NOT INTEND TO TURN LEFT

The next input

"The car with the right-of-way proceeds.
Who has the right-of-way now?

The first input

(19 PROCEED *CAR4)
(110 HAVE-RIGHT-OF-WAY *CAR4)

causes a new time-frame to be created. The vehicles in the time-frames will be advanced, the final position being determined by the applicable procedures.

The following fact is added to the present time-frame:

(A6 AT *CAR3 ARM2 LEAVE)

giving:

(*TF2 (*CARI *CAR2 *CAR3)(A1 A2 A4 A5 A6) <*TFI>)

Regarding the question of right-of-way, the same approach as before is taken except that the new time-frame is now used.

Applications of the laws are straightforward and THE CAR WHICH INTENDS TO MAKE A LEFT TURN is given as an answer.

For the next input:

Suppose, when the three cars approached the
intersection the yield sign is replaced by a
flashing yellow light, which car may proceed?

the following input assertions are used:

(111 WHEN ARRIVE-AT *GROUP2 INTERSECTION)
(112 IS *GROUP2 *CAR5 *CAR6 *CAR7)
(113 REPLACE YIELD-SIGN FLASHING-YELLOW)

The following is added by an intersection specialist:

(A7 ARM1 FLASHING-YELLOW)
(A8 ARM3 FLASHING-YELLOW)
(A9 ARM2 FLASHING-RED)
(A10 ARM4 FLASHING-RED)

and the time-frame

(*TF3 (*CAR1 *CAR2 *CAR3)
(AI A2 A3 A5 A6 A8 A9 A10 A11)(*TF1))

is created.

In this case the specialists provide the information that flashing lights come in pairs at intersections. The opposite directions have flashing yellow and red lights, respectively. The flashing red lights will be interpreted as yields. So the right answer is THE CAR AT THE FLASHING YELLOW LIGHT.

Inputing the question "Why?" does not change the time-frame. When goals succeed "reasons" are stored. These are returned as the answer. So, here the reason is THE OTHER TWO DIRECTIONS HAVE YIELD SIGNS. Note that this default condition may not be true at all times, but the inserted assumptions should be true for a majority of the time.

When asking:

If the car proceeds, who has the right-of-way?

the following assertion and time-frame are formed.

(All AT *CAR1 ARM3 LEAVE)
(*TF4 (*CAR1 *CAR2 *CAR3)
(A2 A3 A5 A5 A7 A8 A9 A10 A11)(*TF3))

Again, the goal succeeds and the answer is THE CAR WHICH INTENDS TO GO STRAIGHT. Now asking

If this car proceeds who has the right-of-way?

adds

(A12 AT *CAR3 ARM2 LEAVE)

and

(*TF5 (*CAR2 *CAR3)
(A3 A5 A6 A7 A8 A9 A10 A12)(*TF4))

to the database. There is only one car at the intersection and it has the right-of-way. So, THE CAR WHICH INTENDS TO MAKE A LEFT TURN is the response. Starting with

A truck in front of a car is travelling below
the speed limit. They are on a two-lane high-
way. Can the car pass the truck on the right?

a new scene is required. The following assertions are made:

(A14 IN CAR LANE1)
(A15 IN TRUCK LANE1)
(A 16 IN-FRONT TRUCK CAR)
(A17 TRAVEL TRUCK)
(A18 BELOW A17 SPEED-LIMIT)
(A19 *SCENE2 TWO-LANE-HIGHWAY)
(*TF6 (CAR TRUCK)<A14 A15 A16 A17 A18 A19)(*TF5))

Notice that the model assumes a two-lane highway. Using the lav dealing with passing on the right on two-lane highways we find that the answer is NO. The next question, "Why?" is answered by using the reason, PASSING ON THE SHOULDER IS PROHIBITED.

The next question is

Two cars are approaching an unmarked intersection
from different directions. Who has the right-
of-way?

Because the law is straightforward, the problem is not who has the right-of-way but what are the relative directions of the two cars.

The intersection specialist is called and returns a list of time-frames containing the possible interpretations. In this case there are two which would be considered equally likely (see Fig. 3). The user is asked which of the interpretations is intended. When the information is given, the question can be answered and the proper response, THE CAR PROCEEDING WEST, is given.

V. Concluding Remarks

We have described techniques for semantic modeling in a deductive question-answering system. The system can maintain a dialog and is able to understand situations which can be expressed as a series of sequential time-frames. The system has the ability to fill-in information which corresponds roughly to those assumptions that might normally be made by a person. The system can ask specific relevant questions of the user because it knows which subgoal failed for insufficient information. It can give reasons for answers based upon the tags left by successfully using the laws. The specialists can set up relative scenes, fill in detailed assumptions about situations, and set up scene modifiers.

References:

1. Charniak, E., "Toward a Model of Children's Story Comprehension", Artificial Intelligence Memo AI TR-266, Massachusetts Institute of Technology, Cambridge, Mass., December, 1972.

2. McDerraott, D. V., Susaman, G. J., "The CONNIVER Reference Manual", Artificial Intelligence Memo No. 259, Massachusetts Institute of Technology, Cambridge, Mass., May, 1972.

3. Simmons, R. F., "Answering English Questions by Computer: A Survey", Communications of the ACM, Vol. 8, No. L, Jan., 1965, pp. 53-69.

4. Simmons, R. F., "Natural Language Question-Answering Systems: 1969", Communications of the ACM, Vol. 13, No. 1, January, 1970, pp. 15-36.

5. Sussman, G. J., Winograd, T., Charniak, E., "Micro-Planner Reference Manual", Artificial Intelligence Memo No. 203A, Massachusetts Institute of Technology, Cambridge, Mass., Dec, 1971.

6. Winograd, T., "Procedures as a Representation for Data in a Computer Program for Understanding Natural Language", Report MAC TR-84, Massachusetts Institute of Technology, Cambridge, Mass., February, 1971.

7. Woods, W, A., Kaplan, R. M., "The Lunar Sciences Natural Language Information System", BBN Report No. 2265, Bolt, Beranek and Newman, Inc., Cambridge, Mass., September, 1971.

8. Woods, W. A., Kaplan, R. M., Nash-Webber, G., "The Lunar Science Natural Language Information System: Final Report", BBN Report No. 2378, Bolt, Beranek and Newman, Inc., Cambridge, Mass., June, 1972.