Incorporating Default Inferences into Plan Recognition

Sandra Carberry
Department of Computer and Information Sciences
University of Delaware
Newark, Delaware 19716
carberry@dewey.udel.edu

Abstract[1]

This paper presents a process model of plan inference for use in natural language consultation systems. It includes a strategy that can both defer unwarranted decisions about the relationship of a new action to the user's overall plan and sanction rational default inferences. The paper describes an implementation of this strategy using the Dempster-Shafer theory of evidential reasoning. Our process model overcomes a limitation of previous plan recognition systems and produces a richer model of the user's plans and goals, yet one that can be explained and justified to the user when discrepancies arise between it and what the user is actually trying to accomplish.

Introduction

During task-oriented interactions with an expert consultation system, a user is engaged in seeking information in order to construct a plan for accomplishing a task. A number of researchers have demonstrated the importance of inferring the user's partially constructed domain plan and tracking his focus of attention in the plan structure[1, 9, 12, 13, 14, 17, 21], and models of plan inference have been used to address problems in language understanding and response generation. Unfortunately, current models of plan inference encounter difficulty when there are several possible explanations for an agent's action. They cannot use a priori knowledge about the domain to make default choices among plans that might be inferred from an observed action, nor can they revise incorrect beliefs about an agent's plan. For example, suppose an agent asks how late the Super-Saver Supermarket is open. Current systems are unable to make the default inference that the agent intends to purchase groceries since there are other high-level actions, such as cashing a check, that might motivate his query.

Analysis of naturally occurring dialogue suggests that human information-providers often make default inferences about the plans and goals of the information-seeker, use the resulting beliefs to generate helpful responses, and can explain, rationally justify, and revise their beliefs when they are in error. If natural language consultation systems act in the same manner, their responses will appear intelligent and natural to their human users. If they fail to make these inferences, they will often be forced to engage in lengthy clarification dialogues in order to ascertain with certainty what the user is trying to do, and will therefore appear unintelligent, obtuse, and uncooperative.

We have been investigating how the behavior exhibited by human information-providers can be captured in an intelligent natural language system. This paper presents a process model of plan recognition that is motivated by an analysis of naturally occurring dialogues and by psychological studies of human inference. It includes a strategy for incrementally updating the system's model of the user's plan that can both defer unwarranted decisions about the relationship of a new action to the user's overall plan and sanction rational default inferences. The paper describes an implementation of this strategy using the Dempster-Shafer theory of evidential reasoning. Our process model overcomes a limitation of previous models of plan recognition and produces a richer model of the user's plans and goals, yet one that can be explained and justified to the user when discrepancies arise between it and what the user is actually trying to accomplish.

Intended versus Keyhole Recognition

Default inferencing plays a role in both intended and keyhole plan recognition. Intended recognition is the recognition of those goals and plans that an agent intends to convey and is essential in identifying the intended meaning of a speaker's utterance[3]. Allen[13] was the first to model intended recognition. When n mutually exclusive higher-level goals could be inferred from a given subgoal, he used a branching heuristic that reduced the ratings of the alternative inferred partial plans to $1/n$ of the rating of the partial plan from which they were inferred; he justified this heuristic by arguing that the speaker would have realized that the hearer would have difficulty identifying which

[1]This material is based upon work supported by the National Science Foundation under Grant No. IRI-8909332. The Government has certain rights in this material.

inference he was intended to choose. In [1] and [20], upward inference of higher-level goals was terminated once a choice among several possible but mutually exclusive plans had to be made. We contend that if the speaker believes that the mutual domain knowledge suggests that one goal is substantially more likely than the other possible alternatives to be the motivation for his actions, he may *intend* the listener to recognize it by virtue of default inferencing.

Keyhole recognition is the inference of an agent's goals and plans by unobtrusively observing the agent, as if through a keyhole, and is useful in generating cooperative, helpful responses[3]. A natural language consultation system's ability to provide useful advice will be directly related to how well it is able to recognize what the user is trying to do. Since an action can generally play a role in many different plans, some of them a priori far more likely than others to be the motivation for the action, the system must be able to make appropriate default inferences if it is to develop a rich and tractable model of the user's plans and goals. Thus default inferencing is important in both intended and keyhole plan recognition. Although a number of researchers have studied default reasoning and developed formal models of default inference[4, 15, 16], little attention has been given to incorporating default inferencing into incremental plan recognition in natural language advisement systems.

Motivation for Our Process Model

Our objective is a plan inference framework that will produce a rich model of the user's underlying task-related plan. What requirements should be placed on this model? We contend that rather than be the *best* model in the sense of being the model with the greatest mathematical probability of representing the user's intended plan, the model should represent intelligent, rational decisions about the user's intentions — decisions that can be explained and justified to the user when questions about them arise. Although various strategies could be devised for constructing a model with the highest probability of occurrence[2, 7, 8], these schemes require that the system perform a great deal of time-consuming and complicated computation. In addition, even if the system's model of the user's plan is the one *most likely* to be correct, it *can* still be wrong. Analysis of naturally occurring dialogue indicates that although human information-providers occasionally make incorrect inferences during a dialogue, they can rationally justify them to the information-seeker when errors arise. Unfortunately, complex probability computations are extremely difficult to explain and justify to a lay person. Yet if an information system is to be viewed by users as intelligent and cooperative, it must appear rational and must be able to explain its reasoning; otherwise its users will not have confidence in the system's ability to intelligently assimilate a dialogue and provide helpful advice.

We contend that an advisement system should model the user's plans and goals by following strategies that not only produce a rich set of beliefs about what the user is trying to do but also capture the kind of behavior exhibited by intelligent human information-providers. If the system does this, then when it finds that its model of the user's plan is in error, it can explain its reasoning to the user and the user is likely to accept as reasonable the system's decisions and bases for making them. The user will come to expect the system to make the kinds of inferences human information-providers generally make and a naturally appearing dialogue can ensue. Note that our intent is not to simulate a particular individual or group of individuals, but rather to produce behavior that can be justified to a human observer and which the observer would regard as intelligent, rational, and natural.

Research by psychologists has provided insight on prediction and inference by humans. In [18], it was found that humans tend to develop a hypothesis explaining a planning agent's actions, expand their beliefs about the agent's plan as much as possible without making unwarranted assumptions about subactions and parameter bindings, and then revise this hypothesis as necessary to accommodate new actions. In [6], it is argued that humans do not reason with large numbers of alternative scenarios and make inferences from these possibilities by complex combinations of uncertain information. Instead, the *As-if* model[6], proposed to explain human behavior in multi-stage inferencing, hypothesizes that humans gather additional information until their confidence in an intermediate conclusion exceeds a threshold, then adopt this conclusion as certain evidence in the next stage of inferencing.

Our model of plan inference incorporates these aspects of human inferencing. It develops a rational hypothesis about an agent's plan by both sanctioning appropriate default inferences and deferring unwarranted decisions until further evidence is available. It reasons about the support that evidence about the user's intended actions gives to alternative conclusions about his goals, has a confidence threshold at which a conclusion is accepted and added to the system's beliefs about the user's plan, views actions that are representative parts of performing a higher-level action as confirming the latter's presence in the system's beliefs about the user's plan, and revises the model when contradictions are detected.

The Process Model

System Overview

Dynamic plan recognition requires that the user's plan be built incrementally as the dialogue progresses. We use a tree structure called a context model[1] to represent the system's beliefs about the user's plan as inferred from the preceding dialogue. Each node in the tree represents an action that the system believes

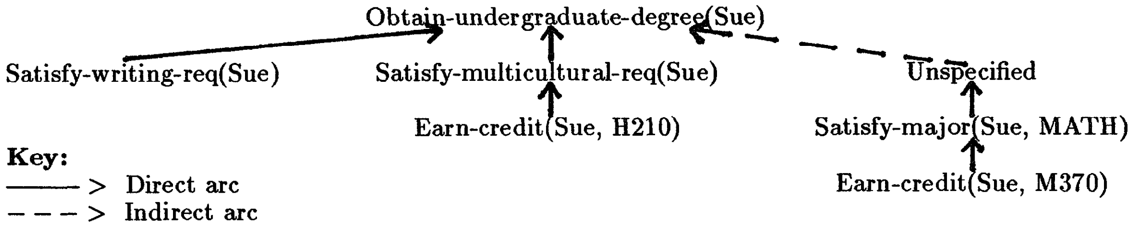

Figure 1: A sample context model

the user wants to perform, along with a set of beliefs that have been ascribed to the user in inferring the action[14]. The arcs connecting nodes are of two types: 1) a direct arc from action A_i to A_j indicates that A_i plays a direct role in performing A_j — ie., A_i is part of an operator description for A_j in the system's plan library, and 2) an indirect arc from A_i to A_j indicates that A_i is part of an expanded plan for A_j, but exactly what role it plays in that plan is as yet ambiguous — ie., there is more than one plausible way to expand the operator for A_j so that it includes A_i and a decision about which expansion represents the user's intentions is unwarranted at the current time. In addition, direct arcs are divided into three classes, representing a generation[14], enablement, or subaction relationship between the actions they connect. Figure 1 illustrates a sample context model containing both direct and indirect arcs. It indicates that earning credit in H210 (an Asian history course) is a subaction in Sue's plan for satisfying the multicultural requirement for getting an undergraduate degree. In addition, it indicates that the system believes that Sue wants to take M370 as part of fulfilling a math major, but that there is insufficient evidence for deciding exactly how satisfying the requirements for a math major fits into her overall plan (ie., whether it is part of obtaining a BA, BS, or Liberal Studies degree).

As each new utterance occurs, the context model must be adjusted and expanded to reflect an updated hypothesis about the information-seeker's plans and goals. We use plan identification heuristics to hypothesize domain actions that might motivate the user's new utterance and focusing heuristics to identify the most coherent relationship between a hypothesized action and the user's current focus of attention in the context model. These are described in [1] and will not be repeated here. If there is only one expansion of the context model that captures the relationship identified by the focusing heuristics, that expansion becomes the updated context model. However, there may be more than one way to expand the context model to include the new action in the specified way. Dynamically computed preference rules are used to represent alternative inferences that might be made from given evidence and order them according to plausibility. The decision to sanction a default inference is based on the conclusion's plausibility with respect to the plausibility of alternative competing conclusions, as estimated by their respective plausibility factors.

Preference Rules

Although our model of plan recognition does not rely on any particular representation scheme, we do need a means for representing how individual pieces of evidence lend support to alternative hypotheses about the user's plans and goals and for reasoning about the combined support offered by several items of evidence. We are using the Dempster-Shafer theory[10, 19] for this purpose. One advantage of the Dempster-Shafer theory over other approaches to evidential reasoning is that belief is distributed over subsets of the possible hypotheses, thereby allowing belief to be attributed to a hypothesis H without all remaining belief being attributed to $\neg H$. In addition, the result of combining evidence in the Dempster-Shafer theory models how the set of plausible hypotheses shrinks as additional evidence is evaluated and taken into consideration[10]. These features are useful for incorporating default inferencing into incremental plan recognition. They will allow us to view several competing conclusions as totally plausible without completely ruling out other conclusions and enable us to model how several pieces of information incrementally gleaned from an ongoing dialogue tend to single out one conclusion as far more plausible than the other possibilities.

In the Dempster-Shafer theory, the set of mutually exclusive alternative possible conclusions is called the *frame of discernment* Θ. A basic probability assignment m represents the impact of given evidence on belief in the alternative conclusions by distributing probability mass among the subsets of Θ. Dempster's rule of combination provides a method for combining evidence by computing a new basic probability assignment from the basic probability assignments associated with individual pieces of evidence.

We are using Dempster-Shafer basic probability assignments to represent the support that evidence about the user's actions gives to alternative conclusions about his goals. The frame of discernment Θ is the set of mutually exclusive higher-level goal actions that might be inferred as the reason that a user wants to perform the actions comprising the evidence E. We are assuming that an agent has only one immediate primary goal as his reason for wanting to perform the actions in E; thus the elements of Θ are mutually exclusive as motivation for the actions in E. *Others* is included as a subset of Θ in order to make it exhaustive and to account for user misconceptions and novel plans. The semantics of an entry $m(X) = k$ in a basic probability assignment is that evidence E commits a portion k of belief to the hypothesis that the user wants to perform some higher-level action in the set X; however, other entries may distribute extra portions of belief to subsets of X.

We contend that default recognition of user goals should be based on a goal's plausibility with respect to other competing goals. In order to do this, we need the notion of a default inference rule in which alternative possible conclusions are ordered by levels of plausibility. We capture this in *preference rules*[2] which have the form **IF** $< E >$ **P-THEN** $<$ *p-list* $>$ where

$< E >$	::=	actions comprising the evidence
$<$ *p-list* $>$	::=	$<$ *p-pair* $>$ $\vert$ $<$ *p-pair* $><$ *p-list* $>$
$<$ *p-pair* $>$	::=	$< A_g >< PF_p(A_g) >$
$< A_g >$	::=	an element of Θ
$< PF_p(A_g) >$	::=	plausibility factor for $< A_g >$

The preference rule P associated with evidence $< E >$ is constructed from the basic probability assignment m associated with $< E >$ by using its frame of discernment Θ and computing plausibility factors $PF_p(A_g)$ as $Plausible_m(\{A_g\}) = 1 - \sum_{(Y \subseteq \Theta) \wedge (Y \cap \{A_g\} = \emptyset)} m(Y)$. This last formula measures the extent to which the inconclusive evidence accumulated thus far makes A_g plausible by failing to refute it[19]. Thus the ***plausibility factor*** for a goal action A_g captures what Reiter[16] calls the intuitive plausibility of a default. Figure 2 illustrates several basic probability assignments (bpa's) and the preference rules compiled from them.

Sanctioning Default Inferences

If the system believes that the user wants to perform a set of actions E and that there is more than one action A_g whose performance would include all the actions in E, then the system must be able to decide whether one of these actions should be identified as the user's goal. We contend that this decision should be based on the plausibility of a goal action with respect to the alternative possible goal actions — namely, if one action is extremely plausible as the user's goal and far more plausible than the other possibilities, then it should be recognized as part of the user's plan. We are modeling this decision-making by maintaining a *threshold plausibility level* ϵ_{pl} and a *threshold difference level* ϵ_d, and sanctioning inference of a goal action A_g by default from the actions comprising the evidence E if A_g is the most plausible goal that might motivate the actions in E, A_g's plausibility factor exceeds the threshold plausibility level ϵ_{pl}, and no other action suggested by the evidence E has a plausibility factor within the threshold difference level ϵ_d of A_g. If A_g can be inferred from E, either with certainty or by default, then we say that $Infer(E) = A_g$. More formally, if P is the preference rule associated with evidence E, then

$$
\begin{aligned}
&Infer(E) = \\
&A_g \quad \text{if } A_g \text{ is the only action whose associated} \\
&\qquad \text{operator contains the actions in E} \\
&A_g \quad \text{if } (A_g \in \Theta_p) \wedge (PF_p(A_g) > \epsilon_{pl}) \\
&\qquad \wedge[\neg \exists A_k \text{ s.t. } (PF_p(A_k) > PF_p(A_g) \\
&\qquad\qquad \vee PF_p(A_g) - PF_p(A_k) < \epsilon_d)] \\
&\emptyset \quad \text{otherwise}
\end{aligned}
$$

The threshold settings are determined by the criticality of the interaction (medical versus travel domain).

Building the Context Model

Although our processing strategy and heuristics are domain-independent, the system must be provided with domain knowledge representative of that required by a capable human information-provider. This knowledge consists of the set of actions that a user might pursue in the domain, operators that describe how to perform these actions, and basic probability assignments representing how individual actions lend support to alternative conclusions about the user's goals. Preference rules are computed dynamically from combinations of one or more basic probability assignments. Since operator descriptions contain subactions which also have associated operators, a plan for an action can be expanded to any desired degree of detail by starting with the operator for the action and repeatedly replacing subactions with their own operator descriptions.

Each new utterance must be assimilated into the context model to produce an updated hypothesis about the user's plans and goals. As described earlier, plan identification heuristics are used to hypothesize domain actions that might motivate a user utterance and focusing heuristics are used to identify the most coherent relationship between a hypothesized action and the user's current focus of attention in the context model. The preceding sections showed how preference rules rank alternative conclusions by plausibility and can be used to sanction default inferences. This section presents some of our rules for incorporating default inferences into a model of plan recognition, along with examples illustrating the rules. The examples are taken from a student advisement domain; the relevant basic probability assignments are shown in Figure 2. We will assume a threshold plausibility level of $\epsilon_{pl} = .9$ and a threshold difference level of $\epsilon_d = .7$.

If there is only one expansion of the context model that captures the relationship identified by the focusing heuristics, the context model should be updated to include it. When more than one expansion satisfies the constraints of the focusing heuristics, Rule-D1 captures the notion of making default inferences that

[2]The initial work on preference rules was done with Kathy Cebulka.

Basic Probability Assignments			Preference Rules		
1) Evidence:	{Earn-credit(_user, M370)}		IF {Earn-credit(_user, M370)}		
	$m_1(\{A_3\})$	= .85	P-THEN	A_3	1.000
	$m_1(\{A_2, A_3\})$	= .12		A_2	.150
	$m_1(\Theta_1)$	= .03		*Other*	.030
2) Evidence:	{Earn-credit(_user, EE202)}		IF {Earn-credit(_user, EE202)}		
	$m_2(\{A_1\})$	= .15	P-THEN	A_1	1.000
	$m_2(\{A_1, A_2\})$	= .80		A_2	.850
	$m_2(\{A_1, A_2, A_3\})$	= .04		A_3	.050
	$m_2(\Theta_1)$	= .01		*Other*	.010
3) Evidence:	{Earn-credit(_user, CS321)}		IF {Earn-credit(_user, CS321)}		
	$m_3(\{A_2, A_3\})$	= .95	P-THEN	A_2	1.000
	$m_3(\{A_1, A_2, A_3\})$	= .02		A_3	1.000
	$m_3(\Theta_1)$	= .03		A_1	.050
				Other	.030
4) Evidence:	{Satisfy-major(_user, CS)}		IF {Satisfy-major(_user, CS)}		
	$m_4(\{A_4\})$	= .74	P-THEN	A_4	1.000
	$m_4(\{A_4, A_5\})$	= .25		A_5	.260
	$m_4(\Theta_2)$	= .01		*Other*	.010
5) Evidence:	{Satisfy-major(_user, MATH)}		IF {Satisfy-major(_user, MATH)}		
	$m_5(\{A_5\})$	= .15	P-THEN	A_5	1.000
	$m_5(\{A_4, A_5\})$	= .84		A_4	.850
	$m_5(\Theta_2)$	= .01		*Other*	.010

A_1 = Satisfy-major(_user, EE)
A_2 = Satisfy-major(_user, CS)
A_3 = Satisfy-major(_user, MATH)
A_4 = Earn-degree(_user, BS)
A_5 = Earn-degree(_user, BA)

Figure 2: Sample bpa's and Preference Rules for a University Advisement Domain

coherently mesh with the user's current focus of attention in his plan, while deferring unwarranted decisions until further evidence is accumulated. Note that the context model can now contain indirect arcs, indicating an incompletely specified relationship between actions.

Rule-D1: Suppose that the focusing heuristics have determined that the new action A_{new} associated with the user's utterance is part of a plan for performing an action A_c in the context model, but that there is more than one way of constructing a plan for A_c that includes A_{new}. If $Infer(\{A_{new}\}) = A_j$ and A_j can play a role in a plan for A_c, then add A_j to the context model, with a direct arc from A_{new} to A_j, and repeat Rule-D1 with A_j in place of A_{new}; otherwise add an uninstantiated node with a direct arc from A_{new} to this new node and an indirect arc from the new node to A_c.

Example-1:
Suppose that Sue has asked about satisfying such university requirements as a writing project and a multicultural course, leading the system to believe that she wants to get an undergraduate degree, and that Sue then asks about taking M370 (probability theory). The plan identification heuristics identify α = Earn-credit(Sue, M370) as the domain action motivating Sue's query. Since the system's domain knowledge includes an operator for earning an undergraduate degree and that operator can be expanded to produce a plan that includes taking M370, the focusing heuristics determine that Sue wants to take M370 as part of earning an undergraduate degree. However, there are several ways that taking M370 fits into such a plan, including satisfying the requirements for a math or a CS major, or merely filling a free elective. Since the first preference rule in Figure 2 produces Infer($\{\alpha\}$) = Satisfy-major(Sue, MATH) and fulfilling a math major can be part of a plan for getting an undergraduate degree, Rule-D1 results in the default inference that Sue wants to satisfy the requirements for a

math major. Since the fifth preference rule in Figure 2 produces Infer({Satisfy-major(Sue, MATH)})=∅, a decision about precisely how satisfying a math major fits into her plan — ie., about whether she intends to fulfill a math major as part of a BA or a BS degree — is deferred until further evidence is accumulated. The resulting context model was shown in Figure 1.

Since a single piece of evidence may be insufficient to warrant a default inference, a plan recognition system must have the ability to combine individual pieces of evidence. The next rule addresses the problem of recognizing actions from an accumulation of evidence.

Rule-D2: Suppose that the context model indicates that one or more actions $A_1, \ldots, A_k$ are part of some as yet unspecified higher-level action and suppose that the focusing heuristics determine that A_{new} is also part of this action. Add A_{new} and a direct arc from A_{new} to this as yet unspecified action, construct a preference rule for $\{A_1, \ldots, A_k, A_{new}\}$ from the bpa produced by combining the bpa's associated with the individual actions $A_1, \ldots, A_k, A_{new}$ (using Dempster's rule of combination), and if $Infer(\{A_1, \ldots, A_k, A_{new}\}) \neq \emptyset$, then instantiate the previously unspecified parent action with $Infer(\{A_1, \ldots, A_k, A_{new}\})$.

Example-2:
Suppose that the system believes that Al wants to take EE202 as part of a plan to perform some as yet unspecified higher-level action A_u. (Note that the second preference rule in Figure 2 produces Infer({Earn-Credit(Al, EE202)}) = ∅.) Now suppose that Al asks about the preconditions for taking CS321 and that Earn-credit(Al,CS321) is identified as the action motivating his new utterance. The focusing heuristics must now determine the most coherent relationship between this new action and the existing context model. Since taking CS321 and taking EE202 are both possible actions for earning majors in math, computer science, and electrical engineering, the focusing heuristics suggest that taking CS321 and taking EE202 are both actions that will be executed as part of a plan to perform the unspecified higher-level action A_u. Rule-D2 applies. Using Dempster's rule of combination[19], the bpa associated with taking EE202 is combined with the bpa associated with taking CS321 (Figure 2) to produce

Combined basic probability assignment

$m(\{A_1\})$	$= .0088$	$m(\{A_2, A_3\})$	$= .0554$
$m(\{A_2\})$	$= .8862$	$m(\{A_1, A_2, A_3\})$	$= .0026$
$m(\{A_1, A_2\})$	$= .0467$	$m(\Theta_1)$	$= .0003$

Preference Rule

IF {Earn-credit(Al,EE202), Earn-credit(Al,CS321)}

P-THEN	A_2	.9912
	A_1	.0584
	A_3	.0583
	Other	.0003

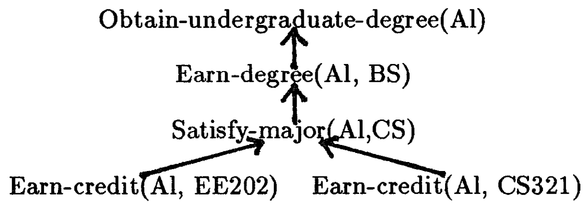

Figure 3: Extended Context Model

In this case, Infer({Earn-credit(Al,EE202),Earn-credit(Al,CS321)}) = Satisfy-major(Al, CS), resulting in the default inference that Al is trying to satisfy the requirements for a CS major. This produces the lower half of the context model in Figure 3.

The next rule addresses the problem of default recognition of higher-level actions motivating the actions that are currently part of the system's beliefs about the user's plan.

Rule-D3: If A_r is the action at the root of the context model and $Infer(\{A_r\}) \neq \emptyset$, then add $Infer(\{A_r\})$ to the context model along with a direct arc from A_r to $Infer(\{A_r\})$; otherwise, add an uninstantiated node with a direct arc from A_r to this new node, indicating that the user may be intending to execute A_r in order to perform some other action whose identify is as yet unknown.

Example-3:
Consider the situation at the end of Example-2, in which the top-level action that has been inferred as part of Al's plan is Satisfy-major(Al, CS). Rule-D3, with this conclusion as evidence, leads to the default inference that he is pursuing a Bachelor of Science degree. Since pursuing a Bachelor of Science degree only plays a role in obtaining an undergraduate degree, this higher-level action is added to the system's beliefs about his overall plan, producing the context model shown in Figure 3.

With the inclusion of default inferences into the context model, new utterances not only can cause the context model to be expanded, but may also confirm, contradict, or call into question components of the model previously added by default. The next two rules accumulate evidence confirming an action added by default to the user's plan; this evidence is being used in our current research to hypothesize the source of disparity when errors in the context model are suspected.

Rule-D4: Suppose that A_d is an action added by default to the context model and that the new utterance either explicitly confirms A_d or is equated with an action that can only be part of a plan for accomplishing A_d — ie., there are no other high-level actions that can account for the domain action motivating the user's utterance. Then mark A_d's status as confirmed and note the type of confirmation.

Rule-D5: Suppose that 1) the focusing heuristics determine that the most coherent relationship between the existing context model and the action A_{new} associated with the current utterance is that A_{new} will be executed as part of a plan to perform an action A_d added by default to the context model, and 2) A_{new} is a plausible action in a plan for A_d (ie., it is representative of what one might expect an agent to do to perform A_d). Then record the occurrence of A_{new} as confirming the presence of A_d in the user's plan.

We contend that representativeness should be used in plan recognition to confirm existing beliefs about what the user wants to do, but that it should not be used as the basis for default inferences. Rule-D5 allows the action of *going to the post office* to confirm a previous default inference that the user wants to *register for the Selective Service*, since going to the post office is representative of (ie., a typical part of a plan for) registering for the Selective Service. However, it avoids the error of using representativeness to make default inferences[11] — ie., when an agent is observed going to the post office, representativeness should not be used to sanction the default inference that he must be registering for the Selective Service.

The user's utterance can also contradict an action A_{d_3} added by default to the context model. Not only must A_{d_3}, along with all higher-level actions inferred from it, be removed, but also a decision must be made about whether to retain the components of the context model that led to inference of A_{d_3}. If A_{d_3} was inferred by default from another action A_{d_2} also inferred by default from an action A_1 and current world knowledge contradicts all plausible conclusions that might result from adopting A_{d_2}, then confidence that A_{d_2} is really part of the user's plan diminishes and it should be removed from the context model. On the other hand, if A_{d_2} is retained in the context model, then A_{d_3} may be replaced with a new action inferred by default once A_{d_3} is removed from consideration. Our method in the case where A_{d_2} has not been confirmed by other evidence is the following:

Rule-D6: Suppose that A_{d_3} was inferred by default from another default action A_{d_2}, A_{d_2} has not been confirmed by other evidence, and A_{d_3} is contradicted and must be removed from the context model.

1. If one of the alternatives to A_{d_3} has at least some minimal plausibility in the preference rule associated with A_{d_2}, then revise this preference rule by recomputing its plausibility factors from the bpa produced by combining $m(\Theta - \{A_{d_3}\}) = 1.0$ with the bpa from which the preference rule was formed. If $Infer(\{A_{d_2}\}) \neq \emptyset$ using this revised preference rule, then add $Infer(\{A_{d_2}\})$ as the new parent of A_{d_2} in the context model.
2. If none of the alternatives to A_{d_3} has at least some minimal plausibility in the preference rule associated with A_{d_2}, then retract A_{d_2} from the context model.

Example-4:
Suppose that the system believes that Al's plan contains the actions shown in Figure 3 and that the system then finds that Al is not pursuing a BS degree.[3] Rule-D6 applies. The default inference that Al is pursuing a BS degree, and therefore an undergraduate degree, is withdrawn. This default inference resulted from the fourth preference rule in Figure 2, using as evidence the belief that Al wanted to major in computer science. Since the alternative conclusion that he is pursuing a BA degree is not implausible, the belief that he wants to major in computer science is retained in the context model and the bpa from the fourth preference rule in Figure 2 is combined with $m(\Theta_2 - \{A_4\}) = 1.0$, producing a revised bpa in which $m(\{A_5\}) = .96$, $m(\{A_4\}) = 0$, and $m(\Theta_2 - \{A_4\}) = .04$. This leads to the revised preference rule

IF {Satisfy-major(USER, CS)}
P-THEN A_5 1.00
Other .04

and Infer({Satisfy-major(Al,CS)})=A_5, thus producing the default inference that Al is pursuing a BA degree instead of a BS degree. From this the system again infers that he wants to get an undergraduate degree.

Future Research

This paper has described a process model for incorporating default inferences into plan recognition and has presented rules for inferring higher-level actions that are the motivation for observed actions. We are also formulating rules that apply when the user is already believed to be pursuing a particular higher-level action. These rules take into account the relative plausibility of alternative possible subactions in the plan for that action and even default inferences about these. In addition, we are working on extending our process model to make generalized inferences[12] about higher-level goal actions that are not the immediate parent of an existing action in the context model.

We are also developing an overall strategy for revising the system's context model when the user's utterances suggest possible disparity between it and what the user is actually trying to accomplish. This strategy will use the system's relative confidence in components of the context model, along with meta-knowledge about how utterances were related to one another using focusing heuristics and how default goals were inferred, to justify and explain the system's beliefs to the user, formulate an intelligent hypothesis about the source of

[3] This might happen in several ways. The simplest case would be a direct statement to this effect by Al. A more realistic scenario would be a query from Al about satisfying a foreign language requirement, where it is mutually believed that BS degrees do not have such a requirement.

error, and guide a negotiation dialogue to remedy the error[5].

Conclusions

This paper has presented a process model of plan inference that is motivated by an analysis of naturally occurring dialogues and by psychological studies of human inference and plan recognition. It includes a strategy for incrementally updating the system's model of the user's plan that can both sanction appropriate default inferences and defer unwarranted decisions until further evidence is available. In this strategy, dynamically computed preference rules are used to rank alternative conclusions according to plausibility, and the decision to sanction a default inference is based on the conclusion's plausibility with respect to the plausibility of alternative competing conclusions. We have presented a set of rules that incorporate appropriate default inferences into the system's model of the user's plan and update and revise the model as the dialogue progresses. Our process model overcomes a limitation of previous plan recognition systems and will produce a rich model of the user's plans and goals that is rational and can be explained and justified to the user when questions about it arise.

References

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

[2] Peter Cheeseman. A method of computing generalized bayseian probability values for expert systems. *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, 1983.

[3] Philip R. Cohen, C. Raymond Perrault, and James F. Allen. Beyond question answering. In W. Lehnert and M. Ringle, editors, *Strategies for Natural Language Processing*, pages 245–274, Lawrence Erlbaum Associates, 1981.

[4] James P. Delgrande. An approach to default reasoning based on a first-order conditional logic. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 340–345, Seattle, Washington, 1987.

[5] Rhonda Eller and Sandra Carberry. A meta-rule approach to dynamic plan recognition. In *Proceedings of the Second International Workshop on User Modeling*, Honolulu, Hawaii, 1990.

[6] Charles F. Gettys, Clinton Kelly III, and Cameron R Peterson. The best-guess hypothesis in multistage inference. In Daniel Kahneman, Paul Slovic, and Amos Tversky, editors, *Judgment Under Undertainty: Heuristics and Biases*, pages 370–377, Cambridge University Press, 1982.

[7] Robert Goldman and Charniak. A probabilistic approach to plan recognition and text understanding. In *Proceedings of the 1989 Workshop on Plan Recognition*, Detroit, Michigan, 1989.

[8] Robert Goldman and Eugene Charniak. A probabilistic atms for plan recognition. In *Proceedings of the AAAI Workshop on Plan Recognition*, Saint Paul, Minnesota, 1988.

[9] Bradley Goodman and Diane Litman. Plan recognition for intelligent interfaces. In *Proceedings of the Sixth Conference on Artificial Intelligence Applications*, 1990.

[10] Jean Gordon and Edward H. Shortliffe. A method for managing evidential reasoning in a hierarchical hypothesis space. *Artificial Intelligence*, 26:323–357, 1985.

[11] Daniel Kahneman and Amos Tversky. On the psychology of prediction. In Daniel Kahneman, Paul Slovic, and Amos Tversky, editors, *Judgment under Uncertainty: Heuristics and Biases*, pages 48–68, Cambridge University Press, 1982.

[12] Henry Kautz and James Allen. Generalized plan recognition. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 32–37, Philadelphia, Pennsylvania, 1986.

[13] R. Perrault and J. Allen. A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics*, 6(3-4):167–182, 1980.

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

[15] D. L. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36(1):27–47, 1988.

[16] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[17] Roger C. Schank and Robert P. Abelson. *Scripts, Plans, Goals and Understanding.* Lawrence Erlbaum Associates, Hinsdale, New Jersey, 1977.

[18] C. F. Schmidt, N. S. Sridharan, and J. L. Goodson. The plan recognition problem: an intersection of psychology and artificial intelligence. *Artificial Intelligence*, 11:45–82, 1978.

[19] G. Shafer. *A Mathematical Theory of Evidence.* Princeton University Press, Princeton, New Jersey, 1976.

[20] Candace L. Sidner. Plan parsing for intended response recognition in discourse. *Computational Intelligence*, 1:1–10, 1985.

[21] Robert Wilensky. *Planning and Understanding.* Addison-Wesley, 1983.