Department of Computer Science,
University of Otago

Technical Report OUCS-2004-05

Multi-agent human-machine dialogue: issues in dialogue management and referring expression semantics

Authors:
A Knott, I Bayard and P Vlugter
Department of Computer Science, University of Otago

Status: This paper is an extended version of a paper to appear at the 8th Pacific Rim Conference on Artificial Intelligence (PRICAI 2004)

Department of Computer Science,
University of Otago, PO Box 56, Dunedin, Otago, New Zealand

http://www.cs.otago.ac.nz/trseries/

Multi-agent human-machine dialogue: issues in dialogue management and referring expression semantics

Alistair Knott, Ian Bayard and Peter Vlugter
Dept of Computer Science, University of Otago
{alik/ibayard/pvlugter}@cs.otago.ac.nz

July 16, 2004

Abstract

This paper describes a human-machine dialogue system which is configured to support dialogue between multiple speakers. The user is one speaker, and the system 'plays' a number of other speakers. We present a number of principles governing dialogue management in such cases, which relate to turn-taking and the identification of the addressees of utterances. We also describe how the syntactic and semantic treatment of first- and second-person personal pronouns, and of addressee terms, need to be extended to deal with the multi-speaker scenario. We conclude by giving some examples.

1 Introduction

Human-machine dialogue systems typically support dialogue between two agents: the human user is one agent, and the system plays the part of another agent. In this scenario, the user and the system take turns at being the speaker, and when one of them is the speaker, the other is the addressee (the agent being spoken to).

However, in real life dialogue, there are frequently more than two participants. Automated dialogue systems can be configured in various ways to operate in a multi-speaker scenario. Firstly, a system can simulate each dialogue participant as a separate autonomous agent (e.g. Padilha and Carletta, 2002; Taylor, 1994). Secondly, a system can play the part of a single agent in a context where there are several human speakers (e.g. Wang, 2002). Finally, the system could support a dialogue between a single human user and several agents, all of which are played by the system. In this case, the agents can either be genuinely autonomous, or they can act in the service of a shared plan, delivering lines given to them by a central controller.

To extend a dialogue system to deal with multi-speaker interactions, whichever of the above scenarios is envisaged, a number of things must be supplied. At the dialogue level, we need a theory of turn-taking, to decide when to make an utterance, and who the addressees of other speakers' utterances are. Secondly, at the level of sentence syntax and semantics, we need to pay special attention to certain constructions: those which are used to refer to dialogue participants (in particular, personal pronouns) and those which are used to control turn-taking (in particular, terms of address).

We have already built a two-speaker dialogue system, which incorporates full sentence parsing and generation using a declarative grammar, and a range of standard dialogue management techniques (see e.g. de Jager *et al*, 2002; Bayard *et al*, 2002). This paper describes how we are extending this system to a multi-speaker environment, focussing on the additional syntactic constructions and dialogue management principles which are required, and on the interactions between these. We begin in Section 2 by describing the core dialogue system. In Section 3, we present the extension to our dialogue management algorithm to deal with multi-speaker dialogue. In Sections 4 and 5, we describe our treatments of the semantics of personal pronouns and addressee terms, which are needed to implement the dialogue management algorithm. We conclude in Section 6 with some examples.

2 Te Kaitito: an English-Māori dialogue system

2.1 Architecture of the system

Our dialogue system, called Te Kaitito[1], supports bilingual human-machine dialogues in English and Māori. The user and the system alternate in generating contributions to a dialogue. When it is the user's turn to contribute, (s)he enters a sentence in English or Māori. The sentence is first parsed, using the LKB system (Copestake *et al.*, 2000), and a set of syntactic analyses is computed. Each analysis is associated with a semantic interpretation. One interpretation is then selected, using various principles (see e.g. Knott and Vlugter, 2003). The dialogue manager then decides what kind of dialogue act is being performed by the utterance, and responds accordingly. If it encounters a problem interpreting the utterance, it responds with a suitable clarification question.

2.2 Presuppositional DRT

When an incoming utterance is parsed, its semantic representation is derived. The grammar associates sentences with representations in the Minimal Recursion Semantics (MRS) formalism (Copestake *et al*, 1999). As a postprocessing stage, we convert these representations to a format called Discourse Representation Structures (DRSs; Kamp and Reyle, 1993), with some additional annotations to deal with presuppositions and dialogue issues.

A DRS is a structure with two fields, one for representing **discourse referents**, and one for representing **conditions** or predications over these referents. DRSs are typically drawn as split boxes, where referents appear at the top, and conditions below. For example, here is the DRS for the sentence *A cat walked*:

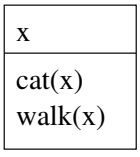

The discourse referent x is created by the indefinite NP *a cat*. This shows that *a cat* has introduced a new discourse referent. The conditions $cat(x)$ and $walk(x)$ were placed in the bottom part by the NP *a cat* and the VP *walked*.

The discourse context is also represented by a DRS structure. This DRS represents the **common ground** of the conversation, in other words, the material which the speakers consider to be common knowledge between them.

A sentence's **presuppositions** are elements of its content which the speaker assumes are already part of the common ground. They are constraints on the kinds of context in which the sentence can be uttered. Here are two examples.

(1) The dog chased a cat.

(2) John's cat slept.

Sentence 1 presupposes that there is a dog in the discourse context (or more precisely, that there is exactly one *salient* dog in the context). Sentence 2 presupposes that there is someone called John, and also that this person has a cat. Presuppositions are triggered by lexical items such as the definite article, proper names, and possessive forms. These triggers determine what is asserted information, and what is presupposed in a given sentence.

As already mentioned, we use a DRT-based treatment of presuppositions as proposed by van der Sandt 1992. A sentence is modelled as an **assertion DRS** and a set of **presupposition DRSs**. The DRSs for Examples 1 and 2 are shown in Figures 1 and 2. Notice that the presupposition DRSs are distinguished by dashed lines.

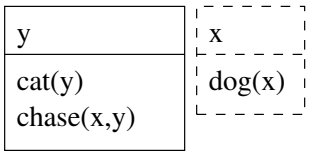

Figure 1: *The dog chased a cat*

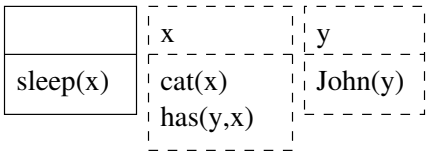

Figure 2: *John's cat slept*

The presuppositions of a sentence need to be **resolved** or **satisfied** in the current discourse

[1]Online demos of Te Kaitito can be found at `http://tutoko.otago.ac.nz:8080/teKaitito/` .

context before its assertional content can be processed. In van der Sandt's DRS-based treatment, this is modelled as a binding operation: the referents in each of the sentence's presupposition DRSs need to be bound to referents in the context DRS which have the properties identified in the presupposition DRS. Once this binding has been done, if the presuppositions of a sentence are not satisfied, referents with suitable properties can be (charitably) assumed to exist, and added to the context DRS, in an operation called **accommodation**.

2.3 Dialogue management concepts

In this section, we describe the main representations and techniques used by the core dialogue manager. These are widely used in managing a dialogue between two agents.

Firstly, it is useful to identify a number of different **dialogue acts**: assertions, questions, answers, acknowledgements and so on. We won't need to buy into any particular theory of these. However, we will distinguish between **forward-looking** and **backward-looking** dialogue acts (c.f. Allen and Core, 1997). A forward-looking act is basically a new initiative taken by a speaker; for example a new assertion made apropos of nothing, or a question about some new topic. A backward-looking act is one which is taken in response to an earlier dialogue act; for example the answer to a question, or the acknowledgement of an assertion. We will refer to a pairing of a forward-looking dialogue act and a backward-looking act as a **subdialogue**.

Another well-known idea in dialogue theory is that subdialogues can be nested inside one another (c.f. e.g. Grosz, 1978; Litman and Allen, 1990). Here is an example of a subdialogue involving a clarification question.

(3) Sid: can you pass the syringe?
Nancy: Which syringe?
Sid: The one on the left.
Nancy: Here you are.

The dialogue context can be thought of as containing a number of forward-looking acts which need to be responded to. These acts are sometimes referred to as the **dialogue stack** (after Grosz and Sidner, 1986), or as a set of **questions under discussion**, unordered except for one **maximal** question, which is what a backward-looking act is assumed to respond to (c.f. Ginzburg and Sag, 2000). In Example 3, after Sid's first utterance, there is one forward-looking act in the dialogue context, and after Nancy's first utterance, there are two, Nancy's being maximal. Sid's second utterance is therefore understood as responding to Nancy's question. After Sid's second utterance, there only remains one forward-looking act, and Nancy's second utterance is therefore assumed to respond to this.

In our system, we model the dialogue stack as a set of DRSs which are kept separate from the common ground.

3 Dialogue management in multi-agent dialogue

When we move from a two-participant to a multi-participant scenario, dialogue management becomes more complex. Principles are needed to determine who the next speaker should be, and when this is decided, who this speaker should address. A huge amount has been written about this topic from a descriptive point of view, beginning with a classic paper by Sacks *et al* (1974). We will focus tightly on what consequences this work should have on the design of a dialogue manager in our scenario, where the system is playing all the participants except one.

3.1 Deciding on the next speaker

Who the next speaker is allowed to be depends on what the previous dialogue act was. A simple rule is that if the previous utterance was a forward-looking dialogue act, the next speaker must be the addressee of that act. However, in nested subdialogues, this rule is not sufficiently general. For instance, in Example 3, Sid's second utterance is backward-looking; however, there is still a forward-looking act on the stack, which Nancy is bound to respond to. The general principle is as follows:

> **Principle 1:** If the dialogue act which is maximal on the stack is forward-looking, the next speaker must be the addressee of that act.

When there is no forward-looking act on the stack—in other words, when the stack is empty—we assume anyone can talk next. Much

has been written about how 'the floor is claimed' in such cases. In our system, we simply hand initiative to the user. The user can cede the floor, simply by pressing 'return' without entering a sentence. If this happens, the system currently selects one of the characters it plays at random to make a new initiative (either asking a new question, or making a new assertion).

3.2 Deciding on the addressee

The next issue is to determine the addressee of an utterance. If the speaker is played by the system, this issue is one of content selection, which relates to what the system wants to achieve. In our case, we keep things very simple, by always talking to the user.

> Principle 2: If the system must make an utterance, the addressee is always the user.

If the speaker is the user, then determining the addresee is an interpretation issue. The remainder of this section sets out some principles which govern how this works.

3.2.1 Structurally-defined addressees

In cases where there is a forward-looking dialogue act on the stack, we can use an analogue of Principle 1 to determine the addressee.

> **Principle 1a:** If a speaker utters a backward-looking dialogue act, the addressee of this utterance is the speaker of the forward-looking dialogue act which is maximal on the stack.

In such a case, the addressee can be identified structurally, without being mentioned explicitly.

3.2.2 Explicit identifications of the addressee

Other methods of indicating the addressee are more explicit. Firstly there are methods which are entirely nonverbal. For instance, when saying an utterance, the speaker can look at a particular person, or can even indicate one or more addressees by actually pointing at them. However, our system does not have the multimodal capabilities to simulate these nonverbal methods.[2] Alternatively, the speaker can be even more explicit, and identify the addressee or addressees linguistically, using what we will call an **addressee term**.

(4) How are you doing, Bill?

(5) Janet, pass me that wrench.

(6) Bill and Janet, you have not been listening to me.

(7) Hello everyone.

(8) Not tonight, Josephine.

These sentences perform different kinds of dialogue act: Examples 4–8 respectively express a question, a command, an assertion, a greeting and an answer. We believe that sentences expressing any type of dialogue act can be modified with an addressee term, though there are restrictions on what form these terms can take for different dialogue acts.

Addressee terms can be used in two dialogue contexts. Firstly, if the addressee is already specified structurally, an addressee term can be given which is consistent with this.

(9) Josephine [to Bert]: Shall we watch a video?
Bert: Not tonight, Josephine.[3] I have a headache.

If an addressee term is used which is inconsistent with the structurally specified addressee, we suggest the result is an incoherent dialogue.

(10) Josephine [to Bert]: Shall we watch a video?
Bert: # Not tonight, Frank. I have a headache.[4]

The second context in which addressee terms can be used is at a point when the speaker is making a forward-looking dialogue act; in other words, where s/he is taking some new initiative.

> **Principle 3:** If the speaker is making a forward-looking dialogue act, (s)he is free to choose any addressee or group of addressees.

[2] Our system does have a talking head interface with a limited ability to track the human user with a camera and identify them visually; see King *et al* (2003). However, the system can only track one user at a time.

[3] We believe that an explicit addressee term in such cases carries connotations of intimacy or of a heightened emotional connection. This seems a good example.

[4] At best, Bert's reply can be understood on the premise that Josephine was asking on Frank's behalf. We do not consider this kind of proxy dialogue move; however, see Section 4.2.1 for a related construction we do cover.

For instance, in Example 11, Bob's first utterance is backward-looking, and has to be understood as being addressed to Sue even though there is no addressee term. But Bob's second utterance is forward-looking; he is thus free to address it to anyone, provided he makes this explicit with an addressee term.

(11) Sue: Shall we go to the cinema tonight, Bob?
Bob: Good idea.
Bob: Do you want to come, Svetlana?

Note that Principle 3 as stated above should apply to forward-looking acts inside nested subdialogues; we expect that nested forward-looking acts should be addressable to any person. Indeed, nested forward-lookig acts do seem to have this property, as the following example shows.

(12) Sue: Shall we go to the cinema tonight, Bob?
Bob: Svetlana, do you want to come?
Svetlana: Good idea.
Bob: Okay.

Note that Svetlana's utterance closes her subdialogue with Bob, and thus that Bob's second utterance (a backward-looking act) has to be understood as being addressed to Sue, by Principle 1a.

3.2.3 Default addressees

One final way of specifying an addressee is by default.

> **Principle 4:** If a forward-looking act F is made without an explicit addressee term, then the addressee is taken to be the set of speakers involved in the previous subdialogue, minus the speaker of F.

Here is an example of Principle 4 in action:

(13) Sue [addressing Bob and Mary]: Shall we go to the cinema tonight?
Bob: Good idea.
Mary: Good idea.
Bob: What film do you want to see?

Sue's first statement, together with Bob and Mary's responses to it, consitute a subdialogue. Bob's second utterance (*What film do you want to see?*) is a forward-looking act. Since there is no explicit addressee term, we assume by Principle 4 that it is addressed to Sue and Mary. Note that the principle also covers the case where the speaker was not involved in the preceding subdialogue. Here is an example of this:

(14) Sue [addressing Bob]: Shall we go to the cinema tonight?
Bob: Good idea.
Mary: Hi there, what's up?

Mary in this case is interpreted as entering a dialogue whose participants are Sue and Bob. In this context, her utterance should be interpreted as addressed to both participants, unless she includes an explicit addressee modifier indicating otherwise.

4 Personal pronouns

Our system needs to be able to generate and interpret utterances in English and Māori. In a multi-speaker context, there are some syntactic constructions which we need to pay special attention to, namely personal pronouns and addressee terms. We consider personal pronouns in this section, and addressee terms in Section 5.

4.1 The syntax of personal pronouns

Personal pronouns are devices which allow a speaker to refer anaphorically to him/herself, to the addressee, and to third parties. In each case, there is provision for the object of reference to be a single person, or a plural entity. The English personal pronoun system is quite simple, comprising first, second and third person pronouns, each of which can be singular or plural.

In other languages, the pronoun system encodes a a richer set of possibilities. For instance, in Māori, the language we are particularly interested in, there is a distinction between singular, dual and plural which is orthogonal to that between first and second person. Moreover, for first person plural, there is a further distinction between 'inclusive' and 'exclusive' pronouns, depending on whether the addressee is included in the set of people including the speaker. First- and second-person Māori pronouns are shown in Table 1.

	First person		Second person
Singular	au/ahau "I/me"		koe "you (by yourself)"
Dual	Inclusive	Exclusive	
	tāua "you and I"	māua "him/her and I"	kōrua "you two"
Plural	Inclusive	Exclusive	
	tātou "us lot (including you)"	mātou "us lot (but not you)"	koutou "you lot"

Table 1: First-person and second-person pronouns in Māori

4.2 The semantics of personal pronouns

All pronouns introduce presuppositions about entities which are already in the discourse context. In the case of singular pronouns, the story is quite simple: first-person pronouns presuppose an object who is the speaker, and second-person pronouns presuppose an object who is the addressee. Consider the following dialogue:

(15) Bob: I love you.
Sue: I love you too.

Both of these sentences can be represented as the DRS in Figure 3. Clearly, in order to interpret such a DRS, the context needs to contain appropriate objects of which `addressee` and `speaker` are predicated. These predicates are unusual in two respects. Firstly, they need to change at each change of either speaker or addressee. For instance, when Bob's utterance in Example 15 is interpreted, the context DRS should be as in Figure 4(a), while when Sue's utterance is processed, it should be as in Figure 4(b). Secondly, there can only be one `speaker` and one `addressee` predicate at any time; we assume there is only one conversation going on.

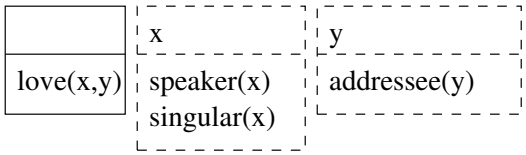

Figure 3: *I love you*

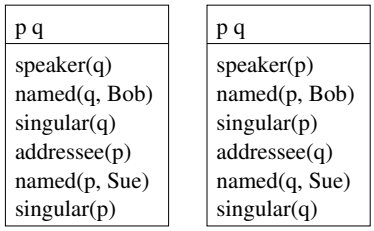

(a) (b)

Figure 4: Dialogue contexts for Example 15

4.2.1 Plural personal pronouns

To represent plural objects in general, we use Kamp and Reyle's (1993) method, in which each plural object is associated with a discourse referent, whose members can be identified with a series of `member` predicates, and about which other special predicates such as `plural-object` and `cardinality` can be asserted.

Plural personal pronouns denote—or more accurately presuppose—groups which include the addressee. For instance, *we* presupposes a group which includes the speaker, and plural *you* presupposes a group which includes the addressee. The group referred to by a plural pronoun must be made salient linguistically, for instance, by using a conjunction of noun phrases, as in Example 16:

(16) Dean: Pearl and I went to the cinema yesterday.
Hank: What did y'all watch?
Dean: We watched 'Casablanca'.

The first utterance here introduces a group entity composed of Pearl and Dean. In Hank's responding utterance, *y'all* presupposes a salient plural entity one of whose members is the addressee (Dean); the group of Pearl and Dean

satisfies this presupposition. In Dean's second utterance, *we* presupposes this same entity.

This general definition of personal pronouns subsumes an interesting case where the speaker or the addressee are *themselves* group entities. Consider this case:

(17) Bob: Sue and Mary, are you ready to go?

The group of people denoted by *you* in this example are not just being referred to, but being addressed; unlike the group denoted by *y'all* in Example 16, *you* in this example actually denotes the group of people who have to respond to the utterance. Note that the second person pronoun can still be defined as presupposing a salient group of people which includes the addressee, provided that (a) we are not talking about strict inclusion, and (b) we assume (reasonably) that the addresse of an utterance is always a salient entity in the discourse context.

According to our dialogue-managment principles, a forward-looking utterance with a group addressee must be responded to by that group. Is it possible for an utterance to have a group speaker? Genuine 'joint utterances' are of course virtually nonexistent (outside the theatre). Our approach is to allow the possibility of group speakers, with a very simple additional dialogue management principle:

> **Principle 5:** utterances made by a group speaker are actually made by an individual member of the group, on behalf of the whole group.

This approach is in fact in keeping with a general assumption in our system that all communication is public, and that there are no disagreements between participants. But in a more realistic situation, clearly much more needs to be said about how members of a group negotiate a response in such cases.

5 Addressee terms

Addressee terms function syntactically as sentence modifiers in English and Māori. To describe the syntax of these constructions, we need to specify firstly the **distribution** of addressee terms, i.e. whereabouts in a sentence they can appear, and secondly the **composition** of addressee modifiers, i.e what kinds of phrase are permissible addressee terms.

5.1 English addressee terms

5.1.1 Distribution

Adressee modifiers in English can appear wherever a sentential adverb can appear; we believe they have the same distribution as conjunctive expressions like *however* or *therefore*. For instance:

(18) John, that is out of the question.

(19) That, John, is out of the question.

(20) That is out of the question, John.

We suggest that commas are fairly mandatory in introducing addressee terms for nearly all dialogue act types. There are a few exceptions: postsentential modification of greetings, answers and acknowledgements seem fine without commas.

(21) Hi John.

(22) No John.

(23) Okay John.

What these seem to have in common is shortness and simplicity.

5.1.2 Composition

English addressee terms are basically like NPs, but with certain restrictions. Firstly, proper names can be addressee terms.

(24) Hit the road, Jack.

Secondly, common nouns can be addressee terms, but only in the context of a 'bare' (determinerless) NP.[5]

(25) Drink up, cats.

(26) Drink up, mate.

(27) ?? Drink up, the cats.

(28) ?? Drink up, a mate.

Notice that adjectives are allowed; so the bare NP is apparently an N-bar.

(29) Drink up, little cats.

[5]Note that many other languages are more permissive of determiners in addressee NPs; c.f. e.g. French *Salut, les gars* (*Hello, the lads*). But here again, not all determiners are allowed; e.g. **Salut, un gars* (*Hello, the lad*.

There is one determiner which can be used in an addressee term: a first-person singular possessive pronoun.

(30) Drink up, my darling.

The word *you* can also be used to introduce a common noun. It seems to behave rather like a determiner in this context.

(31) Get up, you lazybones.

Thirdly, note that ordinary pronouns cannot be used by themselves as addressee terms.

(32) *Drink up, him.

There is a possible exception for second person pronouns:

(33) Drink up, you.

However, we suggest that these must necessarily accompany nonverbal pointing gestures, and hence fall outside the scope of the phenomena we will be looking at in this paper.

5.2 Māori addressee terms

5.2.1 Distribution

In Māori, addressee modifiers are typically found at the beginning of sentences, again introduced by commas.

(34) William, kei te auau ngā kurī. (William, the dogs are barking.)

One important exception: greetings are best modified postsententially, and commas are optional here.

(35) Kia ora William. (Hello William.)

5.2.2 Composition

One significant characteristic of Māori is that if the addressee term begins with a short word, it needs to be prefaced by the particle *e*. Thus:

(36) Kia ora, e Tio. (Hello, Tio.)

(37) Tēna koe, e hoa. (Hello, friend.)

(38) *Kia ora, Tio.

With a long word, *e* is not allowed. Thus:

(39) Kia ora, Tamaterā. (Hello, Tamaterā.)

(40) *Kia ora, e Tamaterā.

What constitutes 'long' and 'short' in this context is defined phonologically; see Bauer (1997):451–5; Harlow (2001:83) for fuller accounts. Note that English proper names are exempt from needing the *e* particle. The reason here is again probably phonological.

Beyond this, the rules for what constitutes an addreessee term are fairly similar to those in English. Proper names are allowed; see Examples 35, 36, 37 and 39. Bare common nouns are allowed:

(41) Kia ora, e hoa. (Hello, friend.)

Determiners are typically not allowed, except first person possessives and definite determiners:

(42) *Kia ora, e tētahi rangatira. (Hello, a chief.)

(43) Kia ora, e taku hoa. (Hello, my friend.)

(44) Tēnā koe, e te rangatira., the chief.)

(Note that using a definite determiner is used to add formality to the statement.) Pronouns are not allowed, except for second-person pronouns.

(45) E koe, kei te auau tētahi kurī. ([Hey] you, a dog is barking.)

This example is analogous to the English Example 33; the pronoun is best understood as accompanying a nonverbal pointing gesture, so we will not be including it in our system.

Finally, it is possible to combine a common noun and a proper noun, as in the following example:

(46) E hoa Tāwhaki, kare he kai mau. (Friend Tawhaki, there's no food for you. Mead, 1959:24)

In greetings, it may appear that the second person pronoun can be used as an addressee modifier. This is because greetings typically incorporate a second-person pronoun already, either implicitly or explicitly. There are two common greetings: *kia ora (koe/kōrua/koutou)* and *tēna koe/kōrua/koutou*. The pronoun is optional for *kia ora* but obligatory for *tēna*. But these pronouns are better thought of as appearing within the sentence, rather than within an addressee term. Firstly, they are obligatory for *tēna*, as just mentioned. Secondly, ordinary addressee modifiers can be used with greetings:

(47) Kia ora koe, e Hone. (Hello, Hone.)

(48) Tēna koutou, e ngā rangatira. (Hello, the chiefs.)

(There is one wrinkle to the above account: it is also possible to say *kia ora e koe*. *Koe* here behaves more like an addressee modifier, since it is prefaced by the particle *e*. But it is better to handle this as an exception than to try and build a general account around it.)

5.3 The semantics of addressee terms

What does an addressee term contribute semantically to a sentence? Our suggestion is that it contributes something very like a presupposition about the addressee, just as second-person pronouns do. For an addressee term, the presupposition has additional content as well, namely, all the properties which it mentions. For instance, here is an extract from Example 15 with heightened passion:

(49) Bob: I love you, Sue.

The DRS for this example is given in Figure 5. There are two addressee presuppositions here,

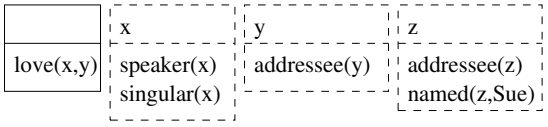

Figure 5: DRS for *I love you, Sue*

one contributed by *you*, and one by *Sue*. Note that the latter presupposition does not bind to anything in the assertion DRS, but that because there is only ever one addressee entity in the context, the two addressee presuppositions are constrained to corefer in this case.

There are nonetheless some important differences between the semantics of an addressee term and that of a second-person pronoun. Firstly, a plural addressee term squarely presupposes the addressee entity, not simply an entity including the addressee. Secondly, as emphasised in Section 3.2.2, addressee terms can function to *change* the default addressee, not simply to refer. This process needs to be described in some detail.

Assume we are dealing with a new incoming utterance from the user. Before processing this utterance, the system the system deletes all the existing `speaker` and `addressee` predicates from the context, sets the new `speaker` to be the user, and sets the `addressee` predicate according to the defaults given in Section 3. For instance, consider Example 50:

(50) User: I like movies, Bob.
Bob (played by the system): That's interesting.
User: Do you like movies, Sue?

After Bob's utterance, the system will set `speaker` to User, and `addressee` to Bob, by Principle 4. The user's second utterance presupposes that Sue is the addressee. The only way of allowing this is by accommodation of an `addressee` predicate which holds of the object named Sue. Since there can only ever be one `addressee` predicate, this also means deleting the existing `addressee` predicate holding of Bob. Accommodation of presuppositions due to addressee terms is thus non-monotonic. One problem this raises is in how to handle the presupposition due to the pronoun *you* in the utterance. If it was processed before that of the addressee term, it ends up being bound to an object which is no longer the addressee! These complications make it important to handle presuppositions due to addressee terms before any other presuppoitions. We therefore augment the DRS language we use to represent utterances, to include a special sub-DRS for addressee terms.

6 An example dialogue for practicing pronouns in conversational Māori

The main application we have in mind for our dialogue system is a computer-aided language learning system for Māori which uses dialogue as its paradigm. To teach the Māori pronoun system, it is useful to support more than two speakers. We have implemented all of the syntax and semantics of pronouns and addressee terms, and many of the dialogue management principles, but we still have some work to do in combining these into a unified framework. In the meantime, Figure 6 gives an example of the kind of multi-speaker dialogue we have in mind. (An English translation is given in Figure 7.) The student is Jason; Hone and Piri are characters whose utterances are generated by the dia-

logue system. In Utterance 1, Hone and Piri

1	Hone/Piri	Kia ora, Jason!
2	Jason	Tēna kōrua.
3	Hone	Kei te hiakai ahau.
4	Piri	Kāore ahau i te hiakai.
5	Piri	Kei te hiakai koe, Jason?
6	Jason	Kāore.
7	Hone	Jason, kei te hiakai kōrua ko Piri?
8	Jason	Kāore māua i te hiakai.

Figure 6: A sample teaching dialogue in Māori

1	Hone/Piri	Hello, Jason!
2	Jason	Hello (you two).
3	Hone	I'm hungry.
4	Piri	I'm not hungry.
5	Piri	Are you hungry, Jason?
6	Jason	No.
7	Hone	Jason, are you (dual) and Piri hungry?
8	Jason	No, we (dual, exclusive) aren't hungry.

Figure 7: An English translation of Figure 6

are the speakers. They address Jason explicitly to start with. When Jason responds, Hone and Piri are by default the addressees, and Jason's greeting uses the appropriate dual pronoun. In 3 and 4, Hone and Piri provide some information about themselves. (In reality, of course, there would be a lot more of this.) In 5, Piri asks a question of Jason. Since Hone is the default addressee of this utterance, Piri has to identify Jason explicitly with an addressee term. In 6, Jason answers Piri. In 7, Hone jumps in; by default, his utterance therefore has Jason and Piri as its addressees. However, Hone's utterance has an explicit addressee term (*Jason*) which overrides this. Finally, in 8, Jason replies, using an appropriate dual and exclusive first-person pronoun to refer to himself and Piri. The utterance is assumed to be addressed to Hone by default.

Implementing a dialogue system which supports this kind of interaction will allow a student great flexibility in practicing Māori questions, answers, assertions, all in a context where distinctions between the alternative personal pronouns are clearly motivated.

References

Allen, J. and Core, M. (1997). Coding dialogs with the damsl annotation scheme. In *Working notes of the AAAI Fall 1997 Symposium on Communicative Action in Humans and Machines*.

Bauer, W. (1997). *The Reed reference grammar of Māori*. Reed books, Auckland, NZ.

Bayard, I., Knott, A., and de Jager, S. (2002). A unification-based grammar for a fragment of Māori and English. In *Proceedings of the 2nd Australasian Natural Language Processing Workshop (ANLP 2002)*.

Copestake, A. (2000). The (new) LKB system. CSLI, Stanford University.

Copestake, A., Flickinger, D., Sag, I., and Pollard, C. (1999). Minimal Recursion Semantics: An introduction. Manuscript, CSLI, Stanford University.

de Jager, S., Knott, A., and Bayard, I. (2002). A DRT-based framework for presuppositions in dialogue management. In *Proceedings of the 6th workshop on the semantics and pragmatics of dialogue (EDILOG 2002)*, Edinburgh.

Ginzburg, J. and Sag, I., editors (2000). *Interrogative investigations*. CSLI Publications, Stanford.

Grosz, B. J. (1978). Discourse analysis. In D. Walker, editor, *Understanding spoken language*, pages 235–268. North-Holland, New York.

Grosz, B. J. and Sidner, C. L. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, pages 175–203.

Harlow, R. (2001). *A Māori reference grammar*. Longman/Pearson Education New Zealand Ltd, Auckland.

Kamp, H. and Reyle, U. (1993). *From discourse to logic*. Kluwer Academic Publishers, Dordrecht.

King, S., Knott, A., and McCane, B. (2003). Language-driven nonverbal communication in a bilingual conversational agent. In *Proceedings of the 16th International Conference on Computer Animation and Social Agents (CASA)*.

Knott, A. and Vlugter, P. (2003). Syntactic disambiguation using presupposition resolution. In *Proceedings of the 4th Australasian Language Technology Workshop (ALTW2003)*, Melbourne.

Litman, D. and Allen, J. (1990). Discourse processing and commonsense plans. In P. Cohen, J. Morgan, and M. Pollack, editors, *Intentions in Communication*. MIT Press.

Mead, H. (1959). Ka haere a tawhaki ki te tangi. Te Ao Hou. No. 28.

Padilha, E. and Carletta, J. (2002). A simulation of small group discussion. In *Proceedings of the 6th workshop on the semantics and pragmatics of dialogue (EDILOG 2002)*, pages 117–124, Edinburgh.

Sacks, H., Schegloff, E., and Jefferson, G. (1974). A simplest systematics for the organization of turn-taking for conversation. *Language*, **50**, 696–735.

Taylor, J. (1994). *A Multi-Agent Planner for Modelling Dialogue*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex.

Van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics*, **9**, 333–377.

Wang, H.-C., Huang, C.-Y., Yang, C.-H., and Wang, J.-F. (2002). Multi-speaker dialogue for mobile information retrieval. In *Proceedings of the international symposium on Chinese spoken language processing (ISCSLP)*. National Cheng-Kung University, Tainan, Taiwan.