

# Goal Inference in Information-seeking Environments

Bhavani Raskutti & Ingrid Zukerman

Computer Science Department

Monash University

Clayton, Victoria 3168

Australia

## Abstract

In cooperative information-seeking environments, we have observed that the dialogues have the following characteristics: (1) they contain sufficient relevant information, (2) they are coherent, and (3) they are well-structured. In this paper, we describe a mechanism for plan inference which takes advantage of these observed features to reduce the number of alternate interpretations of a user's statements. This reduction is achieved as follows: initially, we take advantage of the relevant information trait by using guiding principles and meta predicates to constrain the number of possible interpretations of a single statement. *Discourse coherence* considerations are then applied to integrate subsequent statements and drop incoherent interpretations. The retained interpretations are evaluated using a measure of *information content*, which is used to prefer the interpretations that have more relevant information. The entire mechanism is based on an approach that takes advantage of the well-structured nature of information-seeking dialogues to arrive at the intended interpretation as efficiently as possible.

## Introduction

In this paper, we present a mechanism which infers the plans and goals of a user in cooperative information-seeking settings, e.g., travel agency. The inference is done by generating the possible interpretations of a user's statements, evaluating them and then selecting the more likely interpretations. An interpretation of a user's statements consists of a sequence of plans that the user proposes to carry out, and a plan consists of an action defined by a number of parameters. For instance, in the travel domain, the proposal to fly from Melbourne to Sydney on December 1st, 1990, is a plan, where *flying* is the action, and the parameters *origin*, *destination* and *departure date* are instantiated.

Our inference mechanism operates in the framework provided by a Natural Language Interface (NLI) and a planner. However, in this paper, we focus on the plan inference mechanism, and the only references to the

NLI and planner concern our assumptions about them.

In cooperative information-seeking settings, we have observed that the dialogues have the following features: (1) they contain enough relevant information so that the information seeker's plans may be understood easily by the information provider, (2) they are coherent, and (3) they are well-structured, i.e., the dialogue consists of an initial discourse, followed by a clarification interaction. For instance, the following dialogue<sup>1</sup> at a Melbourne travel agency displays these features:

Traveler: "I need a ticket to Sydney tomorrow.

I am going to Hawaii on the 11:00 am flight.

By the way, I'll be leaving from Adelaide."

Agent: "What about your ticket to Adelaide?"

Traveler: "That is all arranged. Thank you."

Our plan inference mechanism takes into consideration these traits of the dialogue while interpreting a user's statements. It differs from earlier mechanisms for plan inference (Grosz 1977, Schmidt, Sridharan & Goodson 1978, Allen & Perrault 1980, Sidner & Israel 1981, Carberry 1983, Litman & Allen 1987, Pollack 1990) in that it can handle multiple interpretations. The problem of multiple interpretations has been handled by Carberry (1990) by using plausibility factors of alternate hypotheses. However, in domains such as travel, where the plans in an interpretation are defined by a number of parameters, this approach alone does not cope with the problem of multiple interpretations. The mechanism presented in this paper addresses this problem by using the observed features of cooperative dialogues, namely, *relevance*, *coherence* and *structure*.

Initially, we take advantage of the relevance trait by using guiding principles and meta predicates to constrain the number of interpretations of a single statement. In the final stages, the amount of relevant information in an interpretation is measured by its information content, and is used to prefer interpretations with more relevant information.

The coherence of an interpretation is determined by

<sup>1</sup>This dialogue is similar to those found in the transcripts of telephone conversations at travel agencies, which were provided by Prof. J. Roach and D. Sanford from the Virginia Polytechnic Institute and State University.

the relation between the interpretation of a statement and the interpretation of the previous discourse. Litman and Allen (1987) have applied ideas related to discourse coherence in order to select a preferred plan. However, since they maintain only one preferred option, they may choose an interpretation that is locally coherent, but does not fit in with the entire dialogue. In this research, this situation is prevented by maintaining other interpretations with less coherent relations while pursuing the more likely interpretations.

All dialogues, whether well-structured or otherwise, can be handled by a single approach, which we call the *cognitive* approach. In this approach the system draws both direct inferences based on the user's statement, and indirect inferences based on domain and world knowledge from each statement as it is uttered. For instance, upon receiving the first statement in our sample text, the system, after making all the direct inferences, will go on to make the indirect inference that the speaker is probably leaving from Melbourne, where the conversation is being conducted. However, some of the system's indirect inferences may have to be retracted due to information provided later on, as when the traveler says "By the way, I'll be leaving from Adelaide," in our sample text. Evidence provided by written and spoken text, where speakers/writers address inferences presumably performed by their hearers/readers, leads us to believe that this approach models human cognition. For instance, meta comments such as "however," "indeed" and "in fact" indicate violation of expectations, realization of expectations and violation of implicatures, respectively.

The cognitive approach is computationally expensive for the initial discourse, due to the retractions that may have to be performed. These retractions may be avoided by using another approach, namely the *implementational* approach. In this approach, the system first draws all the direct inferences from all the statements in the initial discourse, and only then draws the indirect inferences. In this case, there are less chances of early unwarranted inferences, and hence, from the viewpoint of implementation, this approach is attractive. However, this approach alone cannot cope with a clarification interaction after the initial discourse.

In this paper, we argue for a hybrid approach that takes advantage of the benefits of both approaches by taking into account the observed structure of information-seeking dialogues. In the hybrid approach, the implementational approach is applied to the initial discourse, and the cognitive approach is applied to subsequent statements uttered either after a pause or in response to a query posed by the system. The rest of this paper describes our mechanism for plan inference based on the hybrid approach, with particular reference to the strategies that take into account the characteristics of cooperative information-seeking dialogues in order to reduce the number of alternate interpretations of a user's statements.

## The Main Mechanism

The inference mechanism operates on input provided by the NLI, consisting of predicates, such as FLY and GO, and meta predicates that indicate the modality of the statements, such as WANT. Based on this input, it generates the intended plans of the user. The mechanism is based on the hybrid approach and it consists of two parts: (1) the implementational part for handling the initial discourse  $S_i$ , and (2) the cognitive part for processing the clarification statements  $S_c$ . Both parts apply three main processes: (1) Direct inference, (2) Indirect inference, and (3) Evaluation of interpretations. However, they are applied at different times in each part. In the implementational part, direct inference of all statements followed by evaluation precedes indirect inferences followed by another evaluation. In the cognitive part, direct and indirect inferences are interspersed, with the evaluation performed last. These three processes are described below with reference to the two procedures *Implementational-Infer* and *Cognitive-Infer*, which embody the two parts of the mechanism. The set of interpretations,  $\{IS\}$ , is initially set to *nil*.

Procedure *Implementational-Infer* ( $S_i$ )

- 1 For each statement  $s$  in the discourse  $S_i$  do:
- 2     Generate a set of interpretations  $I(s)$  using guiding principles and meta predicates
- 3     Infer the set  $R$  of possible discourse relations between  $I(s)$  and  $\{IS\}$  and order them according to the coherence measure
- 4     Combine  $I(s)$  and  $\{IS\}$  using  $R$  to create a new  $\{IS\}$ , and determine the likelihoods of the interpretations in the new set
- 5     Use these likelihoods to prune unlikely interpretations from  $\{IS\}$
- 6 For each interpretation  $I'$  in  $\{IS\}$
- 7     Use the information content of  $I'$  to revise its likelihood
- 8     Prune unlikely interpretations from  $\{IS\}$
- 9     For each interpretation  $I'$  that is left in  $\{IS\}$
- 10     Repeat until [  $I'$  is fully defined ] or [ the information content of  $I'$  cannot be increased ]
- 11     Infer values for all undefined necessary parameters in all the plans in  $I'$
- 12 Prune  $\{IS\}$  using information content (lines 7-10)

Procedure *Cognitive-Infer* ( $S_c$ )

- 1 For each statement  $s$  in the discourse  $S_c$  do:
- 2     Perform direct inferences as in lines 2-5 in procedure *Implementational-Infer*
- 3     Perform indirect inferences as in lines 9-11 in procedure *Implementational-Infer* to infer the values of all necessary parameters, and update them so that earlier inferences are not refuted by later weaker inferences
- 4 Prune  $\{IS\}$  using information content as in lines 7-10 in procedure *Implementational-Infer*

## Direct Inference

The direct inference stage (lines 2-5 in procedure *Implementational-Infer* and line 2 in procedure *Cognitive-Infer*) consists of: (1) Inferring a set of interpretations from each of the user's statements (line 2 in procedure *Implementational-Infer*); (2) Inferring a set of relations between each interpretation of a user's statement and the previous discourse (line 3); and (3) Generating a set of new interpretations (line 4). In this process, the likelihood of each interpretation is determined, and less likely interpretations are not pursued.

The inference of interpretations from one statement issued by a user consists of inferring a set of possible actions which match this statement, and computing their likelihoods. This inference is done by using a STRIPS-like operator library (Fikes & Nilsson 1971) and plan inference rules (Allen & Perrault 1980). Since the number of interpretations based on just one statement and the operator library can be quite large in a realistic domain, we control the activation of the plan inference rules by means of guiding principles, thus reducing the number of possibilities. In addition, meta predicates are used during this stage to modify the strength of the plan inference rules.

The inference of the discourse relations between the interpretations of one statement and the interpretations of the previous discourse is done so that the coherence of the discourse is increased.

The inference of a new interpretation, based on an interpretation of a new statement, an interpretation of the previous discourse and one of the inferred discourse relations, is done by updating the interpretation of the previous discourse using the relation and the information in the interpretation of the new statement. The likelihoods of the new interpretations are determined using Bayesian theory of probability, and interpretations whose likelihoods are lower than an acceptable threshold are dropped (Raskutti & Zukerman 1991).

## Indirect Inference

The indirect inference stage (lines 9-11 in procedure *Implementational-Infer*, and line 3 in procedure *Cognitive-Infer*) is used to infer information that has not been explicitly stated by the user. In the implementational part, it is performed after direct inference of a chunk of statements. In the cognitive part, it is performed after direct inference of each statement. The strength of indirect inferences is assessed on the basis of their source of information. For instance, the desired mode of transport between Sydney and Hawaii may be inferred by taking into consideration typical assumptions about the domain. This type of inference is stronger than an inference based on general 'world knowledge', but weaker than a direct inference.

Indirect inferences are applied to infer each of the parameters that are necessary for plan definition, and inferences are drawn until a complete interpretation is inferred or no more inferences can be drawn on the

basis of the existing information. In the implementational approach, only the parameters that remain undefined after direct inference are inferred. In the cognitive approach, however, all the necessary parameters are inferred, and previous inferences may be refuted or corroborated in light of new inferences.

## Evaluation of Interpretations

After a chunk of statements is processed, the generated interpretations are evaluated using a measure of information content, and the interpretations that are better defined according to this measure are retained. This evaluation is performed twice in the implementational approach, first after all the direct inferences have been drawn, and then again after the completion of the indirect inferences (lines 6-8 and 12 in procedure *Implementational-Infer*). In the cognitive approach, on the other hand, the evaluation is deferred until all the statements in a chunk are fully processed (line 4 in procedure *Cognitive-Infer*).

## The Strategies in Detail

In this section, we describe the strategies that are used to reduce the number of possible interpretations of a user's statements. These strategies take into account two aspects of observed dialogues in information-seeking settings, namely, coherence and relevance.

### Guiding Principles

Guiding principles restrict the number of interpretations that match a single statement of a user by controlling the search through the operator library that defines the basic actions in the domain. Each operator is defined in terms of its preconditions, effects and composition (also called the body of the operator).

During the inference of interpretations of a statement that has been parsed into an input predicate, all those operators that have the predicate in the definition of their precondition, effect or body, are chosen as possible matches. This can lead to the selection of a large number of operators as possible matches. Hence, we use the following guiding principles, which are based on the relevance trait, to restrict the inference process, so that only justifiable inferences are drawn.

1. *Least Complication* – requires that the most primitive of the matched operators must be selected. For instance, when a user says "I want to go to Sydney," then while both one-way and two-way travel are possibilities, one-way travel is chosen.
2. *Least Commitment* – requires that the operator that needs the least number of additional assumptions be chosen (Sacerdoti 1977). In practice, this means that more general operators are chosen, unless there is an explicit reference to a specific operator. For instance, unless there is an explicit reference to a means of transport, those operators that are uncommitted about the means of transport are chosen.

## Meta Predicates

Meta predicates, such as CAN and WANT, represent the modality or propositional attitude of the statements. They are returned by the NLI in order to retain information that may be lost when parsing from natural language to predicates. The presence of a meta predicate enables us to resolve ambiguity when the input predicate appears in both the precondition and the effect of an operator. For instance, if a user's request about BEING at a place is expressed as "I want to be ... " or "I can be ... ," but the NLI returns it as BE(...), this predicate can refer either to the precondition or the effect of an operator with equal likelihood. By taking into account the presence of a meta predicate, we increase the bias towards the appropriate inference rule, e.g., if the user had said "I can be ... " then the likelihood of the precondition rule is increased.

## Discourse Coherence

Discourse coherence considerations are used during direct inference to bias the inference process towards interpretations that make the discourse more coherent. To this effect, the inferred relations are assigned likelihoods that prefer normal patterns of discourse, e.g., the elaboration of the latest plan is preferred to the elaboration of older plans; the likelihood of an introduction is reduced if elaboration is possible; and if there are cue words to indicate a relation, then these are used during the inference of relations (Raskutti & Zukerman 1991).

A new statement of the user can either *elaborate* on an earlier topic, *introduce* a new topic, *digress* from the last topic, or *correct* some information specified earlier. The different possible relations coupled with the different topics that may be corrected or elaborated give rise to a number of different interpretations. The likelihoods of these interpretations depend on the likelihoods of the inferred relations, which in turn depend on whether the NLI has recognized a discourse-statement relationship and/or the referred topic. Thus, there are four possibilities:

1. If there is an explicit reference to an earlier topic, and a discourse relation is specified, then the system has no uncertainty about the topic or the relation.
2. If only a topic is specified, then we assume an elaboration relation, e.g., "About the Hawaii trip, I'll be flying QANTAS."
3. If only a discourse-statement relation is specified, e.g., "On second thought, make it 10:00 am" (correction), the system needs to determine which of the earlier topics is being referred to. Notice, however, that the only possible discourse relations pertaining to previous topics are Elaboration and Correction. This is due to the fact that in Introduction, the introduced topic is unrelated to previous topics, and Digression is considered as a special case of Elaboration, where the likelihood of elaborating on the

latest topic is substantially reduced. The determination of the topics which are elaborated or corrected is performed by following common patterns of discourse, where the elaboration or correction of topics mentioned later in the discourse is preferred to that of topics mentioned earlier.

4. If neither a topic nor a discourse relation is specified, then the inference of discourse relations is performed by first determining which relations are possible and then calculating the likelihood of each postulated relation. We do not postulate that the new statement is a correction or a digression unless there is an explicit indication from the NLI. Hence, the only possible relations are Elaboration of each topic in the interpretation of the earlier discourse, and Introduction of a new topic. Like above, elaboration of topics mentioned later in the discourse is preferred to that of topics mentioned earlier. Introduction of a new topic is always considered a possibility. However, its likelihood is reduced if Elaboration is possible.

## Strength and Categorization of Inferences

The strength of an inference is directly proportional to the reliability of the information source on which the inference is based. Hence, we categorize different information sources that are used to draw inferences and list them below in decreasing order of reliability.

1. *User's Statements* – Direct inferences from what is explicitly stated. While these inferences can be presumed correct, there is still a degree of uncertainty in relating a new statement to the earlier ones due to the different discourse relations possible.
2. *Domain Knowledge* – Indirect inferences that are derived by using the system's beliefs about the user's domain knowledge. A typical example is the inference of the arrival time at the destination once the departure time is known.
3. *Domain Assumptions* – Indirect inferences that are derived by assuming what is normal in the domain. For example, when the mode of travel is not specified, it is possible to derive this information from the usual mode of transport between two places.
4. *User Model* – Indirect inferences that are made on the basis of the system's model of the user. The user model may be a default model describing a typical user, or it may be more specific. In our system, we have adopted a default model based on the assumption that, typically, the information provider cannot form an extensive user model.
5. *Common-Sense* – Indirect inferences that are derived using common notions outside the domain of interest. Typically, such notions are used when we postulate return journeys based on the assumption that people usually do not move from their residence.

The inference types are assigned a strength in the (0,1] range, and this strength is used during the computation of the information content of a parameter. The inferences derived from the user's statements have

a strength of 1 and all other inference types have a progressively decreasing strength, according to the reliability of their source of information. The undefined parameters are assigned a minimum strength. This assignment enables us to distinguish between parameters that are defined inexactly by the user and parameters that are left undefined, and assign less information content to undefined parameters.

During the process of inferring the value of a parameter, we emulate one aspect of human behavior whereby once a conclusion is accepted with a particular degree of confidence, people consider it to be certain when drawing subsequent conclusions (Gettys, Kelly & Peterson 1982). To this effect, each parameter in the plans in each interpretation is tagged with the type of inference that gave rise to the value of the parameter, without taking into account the inference types of other parameters that were used for computing the parameter in question. Thus, like Carberry (1990), we do not compound the uncertainty in chains of inferences.

The strength of the inference type in a parameter's tag is used in the cognitive approach to determine whether a particular parameter should be revised by a new inference. If the strength of the new inference is the same or higher than the strength of the inference type in the tag, the new inference replaces the old one. Otherwise, the old inference is retained. In this manner, a weaker inference is prevented from refuting the results obtained from a stronger inference.

### Information Content

We define the information content of an interpretation as the sum of the information content of all the plans in the interpretation, and the information content of a plan as the sum of the information content of all the parameters that are necessary for the definition of the plan. The information content of a parameter, in turn, depends on two factors: (1) its specificity, which is defined as the reciprocal of the number of possible values assigned to this parameter; and (2) its strength, which depends on the source of information from which this parameter was obtained. Thus, both a parameter with multiple values assigned to it and a parameter derived from an unreliable source of information are deemed to have a low information content (Raskutti & Zukerman 1991).

The information content measure is used to update the likelihood of each interpretation in the set of interpretations. This set is then pruned by dropping those interpretations whose likelihoods fall below a relative rejection threshold. In this manner, the interpretations with more relevant information, i.e., those with a higher information content, are chosen.

### Conclusions

We have offered a mechanism which directs inferences towards more likely interpretations of a user's statements by using strategies which take into account

the observed traits of information-seeking dialogues, namely, relevance, coherence and structure. Currently, the algorithm for the implementational part is fully operational, with the indirect inference restricted to a few rules from each category. The cognitive part is in the last stages of its implementation. Finally, our experiments with a few discourse samples have indicated that our algorithm chooses the same interpretation that people choose. For instance, in the example presented in the Introduction, our mechanism considers four possible itineraries, and then chooses the intended one, i.e., *Melbourne* → *Adelaide* → *Sydney* → *Hawaii* → *Melbourne*, as the best interpretation.

### References

- Allen, J.F. and Perrault, C.R. 1980. Analyzing Intention in Utterances. *Artificial Intelligence* 15:143-178.
- Carberry, S. 1983. Tracking User Goals in an Information-seeking Environment. In Proceedings of the National Conference on Artificial Intelligence, 59-63.
- Carberry, S. 1990. Incorporating Default Inferences into Plan Recognition. In Proceedings of the National Conference on Artificial Intelligence, 471-478.
- Fikes, R.E. and Nilsson, N.J. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence* 2:189-208.
- Gettys, C.F.; Kelly III, C. and Peterson, C.R. 1982. The Best-guess Hypothesis in Multistage Inference. In Kahneman, D., Slovic, P. and Tversky, A. eds., *Judgment under Uncertainty: Heuristics and Biases*, 370-377. Cambridge University Press.
- Grosz, B.J. 1977. The Representation and Use of Focus in Dialogue Understanding. Technical Note 151, SRI International, Menlo Park, CA.
- Litman, D. and Allen, J.F. 1987. A Plan Recognition Model for Subdialogues in Conversation. *Cognitive Science* 11:163-200.
- Pollack, M. 1990. Plans as Complex Mental Attitudes. In Cohen, P., Morgan, J. and Pollack, M. eds., *Intentions in Communications*, MIT Press.
- Raskutti, B. and Zukerman, I. 1991. Generation and Selection of Likely Interpretations during Plan Recognition in Task-Oriented Consultation Systems. To appear in *User Modeling and User Adapted Interaction*, an International Journal.
- Sacerdoti, E. D. 1977. *A Structure for Plans and Behavior*. American Elsevier, New York.
- Schmidt, C.F.; Sridharan, N.S. and Goodson, J.L. 1978. The Plan Recognition Problem: An Intersection of Artificial Intelligence and Psychology. *Artificial Intelligence* 10:45-83.
- Sidner, C.L. and Israel, D.J. 1981. Recognizing Intended Meaning and Speakers' Plans. In Proceedings of the Seventh International Joint Conference on Artificial Intelligence, 203-208.