

Hypothetically Speaking

Default Reasoning and Discourse-Structure

Andrew Csinger
csinger@cs.ubc.ca

David Poole*
poole@cs.ubc.ca

Department of Computer Science
University of British Columbia
Vancouver, Canada V6T 1Z2

Abstract

We consider the problem of how to represent domain knowledge for user-modelling applications. The problem yields insight into methodology for representing knowledge in hypothetical reasoning systems. In general, there are not enough constraints in a domain to uniquely determine the approach that a reasoning system should take in formalizing its characteristics, but domain-independent methodological principles can help guide the designer. We present the *shared information constraint* on the domain of inter-agent communication, which implies an interesting representational duality between speakers (designers of utterances) and hearers (recognizers of mental state) in a simple communication environment. This work provides insight into default reasoning applications development.

1 Introduction

The theory of non-monotonic reasoning is young (see, e.g., [Reiter, 1980]), and serious implementations are few. Not surprisingly, there is very little methodological background for would-be practitioners to draw upon.¹ There does not yet exist a handbook for the effective use of a default reasoning framework, and we have often wished for one in our efforts to construct applications. Such a handbook would include a chapter on good methodologies for building default reasoning applications, just as the informative and complete Pascal textbook of today includes a chapter on top-down programming.

We are trying to build a body of methodological principles which may be of use to future practitioners of default reasoning.

Although we illustrate some of the issues with an example that deals with knowledge representation problems arising in natural language understanding and user modelling, the focus of this paper is the development and application of (default) reasoning methodologies.

The other issues have been dealt with elsewhere (see e.g. [Joshi *et al.*, 1981] and [Csinger, 1991]).

We have considered some of the elements that might participate in a theory of communications, and have argued that these can be represented within a default reasoning framework [Csinger and Poole, 1989].² Through the following example, we examine some general design issues we have faced in implementing our theory.

Example 1— Gas or Cigarettes?: Someone stops you on the road and asks you for the location of the nearest gas station. What do you say or do? It depends, of course, on the reasons for which *you believe* you are being asked, and upon what *you happen to know* at the time.³ For instance, in addition to knowing the location of the gas station (“around the corner...”), you might know that it is closed. Under reasonable assumptions of cooperativeness [Grice, 1975], you might want to include this knowledge in your response (“...but it’s closed”) [Mercer and Rosenberg, 1984]. If, however, you think the question arises not from the need for gas but from the desire for cigarettes or cold Beer, you might want to explain that there is a 24 hour shop in another direction...

This example is intended to show that the mental states of interlocutors should participate in a theory of communication. It also provides a useful excursion into representation issues for hypothetical reasoning systems.

The rest of this paper is structured as follows: Section 2 is a brief review of some existing work on methodology, introducing the terms by which our method will be described. Section 3 considers the alternatives for formalizing our domain, in terms of a search for appropriate constraints on the communication process. We conclude

²That paper showed how to use a hypothetical reasoning system in a natural language environment to determine the *presuppositions* and *implicatures* of utterances as part of the process by which to derive the *beliefs* of the speaker.

³The example is presented from the point of view of a fictitious hearer, but similar considerations apply to the dual of the example, presented from the point of view of the speaker. Such a dual rendition might begin as follows: *You are out of cigarettes and don’t know what to do. You know that in this part of the world, gas stations generally sell cigarettes...*

*Scholar, Canadian Institute for Advanced Research.

¹But see [Etherington, 1983; Poole, 1989b; Poole, 1990; Kautz, 1987].

2 Default-Programming Methodology

To make the following presentation more precise, we use the simple hypothetical reasoning framework of Theorist [Poole *et al.*, 1987], where F is a set of closed formulae, called the “facts”, and H is a set of (possibly open) formulae called “possible hypotheses.” The following definitions are relevant:

Definition 1 (Scenario) A scenario is $F \cup D$ where D is a set of ground instances of elements of H such that $F \cup D$ is consistent.

Definition 2 (Explanation) If g is a closed formula, an explanation of g is a scenario that implies g . Such a g is referred to here as an explanandum (the plural being, of course, explananda).

Definition 3 (Extension) an extension is the set of logical consequences of a maximal (with respect to set inclusion) scenario.

There's more than one way to use a hypothetical reasoning formalism. You can use it at least for prediction and for abduction, often in a single domain or problem.

Theorist is particularly of interest because the same formal definition allows for both default and abductive reasoning [Poole, 1989a]. It is also implemented; the examples provided in this paper have been tested on a running version of the program.

These different uses of Theorist can be characterized along two dimensions:

- Status of Explananda, and
- Status of Assumptions

2.1 Status of Explananda

The first dimension concerns whether the explanandum is known or not. This distinction corresponds to a choice between the following:

Abduction: The system regards the explanandum (the observation of the world or the design objective) as given, and needs to find an explanation for it. The idea is to find assumptions that imply the goal. We consider all explanations of the goal as possible descriptions of the world.

Prediction: The system does not know if the explanandum is true, and the idea is to determine what can be predicted from the facts (the general knowledge and the observation or design objective).

The issue is whether the explanandum is known to be true or whether it is something that has to be determined. For instance, if a reasoning agent knows (or has as defaults) that $a \rightarrow b$ and that a , the agent can predict b from its knowledge. Another agent who also knows that $a \rightarrow b$, as well as b , might be able to assume in the absence of contradictory evidence that a . The first agent is using prediction, while the second agent is using abduction.

One interesting difference between abduction and prediction is in the relevance of counter-arguments. For instance, when predicting g , it is important to know if $\neg g$

		Explanandum	
		Known Abduction	Unknown Prediction
Who	Design User		
	Recognition Nature		

Table 1: Domain-Formulation

can also be explained. In abduction, however, an explanation of $\neg g$ is irrelevant [Poole, 1989a].

2.2 Status of Assumptions

Along the other dimension we can distinguish between the two types of tasks:

Design tasks [Finger and Genesereth, 1985] are those in which the system can choose any hypotheses it wants. For example, a system can choose the components of the design in order to fulfil its design objective, or choose utterances to make in order to achieve a discourse goal.

A consistency check is used to rule out impossible designs. All other sets of components that fulfil the goal are possible, and the system can choose the “best design” to suit its goal.

Design can be done abductively to try to hypothesize components in order to imply a design goal [Finger and Genesereth, 1985]. Alternatively, design can be done predictively to derive a design from goals and any hypotheses we care to choose.

Recognition tasks are those in which the underlying reality is unknown, and all we can do is to guess at it based on the observations we make about it. This definition includes diagnosis, scene recognition and plan recognition.

Recognition can also be performed abductively or predictively [Poole, 1989b]. In an abductive framework, each explanation is a possible description of the world, while the disjunction of all explanations is the description of the world. In the predictive framework, an appealing strategy is to predict something only if it is explained from the observations even when an adversary chooses the hypotheses [Poole, 1989a], which corresponds to membership in all extensions.

This distinction turns on whether the system is free to choose any hypothesis that it wants or whether it must try to “guess” some hypothesis that “nature” or an adversary has already chosen.

For example, agents who know only that there is to be a meeting on the hour, sometime between 09:00 and 12:30, but not at 10:00, are able to make only disjunctive statements about when it will take place; they *recognize* that the meeting is at 09:00 or 11:00 or 12:00. In contrast, the agent organizing (*designing*) the meeting is free to pick any time that is consistent with its own knowledge. The planning agent is free to choose that the meeting will take place at 11:00.

Referring to our gas-station example, the Speaker is free to design any utterance she thinks will satisfy her goal. The Hearer, on the other hand, may be able to derive only that the Speaker wants gas or cigarettes or beer. The Speaker hopefully knows what she means; the Hearer is never certain.

Note that these frameworks are different ways to use the same formal system for different purposes. In order to use the system we have to choose one way to implement our domain.

In general, there are not enough constraints in a domain to uniquely determine the approach that the reasoning system should take in formalizing its characteristics [Poole, 1989b]. The ‘causality’ in the domain does not uniquely constrain its default-reasoning axiomatization.

3 Domain Formulation

We now turn to formulating the simple gas station domain within the default reasoning framework. The problem of finding the right constraints on the domain breaks down into the problem of where to place the interlocutors of the Speaker–Hearer pair on the domain-formulation grid of Table 1.⁴

As we have presented the domain, there are essentially two kinds of information available to, and distributed between, the interlocutor pair: the Speaker *knows* mental-state⁵ and *designs* utterances, while the Hearer *knows* utterances and *recognizes* mental-state. These aspects of the domain allow us to conclude that it is the Speaker-agent that occupies the first row of the domain formulation grid, and that the Hearer-agent will occupy the second. For convenience, we have labelled the agents in Table 2 with the coordinates of the box they occupy in it.

3.1 Axiomatizing the Domain

Let us consider each of the four blocks in the Domain-Formulation Grid, beginning with the top-left corner, which represents a speaker performing abduction for design of the utterance.

Speaker₁₁ When the Speaker is doing design by abduction, she seeks to find explanations of her goals in terms of suitable utterances. These considerations force an axiomatization of the domain according to the schema: *utterance*⇒*mental-state*.

Speaker₁₂ When the Speaker is doing design by prediction, she attempts to derive utterances from knowledge of her goals, which enforces an axiomatization that adheres to: *mental-state*⇒*utterance*.

Hearer₂₁ When the Hearer is doing recognition by abduction, he makes assumptions about the mental

⁴The grid merely summarizes some of the different possible uses of the hypothetical reasoning formalism; different problems/domains fall into different boxes, corresponding to different uses of the reasoning system.

⁵We do not attach ontological significance to the use of terms like ‘belief’ and ‘goal’; they are merely components of the [mental] states of the agents participating in a discourse. These agents may or may not be human, and the anthropomorphism is purely illustrative.

		Explanandum (x)	
		Known Abduction	Unknown Prediction
Who	Design User	I. Speaker ₁₁ <i>x</i> = <i>bel_s</i>	II. Speaker ₁₂ <i>x</i> = <i>utt_s</i>
	Recognition Nature	II. Hearer ₂₁ <i>x</i> = <i>utt_s</i>	I. Hearer ₂₂ <i>x</i> = <i>bel_s</i>

Table 2: Communication Domain Formulation

state of the speaker to explain what he knows (the utterance). Such a process is supported by an axiomatization of the form: *mental-state*⇒*utterance*.
Hearer₁₂ When the Hearer is doing recognition by prediction, he reasons from what he knows (the utterance) to find possible speaker mental states that might be responsible. This process requires an axiomatization of the form: *utterance*⇒*mental-state*.

The domain can be implemented in at least four different ways, corresponding to the four different possible combinations of Speaker and Hearer, as represented in the domain formulation grid.

3.2 The Shared-Information Constraint

When an agent alternates between speaking and hearing, we want to support both of these roles with a single, minimal knowledge base. Note that this is not possible if we insist that a single agent use only one reasoning strategy: if we implement an agent with only abductive reasoning capabilities, we force different representations upon the knowledge required for design and recognition (as described in Section 3.1).

Regardless of which view of the communication domain ‘causality’ we adopt as the developer of a hypothetical reasoning system (*i.e.*, whether we adopt the view that mental state influences utterance content — a speaker’s view, or that utterances have perlocutionary effects upon mental state — a hearer’s view), the domain-formulation grid encourages an interesting duality: for the Speaker and Hearer to use the same representation, they must use different reasoning strategies. *I.e.*, one agent will be an abductive one, and the other predictive.

Of the four possible implementations, only two permit the sharing of information in the aforementioned, desired fashion. These are I) Speaker₁₁ with Hearer₂₂, and II) Speaker₁₂ with Hearer₂₁. These two combinations of Speaker and Hearer are referred to from now on as Case I and Case II, respectively.

3.3 The Two Alternative Strategies

Accepting the argument for mutually represented information, and adopting the *Shared Information Constraint* in the formulation of the domain, there remain two alternatives: any domain is likely to admit of this kind of ‘vagueness’, which is not unlike the problem of choosing an algorithm in a conventional programming language.

In both Case I and Case II, some beliefs are assumable by both the Speaker and the Hearer agents. These

have been provided to Theorist by setting H , the set of hypotheses, to the following:

$$H = \{ \text{hm, go, kw, atk, dkw, gg} \}$$

The following interpretations are intended⁶:

- hm (hearer believes) speaker believes speaker has money,
- go (hearer believes) speaker believes the gas station is open,
- kw (hearer believes) speaker believes hearer knows where the gas station is located,
- atk (hearer believes) speaker is asking to know (relevance),
- dkw (hearer believes) speaker doesn't know location of station,
- gg (hearer believes) speaker wants to know where the gas station is in order to get gas.
- goal(get-gas) (hearer believes) it is a goal of speaker to go and get gas,
- subgoal(know-where) (hearer believes) it is a subgoal⁷ of speaker to know where the nearest gas station is,
- subgoal(go-to-gas) (hearer believes) it is a subgoal of speaker to go (now) to the nearest gas station,
- utt(ask-where-gas) speaker asks where the nearest gas station is.

3.3.1 Case I: Design by Abduction, Recognition by Prediction

Choosing the axiomatization of Case I means the Hearer agent uses prediction, while the Speaker agent uses abduction, and that the facts will be of the form: *utterance* $\Rightarrow$ *mental-state*. For example:⁸

$$F_I = \{ \text{go} \wedge \text{hm} \wedge \text{gg} \wedge \text{subgoal}(\text{know-where}) \Rightarrow \text{goal}(\text{get-gas}),$$

$$\text{atk} \wedge \text{kw} \wedge \text{utt}(\text{ask-where-gas}) \Rightarrow \text{subgoal}(\text{know-where}) \}$$

The Speaker's View: In adopting the abductive approach for the Speaker, we add to H_I the predicate $\text{utt}(X)$, for some X , indicating that the Speaker is free to use any utterance available to her in achieving her goal. We engage the abductive mechanism of Theorist by asking it to explain $\text{goal}(\text{get-gas})$, for which it provides the explanation:

$$D_{ask}^I = \{ \text{hm, go, kw, atk, gg, utt}(\text{ask-where-gas}) \}$$

⁶Inferences with these predicates actually take place within modal contexts, but such details would only obfuscate the message we are trying to get across in this paper. *hm*, for instance, should read for the speaker: $\text{bel}(s, \text{hm}(s))$, and for the hearer: $\text{bel}(h, \text{bel}(s, \text{hm}(s)))$.

⁷The distinction between *goals* and *subgoals* is purely syntactic. The subgoals are those goals of the speaker that are not appropriate for the hearer to refer to as *the goals* of the speaker.

⁸The hypothesis *gg* can be seen as the name of the default [Poole, 1989a].

In other words, Theorist has licensed, via the stated assumptions, the design of $\text{utt}(\text{ask-where-gas})$. This is to be interpreted as utterance of the question *Where is the (nearest) gas station?*

The Hearer's View: In adopting the predictive approach for the Hearer, we augment F_I with the predicate $\text{utt}(\text{ask-where-gas})$ and ask Theorist to predict $\text{goal}(X)$, for some X ; in other words, we ask for the goal of the Speaker given the utterance. There is a single extension containing:

$$D_{gas}^I = \{ \text{hm, go, gg, kw, atk, goal}(\text{get-gas}) \}$$

From this extension we predict $\text{goal}(\text{get-gas})$, i.e., that it is a goal of the Speaker to get gas.

3.3.2 Case II: Design by Prediction, Recognition by Abduction

Choosing the axiomatization of Case II means the Hearer agent uses abduction, while the Speaker agent uses prediction, and that the facts will be represented in the form: *mental-state* $\Rightarrow$ *utterance*.

$$F_{II} = \{ \text{goal}(\text{get-gas}) \wedge \text{hm} \wedge \text{go} \Rightarrow \text{subgoal}(\text{go-to-gas}),$$

$$\text{dkw} \wedge \text{subgoal}(\text{go-to-gas}) \Rightarrow \text{subgoal}(\text{know-where}),$$

$$\text{kw} \wedge \text{atk} \wedge \text{subgoal}(\text{know-where}) \Rightarrow \text{utt}(\text{ask-where-gas}) \}$$

The Speaker's View: To use Theorist to predict the Speaker's discourse goal, we add to F_{II} what the Speaker knows: $\text{goal}(\text{get-gas})$, and we instruct Theorist to explain $\text{utt}(X)$ for some X ; i.e., we see what utterance is in an extension given the goal. There is one extension containing the following explanation for $\text{utt}(\text{ask-where-gas})$:

$$D_{ask}^{II} = \{ \text{hm, go, kw, dkw, atk} \}$$

So the agent should ask where the gas station is.

The Hearer's View: The Hearer needs to arrive at an explanation of his observation of the Speaker's utterance. To use Theorist in this way, we add to H_{II} a predicate licensing assumptions about the Speaker's mental state: $\text{goal}(X)$. We then ask the system to explain the observation $\text{utt}(\text{ask-where-gas})$. Theorist provides the following explanation:

$$D = \{ \text{hm, go, kw, dkw, atk, goal}(\text{get-gas}) \}$$

3.4 Multiple Extensions

So far, we have ignored the myriad other reasons why the speaker agent of our example might have uttered her query; there are many reasons to ask for the location of the nearest gas station. There might also be other ways to find gas stations besides asking potentially unreliable strangers. To complete the picture, we add to the representations new facts and defaults to express the ideas that another reason for going to gas stations is to buy cigarettes, and that just driving around is also a way to find gas stations.

H is first of all augmented with the following:

{ gc, dtk, gggc }

The following interpretations are intended:

- gc (hearer believes) speaker wants to know where the gas station is in order to get cigarettes,
- dtk (hearer believes) speaker is going to drive around in order to find out where a gas station is,
- gggc (hearer believes) speaker wants to go to the gas station in order to get cigarettes (as opposed to going to the corner store, for example),
- do(drive) agent (who would otherwise be the speaker) drives car around,
- goal(get-cigs) (hearer believes) it is a goal of speaker to go and get cigarettes.

3.4.1 Case I Revisited

To F_I we add:

{ subgoal(know-where) $\wedge$ go $\wedge$ gc $\wedge$ hm
 $\Rightarrow$ goal(get-cigs),
 dtk $\wedge$ do(drive)
 $\Rightarrow$ subgoal(know-where) }

Speaker: With the addition to H_I of $\text{utt}(X)$ and $\text{do}(X)$, there are two designs (two explanations) for either goal(get-cigs) or goal(get-gas). For example, goal(get-gas) has explanations:

$D_{ask}^I = \{ \text{hm, go, kw, atk, gg, utt(ask-where-gas)} \}$
 $D_{drive}^I = \{ \text{hm, go, dtk, gg, do(drive)} \}$

So the Speaker can ask where the gas station is or drive around to find one. Suppose that she asks.

Hearer: We once again augment the set of facts with what the Hearer knows, namely $\text{utt(ask-where-gas)}$,⁹ and ask Theorist to predict goal(X).

There are two extensions, corresponding to the following explanations, from which we predict goal(get-cigs) $\vee$ goal(get-gas).

$D_{gas}^I = \{ \text{hm, go, atk, kw, gg, goal(get-gas)} \}$
 $D_{cigs}^I = \{ \text{hm, go, atk, kw, gc, goal(get-cigs)} \}$

3.4.2 Case II Revisited

To F_{II} we add¹⁰:

{ goal(get-cigs) $\wedge$ go $\wedge$ hm $\wedge$ gggc
 $\Rightarrow$ subgoal(go-to-gas),
 dtk $\wedge$ subgoal(know-where)
 $\Rightarrow$ do(drive),
 atk $\wedge$ dtk $\Rightarrow$ false }

⁹For technical reasons we also need to add an integrity constraint $\text{gg} \wedge \text{gc} \Rightarrow \text{false}$. The recognizer does not want to assume both that the speaker is going for gas and that the speaker is going for cigarettes. This does not mean that the speaker is not going for both, but rather that conservatism will not allow the assumption of both. This is a general technique for making conservative predictions [Kautz, 1987].

¹⁰The integrity constraint (rule implying false) is again used here to make two extensions. It is up to the speaker to choose one of these two assumptions (i.e., whether to ask to know where the gas station is, or drive to know where it is).

Speaker: For either of goal(get-cigs) or goal(get-gas), added as facts, there are two extensions, corresponding to two designs. For example, for goal(get-gas), the two extensions correspond to:

$D_{ask}^{II} = \{ \text{hm, go, atk, kw, dkw, utt(ask-where-gas)} \}$
 $D_{drive}^{II} = \{ \text{hm, go, dtk, kw, dkw, do(drive)} \}$

Hearer: We once again augment the set of allowable assumptions with goal(X), and ask the system to explain $\text{utt(ask-where-gas)}$. Theorist provides the following explanations:

$D_{gas}^{II} = \{ \text{hm, go, atk, kw, dtk, goal(get-gas)} \}$
 $D_{cigs}^{II} = \{ \text{hm, go, atk, kw, dtk, gggc, goal(get-cigs)} \}$

3.5 Summary

Formalisms for hypothetical reasoning can be used abductively or predictively. Theorist is one such formalism; it has been applied in this section to an illustrative example in the domain of inter-agent communication. Agents in such domains rely heavily upon shared knowledge [Joshi *et al.*, 1981] [Brennan, 1991] which should be represented in a form that supports the inferences of both speakers (as utterance designers) and hearers (as goal or belief recognizers).¹¹ These considerations motivate the modelling of agents that are dual in the sense that to share a common representation of principles which govern communication, interlocuting agents can not reason with the same strategy. In the example presented, we showed two obvious ways in which the domain might be axiomatized, and then demonstrated this duality using Theorist; in each axiomatization, the shared-information constraint results in a Speaker-Hearer pair of homogeneous agents where each agent uses a different strategy (or, equivalently, a single agent uses a different strategy for design and for recognition, respectively).

4 Conclusions

Finding enough constraints in a domain to uniquely define its default axiomatization is not usually possible. Default implementations can be classified along (at least) two dimensions: the assumption and *explananda* status dimensions, which we have represented as the rows and columns of the domain-formulation grid shown in Tables 1 and 2. The domain formulation task can be superficially regarded as one of finding how the domain fits into the grid's representation framework. The example of the simple communication domain demonstrates this process, and is particularly illustrative for a number of reasons:

- it exhibits aspects along both dimensions of the space defined in Section 2,
- it is characteristic of the plurality of implementation strategies of which many domains will admit,

¹¹These considerations are not restricted to natural language environments. Another example of shared principles or convention is social law, highly pertinent, for instance, in the design of societies of interacting robots [Shoham, 1990], [Shoham and Tennenholtz, 1992]. The observations in this paper may be also applicable to such domains.

- it brings to light a general constraining principle, the *shared-information constraint*, whose consideration led to an implementation that allows the speaker and hearer agents to share an axiomatization of the domain,
- it highlights a duality between interlocuting agents; applying the shared-information constraint forces homogeneous speaker and hearer agents to use different reasoning strategies, necessarily placing these agents at different points along the *explananda* dimension defined in Section 2.

A general observation stemming from this work and which may lead to further research is preliminary evidence for the claim that extant logical databases designed for planning applications might be pressed into service as databases for recognition applications, and vice-versa. An important application area which springs to mind is on-line user help systems that model the user to anticipate errors as well as to provide pro-active advice [Thies and Berger, 1992]; such a system is both hearer and speaker, as described in this paper. An important question to be answered in this regard concerns the amount of attention that must be devoted during the design of the database to ensuring that this dual usage will be possible.¹²

The cognitive ramifications for designing societies of heterogeneous agents may be of interest as well. It is noteworthy that in the context of the work presented in this paper, there is no way to decide externally, whether an agent engaged in communication is employing a particular axiomatization (e.g., Case I or Case II, etc.), or using a particular inference strategy (i.e., abduction or prediction).

Acknowledgements

This research was supported by Canadian National Sciences and Engineering Research Council (NSERC) grant OGPOO44121, and by Project B5 of the Institute for Robotics and Intelligent Systems. Thanks to Michael Horsch and Verónica Becher for timely and considered criticism.

References

- [Brennan, 1991] Susan Brennan. Conversation with and through computers. *User Modeling and User-Adapted Interaction*, 1(1), 1991.
- [Csinger and Poole, 1989] Andrew Csinger and David Poole. From utterance to belief via presupposition: Default reasoning in user-modelling. In *Proceedings of the Conference for Knowledge Based Computing Systems, KBCS-89*, pages 408-419, Bombay, India, December 1989. Also appeared in Springer-Verlag *Lecture Notes in Artificial Intelligence*, number 444, pages 407-417.
- [Csinger, 1991] Andrew Csinger. Implementing a theory of communications in a default reasoning framework. TR 91-30, University of British Columbia, 1991.
- [Etherington, 1983] David W. Etherington. Formalizing non-monotonic reasoning systems. Technical Report 1, University of British Columbia, Vancouver, Canada, V6T 1W5, 1983.
- [Finger and Genesereth, 1985] J. J. Finger and M. R. Genesereth. Residue: A deductive approach to design synthesis. Technical Report STAN-CS-85-1035, Department of Computer Science, Stanford University, Stanford, Cal., 1985.
- [Grice, 1975] H.P. Grice. Logic and conversation. In P. Cole and J.L. Morgan, editors, *Syntax and Semantics: Speech Acts*, vol 3, pages 47-58. Academic Press, New York, 1975.
- [Joshi et al., 1981] Joshi, Webber, and Sag, editors. *Elements of Discourse Understanding*. Cambridge University Press, 1981.
- [Kautz, 1987] Henry A. Kautz. A formal theory of plan recognition. Technical Report 215, Dept. of Comp. Sci., University of Rochester, Rochester, NY 14627, 1987.
- [Mercer and Rosenberg, 1984] Robert Mercer and Richard Rosenberg. Generating corrective answers by computing presuppositions of answers, not of questions. In *Proceedings of the 1984 Conference*, pages 16-19, University of Western Ontario, London, Ontario, May 1984. Canadian Society for Computational Studies of Intelligence.
- [Poole et al., 1987] David Poole, Randy Goebel, and Romas Aleliunas. A logical framework for default reasoning. In *The Knowledge Frontier: Essays in the Representation of Knowledge*, pages 331-352. Springer Verlag, New York, NY, 1987.
- [Poole, 1989a] David Poole. Explanation and prediction: an architecture for default and abductive reasoning. *Computational Intelligence*, 5(2):97-110, 1989.
- [Poole, 1989b] David Poole. Normality and faults in logic-based diagnosis. In *IJCAI*, pages 1304-1310, Detroit, MI, August 1989.
- [Poole, 1990] David Poole. A methodology for using a default and abductive reasoning system. *International Journal of Intelligent Systems*, 5:521-548, 1990.
- [Reiter, 1980] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1,2):81-132, 1980.
- [Shoham and Tennenholtz, 1992] Yoav Shoham and Moshe Tennenholtz. On the synthesis of useful social laws for artificial agent societies. In *Proceedings NCAI-92*, 1992.
- [Shoham, 1990] Yoav Shoham. Agent-oriented programming. STAN-CS 1335-90, Stanford University, 1990.
- [Thies and Berger, 1992] Markus A. Thies and Frank Berger. Intelligent user support in graphical user interfaces: Plan-based graphical help in object-oriented user-interfaces. Technical report, Deutsches Forschungszentrum fuer Kuenstliche Intelligenz, Stuhlsatzenhausweg 3, D-6600 Saarbruecken 11, Germany, March 1992.

¹²Similar considerations motivate the concept of code reuse in software engineering.