# Goals for Future Computational Models of Conversational Implicature

**Nancy Green and Jill Fain Lehman***
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15231-3891, USA
(Internet: nancy.green@cs.cmu.edu)

**Abstract**

We propose research goals for future models of conversational implicature, and describe how NL-Soar, a computational model of dialogue interpretation and generation implemented within the Soar cognitive architecture, serves as a step towards achieving these goals in a unified framework.

One avenue of research on conversational implicature (Grice 1975) in non-figurative discourse which should be addressed in the immediate future is the construction of computational models of agents who generate and comprehend conversational implicatures in the course of real-time dialogue. In the first section of this paper, we propose a set of research goals for such models. In the second section, we describe an earlier plan-based approach to conversational implicature which motivates our current research. Finally, we describe how our current work in NL-Soar, a computational model of dialogue interpretation and generation implemented within the Soar cognitive architecture, serves as a step towards achieving these goals in a unified framework.

## Goals for Future Models

We believe that future computational models of conversational implicature (CI) should aspire to meet the following goals:

- To derive all and only the conversationally licensed inferences which a speaker intends for the hearer to recognize as intended in the current discourse context. To model the contextual dependence of CI. To model cancelability phenomena. To be capable of handling multiple types of CI without unwarranted conflict.
- To provide tractable procedures for generating and interpreting CI. To model a speaker's motivation for use (or avoidance) of discourse conveying CI, and to model the processes of selecting *what* to say and *how* to say it. To model a hearer's process of recognizing that the speaker must have meant "more than what he said" and to model the process by which CI are derived.
- To be a cognitively plausible model. That is, to be subject to real-time performance constraints and human resource limitations.
- To learn from experience how to interpret novel CI. That is, new classes of CI can be interpreted without requiring the model's builder to provide new licensing rules specific to those classes.
- To model diachronically how certain CI become semi-conventional. E.g., to model the association of particular linguistic forms with certain CI, making it problematic to use those forms to convey just their original meaning.
- To lead to new predictions about and insights into use of language and how it interacts with other cognitive capabilities (e.g., planning and reasoning) of embodied agents capable of interacting with their environment through linguistic and non-linguistic means.

To elaborate, the first of the above goals suggests what we mean by an adequate model of conversational implicature as opposed to other aspects of discourse interpretation, e.g., entailment, presupposition, and ambiguity. Although it has been notoriously difficult to formulate a precise definition of CI, the notions of speaker meaning (Grice 1957), mutual belief (Clark & Marshall 1981), and speaker plans have been shown to be important (Hirschberg 1985; McCafferty 1987; Thomason 1990). In addition to accounting for cancellation by previous and subsequent context, models of CI should account for the change in a potential conversational implicature's status from cancelable to

*This research was supported in part by the Wright Laboratory, Aeronautical Systems Center, Air Force Materiel Command, USAF, and the Advanced Research Projects Agency under grant number F33615-93-1-1330. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Wright Laboratory or the U.S. government.

noncancelable as a discourse progresses(Gunji 1981). Moreover, this "window" of cancelability is determined by a discourse's intentional structure (Green 1994).

The second goal describes what we mean by a computational model of CI. The goal of cognitive plausibility can be justified not only on scientific but also on practical grounds. First, since human conversationalists comprehend and generate CI, searching for a cognitively plausible model is a useful research strategy. Second, if the model is to be used in human-computer interaction, it is desirable for it to meet the same real-time performance requirements as its human conversational partners.

The goals of learning to interpret novel CI and of modeling how certain CI become semi-conventional are desirable since it is unclear whether human model builders can anticipate all the situations in which CI may arise. Finally, the last goal seems desirable, not only for the sake of science, but also to enable our models to better approximate human competence in the long term. For example, for a first approximation, a model of a certain type of CI might be built which presupposes that the agent has acquired knowledge of principles of politeness. A more powerful model is one that explains how such principles arise, e.g., by learning through social interaction with other agents.

Traditionally, most work on CI has focused on the first two of the above goals. In contrast, our research strategy is to begin by building a foundational model of dialogue processing which will enable all of the above goals to be addressed in a unified framework. In the next section, we provide some motivation for this strategy.

## A Plan-based Approach to Indirect Answers

As for the goal of interpreting novel types of CI without requiring the model builder to provide licensing rules specific to that class of CI, first, let us consider the first author's computational model for interpreting and generating indirect answers to Yes-No questions (Green 1994). By use of an indirect answer as in (1), R conversationally implicates the intended answer (*no*) to Q's question.

(1) Q: Are you going to rent a car in Chicago?
R: I can't drive.

Interpretation and generation of indirect answers are modeled, respectively, as recognition and construction of the responder's discourse plan specification for a full answer. A full answer consists of the requested information, i.e. the direct answer, and possibly extra relevant information. A discourse plan specification relates a speaker's beliefs and discourse goals to his program of communicative actions. An indirect answer is the result of the responder providing only part of the planned response, but expecting that his full discourse plan specification will be recognized by the hearer.

An important knowledge source in that model is a shared set of discourse plan operators encoding generic programs for conveying full answers. A wide range of full answer types are characterized by the operators, which encode generalizations about *subject-matter* coherence relations (Mann & Thompson 1987) in responses, as well as about general discourse strategies motivating the selection of the "extra" information in responses. These operators play a key role in recognizing and constructing the speaker's discourse plan specification. Note that the model does not employ a licensing rule specific to indirect answers. Instead, an indirect answer is comprehended as a result of attempting to recognize the responder's discourse plan specification for a full answer. (When a full answer is not supplied by a speaker, the model uses a combination of hypothesis generation and plan inference to recognize his discourse plan specification.)

The discourse plan operators used in that model were constructed by the model's builder based on her analysis of examples of uses of indirect answers. An interesting question arises as to whether such a discourse-plan-based approach has any cognitive plausibility. Human conversationalists seem to be able to comprehend many CI in real-time, suggesting an approach that does not require reasoning from first principles in all cases, i.e. an approach that uses some type of "compiled" knowledge. As a model of natural language processing situated within a cognitive architecture, NL-Soar provides us with a starting point for investigating this and related questions, e.g., if people do use discourse operators, how could they learn them?, what sorts of discourse operators are possible?, and could the same discourse operators be used in interpretation as well as in generation? Thus, we are currently investigating how NL-Soar's learning mechanism can be used to acquire discourse operators.

## Dialogue Modeling in NL-Soar

NL-Soar is a computational model of real-time natural language generation and comprehension implemented within the Soar cognitive architecture. Soar is a candidate unified theory of human cognition which has been implemented as a computational system for modeling cognitive tasks.[1] NL-Soar has been implemented using the mechanisms provided by Soar, e.g., short-term and long-term memory, control structure, and basic learning mechanism. Our current research in NL-Soar is directed towards building a dialogue system which can participate in conversation in real-time within a task domain.[2]

[1]For more information about Soar, see (Newell 1990; Lehman, Laird, & Rosenbloom 1996) or *http://www.isi.edu/soar/soar.html* (the Soar home page).

[2]Previous versions of NL-Soar addressed real-time parsing (Lewis 1993) and sentence generation (Rubinoff &

In Soar, cognitive tasks are performed in a problem space by applying operators to states to create new states. The knowledge defining an operator (e.g., when to select or terminate it, and what actions it performs) is represented in long-term memory by means of *productions*, similar to condition-action rules. The condition part of a production specifies a pattern to be matched against the current contents of working (short-term) memory, while the action part specifies a set of actions, such as updates to working memory. Soar's basic control structure is to repeatedly attempt to select an operator and apply it, based upon what productions have fired in response to the current contents of working memory. Note that operator execution time, in Soar's theory of cognition, can be mapped to a constant amount of cognitive time. Thus, it is possible to verify if a Soar model's temporal performance is consistent with data about human performance.

In addition to applying operators, the Soar architecture provides a basic learning mechanism which can be used to acquire new operators. (The learned operators are also represented as productions stored in long-term memory.) We have designed a method whereby discourse operators are acquired by NL-Soar during discourse generation, and are used in both discourse generation and interpretation. (For a detailed description of our approach, see (Green & Lehman 1996).) For example, in generating a speaker's turn over a radio channel, NL-Soar may decide to begin by making a self-introduction using the speaker's radio call sign (e.g., *Dakota 204*) since the addressee might otherwise not be able to identify who is speaking. The result of this planning process automatically results in a learned discourse operator. In subsequent situations with similar features, NL-Soar would be able to generate a self-introduction using the learned operator rather than by repeating the planning process. Also, the same operator can be used during interpretation in similar discourse contexts to recognize the speaker's intention in uttering his radio call sign.

Thus, our current work in investigating the acquisition and use of discourse operators serves as a step towards the goals of processing novel CI in a cognitively plausible computational model. Furthermore, it may be possible to use NL-Soar to model how certain CI become associated with certain linguistic forms. NL-Soar learns syntactic and semantic operators while parsing novel inputs. These operators enable NL-Soar, on encountering similar expressions, to comprehend them much more efficiently using the results of its past experience instead of repeating the problem-solving process. One way of modeling how certain CI become associated with certain linguistic forms would be by use of learned operators combining syntactic, semantic, and pragmatic processing. NL-Soar's bias to use these learned operators before attempting problem-solving from first principles would account for the difficulty of recovering the original meaning of such forms. As for the first two goals, i.e. the goals of identifying the knowledge and processes involved in use of CI, much work remains to be done. However, the results of that work could be incorporated into a model such as NL-Soar. Furthermore, NL-Soar provides a basis for addressing the last goal, by providing a model of natural language processing situated in a general model of human cognition.

Lehman 1994). Other areas of current research include using NL-Soar in modeling second language acquisition and simultaneous translation.

## References


Clark, H., and Marshall, C. 1981. Definite reference and mutual knowledge. In Joshi, A. K.; Webber, B.; and Sag, I., eds., *Elements of discourse understanding*. Cambridge: Cambridge University Press.

Green, N. L., and Lehman, J. F. 1996. An approach to compiling knowledge for dialogue generation and interpretation. unpublished.

Green, N. L. 1994. *A Computational Model for Generating and Interpreting Indirect Answers*. Ph.D. Dissertation, University of Delaware. Department of Computer and Information Sciences, Technical Report No. 95-05.

Grice, H. P. 1957. Meaning. *Philosophical Review* 56:377–388.

Grice, H. P. 1975. Logic and conversation. In Cole, P., and Morgan, J. L., eds., *Syntax and Semantics III: Speech Acts*. New York: Academic Press. 41–58.

Gunji, T. 1981. *Toward a Computational Theory of Pragmatics - Discourse, Presupposition, and Implicature*. Ph.D. Dissertation, The Ohio State University.

Hirschberg, J. B. 1985. *A Theory of Scalar Implicature*. Ph.D. Dissertation, University of Pennsylvania. Department of Computer and Information Science, MS-CIS-85-56, LINC LAB 21.

Lehman, J. F.; Laird, J. E.; and Rosenbloom, P. S. 1996. A gentle introduction to Soar, an architecture for human cognition. In Sternberg, S., and Scarborough, D., eds., *Invitation to Cognitive Science*, volume 4. MIT Press.

Lewis, R. L. 1993. *An Architecturally-based Theory of Human Sentence Comprehension*. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, Pennsylvania. Available from Computer Science Department as Technical Report CMU-CS-93-226.

Mann, W. C., and Thompson, S. A. 1987. Rhetorical structure theory: Toward a functional theory of text organization. *Text* 8(3):167–182.

McCafferty, A. S. 1987. *Reasoning about Implicature: a Plan-Based Approach*. Ph.D. Dissertation, University of Pittsburgh.

Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, Massachusetts: Harvard University Press.

Rubinoff, R., and Lehman, J. F. 1994. Real-time natural language generation in NL-Soar. In *Proceedings of the Seventh International Workshop on Natural Language Generation*, 199–206.

Thomason, R. H. 1990. Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In Cohen, P.; Morgan, J.; and Pollack, M., eds., *Intentions in Communication*. Cambridge, Massachusetts: MIT Press. 325–363.