# Plan Recognition:
# Commitment and Preference

Steve Woods, University of Waterloo, March 1990

In this paper I wish to examine the topic of plan recognition. I will offer several real world examples of planning which are intuitively straightforward, and attempt to bring some of the more important aspects of current plan recognition (PR) work to light in terms of these simple examples. Once these PR concepts have been introduced, I will discuss approaches to PR as described by the authors of three notable PR papers. Once some of the more formal approaches to PR concepts have been introduced, I will attempt to define plan recognition in a more domain independent manner and to bring together some of the wide variety of PR work into something of a comprehensive picture, while still highlighting the differences in some of the notable approaches. Terminology is often a stumbling block when attempting to understand the relationship between the work of many authors, and in an effort to reduce this confusion, I will try to relate the most important ideas from the field into a single complete and understandable description of plan recognition.

Given a clear description of just what plan recognition is, I will describe some of the more recent and promising contributions within this context. In particular, I will describe some approaches to selection of candidate plans which seem to represent new ground in plan recognition, and some of the heuristics which aid in intelligent selection of plans when many plans are available and applicable in certain situations.

At the end of this paper I will offer some general suggestions of my own as to the direction I see PR taking in the future, and discuss what other authors see as important considerations for PR in the future.

## Plan Recognition

Every day people make decisions about how to act and what to say in different situations. In particular, people react to their environment in ways that they believe will bring about desired reactions or results from other people or the world around them. Just crossing the street involves understanding something about the way in which walk signals and other motorists act. In fact, we can view our observations of agents in the real world as our idea of these agents' goals and plans they are likely executing to achieve these goals. This view of the world can perhaps best be seen through example.

Take for instance the case where we wish to cross a street, our walk signal is green, and a car is slowly approaching the intersection. We will infer that the motorist is traveling for some reason, with the most likely goal of getting safely from point A to point B. We also (based on the manner in which the driver is proceeding) may infer that the driver has subgoals in his traveling plan such as not violating traffic signals or not injuring pedestrians. In trying to achieve our goal of crossing the street, we too have a plan, specifically, our plan could be to walk across the street at the nearest intersection.

In choosing to cross the street here and now, we have recognized a portion of the driver's plan, and based on a likely subgoal in this plan we have decided that our own goals of crossing the street, including minimizing our risk, can be satisfied. In this manner, the act of plan recognition has allowed us to act out our plan within our environment or domain. Some authors [All78] describe this close planning/plan recognition interaction strictly within a planning framework, where the actual act of recognizing other agent's plans is an actual subgoal of our own planning process. This approach seems quite logical and consistent with the manner in which people really tend to act and react. Kautz and Allen [Kau86] elaborate on the distinction between planning and plan recognition by describing planning as a process of hypothetical reasoning, while plan recognition is described as a more uncertain process in which actual events that have happened must be represented, and hypothetical explanations of actions must be proposed. PR is described as more uncertain than planning since PR involves recognizing a particular plan an agent is performing, while planning involves merely finding any plan that satisfies certain goals.

In fact, these differences may not be as discerning as Kautz and Allen described them, especially if we consider planning and PR in terms of a hierarchy of specialization. Often in PR, much can be learned about an agent and its plans without fully determining its exact plan. Hierarchical planning and hierarchical PR both allow for various levels of elaboration of plans. For instance, in performing PR, we may determine a general plan that someone is performing, such as cooking dinner without fully specifying the exact meal that is being prepared. Later observations may allow us to either specialize or generalize our current view of the plan in progress. In a very similar manner, plans can be built hierarchically. General plans are formed first, and the postponement of the specializing (filling in the gaps) of a plan may allow us to perform one action or plan immediately without going through the tedious process of fully specifying a plan in advance. This approach is intuitively obvious if one considers how you would go about any daily task. Say you are want to go to the store and pick up some milk. Your high level or general plan might be something like: "Drive car to the store, buy milk, bring it home". Now, while you attempt to execute the first portion of the plan "Drive car to the store", you need to specialize it further (until the actions are at a level where you have pre-built or atomic plans for them). So, you specialize this "Drive" action into something like: "Get the keys, get in the car, start the car, and drive to the store on your favorite route". In effect you are executing the original general plan after specializing the immediate portion of your plan, while the remainder of you plan is still very general.

The point of these hierarchical examples is to demonstrate the high level of similarity in what we commonly describe two separate actions: planning and plan recognition.

Both plan recognition and planning occur in conversation in various forms. In conversation with others, we listen to what someone is saying to us, and formulate our response based on what we understand, not only on the content of the words being spoken. We interpret what is being said to us in a certain context or with respect to knowledge we have in common with the speaker. This knowledge context represents a shared understanding of both the topic we are discussing, and the way in which various utterances relate to the topic of discussion. A valid question to ask here is what constitutes the substance of this common knowledge? One possible answer is to view this substance as a common understanding of what constitutes valid actions, and interactions between goals, and plans in a certain domain.

In order for communication with another agent to be a success, we must share the same (or at least similar) beliefs about the domain in question, and also about the domain of conversation itself. For example, when someone asks a question, both the questioner and the intended respondent realize that an expected response to a question is an answer, or possibly a request for better specification of the question in the form of a "clarification dialogue" [All78]. Both conversation agents realize that any response made is now within the context formed by the initial question, and both understand that the response will be given with this context in mind. This aspect of communicative capability involves the domain of communication, and is independent of the topic of discussion. At a different level in the same exchange, the discussion domain topic itself must be considered. For example if the question was "What time is it?", a valid response might be "The correct time is 11:00", or "It’s 11:00", or even "11:00". In any event, the response is understood within the context of the discussion topic "time" formed by the question.

What we see in this question/answer example is an instance of planning by one agent, and an attempt at recognizing that plan by the other. The questioner is trying to form plans in the communication domain to satisfy the goal of obtaining certain "time" domain information, and the responding agent is trying to infer the questioner’s goals and related plan in order to respond in a manner helpful in the the satisfaction of the questioner’s goals.

The questioning agent’s primary goal is to find out the time. In order to achieve this goal, the agent has formed a plan which might be viewed on a high level as "Get the time information from some other agent who has it". In order to achieve this goal through its plan, the agent must understand something about the domain of discourse as well as the domain of information. In order to receive information, the "curious" agent realizes it must form a question that the information bearing agent will understand, in this case through natural language discourse. In fact, the curious agent wants to satisfy a different type of subgoal here. Specifically, it wants to receive information from the information bearing agent. To this end, the curious agent formulates a plan that might look something like "Formulate a direct question and pose it to the information bearing agent". This plan is no longer in the same domain as the previous plan concerning time, it is in the domain of communication.

The information bearing agent in this example is continually trying to recognize what exactly the curious agent is up to. It may first recognize the initial utterance as a "question" within the set of possible plans or expectations in the discourse domain, and then recognize the time information request within the discussion domain. Armed with this information about the questioner, the information bearing agent must try to find a response which will satisfy the curious agent’s perceived goal of learning the time. The information bearing agent might have several options or plan alternatives which potentially can satisfy this perceived goal, and based on some method of selection, one plan must be selected and attempted. A possible response (as suggested earlier) might be selected within the discussion domain, and the result placed into a satisfactory communication plan for information conveyance in response to a question (ie ANSWER), and execution of the plan (might) result in the satisfaction of the curious agent’s goal. Presumably the information bearing agent tries to recognize the success or failure of its plan execution through further observation and interaction with the curious agent, and will act subsequently based upon this further context.

In plan recognition/planning situations such as these, there is a certain amount of ambiguity as to whether the agents are planning for themselves, or are recognizing the plans of other agents. In fact, the difference may be simply in how one views the problem. Certainly there is a large degree of interdependence between the tasks of planning for one’s own actions and recognizing the needs (goals) of others and attempting (planning) to help achieve these goals (which can also be viewed as one’s own goals in a truly cooperative environment).

In much natural language plan recognition work, the difference between communicative and topic domains is explicitly recognized and discussed [All78, Sid81, Lit86, Car89]. The "time" related planning of the previous example can be seen as a case of domain specific planning for "time" (where knowledge manipulation can be done in the specific domain independent of language), and the natural language or communicative aspect of the conversation depicted can be described as natural language specific planning, and is sometimes referred to as *discourse planning* [Lit86].

It is important to recognize the differences between attempting to learn something about plan recognition in general, attempting to learn something about plan recognition for discourse, and attempting to learn something about plan recognition in some specific domain of discussion. Most authors attempt to keep their work as general as possible, but of course they still must use some domain for examples. In this manner, it often becomes difficult to compare and contrast various works. So, for the remainder of this paper, I will try to point out which PR techniques refer to a domain in particular (such as natural language, or tying shoelaces, etc.), and which are more general. Of course, as mentioned earlier, all applications are in at least one domain, and likely more than one simultaneously. The relationship between these domains are not generally very well defined in any work I have seen.

Sidner & Israel [Sid81] attempt to incorporate language specific understanding features (discourse planning) into their plan recognition framework. They reflect this understanding through what they describe as the *intent* behind an utterance. Sidner states that natural communication is possible between two agents only if each can reason about the intent behind the other's utterances, and about its own responses. Sidner also points out that this type of reasoning involves using both world and domain knowledge known before a conversation begins, and knowledge gained during a discussion. The knowledge shared between interacting agents thus has some domain related starting points, and changes throughout any interaction.

Previous or initial world/domain knowledge can correspond to several aspects of plan recognition, including the proper and expected ways in which to interact in some domain (such as language), and shared domain beliefs. An interesting branch of this concept of world knowledge can also be seen in the selection of one explanation or possible plan during plan recognition. For instance, some action we see someone doing may have one common explanation, and other not so common or likely explanations. World knowledge might imply that we would "guess" that the likely explanation is the one to choose (by default) when there is a lack of information to the contrary. Consider the language domain. A simple sentential recognizer might scan sentences for question marks, and flag sentences ending with one as a query statement in need of a response. This seems reasonable enough, but there is also a possibility that the statement was a rhetorical question, and in this case no response should be given. The point here is simply that in any domain there can be multiple explanations for certain events or actions, and selection of one such explanation is not necessarily a straightforward task. Knowledge we have of the domain in advance (such as a "?" most likely indicates the presence of a query requiring a response) can help us in these selections. Later in this paper we will see some formalization of this line of thinking in the work of Carberry [Car89].

In performing PR, it is important to note that the plan inferring agent can obtain information used to build upon the shared context in which to infer plans in two distinct ways. The observed agent may actively participate with the observer in the PR process, in which case the observing agent tries to makes use of information *intended* by a cooperative observed agent to help in this process. On the other hand, the observing agent may obtain information about a passive observed agent's plans and goals simply through observation of the agent's actions, as if through a keyhole. Carberry [Car89] names these two PR approaches *intended* and *keyhole* recognition. More will be said about these later in a more complete description of Carberry's work.

## Solutions and Approaches

Now that we have seen several examples of some of the interesting aspects of plan recognition and its close relationship with planning, I will specifically describe several different approaches, and try to explain how they build upon one another.

## Allen and Perrault [All78]

This paper constitutes the earliest paper I have included in my evaluation of plan recognition. Many of the concepts presented in this paper evolve into more refined and integrated ideas in later papers.

### The Domain:

This paper describes plan recognition in terms of the domain of a dialogue of advice giving one might find at a train station information booth. This domain itself is quite restricted in that people using the advice giver are interested in only one of two primary goals - meeting a train, or boarding a train. However, the train domain provides a rich basis for dialogue, from which work on plan recognition in the discourse domain can proceed. Basically, Allen et al attempt to show how the fitting of utterances into possible plans can help in the understanding of poorly or incompletely formed utterances made in the train domain.

The Approach:

Allen’s plan recognition system models actions and plans in a manner similar to most other systems, where the plans are described with the following terms: *States* (propositional content of the world), *Operators* (simply a new name for an action template - either a primitive or a plan consisting of a body of more Operators), and *Actions* (instances of Operators). It should be noted that with respect to the construction of plans in this framework, that the ability of an action to be either a primitive or a plan consisting of other actions shows this approach to be utilizing hierarchical planning. Since it is not necessary to fully define a speaker’s plan immediately, much detail and overhead can be omitted when an attempt is being made to match potential plans or actions in the process of plan recognition. Sacerdoti [Sac75] seems to be Allen’s prime reference with respect to plan construction, and although the hierarchical nature of plans is not specifically described in [All78], he does take advantage of the ability of a structure of this type to allow *abstraction*: "We may create a plan involving A without initially considering the details required in its body" [All78]. In a similar way, Sacerdoti recognized the value in structuring plans this way, and specifically referred to the abstraction process as hierarchical planning: "...the hierarchical planning approach used by NOAH can result in the creation of higher-level plans that appear to be feasible at a higher level of abstraction, but whose detailed expansions are invalid".

Actions involving the production of an utterance are known as *Speech Acts*. These actions have a speaker, a hearer, and a propositional content. Speech acts have preconditions and effects which are specified in terms of the beliefs and wants of the speaker. An example of a speech act is INFORM. The precondition governing the use of INFORM is that the speaker (action taker) believes that the propositional content of INFORM is true. The effect of INFORM is that hearer believes that the speaker believes the truth of the propositional content. In Allen’s framework, speech acts deal solely with the domain of discourse I have mentioned previously.

The knowledge agents possess is modeled in terms of *beliefs*, *wants*, and *knows*. A plan exists within a particular agent’s want space. For example, the system is cooperative, and so it has a "want’" in the form of a plan to help the questioning agent. This help plan consists of a subplan to try to recognize the plan of the agent, and other subplans involving helping the agent.

In general, the system contains a limited number of *expected* or possible plans that the agent may be trying to execute. Each of these expectations will contain partial plans describing how these goals are attained. Trying to fit the agent's utterances into these expected plans in terms of speech acts constitutes the plan recognition portion of the system. After the system deduces the agent's plan, it follows through by attempting to execute this plan itself in order to find obstacles which prevent the plan from being executed immediately by the agent. By making these obstacles its own goals, the system can seek a plan for helping the agent (through conversation) based on its own knowledge about the domain.

Inferences are made in this system in two separate ways. First, if a plan action is observed, bottom up inferences can be made, in which the observed actions form the basis for finding a higher level plan of which they are members. This approach is also often described in other work as *generalization*. Allen approaches generalization in a somewhat imprecise manner in that these bottom up inferences are achieved through a set of inference rules describing the (quite domain dependent) cases where the nature of an utterance might indicate a certain type of higher level goal. I claim these rules are at least partially domain dependent since the manner in which they are defined in [All78] is with respect to a Speaker and an agent interpreting the Speaker's utterances. Thus we see that there is no explicit relationship between a low-level plan and a corresponding high-level plan(s), but rather an implicit or inferential link. An example of this type of inference would be when the system recognizes an agent's action as LEARN-TRAIN-LOCATION through an utterance of the agent, realizes that this action constitutes a subgoal of the plan BOARD-TRAIN, and infers that the agent has a plan of boarding a train. The second type of inference occurs when, the system starts with a plan expectation, and tries to simulate the speaker's planning process top-down. For example, if the system believes the agent has a goal LEARN-TRAIN-LOCATION, and the system knows that an action TELL-TRAIN-LOCATION has an effect on the hearer of LEARN-TRAIN-LOCATION, then the system might infer that the agent has a goal of getting the system to TELL-TRAIN-LOCATION. Inferencing in this manner is often referred to as top-down inferencing, or as *specialization* in later works. Once again, the work of Allen [All78] does not explicitly specify these specializing links, but rather creates them inferentially. Inferencing in these two manners suggests possible candidate plans or goals for the system to consider when attempting to recognize the agent's plan.

Based on what the system believes about the agent, such as goals the system believes the agent may have already accomplished or has yet to accomplish, one plan may be selected as a candidate while others are passed over. Of potential interest in this selection process is the fact that it is possible that one or more of these plans may be possible candidates. This aspect of plan selection will be discussed at greater length later in this paper.

The plan recognition or deduction process consists of a search through all possible plan alternatives that are expected or can be constructed, guided by various forms of heuristics. Pairing observed actions with a possible plan expectation forms one alternative hypothesis about what the agent is doing. The plan's "expected portions" or subplans are gradually filled in with observed actions. In this way, alternatives compete with one another and receive ratings. The highest rated task will form the basis for the system's action. Rating of the alternatives is based on factors such as likeliness of goals implied by the alternative (goals already satisfied are not likely), and possible satisfaction of an action's preconditions, among others.

Typically, the plan recognition system described in [All78] behaves in the following manner. First, an observed action is fit into the observed portion of possible alternatives, and bottom up inferences or generalizations are carried out, where possible, using these alternatives. Next, the action's parameters are bound into the alternatives selected. Now, one of the alternatives is accepted as the deduced plan, and any conflicts between alternatives are settled either by rating, or in the case of similarly rated mutually exclusive alternatives, a clarification subdialogue is carried out.

This approach makes use of the fact that the agents are assumed to be cooperative with the system. Allen claims this assumed cooperativeness suggests the use of several heuristics in the plan recognition process. Specifically, similarity between objects in expected and observed plans should suggest that these expected plans are quite likely the ones intended. For example, if it is "common" in some way not explicitly described by Allen for the mention of a city in an utterance to suggest that an expected plan is of the form BOARD-TRAIN(city), then this expected plan should be matched with the object "city" early on in the search for alternatives, rather than later. Also, cheap inferences should be preferred by the system over expensive ones. The rationale behind this heuristic is that the system is designed to fit the domain, and thus obvious inferences are more likely than very complex ones.

Some of the capabilities of Allen's system can be highlighted better through discussion of some simple examples. For a discussion of these examples, please refer to Appendix 1.

## Summary

Allen describes a plan recognition system used in the context of dialogue understanding for the very first time. The insight into viewing speech structures as actual actions in a plan resulted in a great stride towards improving the capability of natural language systems in some of the ways shown in the examples described earlier, such as fragment comprehension. The general recognition process is described in a largely domain independent manner, and the general heuristics apply equally well to any domain where cooperative agents are assumed and a large amount of domain knowledge can be shared between system and agent at the outset of planning.

## Kautz and Allen [Kau86]

In this paper, Kautz and Allen introduce an approach to plan recognition which is more powerful than the work in [All78], primarily due to the generality of the concepts which they describe. The new ideas which I will look at and discuss describe a new view of how to go about the process of plan recognition. Of particular interest are the decisions made about dealing with commitment to candidate plans, and discussion about what makes some inferences "likely". Using some of the examples from [Kau86] as demonstrations, I will try to show what this approach has gained, and later in this paper the cost of this approach will be discussed.

Kautz and Allen offer insight into how to reason about temporally complex situations such as when actions occur concurrently. However, I have omitted this discussion for the purposes of this paper.

## The Domain

As mentioned previously, the work in [Kau86] is quite general in nature, however, examples given are in the domain of cooking, where it is often the case that generalizations and specializations follow quite naturally.

## The Approach

It is easy to see how the concept of non-monotonic reasoning about plans fits into our interpretation of the actions of agents in our world. We observe, and draw conclusions to the best of our ability based on our past experience. If we are not overly stubborn, we are willing to modify our beliefs about actors in our world based on new observations we make, or information we obtain. The difficulty lies in determining exactly when we should rethink our previous judgement and search for new alternatives based on new information. If we delay too long in clinging to our old beliefs, life can pass us by, as in the case of a rapid fire discussion where someone changes topics and leaves us bewildered, still pondering an old point. On the other hand, if we hastily discard the context of our world as we rush on to something new, we can become just as mired in confusion by forgetting the main point of a conversation. This problem is referred to as the *commitment* problem.

Kautz and Allen point out that past plan recognition work has suffered from problems with commitment. Approaches which attempt to fit observations into canned *explanations* suffer from an inability to choose among competing explanations satisfactorily. Arbitrarily selecting one of many possibilities equates to over-commitment to one explanation. A selection criteria of some sort is required. After all, when you feel a drop of rain on your head, are the two explanations "the sky is falling" and "it is starting to rain" deserving of equal consideration? Other approaches attempt to make "likely" inferences based entirely upon heuristic rules, (as alluded to in [All78]). Rules of this type might say "If you observe someone opening his/her purse, then he/sh is likely going to take money out of it". These rules are hard and fast in that they only take into account domain information and not the context in which the rule would apply. Any decision about when to apply these rules seems to be of a domain dependent heuristic nature, and adds little to a general discussion about plan recognition.

What is offered in [Kau86] is a formalization of some of these implicit assumptions about what occurs during plan recognition. Within the bounds of two assumptions, Kautz and Allen propose a PR theory which will supply an agent performing PR with the strongest possible set of conclusions that can be made based upon a set of observations. The first assumption made is that the agent performing PR has a complete plan library. In other words, the known ways of *specializing* or performing an action are the only ways. The second assumption is that all actions occur for a reason, and all reasons for actions are known. For example, if an action RUN is observed, then the set of all more complex actions or plans including RUN (ie RUN-TO-STORE, RUN-AWAY-FROM-DANGER, etc.) contains the more complex action or higher level plan which is actually being carried out. In effect, these two assumption are heuristics in themselves. The first assumption is actually an implicit implementation of a "Closed-World" type of heuristic, and the second is an example of what one might call a "Motive" heuristic. More discussion will focus on these heuristics later in this paper. The completeness assumption allows for reasoning to derive a set of all possible plans containing the given observations. Commitment is handled simply in that no selection is made at all. Reasoning is carried out such that elimination can be performed case by case. The disjunction of the elements of the *explanation set* is considered, and each element or plan is given equal consideration. This approach really ignores all concept of likely or unlikely plans completely. No preference is made based upon likeliness of alternatives. In general perhaps this approach is the most fair, and least susceptible to error, but it is difficult to think of any domain at all in which all possibilities are always equally probable. Carberry [Car89] discusses some of the options that can be utilized instead of this non-approach.

Within the scope of these assumptions, Kautz and Allen essentially reduce plan recognition to a problem of ordinary deductive inference using observations to fit into a complete actions description. Since there is a very high degree of complexity in fitting multiple observations into a coherent plan structure (ie each observation could actually be part of a separate plan), Kautz and Allen introduce what they refer to as *simplicity constraints*. These constraints actually function as a heuristic describing how actions are "typically" observed in some manner. In a hierarchical view of plans, top level plans or actions are done for their own sake, and are not subparts of larger plans. The "Simplicity" heuristic simply states that multiple observations tend to be part of the same top level act, rather than of many top level actions. As few top level actions are recognized as possible.

By considering some of the examples of plan recognition given in [Kau86], the method behind the theory becomes clearer. Refer to Appendix 2 for a discussion of several relevant examples.

Summary

Kautz and Allen succeeded in formally specifying a good deal of plan recognition theory present but not explicitly stated in earlier papers. They have formalized the concepts of specialization and generalization/abstraction in a very rigid framework where complete domain knowledge is assumed and required. Uncertainty is precisely specified in that preference is virtually ignored in cases of ambiguity, and all possibilities are equally considered. It would seem that this approach has the advantage of being rigorous and correct at the expense of flexibility. If the intent of a plan recognition methodology is to mimic an intelligent (human) approach, then this system fails. However, if the intent is to produce results that have little chance of failure and are as correct, or only as correct as deduction would allow, then the system described succeeds. The heuristics mentioned by Kautz and Allen (Closed World, Motive, Simplicity) all seem valuable in plan recognition systems in general, and they appear in various forms in much of the recent work in the literature, such as that done by Carberry [Car89]. However, the cost of strong assumptions such as the Closed World assumptions is large in that one must have complete domain knowledge in advance, including all possible plans leading to any specific high level goal. Relaxation of this assumption is not possible without giving up all *certainty* in the deduction process inherent in the completeness of Kautz's approach. Carberry takes a different approach to PR that is not as strict in order to avoid this and other difficulties.

Carberry [Car89]

Carberry focuses her efforts in her most recent work [Car89] on developing a theory of plan recognition which allows for the existence of default inferences when selecting between possible candidate plans or alternatives during the PR process, and a strategy for assimilating newly obtained beliefs about the world into a belief model. In addition, attention is paid to the problem of justifying inferences or assumptions in an understandable manner in order to lend the PR system credibility in the eyes of the human user.

The Approach

Carberry's previous work [Car83, Car88] has exhibited many of the concepts discussed and introduced in the papers elaborated on earlier in this paper [All78, Kau86]. As in other work, Carberry intends for her PR systems to build a contextual picture or models of the information seeking agent's (IS) plans. Towards this goal, a PR system called TRACK was created which inferred a user's goals in an ongoing, fairly ideal, cooperative dialogue. TRACK utilized hierarchical plans, much like those described in [Kau86], and maintained its picture of the user's plan in a tree-structured context model similar in nature to the generalization/specialization illustrated in Figure 1 of Appendix 2. A notable difference in the construction of the models in [Kau86] and [Car83,88,89] is that the generalization (Up) links of [Kau86] are explicitly defined and reasoned with, while the generalizations are inferred through heuristics in Carberry's model. However, the representation of the intended model of the "world" remains similar, only the construction methodology differs.

Incorporation of new utterances into the context model was achieved through the use of *plan identification* heuristics which identified possible candidate focused plans (CFP) which were intended to represent the task upon which the user was focusing his attention. Plan identification heuristics determine a set of candidate plans through the use of inference rules [All80] which are intended to "guess" the reason that a user issued a particular utterance. An example of an inference rule of this nature might be "If the user wants to know the value of X, where X is a precondition or subaction of some plan P, then P should be a candidate focused plan" [Car83]. As a point of interest, Allen et Kautz [Kau86] fail to specifically discuss these inference rules, and offer nothing as replacement. Perhaps they felt that in a complete, deduction oriented environment, the concepts of generalization and specialization subsumed these inference rules, or equated to the application of all of these inference rules.

Focusing heuristics were used to select and integrate the "best" CFP into the existing context model. These heuristics are based upon the structure which is observed to exist in dialogue. For example, research has shown that people generally ask all their questions relevant to a plan for a single subgoal before moving on to ask questions about another subgoal of an overall task [Car88]. This type of behavior can be reflected in PR systems in terms of the tendency to prefer to continue pursuing the subgoals of a single task before trying to infer new tasks from a new observation. This same type of simple preference logic is seen in the work of Kautz and Allen [Kau86] where they prefer to fit actions into one top level plan rather than into several disjoint top level plans. Another factor which is taken into account in the creation of focusing heuristics is the cooperative nature of the agents in a dialogue. This nature implies that conversation should exhibit something like inertia with respect to topic [Car88, Sid81]. In [Car88] and [Car89], Carberry describes the operation of these focusing heuristics at some length. Basically, the focusing is accomplished by preferring certain relationships of the identified CFP and the existing focus point of the context tree. The relationships are tried in order, and the first one which fits (minimizing the shift of focus of attention) is chosen. Intuitively this approach makes sense, since it is easy to realize that context changes in person to person dialogue do require a bit of mental effort in order to correctly frame the new context of discussion in one's mind. Similarly, changes in topic require some effort on the part of the system, since all *inertial* possibilities must be tried and rejected before a new context can be formed.

### The Bottom-Up Problem

TRACK was unable to make generalizations of actions if a choice needed to be made between mutually exclusive generalizations. While Kautz [Kau86] was able to disjunctively handle this difficulty, it should be noted that this type of inference in general can lead to computational problems. Each action inferred (bottom-up) could potentially have many higher level actions of which it is a part. Kautz heuristically controlled the computational explosion here through the rating method described earlier in this paper. Carberry suggests that rather than stall at these choice points (as TRACK did), or pursue the computationally expensive option of splitting the current context model into many different context models (or equivalently performing a great deal of logic using many disjunctions), that default inferences be used in the action hierarchy in such a manner that certain generalizations would be preferred over others. The problem of likely explanations described earlier in this paper (H20 => Sky is falling or H20 => It is Raining) is handled nicely in this solution, which actually does simplify to "what is the most likely explanation for a certain action?". Carberry [Car89] points out that in the case (in some domains) where one higher level plan is overwhelming more likely than others, in cases of cooperative PR (Intended recognition), the agent may expect the PR system to recognize the most likely higher level inference. Determination of what exactly constitutes "likely enough" for this type of agent-system expectation to exist is unclear, but could suggest a possible heuristic in some domains. As an example of this *likely expectation*, consider the case where Grandma comes over for supper every Sunday, driving her loud jalopy. If your mother tells you to pick up your toys because she hears a loud car in the driveway, you should be able (and in fact are expected) to infer that Grandma has arrived.

## The Top-Down Problem

TRACK also experienced problems when more than one way existed in which to expand the CFP and capture the relationship selected through focusing heuristics since it had no method of preferring one way over another. Kautz's system [Kau86] represented this difficulty in that it could not select between possible specializations of a general plan. In other words, TRACK experienced the same problem all other PR systems do: it had no manner of preferring one possible applicable plan out of many when each are capable of achieving some existing high level goal.

## The Solution : Default Inferencing

Carberry's analysis of natural language dialogues has indicated that in cases of human communication where one person (IS - information seeker) is actively seeking information from another (IP - information provider), the IP will often make default inferences (assumptions) about the IS's plans and goals, and act based on these inferences [Car89]. Since Carberry's primary goal is to attain a system that appears natural and rational in its behavior, it is important that the system should not engage in excessive amounts of clarification dialogue to resolve conflict between mutually exclusive specializations or generalizations since the system will be perceived as being "out of touch". On the other hand, Carberry observes that schemes based upon many levels of probabilities can be difficult to explain or justify to the user, and could result in what I call the "Mr. Spock Syndrome". As an example of this, consider the case where the IP system has inferred via probability logic that the IS agent is trying to find out how to board a train to Kokomo, and utters "To board the Kokomo train, go to gate 99.", to which the IP replies "Why do you think I want to board the Kokomo train", a request by IS for IP to justify its last utterance. A response of the type "85.4 % of all requests in Toronto to board a train at 12:15pm on Saturdays involve boarding the train to Kokomo" might not inspire confidence on the part of the IS that he can relate to the system successfully.

What seems to be required then is an approach to "likely" or "common" selections that are handled in a way similar to the way humans make such inferences. Carberry [Car89] cites references which indicate that people people often select solutions most *representative* of the evidence at hand, rather than depending on strict probability assessment. Consider the following scenario as an example of representativeness. If some agent is believed to be pursuing some goal G, and is observed performing several subgoals typically found in (or representative of) plans for G, then this evidence should confirm (perhaps in the manner of a rating system) the belief in G as a goal of the agent.

Carberry also notes that other researchers have discovered that people tend to develop a rather complete, workable hypothesis to explain some agent’s plans, and revise this hypothesis only when evidence emerges to contradict it. Rather than maintaining many possibilities with associated probabilities for each, there seems to be evidence that people maintain or discard a hypothesis based upon some threshold which indicates *skepticism* or the level of certainty one has in that hypothesis. This concept of a threshold could be seen as some measure of the level of conservativeness of a person. The higher the threshold required for discarding a hypothesis, the more conservative a person is; the lower the threshold, the more likely a person is to jump to new conclusions.

The *plausibility factor* which Carberry uses in conflict resolution parallels Reiter’s concept of "intuitive plausibility", where the factor simply represents the extent to which evidence accumulated so far makes some goal G plausible. Bottom-Up (or generalization) conflict resolution would be settled through *preference rules* such that of the possible generalizations G of subgoals E, we select the Gi with the highest plausibility level of all the G, such that the plausibility of Gi exceeds the predetermined threshold or skepticism level, and no other Gj is nearly as plausible within a predetermined difference/similarity level. It is important to note that the level of plausibility of some goal is used only in preference selection, and goals selected in this way are treated as any other inferred goal. The values of the skepticism and difference levels control how much default inferencing is performed. Although settings of these values are not discussed at any length in [Car89], their values determines the "personality" of the system and deserve some research for different domains and situations.

Through the discussion of a simple example, the application of the results of Carberry's work can be clarified. Refer to Appendix 3 for this discussion.

## Summary

Carberry's work with default inferences seems to shed some light on a practical way to help the process of PR through use of heuristics which attempt mimic the method of PR undertaken by people. Her described application of preference or plausibility values for explanations or generalizations manages to capture the concept of "likely" explanations discussed in [All78] and [Kau86], and also in many other papers discussing PR. Carberry's method exhibits a great deal of flexibility in that different settings of her skepticism value can produce a system that would produce inferences over the range of ways people do, from the very conservative or skeptical, to the quick to jump to conclusions. By reasoning in a manner similar to the way people actually reason, users of the system will have a greater level of confidence in the system since they can understand the reasoning process undertaken. No problems arise in explaining to the relatively naive user the complex probabilistic formulas and computations required as in deeply probabilistic systems, or the arbitrary selection of "any old" explanation for some actions set, such as sticking to one conclusion (commitment) until it becomes questionable due to the discovery of new evidence.

Carberry discussed the need to incorporate the plausibilistic inferencing in a Top-Down or specialization manner also, suggesting the case where certain goals are achieved by likely plans more than by other less likely plans (in some circumstances). Of additional interest here is the cases where goals are achieved by plans which have not been used or seen before (unique). Recognition in this type of non-closed world is not discussed at any length in [Car89], although it is certainly a problem that people can handle through analogy to other similar situations, or through straightforward deduction based on a deeper (?) understanding between the relationships and objects in the domain, and in the world in general.

## On Plan Commitment and Preference Heuristics

What we have seen in these three "indicator" papers is three different levels of plan commitment, and three different levels of approach to candidate plan preference.

Initially, PR approaches suggested tended to overcommit to the first possible plan for a goal, or explanation of a set of actions. This approach can be computationally expensive in cases where an "unlucky" selection is made, or the domain is structured in such a way that the desired or correct solutions are not discovered until much search through the space of possibilities. In cases where this expense is incurred, the system performs and interacts stupidly in terms of much reliance upon clarification subdialogues and utterances to the user which do not portray the system as understanding the domain in question. The result of these difficulties can be a loss of user confidence and hence an unwillingness to depend on or utilize the system.

Kautz and Allen developed an approach to PR quite different from the approach discussed previously. This approach exhibited no problem with overcommitment to a plan based on inconclusive evidence because no commitment was given to a "possible" solution at all until there was no question as to the correctness of the solution. Another way of viewing their approach is to consider it overcommitment to more than one possible explanation, in that all possibilities are treated equally until some are logically excluded. The disjunctive acceptance of all possibilities produces a scheme that has the expense of computational complexity spread out over all inferencing, as opposed to the previous approach which only incurred the expense if it was unlucky. In any situation where more than one generalization or specialization existed, all would be accepted and offered. The advantage of this approach is that it is never "wrong" in that a guess has been made incorrectly, but its disadvantages include: 1) being expensive with respect to inferencing; 2) being hesitant in that a request for a "likely" solution would not indicate one over another, despite the fact that the domain might suggest one as far more likely than another; and 3) having an inability to reason in a manner consistent with the reasoning of people. Clearly no one considers all possibilities logically possible at all times when performing PR — "common sense" reasoning takes a large part in the trimming of the possible search space.

The more recent work of Sandra Carberry shows a tendency to commit only to a plan or explanation that is likely in the context of the domain. In this manner, likely options are exhausted before unlikely possibilities are examined. The expense of this approach is controllable in that the factor of likelihood and the personality of the system performing the reasoning on a plausible level can be set to match a particular need or domain. Either of the two previous approaches can be modeled in Carberry's suggested framework of plausible Top-Down and Bottom-Up reasoning. It is also important to note that the computational expense incurred through much search through possible plans, and the expense of performing subdialogues for clarification can be reduced, since these expenses occur only in the cases which are, by definition in the domain, less likely. Since the domain knowledge (ahead of execution) includes some measure of likelihood, we can profit by this in our restriction of the search space, expanding it only when likely options are exhausted. The obvious problem inherent to this approach is the discovery of the levels of plausibility to use in the domain definition, and although there is no cut and dried answer, Carberry's approach can live without these if necessary by reverting to one of the earlier approaches, and still benefit in cases where likely knowledge exists.

In a manner paralleling the establishment of commitment in the three examples above, different preference approaches for candidate plans have been used. The use of the *simplicity* constraint discussed earlier in this paper (prefer explanations/generalizations which involve a minimum of top level or higher level goals) may be applicable only in certain domains. There is no establishing evidence suggesting that this heuristic applies equally well to all planning or PR domains. To a large extent at least, this principle would hold only if the domain's action hierarchy structure was created with this principle in mind.

Bottom-Up and Top-Down preference heuristics were evidenced to some extent in each of the three approaches we have seen. Allen [All78] utilized a rating approach to preference inferencing that was only partially (and vaguely) defined. This approach seemed to be meant to work only in cases of generalization (Bottom-Up), and a similar method of performing specialization (Top-Down) inferencing is not suggested.

Kautz's approach implied that certain candidate plans were preferred only if they were known to be certain. In cases of uncertainty, no preferences were used at all, all candidates were treated equally.

Carberry includes the concept of preference in the explicitly defined domain information, in the same way as the action hierarchy is defined for reasoning under the closed world assumption. No method is suggested for preferring unique plans less or more than non unique plans, just as no method is suggested for inferring unique plans as valid explanations. Problems with this approach include establishing the values of plausibility in a particular domain, and utilizing this approach with Top-Down inferencing. In the case of (T-D) specialization inferencing it will be necessary to define explicitly what is meant by plausibility, and this definition may need to be quite domain specific.

## Possible PR Directions Using Preference and Commitment

In general, any approach to candidate plan preference should attempt to prefer certain generalizations and certain specializations, perhaps only those that make most sense in a particular domain. One consideration might be developing a scheme which allowed for the dynamic assignment of likeliness values over time. In this manner, the PR system might be able to start with what the domain action hierarchy designer believed to be accurate plausibility values, and modify these over time based on statistics it is able to keep on known explanations and plans, or perhaps change these plausibility values to account for the possibility of discovery of new explanations and plans in some manner.

It would be interesting to consider the case of imposing an ordering on the preferences for explanation selection and plan inferencing, while at the same time discarding the concept of using mathematical plausibility and skepticism levels described by Carberry. A system which exhibited this kind of preference, with no expense other than the ordering of the domain when it is defined, could perhaps gain some of the advantages of Carberry's work while at the same time remaining simplistic in its explanation selection. For a diagram of how this type of selection might occur, please refer to Appendix 4, page 8 of the overheads.

In addition, all of the approaches mentioned exhibit quite different degrees of limitation of the levels of inferencing allowed. Carberry limits the inferencing levels through the fact that levels of plausibility of explanations decrease with each successive level. Allen's paper suggests that inferencing generally involves few levels of inference, and limits the levels implicitly in the definition of inference rules. Kautz takes the most general view where inferencing can be many levels deep - commitment is made completely to all possibilities, generating all possible results at all possible levels of inferencing. Investigation into what "makes sense" in a particular domain in terms of depth of inferences allowed could impact an implementation of a system of this type dramatically. If perhaps only one level of inferencing were allowed for each observation, and that level confirmed in some fashion before another level was allowed, a much more skeptical system would result. Certainly there are cases where this approach would be undesirable in that it might be difficult to arrive at a very high level, or general explanation, but perhaps this would be beneficial in some domains where concern on the planning nearest the focus is of primary concern. An example of this might be a situation where there simply is no time available to perform a high level of inferencing, such as a very demanding real time environment. Of course, the speed with which such a decision is made through few inferencing levels might provide only limited contextual information, but in the same fashion it is less susceptible to errant complex assumptions about the context.

## Closing Comments

The work examined in this paper constitutes only three of many papers currently of interest in the PR field, however they each represent some different aspect of approaches commonly taken towards PR. Each has a different view towards what is important — either completeness and correctness of inferences, efficiency and application, or some combination of these factors. Perhaps the most valuable separation of approaches lies in the attention Carberry pays to trying to duplicate the PR process of people. Although no algorithm or machine has yet demonstrated the quick and comprehensive ability to perform multiple domain PR in large, unwieldy and complex environments, it is a fact that that humans succeed in this endeavor every minute of every day. Just as other researchers try to duplicate the inner workings of the human brain using current biological research insights in an attempt to achieve something closer to intelligent machines, Carberry is trying to duplicate the process of PR at a macro level based on some very simple yet fundamental observations about how people tend to infer the plans occurring in their environment. Heuristics such as those suggested in [All78], and used by others in identifying possible plans (inference rules), and many other heuristics used in PR are simply expressions of what seems to make sense, of what a person might infer, in his/her attempt to recognize the plans of agents in his/her environment. Carberry is simply taking these heuristics a bit further in her approach using the concept of plausibility. As further research yields new information about the human action of planning and plan recognition, new insights will be gained into how work in plan recognition should proceed.

## References

[All78] James F. **Allen**, C. Raymond **Perrault**, Participating in Dialogues: Understanding via Plan Deduction, *CSCSL*, 1978.

[All80] James F. **Allen**, C. Raymond **Perrault**, Analyzing intention in utterances, *Artificial Intelligence*, 15:143-178, 1980.

[Car83] Sandra **Carberry**, Tracking user goals in an information-seeking environment, *Proceedings of The National Conference on Artificial Intelligence*, pages 59-63, Washington, D.C., August 1983.

[Car88] Sandra **Carberry**, Modeling the user's plans and goals, *Computational Linguistics*, 14(3) : 23-37, 1988.

[Car89] Sandra **Carberry**, 1989, *A New Look at Plan Recognition in Natural Language Dialogue*, Technical Report No. 90-08, Department of Computer and Information Sciences, University of Delaware, Newark, Delaware.

[Cha87] David **Chapman**, Planning for conjunctive goals, *Artificial Intelligence*, vol 32, p.333, 1987.

[Kau86] Henry **Kautz**, James **Allen**, Generalized Plan Recognition, *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 32-37, Philadelphia, Pennsylvania, 1986.

[Pol86] Martha **Pollack**, A model of plan inference that distinguishes between the beliefs of actors and observers, *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, pages 207-214, New York, New York, 1986.

[Lit86] Diane J. **Litman**, Linguistic Coherence: A Plan-Based Alternative, *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, pages 215-223, New York, New York, 1986.

[Sac75] E.D. **Sacerdoti**, The nonlinear nature of plans, Advance Papers, IJCAI 1975, Tbilisi, USSR, vol. 1, pp. 206-214.

[Sac77] E.D. **Sacerdoti**, *A Structure for Plans and Behavior*, American Elsevier, New York, 1977.

[Sid81] Candace L. **Sidner**, David J. **Israel**, *Recognizing Intended Meaning and Speakers' Plans*, IJCAI, Vancouver, 1981, 203-208.

[Smi88] David E. **Smith**, 1988, *A decision-theoretic approach to the control of planning search*, Stanford Logic Group Report LOGIC-87-11, Department of Computer Science, Stanford University, Stanford, CA.

[Tat75] Austin **Tate**, Interacting goals and their use Advance Papers, IJCAI 1975, Tbilisi, USSR vol. 1.

[Wile83] Robert **Wilensky**, *Planning and Understanding: A Computational Approach to Human Reasoning*, Addison-Wesley, 1983.

[Wilk84] D.E. **Wilkins**, Domain independent planning: representation and plan generation, *Artificial Intelligence*, vol. 22, no. 3, April, 1984.

Appendix 1 : Examples of Allen’s approach to PR

Sentence fragments can be resolved by taking the context of their utterance in some domain into account. An example of a possible fragment might be: "The 3:15 train to Windsor?". Trying to interpret this utterance without knowledge of the possible or expected domain of plans under which it is uttered is very difficult, and leads to many possible interpretations, such as "What color is the 3:15 train to Windsor?", "Who is driving the 3:15 train to Windsor?", etc. However, it is quite easy to take advantage of the limited plan domain we are assuming. Since we know that the agent making the utterance is interested in only boarding a train or meeting a train, we can interpret this utterance much more specifically. The system is located in some arbitrary location (Toronto in this case), and the system knows this when it makes inferences within its domain. The use of the *clue word* "to" in the utterance with respect to the location "Windsor" can serve as enough information for the system to infer that the querying agent is interested in a train from Toronto to Windsor. The system can now infer that the querying agent’s plan is to board a train traveling from Toronto to Windsor since no other possible plan expectation fits the context. Subgoals of the agent’s plan might include knowing where the train is (ie gate) and knowing when to board the train. The system might recognize the time reference in the query ("3:15") and conclude that the agent has already satisfied the time subgoal. Based on the fact that the agent wishes to board a train to Windsor, and has one of two subgoals satisfied, the system could infer that for the agent’s high level goal to be satisfied, the agent must be given information about the remaining subgoal, the location of the train to board. A reasonable response could now be given as succinctly as "Gate 10". In certain cases similar to this one, a plan might be recognized with more than one possible unsatisfied subgoal. In cases such as this, the system could easily determine that it could *help* by offering information to satisfy as many subgoals as possible. Clearly then, given a well restricted domain, this system is able to recognize a possible plan, based upon only fragmentary information. Even more, it is able to respond in a much more satisfactory manner than if no plan recognition had been attempted.

It can often occur in dialogues such as the one described above, that not enough information is available at a particular point for the system to ascribe just one possible goal or subgoal to the querying agent. This would be the case if, in the previous example, the query were formed without the clue word "to" as "The Windsor train at 3:15?". Now, if it were the case that a train arrived from Windsor at 3:15 and one left for Windsor at 3:15, then the statement is ambiguous with respect to which of the MEET or BOARD goals the agent has in mind. The system would realize that more information is needed about the agent’s plan in order to narrow it to one of the possibilities. A question of the form "Are you traveling TO

Windsor?" and a response by the agent in either the affirmative or negative would be enough to conclude which plan the agent is executing, and an appropriate response can be made to help the agent satisfy his remaining subgoal. Responses might be of the form "The 3:15 train TO Windsor leaves from Gate 10", or "The 3:15 train FROM Windsor arrives at Gate 8". This type of a system initiated sub-dialogue is referred to in [All78] as a *clarification dialogue*.

## Appendix 2 : Examples of Kautz's approach to PR

Consider the cooking domain for the following examples. In the following diagram the domain is explained, with broad arrows indicating specializations such as "MakePastaDish is a specialization of PrepareMeal", and narrow arrows indicating the hierarchy of subactions such as "Boil and MakeNoodles are two subactions of MakePastaDish".

If disjunctive observations could be made in some manner where it is not quite possible to distinguish between two actions, the deductive approach can still provide definite information about the plan in question, taking advantage of the hierarchical nature of plans. Consider the case where an agent is observed to be either MakeFettucini(M-F) or MakeSpaghetti(M-S). Since both M-F and M-S are specializations of MakeNoodles(M-N), we can abstract the fact that the disjunction of the lower level actions is an instance of the higher level action and conclude that M-N is part of the agent's plan. Since M-N is a subaction of just one plan, MakePastDish(M-P-D), we can infer that M-P-D is part of the agent's plan, and since M-P-D is a specialization of Prepare-Meal(P-M) that P-M is the top-level act that the agent is performing. Based on this knowledge we have built up, we can anticipate future observations such as MakeSauce as a subaction of M-P-D, and can realize that certain actions must have occurred outside our observation, such as GoToKitchen, a subaction of P-M that must occur before M-P-D.

Suppose that a second observation was made of the action MakeMarinara(M-M). Since M-M is a subaction of two higher level actions MakeFettuciniMarinare(M-F-M) and MakeChickenMarinara(M-C-M), we can infer that the agent is performing one of these two higher level actions. More simply, the disjunction of these two alternatives holds. M-F-M is a specialization of the action M-P-D, and M-C-M is a specialization of the action MakeMeatDish(M-M-D), so we know that one of these two holds also. Since both of these actions M-P-D and M-M-D are specializations of the P-M action, we can infer P-M as part of the agent's plan as well.

In each of the previous two examples we have derived explanations of a single observation. In order to combine these two explanations, we first appeal to the simplicity constraint described previously, which states that a single top level goal is preferred over more than one top level goal. The assumption basically says that it makes sense to try and interpret two actions together first, before trying to interpret them separately. Based on this, we attempt to combine the observations. The disjunction inferred in the second example between M-F-M and M-C-M can now be resolved. Since the first example established that M-P-D was the (mutually exclusive) specialization of P-M, and M-F-M is a specialization of M-P-D, we can conclude that M-F-M is in the agent's plan, and that M-C-M is not.

Appendix 3 : Examples of Carberry's approach to PR

A simple example of the approach taken by Carberry [Car89] can be explained using the utterance "Who is teaching EE202 next term?" by an IS in the domain of student advising. The *plan identification* heuristic might give the corresponding CFP as LEARN-MATERIAL(User, EE202), based on the fact that a student asking a question about the EE202 class could very likely have a higher level goal in mind of learning the material in that course. Performing generalizing inferences in this domain, based on the identified CFP, the system attempts to explain the CFP. In this domain, two higher level goals exist of which Learn-Material is a subaction. These goals are AUDIT(User, EE202) indicating the student may wish to audit EE202, or TAKE-CREDIT(User, EE202), indicating the student may wish to take EE202 for credit. If these generalizations have *plausibility* levels associated with them of .98 and .015 respectively as explanations of Learn-Material, then the most plausible one will be selected provided it has a plausibility greater than the level of skepticism of the system. Say the system skepticism was set at .90, indicating the minimum level of plausibility required by an explanation for it to be inferred. We would then see AUDIT(User, EE202) inferred into the context model. If there were no plausible enough explanation of Audit in the domain, then inferencing would stop at this point. It is important to note that another heuristic is used to determine if there is grounds to select between two explanations that each may exhibit a high degree of plausibility. This heuristic is known as the *difference* level, and simply states the amount two explanations must differ by in order for one to be selected before another as the current explanation in the context model. The context model look like the following.

```
AUDIT(User, EE202)
         |
LEARN-MATERIAL(User, EE202)  <- Focus
```

If the skepticism level were set at .99 for this example, no generalization would have been selected to add to the context model, since the system would have been "too skeptical" of the possible explanations. A setting of 1.00 indicates that only certain explanations are adopted, and a setting of 0.00 indicates that any explanation offered would be adopted.

Appendix 4 : Overheads of presentation of this paper

The following pages are photocopies of the slides used in the presentation of the main ideas in this paper. The presentation is approximately 20 minutes in length