

# REPAIRING CONVERSATIONAL MISUNDERSTANDINGS AND NON-UNDERSTANDINGS\*

Graeme Hirst, Susan McRoy,<sup>†</sup> Peter Heeman,<sup>‡</sup> Philip Edmonds, and Diane Horton

Department of Computer Science  
University of Toronto  
Toronto, Ontario M5S 1A4  
Canada

## ABSTRACT

Participants in a discourse sometimes fail to understand one another, but, when aware of the problem, collaborate upon or negotiate the meaning of a problematic utterance. To address nonunderstanding, we have developed two plan-based models of collaboration in identifying the correct referent of a description: one covers situations where both conversants know of the referent, and the other covers situations, such as direction-giving, where the recipient does not. In the models, conversants use the mechanisms of refashioning, suggestion, and elaboration, to collaboratively refine a referring expression until it is successful. To address misunderstanding, we have developed a model that combines intentional and social accounts of discourse to support the negotiation of meaning. The approach extends intentional accounts by using expectations deriving from social conventions in order to guide interpretation. Reflecting the inherent symmetry of the negotiation of meaning, all our models can act as both speaker and hearer, and can play both the role of the conversant who is not understood or misunderstood and the role of the conversant who fails to understand.

## 1 INTRODUCTION

A common attitude in artificial intelligence research is that the tasks that are so difficult for our computers to perform—often, seemingly impossible—are trivial for people. Language, for example, seems effortless for people, and yet after 40 years of research in computational linguistics, we are still far from a complete solution.

But in fact, language understanding is often difficult for people too. Much that is written or said is not understood. Sometimes this is due to inattentiveness or difficulty in hearing. But sometimes the causes are more fundamental. A scientific or technical paper that expresses complex ideas can be hard to understand simply because of the difficulty of the material. And even simple ideas can be hard to understand if they are poorly written or expressed; high-quality language generation is very difficult for people.

Nevertheless, people are, in general, quite successful in their use of language. That's because they have strategies

for coping with their linguistic limitations. If they can't understand what is being said to them, they seek clarification and try to work things out. And people are very good at noticing when a conversation has gone awry as a result of one party misunderstanding the other, figuring out the exact nature of the problem, and saying the right thing to recover from the error.

It would be unreasonable, then, to expect computers to always understand language perfectly. But what we should expect is that computers have the flexibility to deal with the consequences of this imperfection. Like people, they should have the ability to recognize problems in understanding and to correct them.

## 2 NOT UNDERSTANDING AND MISUNDERSTANDING

Participants in a dialogue bring to it different beliefs and goals. These differences can lead them to make different assumptions about one another's actions, construct different interpretations of discourse objects, or produce utterances that are either too specific or too vague for others to interpret as intended. This may lead to not understanding or to misunderstanding.

By *not understanding*, we mean a participant's failure to find any complete interpretation of an utterance. This could mean finding no interpretation at all for some or all of the utterance or it could mean finding more than one interpretation and not being able to choose between the alternatives. An important aspect of not understanding is that the participant is *aware* that it has happened.

By contrast, the participant is not aware, at least initially, when *misunderstanding* has occurred. In misunderstanding, the participant obtains an interpretation that she believes is complete and correct, but which is, however, not the one that the other speaker intended her to obtain.<sup>1</sup> It is possible that a misunderstanding will remain unnoticed in a conversation and the participants continue to talk at cross-purposes. Alternatively, the conversation might break down, leading one participant or the other to

---

\*This paper summarizes work reported in greater detail in references [15, 20, 21, 22, 27, 28].

<sup>†</sup>Susan McRoy is now at the Department of Electrical Engineering and Computer Science, University of Wisconsin–Milwaukee, Milwaukee, WI 53201, U.S.A.

<sup>‡</sup>Peter Heeman is now at the Department of Computer Science, University of Rochester, Rochester, NY 14627, U.S.A.

---

<sup>1</sup>Misunderstanding should not be confused with *misconception*. A misconception is an error in the prior knowledge of a participant. McCoy [26], Calistri-Yeh [5], Pollack [33, 34], and others have studied the problem of how one participant can, in conversation, determine the misconceptions of another in order to correct them.

determine that a misunderstanding has occurred.

It is thus useful to divide misunderstanding into two types: *self-misunderstandings* are those that are both made and detected by the same participant, and *other-misunderstandings* are those that are made by one participant but detected by another. Self-misunderstandings arise when a participant finds that he cannot incorporate an utterance into the discourse consistently, unless he interprets one of the other's earlier utterances differently. Other-misunderstandings occur when a participant recognizes that if one of his own acts had been interpreted differently, the other's utterance would have been the expected response to it. The participant might then attempt to change the other's interpretation. For example, he might restate his message, or explicitly tell the other that she has misunderstood; or he might do nothing (*cf* [4]), perhaps in order to avoid social awkwardness.

In our research, we have considered both misunderstanding and not understanding, trying a somewhat different approach to each.

### 3 NOT UNDERSTANDING A REFERENT

#### 3.1 Referring as collaboration

The linguistic task of referring to some object or idea can involve a *collaboration* between the speaker and the hearer. The speaker has the goal of having the hearer identify the object that the speaker has in mind. The speaker attempts to achieve this goal by constructing a description of the object that she thinks will enable the hearer to identify it. But since the speaker and the hearer will inevitably have somewhat different beliefs about the world, the hearer might not be able to identify the object from this description—that is, not understand the reference. Often, in such cases, the speaker and hearer will collaborate in making a new referring expression that accomplishes the reference.

In an important series of experiments, Clark and his colleagues—especially Wilkes-Gibbs—have shown that conversants will often engage in a kind of *negotiation* in order for one of them to understand a reference that the other wishes to make [6, 8]. In their fundamental experiment, Clark and Wilkes-Gibbs gave pairs of subjects each a copy of a set of hard-to-describe tangram figures. The subjects' task was to arrange their sets in the same order, and to do so by conversation alone; neither could see the other's set. The subjects were thus obliged to construct descriptions for each tangram that they hoped the other could interpret correctly; for example, *the one that looks like an angel with a stick*.

Clark and Wilkes-Gibbs found that typically the participant trying to describe a tangram pattern would present an initial referring expression. The other participant would then pass judgment on it, either *accepting* it, *rejecting* it, or *postponing* his decision. If it was rejected or the decision postponed, then one participant or the other would

*refashion* the referring expression. This would take the form of either *expanding* the expression by adding further qualifications or *replacing* the original expression with a new expression. The referring expression that resulted from this was then judged, and the process continued until the referring expression was acceptable enough to the participants for their current purposes. The final expression was contributed to the participants' common ground.

This excerpt from Clark and Wilkes-Gibbs's data illustrates rejection, replacement, and acceptance:

#### Example 1

- 1 **A:** Okay, and the next one is the person that looks like they're carrying something and it's sticking out to the left. It looks like a hat that's upside down.
- 2 **B:** The guy that's pointing to the left again?
- 3 **A:** Yeah, pointing to the left, that's it! [*laughs*]
- 4 **B:** Okay.

In this dialogue, B implicitly rejects A's initial presentation by replacing it with a new referring expression in line 2, *the guy that's pointing to the left again*. A then accepts the refashioned referring expression in line 3.

This kind of reference negotiation is not found only in laboratory settings. A particularly clear instance can be seen in the following example from the London-Lund *Corpus of English conversation* [42, S.2.4a:1–8], in which the conversants collaborate simultaneously on the phrases *that weird creature* and *over there*.

#### Example 2

- 1 **A:** What's that weird creature over there?
- 2 **B:** In the corner?
- 3 **A:** [*affirmative noise*]
- 4 **B:** It's just a fern plant.
- 5 **A:** No, the one to the left of it.
- 6 **B:** That's the television aerial. It pulls out.

#### 3.2 A model of collaboration on referring

Heeman and Hirst have presented a computational model of this kind of collaboration; it covers the agent who makes the initial referring expression, that is, the *initiator*, and the agent who is to understand the expression, the *recipient*. In this model, the initiator has the goal of referring to something, and constructs a plan, in the form of a sequence of surface speech actions, to try to achieve it, given a set of beliefs about what the recipient believes. The recipient, seeing only the surface speech actions, tries to infer the plan in order to understand the reference. Thus, referring expressions are represented by plan derivations, and an unsuccessful referring expression is an invalid plan in whose repair the agents collaborate. This collaboration takes place through the use of plans that judge and refashion the expression.

Because the model uses plans to represent referring expressions and for the agents' collaboration itself, two

copies of the model can converse with one another, each copy alternating between the roles of speaker and hearer.<sup>2</sup> Acting as a hearer, the system performs plan inference on each set of actions that it observes, and then updates the state of the collaboration. It then switches to the role of speaker in order to reply. As the new speaker, the system looks for a goal that it can adopt, and then constructs a plan to achieve it. Next, presupposing the other participant's acceptance of the plan, it updates the state of the collaboration. It repeats this until there are no more goals to adopt. The surface actions of the constructed plans form the response of the system. The system then switches back to the role of hearer and waits for a response from the other copy.

In the rest of this section, we give a brief overview of the model; details may be found in references [20, 21].

### 3.3 Plans for referring

We extend the earlier approaches of Cohen [9] and Appelt [2] in planning not only the occurrence of a referring expression but also, at the same level, its *content*. We use a surface speech action, *s-refer*, to express the speaker's intention to refer, and a surface speech action, *s-attrib*, for each attribute that the referent can be described in terms of. Constraints on these actions express the conditions under which they can be used; for instance, that it be mutually believed that the object has a certain attribute [7, 31, 30].

These speech actions are the building blocks that referring expressions are made from. Acting as the mortar are intermediate plans that encode the knowledge of how a description can allow a hearer to identify an object, and these ensure that the referring expression includes sufficient descriptors that the hearer can (in the speaker's opinion) identify the referent (*cf* Dale [12], Reiter [37]). The intermediate plans do this by having *mental actions* as steps in their decomposition. These mental actions determine which objects could be believed to be the referent of the referring expression. The mental actions are performed on the potential referents, and the constraints are evaluated, and so the referent can be determined in a manner analogous to constraint satisfaction.

Following Pollack [34], our plan inference process can infer plans in which, in the hearer's view, either a constraint does not hold or a mental action is not executable. In inferring a plan derivation, the system as hearer first finds the set of plan derivations that account for the primitive actions that were observed, reasoning about the other participant's (believed) beliefs. Second, it evaluates each of these derivations by attempting to find an instantiation for the variables such that all of the constraints hold and the mental actions are satisfiable with respect to the

hearer's beliefs about the speaker's beliefs. If the hearer is able to satisfy the constraints, then he will have understood the plan and be able to identify the referent, since a term corresponding to it will have been instantiated in the inferred plan. Otherwise, he has a constraint that is unsatisfiable, making this derivation ill-formed and the referent unresolvable. After all derivations have been evaluated, if there is just one valid plan then the hearer will believe that he has understood (and will have identified the referent). Otherwise, if there is one invalid derivation, the constraint or mental action that is the source of the invalidity is noted, and will be used by the hearer in his attempt to fix the non-understanding.<sup>3</sup>

### 3.4 Plans for collaborating

If the initial referring expression is not understood, then the conversants will collaborate in its repair. Plan repair techniques can be used to refashion the expression and the parts of the plan that are removed or added can be communicated to the other participant using discourse plans. Thus we model a collaborative dialogue in terms of the changes that are being made to the plan derivation.

The first step in repairing an ill-formed plan is to communicate the source of the error to the other participant. So, this goal is given to the plan constructor. If the failure is because the hearer can find no objects that he believes match the referring expression, then it is overconstrained and the goal is *reject-plan*. So if the referring expression were *the weird creature*, and the hearer couldn't identify anything that he thought weird, he might say *What weird thing?*, thus indicating that problems arose at the surface speech action corresponding to *weird*. If several objects seem to match, the referring expression is underconstrained, and the goal will be *postpone-plan*, which might be expressed as a tentatively spoken *Okay?*

In either case, the hearer or the other will *refashion* the expression in the context of the rejection or postponement. In keeping with Clark and Wilkes-Gibbs, we use two discourse plans for refashioning: *replace-plan* and *expand-plan*. The first is used to replace some of the actions in the referring expression plan with new ones (as in line 2 of Example 1 above), and the second is to add new actions. Replacements can be used if the referring expression either overconstrains or underconstrains the choice of referent, while the expansion can be used only in the latter case.

The decomposition of the refashioning plans encodes how a new referring expression can be constructed from the old one. This involves three tasks: first, a single candidate referent is chosen; second, the referring expression is refashioned; and third, this is communicated to the hearer by way of *s-actions*, which was already discussed.<sup>4</sup> The

<sup>2</sup>The model is implemented in Prolog. Input and output are in the form of surface-level descriptions of speech acts; for example, *s-accept* represents the acceptance of a referring expression, and might be realized in a complete natural language system as *okay*.

<sup>3</sup>We have not explored ambiguous situations, those in which more than one valid derivation remains, or, in the absence of validity, more than one invalid derivation.

<sup>4</sup>Another approach would have been to separate the communicative task

first step involves choosing a candidate. If the speaker of the refashioning is the person who initiated the referring expression, then this choice is obviously pre-determined. Otherwise, the speaker must choose a possible candidate. Goodman [18] has addressed this problem for the case of when the referring expression overconstrains the choice of referent. He uses heuristics to relax the constraints of the description and to pick one that *nearly* fits it. This problem is beyond the scope of this research, and so we simply choose one of the referents arbitrarily (but see Heeman [20] for how a simplified version of Goodman’s algorithm that only relaxes a single constraint can be incorporated into the planning paradigm).

The second step is to refashion the referring expression so that it identifies the candidate chosen in the first step. This is done by using plan repair techniques [19, 45, 46]. Our technique is to identify a part of the plan that includes the constraint in error, to construct a replacement for it, and then to substitute the replacement into the old plan. This substitution undoes any decisions that were in the removed part that affect other parts of the old derivation. This technique has been encoded into our refashioning plans, and so can be used for both constructing repairs and inferring how another agent has repaired a plan.

Once the refashioning plan is accepted, the common ground of the participants is updated with the new referring expression. So the effect of the refashioning plan is that the hearer will believe that the speaker wants the new referring expression plan to replace the current one. This is done regardless of whether the referring expression plan is in fact valid. If it is, however, valid, then the referent can be understood, prompting the hearer to adopt the goal to communicate this with the speaker, leading to the surface speech action of *accept-plan*. Otherwise, the process will repeat, but this time with the new referring expression.

Table 1 shows two copies of the system engaging in a (simplified!) version of Example 2.

## 4 COLLABORATION ON REFERRING TO OBJECTS THAT ARE NOT MUTUALLY KNOWN

### 4.1 Referring in direction giving

A crucial assumption of Clark and Wilkes-Gibbs’s work—and of Heeman and Hirst’s model—is that the hearer of the initial referring expression already has some knowledge of the referent in question. In Clark and Wilkes-Gibbs’s experiments, it is one of the given tangram figures; in Example 2, it is an object in the room that both speakers can see. However, a speaker sometimes has to refer to an object that is not previously known to the hearer. One particular situation in which this arises is in giving directions. For example, the speaker might give a direction like the following:

### Example 3

1 **A:** Go straight ahead until you get to *a funny-looking building*.

The recipient has to understand the reference well enough that when he later reaches the building, he will recognize it as the intended referent. Although this type of reference is different from the kind of referring action that Heeman and Hirst modelled, conversants can nevertheless collaborate to achieve an understanding of them. This can be seen in the following portion of a telephone conversation recorded by Psathas [36, p. 196].

### Example 4

- 1 **A:** Ya just stay on 2A, until ya get to Lowell Street.  
 2 **B:** Is it marked?  
 3 **A:** Yeah, I think there’s a street sign there, it’s an intersection with lights.  
 4 **B:** Okay.

In this dialogue, B has not understood the reference to the intersection at Lowell Street, and so suggests that the intersection might be marked. A replies with an elaboration of the initial expression.

Edmonds [15] has presented a computational model of this type of collaboration that draws from Heeman and Hirst’s model. The domain is that of giving directions for someone unfamiliar with an area to get to a particular place. In this section, we give an overview of Edmonds’s model.

The basis of the model is that the hearer can accept a referring expression plan if (1) the plan contains a description that is *useful* for making an *identification plan* that the hearer can execute to identify the referent, and (2) the hearer is *confident* that the identification plan is *adequate*.

The first condition, originally described by Appelt [1], is important because the success of the referring action depends on the hearer formulating a useful identification plan. We take the referring expression plan itself to be the identification plan. The mental actions in the intermediate plans encode only useful descriptions. For the second condition to hold, the hearer must believe that the identification plan is good enough to uniquely identify the referent when it becomes visible. This involves giving enough information by using the most visually prominent or *salient* attributes of the referent.

Each agent associates a numeric *confidence value* with each of the attributes in the referring expression, and by composing these,<sup>5</sup> computes a level of confidence in the adequacy of the complete referring expression plan that can be interpreted as ranging from low confidence to high confidence. If the overall confidence value exceeds some set value, the agent’s *confidence threshold*, then the agent

<sup>5</sup>The present composition function is simple addition. One could envision more complex systems to compute confidence such as an algebra of confidence or a non-numeric system.

Table 1: Example of referent negotiation.

---

1	<b>A:</b> See the weird creature. <code>s-refer(Entity)</code> <code>s-attrib(Entity, λX·assessment(X, weird))</code> <code>s-attrib(Entity, λX·category(X, creature))</code>
2	<b>B:</b> In the corner? <code>s-postpone(p1)</code> <code>s-actions(p1, [s-attrib(Entity1, λX·category(X, corner)),</code> <code>s-refer(Entity1),</code> <code>s-attrib-rel(Entity, Entity1, λX·λY·in(X, Y))])</code>
3	<b>A:</b> No, on the television. <code>s-reject(p56, [s-attrib-rel(Entity, Entity1, λX·λY·in(X, Y))])</code> <code>s-actions(p56, [s-attrib-rel(Entity, Entity2, λX·λY·on(X, Y)),</code> <code>s-refer(Entity2),</code> <code>s-attrib(Entity, λX·category(X, television))])</code>
4	<b>B:</b> Okay. <code>s-accept(p142)</code>

---

believes the plan is adequate. That is, if the agent is the initiator, she believes that the other will be able to understand the reference; if the agent is the other, he believes that he has understood the reference.

Now, the confidence value of each attribute is equivalent to its *salience* within the context of the referring expression. Salience, for our purposes in direction-giving, is primarily visual prominence, but can also involve identifiability, familiarity, and functional importance [14, 25]. One approach is to encode the salient properties in a static hierarchy as Davis [13], and Reiter and Dale [38] have done. But, ideally, salience should depend on the context surrounding the referent. For example, the height of a tall building would normally be salient, but not if it were surrounded by other tall buildings. This computation would be quite complex, so we have adopted a middle ground between the simple context-independent approaches, and a full-blown contextual analysis. The middle ground involves taking the type of object into account when choosing attributes and landmarks that relate to it. For example, height and architectural style can be very salient features for describing a building, but not for describing an intersection for which having a sign or traffic lights is important. This approach still allows us to encode salience in a hierarchy, but it is dependent on the referent.

Table 2 is an example of a simple salience hierarchy that an agent might have. The hierarchy is actually a set of partial orderings of attributes, represented by lambda expressions, indexed by object type. In the table, the confidence value of using architectural style to describe a building is 4. The confidence value of a tall building is 3, and so this attribute is less salient than architectural

style. The other rows (for describing intersections) follow similarly.

## 4.2 Construction and inference of referring plans

The salience hierarchy is used in both plan construction and plan inference.

In plan construction, salience is used for constructing initial referring expression plans, elaborating on inadequate plans, and for suggesting possible elaborations to plans by allowing an agent to choose the most salient properties of the referent first. The agent constructs an initial referring expression plan in almost the same way as in Heeman and Hirst’s system. Mental actions in the intermediate plans of a referring expression plan allow the speaker to choose the most salient attributes that have not yet been chosen, and constraints in the surface speech actions make sure the speaker believes that each attribute is true.<sup>6</sup> For example, to construct the reference to the *building* in Example 3, the speaker consulted her salience hierarchy (in table 2) and determined that architectural style is salient. Hence, she described the building as *funny-looking*. This single attribute was enough to exceed her confidence threshold.

During plan inference, the salience hierarchy is used when judging a recognized plan. Mental actions in the intermediate plans determine the confidence values of each attribute, and add them up. A final constraint in the plan makes sure the overall confidence sum exceeds the confidence threshold of the agent. This means that judging

<sup>6</sup>In Heeman and Hirst’s model, an attribute has to be mutually believed to be used. Here, mutual belief is not possible because the hearer has no knowledge of the referent, but mutual belief is an intended effect of using this plan.

Table 2: A salience hierarchy.

---

salient-attribute(4, building,	$\lambda X \cdot \text{architectural-style}(X, \text{Style})$ ).
salient-attribute(3, building,	$\lambda X \cdot \text{height}(X, \text{tall})$ ).
salient-attribute(3, intersection,	$\lambda X \cdot \text{called}(X, \text{Name})$ ).
salient-attribute(2, intersection, sign,	$\lambda X \cdot \lambda Y \cdot \text{has}(X, Y)$ ).
salient-attribute(2, intersection, traffic-lights,	$\lambda X \cdot \lambda Y \cdot \text{has}(X, Y)$ ).

---

the adequacy of a referring expression plan falls out of the regular plan evaluation process. For example, after recognizing the reference in Example 3, the hearer evaluates the plan. Assuming he believes the salience information in table 2, he computes the confidence value of 4. If this value exceeds his confidence threshold, then he will accept the plan. If not, he will believe that there is an error at the constraint that checks his confidence threshold.

### 4.3 Suggestion and elaboration

If the hearer is not confident in the adequacy of the plan, he uses an instance of `postpone-plan` to inform the initiator that he is not confident of its adequacy (which causes the initiator to raise her own confidence threshold). Now, although he cannot refashion the expression himself, he does have the ability to help the initiator by *suggesting* a good way to expand it; *suggestion* is a conversational move in which an agent suggests a new attribute that he deems would increase his confidence in the expression’s adequacy if the expression were expanded to include the attribute. Continuing with the example, if the hearer were not confident about the adequacy of *the funny-looking building*, he might suggest that the initiator use height (as well as architectural style), by asking *Is it tall?*. From this suggestion the initiator might expand her expression to *the tall funny-looking building*. So, in our sense, a suggestion is an illocutionary act of questioning; along with actually suggesting a way to expand a plan, the agent is asking whether or not the referent has the suggested attribute.

To decide what suggestion to make, the agent uses an instance of `suggest-expand-plan`, which has a mental action in its decomposition that chooses the most salient attribute that has not been used already.

However, only the initiator of the referring expression can actually *elaborate* a referring expression, because only she has the knowledge to do so. Depending on whether the hearer of the expression makes a suggestion or not, the initiator has two options when elaborating a plan. If no suggestion was made, then she can expand the plan according to her own beliefs about the referent’s attributes and their salience. On the other hand, if a suggestion was made, she could instead attempt to expand the plan by using the attribute suggested.

The decomposition of `expand-plan` calls the plan constructor with the goal of constructing a `modifiers`

schema, and with the suggested attribute as input. The plan constructor attempts to find a plan with the surface speech actions for the attribute in its yield, but this might not be possible. In any case, the speaker constructs an expansion that will make the plan adequate according to her beliefs.<sup>7</sup>

The response to a suggestion depends, obviously, on whether or not the suggestion was used to expand the plan. The speaker can (1) affirm that the plan was expanded with the suggestion by using the `s-affirm` speech act; (2) affirm that the suggestion was used, along with additional attributes that weren’t suggested, by using `s-affirm` and `s-actions`; or (3) deny the suggestion with `s-deny`, and inform the other by `s-actions` as to how the plan was expanded.

By using the expansion and suggestion moves, the two agents collaborate on refashioning the referring expression until the recipient of the directions is confident that it is adequate.

### 4.4 Example

We have implemented the model in Prolog. Table 3 shows two copies of the system engaging in a simplified version of Example 4.

## 5 MISUNDERSTANDING

### 5.1 Misunderstanding and repair

Participants in a conversation rely in part on their expectations to determine whether they have understood each other. If a participant doesn’t notice anything unusual, she may assume that the conversation is proceeding smoothly. But if she hears something that seems inconsistent with her expectations, she may hypothesize that there has been a misunderstanding, either by herself or the other, and produce a *repair*—an utterance that attempts to correct the problem.

One common type of repair involves correcting another speaker’s interpretation of the discourse. In the simplest case, a speaker makes an utterance displaying her misunderstanding in the turn immediately following the one she misunderstood. If the other speaker then recognizes the misunderstanding and initiates a sequence to resolve the

<sup>7</sup>Recall that she raised her confidence threshold as a result of the hearer’s postponement move, so now she must meet the new threshold.

Table 3: Example of suggestion and elaboration.

---

1	<b>A:</b> Go to the Lowell Street intersection. <code>s-goto(Entity)</code> <code>s-refer(Entity)</code> <code>s-attrib(Entity, λX.category(X, intersection))</code> <code>s-attrib(Entity, λX.called(X, 'Lowell Street'))</code>
2	<b>B:</b> Does it have a sign? <code>s-postpone(p1)</code> <code>s-suggest(p1, [s-attrib-rel(Entity, Entity2, λX. λY.has(X, Y)),</code> <code style="padding-left: 20px;">s-refer(Entity2),</code> <code style="padding-left: 20px;">s-attrib(Entity2, λX.category(X, sign))])]</code>
3	<b>A:</b> Yes, it does, and it also has traffic lights. <code>s-affirm(p1, [s-attrib-rel(Entity, Entity2, λX. λY.has(X, Y)),</code> <code style="padding-left: 20px;">s-refer(Entity2),</code> <code style="padding-left: 20px;">s-attrib(Entity2, λX.category(X, sign))])]</code>  <code>s-actions(p1, [s-attrib-rel(Entity, Entity3, λX. λY.has(X, Y)),</code> <code style="padding-left: 20px;">s-refer(Entity3),</code> <code style="padding-left: 20px;">s-attrib(Entity3, λX.category(X, traffic-lights))])]</code>
4	<b>B:</b> Okay. <code>s-accept(p123)</code>

---

misunderstanding, this is a *third-turn* (or *third-position*) repair, so called because the repair is initiated in the third turn of the top-level sequence, counting from the misunderstood utterance. Consider Example 5 from Coulthard and Brazil [11]. In this example, B has responded to line 1 with an *acknowledgement*, interpreting line 1 as an *inform*.

**Example 5**

- 1 **A:** So the meeting's on Friday.
- 2 **B:** Thanks.
- 3 **A:** No, I'm *asking* you.

However, A intended line 1 to be *yes-no* question (presumably with an *inform* as the expected reply). Recognizing B's misunderstanding, A produces a third-turn repair in line 3, telling B what action A had intended in line 1. A could have also told B the intended goal (e.g., "No, I want you to tell me.")

Another type of repair involves producing a new reply to a turn that one has apparently misunderstood. If a conversant hears an utterance that seems inconsistent with her expectations (perhaps because she has misunderstood some previous utterance) and the inconsistency leads her to reinterpret an earlier utterance and produce a new response to it, this is a *fourth-turn* (or *fourth-position*) repair [39]. Such repairs not only display the alternative interpretations, but also indicate some of the information that may underlie a participant's decision to favor one of them over another. Consider the fragment of conversation between a mother and her child (named Russ), shown in Example 6 [43].

**Example 6**

- 1 **Mother:** Do you know who's going to that meeting?
- 2 **Russ:** Who?
- 3 **Mother:** I don't know.
- 4 **Russ:** Oh. Probably Mrs McOwen and probably Mrs Cadry and some of the teachers.

In this dialogue, Russ initially interprets line 1 as expressing Mother's desire to tell, that is, as a *pretelling* or *preannouncement*, but finds this interpretation inconsistent with her next utterance. In line 3, instead of telling him who's going (as one would expect after pretelling), Mother claims that she does not know (and therefore could not tell). Russ recovers by reinterpreting line 1 as an indirect request, which his line 4 attempts to satisfy. This example also demonstrates agents' reluctance to repair the problems in the utterances of others [41]; although Mother might have produced a third-turn repair at line 3, the manifestation of a misunderstanding provided her with an expectable option that allowed her to avoid having to produce an explicit repair.

**5.2 The need for both intentional and social information**

Any dialogue system must account for the detection and repair of misunderstandings. But to consider possible misunderstandings in addition to intended interpretations would explode the number of alternatives that an interpreter would need to consider, unless there were adequate

constraints. However, predominant computational approaches to dialogue, which are based on intention, already have difficulty constraining the interpretation process. Proposed sociological accounts are more constrained, but none are computational. Some synthesis of intentional and social accounts of discourse is required.

In intentional accounts, speakers use their beliefs and goals to decide what to say; when hearers interpret an utterance, they try to identify goals that might account for it. This sort of reasoning is difficult to constrain because, although beliefs can narrow the search for an interpretation, there is no principled way of constraining the depth of the search. For each motivation that a hearer considers, he must also consider any higher-level motivations that it might support. To make such an approach workable, many simplifying assumptions have to be made, including the assumption that previous parts of the conversation have been understood correctly. However, there is another way to address misunderstanding that avoids this unconstrained inference of goals: use expectations deriving from social conventions (rather than intention) to guide interpretation.

In sociological accounts provided by Ethnomethodology, both coherent discourse interactions and repairs of misunderstandings are normal activities guided by social conventions [17, 40]. There are conventions regarding the expected range of responses to every action, for example. People then can assume that others are behaving as expected, unless they have reason to believe otherwise. In this way, the conventions give speakers a guide to possible interpretations. Reasoning is also limited, because conventions do not depend on the psychological characteristics of particular participants. What these accounts lack that computational accounts provide is an explanation of how people can identify the convention that is relevant, especially when there is no pre-existing expectation.

### 5.3 A synthesis

In our work (described more fully in references [27, 28]) we have developed a model of interaction that addresses the possibility that the participants might differ about the speech act that is performed by some utterance, without requiring extended reasoning about the speaker's goals. According to the model, speakers form expectations on the basis of what they hear, and thus monitor for differences in understanding. If necessary, they also reinterpret utterances in response to new information and generate repairs. Beliefs about the discourse context and conventions for interaction are used to select speech acts that are appropriate for accomplishing the speakers' goals. Interpretation and repair attempt to retrace this selection process abductively—when speakers attempt to interpret an observed utterance, they try to identify the goal, expectation, or misunderstanding that might have led the other agent to produce it.

The model uses both intentional and social sources of knowledge. Intentional information is captured by two re-

lations: one between utterances (input forms) and speech acts, and one between utterances and the attitudes that they express. These relations are the basis for deciding whether a set of utterances is consistent. To capture socially-derived expectations, the theory includes a relation on the speech acts—for each act, which acts are expected to follow. It also contains an axiomatization of speakers' knowledge for generating appropriate utterances and for detecting and repairing misunderstandings. The model demonstrates how these decisions depend on interactions among discourse participants' beliefs, intentions, previously expressed attitudes, and knowledge of social conventions.

The key features of the model that distinguish it from previous ones are the following:

- An account of the detection and repair of speech act misunderstandings and its relation to generation and interpretation. Although there has been work on identifying *potential* sources of misunderstanding, none of it addresses the problem of identifying and repairing *actual* misunderstandings. Also, unifying these tasks requires that linguistic knowledge and processing knowledge be kept distinct, improving the clarity of the model and permitting general knowledge about language to be reused.
- An integration of the socially-determined, structural conventions that have been identified by Ethnomethodology with the use of belief and intention that has been popularized within Artificial Intelligence. As a result, the model does not do extended inference about goals when it is not necessary.
- An account of the nonmonotonicity of discourse reasoning. In particular, the interpretation of utterances and the detection of misunderstandings are both characterized as abduction problems; speech act generation is characterized as default reasoning. As a result, all three processes can be specified within a single theory of communicative interaction.
- A reification of expectation. According to the model, agents form expectations on the basis of social conventions. They filter these expectations by considering the consistency of the Gricean intentions that they have expressed. By contrast, previous models of discourse attempt to eliminate interpretations by using some (necessarily incomplete) set of felicity conditions.
- An axiomatization in Prioritized Theorist [35]. Theorist is a declarative framework for default and abductive reasoning. Thus, linguistic knowledge and processing knowledge are kept distinct.

## 5.4 The architecture of our model

In the architecture that we have formulated, producing an utterance is a default, deductive process of choosing both a speech act that satisfies an agent’s communicative and interactional goals and a utterance that will be interpretable as this act in the current context. Utterance interpretation is the complementary (abductive) process of attributing communicative and interactional goals to the speaker by attributing to him or her a discourse-level form that provides a reasonable explanation for an observed utterance in the current context. Expectations deriving from social norms delimit the range of responses that can occur without additional explanation. The attitudes that speakers express provide additional constraints, because speakers are expected not to contradict themselves. We therefore attribute to each agent:

- A set  $\mathcal{B}$  of prior assumptions about the beliefs and goals expressed by the speakers (including assumptions about misunderstanding).
- A set  $\mathcal{M}$  of potential assumptions about misunderstandings and meta-planning decisions that agents can make to select among coherent alternatives.
- A theory  $\mathcal{T}$  describing his or her linguistic knowledge, including principles of interaction and facts relating linguistic acts.

**DEFINITION 1:** An *interpretation* of an utterance  $\mathbf{u}$  to hearer  $\mathbf{h}$  by speaker  $\mathbf{s}$  in discourse context  $\mathbf{ts}$  is a set  $M$  of instances of elements of  $\mathcal{M}$ , such that

1.  $\mathcal{T} \cup \mathcal{B} \cup M$  is consistent
2.  $\mathcal{T} \cup \mathcal{B} \cup M \models \text{utter}(\mathbf{s}, \mathbf{h}, \mathbf{u}, \mathbf{ts})$
3.  $\mathcal{T} \cup \mathcal{B} \cup M$  is not in conflict with any stronger defaults that might apply.<sup>8</sup>

**DEFINITION 2:** It would be *coherent* for  $\mathbf{s}$  to utter  $\mathbf{u}$  in discourse context  $\mathbf{ts}$  if  $\mathbf{u}$  is a solution to the following default reasoning problem:

$$\mathcal{T} \cup \mathcal{B} \cup M^{\text{meta}} \vdash (\exists u) \text{utter}(\mathbf{s}, \mathbf{h}, u, \mathbf{ts})$$

where  $M^{\text{meta}}$  is a set of assumptions about meta-planning decisions in  $\mathcal{M}$ , such that

1.  $\mathcal{T} \cup \mathcal{B} \cup M^{\text{meta}}$  is consistent
2.  $\mathcal{T} \cup \mathcal{B} \cup M^{\text{meta}} \models \text{utter}(\mathbf{s}, \mathbf{h}, \mathbf{u}, \mathbf{ts})$
3.  $\mathcal{T} \cup \mathcal{B} \cup M^{\text{meta}}$  is not in conflict with any stronger defaults that might apply.

**DEFINITION 3:** A speaker  $SI$  is *expected* to do action  $R$  in dialogue  $TS$  whenever there is an action  $A$  that is active in  $TS$  (because it was performed earlier),  $R$  as the normal expected reply to  $A$ , and the linguistic intentions of  $R$  are compatible with the active suppositions of  $TS$ .

<sup>8</sup>More precisely,  $\mathcal{T} \cup \mathcal{B} \cup M$  satisfies the priority constraints of Prioritized Theorist.

In addition, acts of interpretation and generation update the set of beliefs and goals assumed to be expressed during the discourse. The current formalization focuses on the problems of identifying how an utterance relates to a context and whether it has been understood. The update of expressed beliefs is handled in the implementation, but outside the formal language.<sup>9</sup>

The following formulates our characterization of misunderstanding and repair:

### Other-misunderstanding

Speaker  $s_1$  might be attempting action  $a_{\text{new}}$  in discourse  $ts$  if:

1. Earlier, speaker  $s_2$  performed act  $a_{\text{intended}}$ ;
2. Actions  $a_{\text{intended}}$  and  $a_{\text{similar}}$  can be performed using a similar surface form;
3. If  $s_2$  had performed  $a_{\text{similar}}$ , then  $a_{\text{new}}$  would be expected;
4.  $s_1$  may have mistaken  $a_{\text{intended}}$  for  $a_{\text{similar}}$ .

### Self-misunderstanding

Speaker  $s_1$  might be attempting action  $a_{\text{new}}$  in discourse  $ts$  if:

1.  $s_1$  has performed action  $a_{\text{observed}}$ ;
2. But, the linguistic intentions of  $a_{\text{new}}$  are inconsistent with the linguistic intentions of  $a_{\text{observed}}$ ;
3.  $a_{\text{observed}}$  and action  $a_{\text{intended}}$  can be performed using a similar surface-level speech act; and
4.  $s_2$  may have mistaken  $a_{\text{intended}}$  for  $a_{\text{observed}}$ .

### Third-turn repair

Speaker  $s_1$  should tell speaker  $s_2$  that she intended to perform  $a_{\text{intended}}$  in discourse  $ts$  if:

1.  $s_2$  has apparently mistaken an instance of act  $a_{\text{intended}}$  for act  $a_{\text{observed}}$ ; and
2.  $s_1$  may perform a third-turn repair (*i.e.*, the act that  $s_1$  would expect to follow her intended action has not already been done and it would be consistent for  $s_1$  to tell  $s_2$  that she intended to perform  $a_{\text{intended}}$ ).

### Fourth-turn repair

Speaker  $s_1$  should do action  $a_{\text{reply}}$  in discourse  $ts$  when:

1.  $s_1$  has mistaken an instance of act  $a_{\text{intended}}$  as an instance of act  $a_{\text{observed}}$ .
2. A reconstruction of the discourse is possible.
3.  $s_1$  would expect to do  $a_{\text{reply}}$  in this reconstruction.
4. And,  $s$  may perform a fourth-turn repair.

<sup>9</sup>A related concern is how an agent’s beliefs might change after an utterance has been understood as an act of a particular type. Although we have nothing new to add here, Perrault [32] shows how default logic might be used to address this problem.

## 5.5 Example

We have implemented the model in Prolog and the Theorist [35, 44] framework for abduction with prioritized defaults. Table 4 shows two copies of the system engaging in Example 6.

## 5.6 Other kinds of misunderstanding

Although the framework that we have developed is very general, we have developed within it the knowledge for dealing with only certain types of misunderstanding at the level of sentences and speech acts. Many other kinds of misunderstanding occur when the conversants differ in their beliefs as to what the plans of the other are, regarding both the domain of the discourse and its structure.

We have characterized the types of discourse-model discrepancy that can arise, and for most, have constructed or found natural instances of examples [22]. In example 7, in which the conversants are planning a party, B misunderstands A's ellipsis in line 7, (as *do you want to ask Karin to come to the party?* instead of *do you want to ask Karin for the recipe?*) and thinks that a topic shift has occurred.

### Example 7

- 1 A: I wonder if we've forgotten anyone.
- 2 B: Did you invite the Mooneys?
- 3 A: Mm hm.
- 4 B: Can you get me the cake recipe?
- 5 A: It's on the shelf above the stove.
- 6 B: It's not there.
- 7 A: Do you want to ask Karin?
- 8 B: Isn't she going to be out of town?
- 9 A: No, I mean ask her for the recipe.

To handle conversations, such as these, in which miscommunication occurs due to discourse-model discrepancies, a plan recognition scheme must be able to both detect these discrepancies and make any necessary revisions. Such abilities would also enable the modeling of miscommunication arising from discrepancies in the *domain* plans conversants attribute to each other. In the following exchange, a discrepancy arises concerning what domain plan A's question is intended to help her fulfill:

### Example 8

- 1 A: Where's CSC104 taught?
- 2 B: Sidney Smith Building, room 2118, but the class is full.
- 3 A: No, I teach it.

B assumes that A's domain plan is to take CSC104; in fact, her plan is to teach it.

Unfortunately, most existing plan recognition schemes cannot be employed as the foundation of a model of these sorts of miscommunication. Although techniques have been developed to handle situations in which agents' plan

libraries differ in some respects, *e.g.*, Pollack [34], Calistri-Yeh [5], most current schemes have no mechanism for detecting discrepancies in, or for revising, the plans inferred. Appelt and Pollack [3] suggested the use of weighted abduction to model the nonmonotonic aspects of plan inference. Although the weighted axioms they define provide limited coverage, the method itself is interesting. Eller and Carberry [16] proposed another mechanism that performs detection and revision, based on the insight that dialogues requiring plan inference revision are analogous to semantically or syntactically ill-formed input. However, their means of detecting a need to revise<sup>10</sup> is limited; it does not make use of clues in the conversation itself and revision initiated by such clues would appear to require another mechanism. Further, the meta-rules they propose for revision are procedural, and vary greatly in their specificity. A knowledge-level account of revision would be preferable.

But, more generally, the question arises as to the extent to which it is necessary to infer the other conversant's goals at all. Our abductive model synthesized intentional and social accounts. However, sociolinguists maintain that expectation and social accountability are sufficient by themselves. This question is important because, if it were possible to eliminate goal inference, then we would have succeeded in finding a principled way to greatly constrain the interpretation of utterances. If not, it would still be the case that social expectations could be used in constraining and directing the search within the goal inference procedures.

Our current model does not resolve the question of whether goal inference is ever required, because it does not cover non-literal utterances such as irony or sarcasm, nor partially expressed propositions such as ellipsis and fragmentary utterances. We are presently investigating the generality of our approach by extending the model to handle such cases, which many have claimed require extended goal inference.

## 6 CONCLUSION

Participants in a dialogue are necessarily limited in the amount of information that they can make explicit. Discourse participants compensate for this limitation by using the evidence provided by their utterances to verify each other's understanding of the conversation as it progresses. To show his understanding and acceptance of an utterance, a hearer may reply with an utterance that is consistent with the speaker's expectations. Alternatively, if he disagrees with the speaker's displayed interpretation, he can initiate a repair. In this way, participants negotiate the meaning of utterances. Reflecting the inherent symmetry of this negotiation, all our models can act as both speaker and hearer, and can play both the role of the conversant who

<sup>10</sup>Specifically, revision occurs when their plan inference scheme is unable to accommodate the next utterance without violating one of their constraints on well-formed dialogue.

Table 4: Example of repair of misunderstanding.

---

1	<b>‘Mother’:</b> Do you know who’s going to that meeting? surface-request(m, r, informif(r, m, knowref(r, whoIsGoing)))
2	<b>‘Russ’:</b> Who? surface-request(r, m, informref(m, r, whoIsGoing))
3	<b>‘Mother’:</b> I don’t know. surface-inform(m, r, not knowref(m, whoIsGoing))
4	<b>‘Russ’:</b> Oh. Probably Mrs McOwen and probably Mrs Cadry and some of the teachers. surface-informref(r, m, whoIsGoing)

---

is not understood or is misunderstood and the role of the conversant who fails to understand.

To address nonunderstanding, we have developed two plan-based models of collaboration in identifying the correct referent of a description: one covers situations where both conversants know of the referent, and the other covers situations, such as direction-giving, where the recipient does not. In the models, conversants use the mechanisms of refashioning, suggestion, and elaboration, to collaboratively refine a referring expression until it is successful.

To address misunderstanding, we have developed a model that combines intentional and social accounts of discourse to support the negotiation of meaning. The approach extends intentional accounts by using expectations deriving from social conventions in order to guide interpretation.

### ACKNOWLEDGEMENTS

Our work is supported by the Natural Sciences and Engineering Research Council of Canada and the University of Toronto. For discussions that influenced the course of this work, we are grateful to James Allen, Hector Levesque, Ray Reiter, Jan Wiebe, Mike Gruninger, Sheila McIlraith, Paul van Arragon, Linda Peto, Javier Pinto, Steven Shapiro, Stephen Green, and Suzanne Stevenson.

### REFERENCES

- [1] Appelt, Douglas E. “Some pragmatic issues in the planning of definite and indefinite noun phrases.” *Proceedings of the 23rd annual meeting of the Association for Computational Linguistics*, Chicago, 1985, 198–203.
- [2] Appelt, Douglas E. “Planning English referring expressions.” *Artificial Intelligence*, **26**(1), 1985, 1–33.
- [3] Appelt, Douglas E. and Pollack, Martha E. “Weighted abduction for plan ascription.” *User Modeling and User-Adapted Interaction*, **2**(1–2), 1992, 1–25.
- [4] Blum-Kulka, S. and Weizman, E. “The inevitability of misunderstandings: Discourse ambiguities.” *Text*, **8**(3), 1988, 219–241.
- [5] Calistri-Yeh, Randall J. “Utilizing user models to handle ambiguity and misconceptions in robust plan recognition.” *User Modelling and User-Adapted Interaction*, **1**(4), 1991, 289–322.
- [6] Clark, Herbert H. *Arenas of language use*. The University of Chicago Press, Chicago, and Center for the Study of Language and Information, Stanford, 1993.
- [7] Clark, Herbert H. and Marshall, Catherine R. “Definite reference and mutual knowledge.” In [23], 10–62 and [6], 9–59.
- [8] Clark, Herbert H. and Wilkes-Gibbs, Deanna. “Referring as a collaborative process.” *Cognition*, **22**, 1986, 1–39. Reprinted in [10], 463–493 and in [6], 107–143.
- [9] Cohen, Philip R. 1981. “The need for referent identification as a planned action.” *Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI-81)*, Vancouver, 1981, 31–36.
- [10] Cohen, Philip R.; Morgan, Jerry; and Pollack, Martha E. (editors). *Intentions in communication*. Cambridge, MA, The MIT Press, 1990.
- [11] Coulthard, Malcolm and Brazil, David. “Exchange structure.” In Coulthard, Malcolm and Montgomery, Martin (editors), *Studies in discourse analysis*. Routledge & Kegan Paul, London, 1984, 82–106.
- [12] Dale, Robert. “Cooking up referring expressions.” *Proceedings of the 27th annual meeting of the Association for Computational Linguistics*, Vancouver, 1989, 68–75.
- [13] Davis, J. R. *Back Seat Driver: Voice assisted automobile navigation*. PhD thesis, Massachusetts Institute of Technology, 1989.
- [14] Devlin, A. S. “The ‘small town’ cognitive map: Adjusting to a new environment.” In Moore, G. and Golledge, R. (editors), *Environmental knowing: Theories, research and methods*. Dowden, Hutchinson and Ross, Stroudsburg, PA.
- [15] Edmonds, Philip Glenny. *A computational model of collaboration on reference in direction-giving dialogues*. MSc thesis, to be published as a technical report, Department of Computer Science, University of Toronto, September 1993.
- [16] Eller, Rhonda and Carberry, Sandra. “A meta-rule approach to flexible plan recognition in dialogue.” *User Modeling and User-Adapted Interaction*, **2**(1–2), 1992, 27–53.
- [17] Garfinkel, Harold. *Studies in Ethnomethodology*. Prentice Hall, Englewood Cliffs, NJ, 1967. (Reprinted by Polity Press, Cambridge, England, 1984.)
- [18] Goodman, Bradley A. “Repairing reference identification failures by relaxation.” *Proceedings of the 23rd annual meeting of the Association for Computational Linguistics*, Chicago, 1985, 204–217.
- [19] Hayes, Philip J. “A representation for robot plans.” *Proceedings of the 4th International Joint Conference on Artificial Intelligence (IJCAI-75)*, Tbilisi, 181–188.

- [20] Heeman, Peter A. *A computational model of collaboration on referring expressions*. MSc thesis, published as technical report CSRI-251, Department of Computer Science, University of Toronto.
- [21] Heeman, Peter A. and Hirst, Graeme. "Collaborating on referring expressions." MS, 1993, submitted for publication (available on request from the authors).
- [22] Horton, Diane and Hirst, Graeme. "Discrepancies in discourse models and miscommunication in conversation." In *Working notes of the AAAI fall symposium: Discourse structure in natural language understanding and generation*, 1991, 31–32.
- [23] Joshi, Aravind K.; Webber, Bonnie L., and Sag, Ivan (editors). *Elements of discourse understanding*. Cambridge University Press, Cambridge, 1981.
- [24] Lambert, Lynn and Carberry, Sandra. "A tripartite plan-based model for dialogue." *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Berkeley, 1991, 47–54.
- [25] Lynch, K. *The image of the city*. The MIT Press, Cambridge, MA.
- [26] McCoy, Kathleen Filliben. "Generating context-sensitive responses to object misconceptions." *Artificial Intelligence*, **41**(2), December 1989, 157–195.
- [27] McRoy, Susan W. *Abductive interpretation and reinterpretation of natural language utterances*. PhD thesis, published as technical report CSRI-288, Department of Computer Science, University of Toronto, 1993.
- [28] McRoy, Susan W. and Hirst, Graeme. "Abductive explanation of dialogue misunderstandings." *Proceedings of the 6th conference of the European chapter of the Association for Computational Linguistics*, Utrecht, The Netherlands, 1993, 277–286.
- [29] McRoy, Susan W. and Hirst, Graeme. "Misunderstanding and the negotiation of meaning." *Working notes of the AAAI fall symposium: Human-computer collaboration: Reconciling theory, synthesizing practice*, Raleigh, NC, November 1993.
- [30] Nadathur, Gopalan and Joshi, Aravind K. "Mutual beliefs in conversational systems: Their role in referring expressions." *Proceedings of the 8th International Joint Conference on Artificial Intelligence (IJCAI-83)*, Karlsruhe, 1983, 603–605.
- [31] Perrault, C. Raymond and Cohen, Philip R. "It's for your own good: A note on inaccurate reference." In [23], 217–230.
- [32] Perrault, C. Raymond. "An application of default logic to speech act theory." In [10], 161–186.
- [33] Pollack, Martha E. "A model of plan inference that distinguishes between the beliefs of actors and observers." *Proceedings of the 24th annual meeting of the Association for Computational Linguistics*, New York, 1986, 207–214.
- [34] Pollack, Martha E. "Plans as complex mental attitudes." In [10], 77–103.
- [35] Poole, David; Goebel, Randy; and Aleliunas, Romas. "Theorist: A logical reasoning system for defaults and diagnosis." In: Cercone, Nick and McCalla, Gordon, editors, *The knowledge frontier: Essays in the representation of knowledge*, Springer-Verlag, New York, 1987, 331–352.
- [36] Psathas, George. "The structure of direction-giving in interaction." In Boden, Deirdre and Zimmerman, Don H. (editors), *Talk and social structure: Studies in ethnomethodology and conversation analysis*. Polity Press, Cambridge, England, 1991, 195–216.
- [37] Reiter, Ehud. "The computational complexity of avoiding conversational implicature." *Proceedings of the 28th annual meeting of the Association for Computational Linguistics*, Pittsburgh, 1990, 97–104.
- [38] Reiter, Ehud and Dale, Robert. "A fast algorithm for the generation of referring expressions." *Proceedings of the 14th International Conference on Computational Linguistics (COLING-92)*, Nantes, 1992, 232–238.
- [39] Schegloff, Emanuel A. "Some sources of misunderstanding in talk-in-interaction." *Linguistics*, **25**, 1987, 201–218.
- [40] Schegloff, Emanuel A. "Repair after next turn: The last structurally provided defense of intersubjectivity in conversation." *American Journal of Sociology*, **97**(5), 1992, 1295–1345.
- [41] Schegloff, Emanuel A.; Jefferson, Gail; and Sacks, Harvey. "The preference for self-correction in the organization of repair in conversation." *Language*, **53**, 1977, 361–382.
- [42] Svartvik, Jan and Quirk, Randolph. *A corpus of English conversation*. Lund Studies in English 56. C.W.K. Gleerup, Lund, 1980.
- [43] Terasaki, A. "Pre-announcement sequences in conversation." Social Science Working Paper 99, School of Social Science, University of California, Irvine. 1976.
- [44] Van Arragon, Paul. *Nested default reasoning for user modeling*. PhD thesis, Department of Computer Science, University of Waterloo, Waterloo, Ontario, 1990. Published by the department as Research Report CS-90-25.
- [45] Wilensky, Robert. "A model for planning in complex situations." *Cognition and Brain Theory*, **4**, 1981.
- [46] Wilkens, David E. "Recovering from execution errors in SIPE." *Computational Intelligence*, **1**, 1985, 33–45.