Reference-based Discourse Structure for Reference Resolution

Helen Seville and Allan Ramsay
Centre for Computational Linguistics
UMIST, PO Box 88, Manchester M60 1QD, England
heleng/allan@ccl.umist.ac.uk

Abstract

The connection between discourse structure and reference is well recognised. In the system of language understanding described, discourse structure is exploited for the purposes of reference resolution and then continuity of reference is exploited to build the discourse structure for the purposes of subsequent reference resolution. The role of discourse structure in dereferencing, pronouns, other NPs, and temporal references is explored, as is the role of these in constraining discourse structure.

1 Introduction

The role of discourse structure in reference resolution, in particular pronoun and temporal reference resolution, is well recognised. However both Rhetorical Structure Theory (Mann and Thompson, 1988) and the intention-based approach to discourse structure associated with Centering Theory (Gross and Sidner, 1986) involve high-level reasoning. Such reasoning itself presupposes the availability of interpretations in which references have already been resolved. This raises the question of the legitimacy of assuming such a discourse segmentation for the purposes of reference resolution. From a practical point of view a number of questions arise. Firstly, to what extent can we obtain the discourse structure we need for reference resolution without recourse to such higher-level reasoning? Secondly, having done our reference resolution, what further higher-level reasoning needs to be done if we are to adequately capture the meaning of the discourse? Being able to *label* our discourse structure with discourse relations will only be useful insofar as these relations have implications for language understanding. Recognising an *explanation* can be seen to contribute to our understanding in concrete ways. It allows us to infer, for instance, the fact that one event caused another together with the fact that the first event preceded the second. Given that this information is implicit in the discourse, it is clear that our understanding would be impaired without it. With other discourse relations, however, the contribution to inference and so discourse understanding is less clear.

We have implemented a unified approach to reference resolution and discourse structure in the system of language understanding described in (Ramsay, 1999). Having used the existing discourse structure to resolve references in the utterance and anchor it in the discourse, we then use information about continuity of reference to attach the new discourse state generated by the utterance to the discourse structure, or *discourse tree*. Attachment is by means of a number of interchangeable attachment rules which we have used to explore how different referential cues contribute to discourse structure (Seville, 1999). The kinds of referential information we use include: *theme*[1] (as this term is used by (Halliday, 1985)), since this is what the discourse at any one point is about; *reference time*[2], since sometimes discourse structure reflects the temporal structure of the events described; and *pronominalisation* in general, since this recognisably contributes to the overall coherence of the discourse. There are further referential cues we have yet to explore, including adverbials like "Then" and "Before that" and bridging descriptions like those in the following examples:

(1) John likes Shirley.
(2) *His mother* doesn't.

(1) The house was grand.
(2) *The door* was carved with gilded cherubs.

[1] This is similar in many ways to the *center*.

[2] Again, this is similar to the notion of *temporal center* (Kameyama, Passoneau, and Poesio, 1993), but with important differences; for instance, *reference time* is an instant, while the *temporal center* is an interval.

The way we evaluate the discourse trees generated by the different attachment rules is with respect to two main criteria. The first of these is pronominalisation. An adequate discourse structure should enable us to do long-distance pronominalisation, and, equally, prevent us from obtaining an antecedent in those cases where the referent would normally not be pronominalised; for example:

(1) Mary$_i$ loves Bill.
(2) He loves John.
(3) John loves Peter.
? (4) Peter loves her$_i$.

Our second criterion is discourse relations. While the discourse trees we generate for reference resolution are unlabelled, we require that they get the attachments or *discourse thread* right, so that, in inferring how events are related, we consider only those events which are actually connected in the discourse.

Below, we describe our approach to reference resolution and discourse structure building. In the case of reference resolution we are concerned with a wide range of referring expressions, including anaphors, pronouns, referential auxiliaries like "had", names and definite descriptions (Seville and Ramsay, 1999). Of particular concern is the relationship between short, non-unique NPs like "the man" and both pronouns like "he" and longer NPs like "the man in the moon". Given that both pronouns and short NPs like "the man" are used to refer to familiar referents, in context, an obvious strategy is to exploit the discourse structure we have available to us for the resolution of the latter as well as the former. However, short NPs may be used to refer to entities for which a pronominal reference is no longer available:

(1) Mary loves Bill.
(2) He loves John.
(3) John loves Peter.
(4) Peter loves Mary/*her.

This suggests a functional division between pronouns and names/definite descriptions, with the former, but not the latter, relying on the discourse structure for reference resolution. Full NPs can, of course, also be used to refer to entities which are pronominalisable, in the appropriate circumstances:

(1) John$_i$ loved Bill$_j$.
(2) He$_i$ adored him$_j$.

(1) John$_i$ loved Bill.
(2) Bill adored him$_i$.

Given that we are not treating definite descriptions as dependent on the discourse structure for their reference, we still have to consider how a short NP like "the man" may successfully refer when there are apparently several potential referents in the context:

"Alex$_i$ was *the man*$_i$ conducting the interviews. He$_i$ arrived on time but couldn't find the entrance. A man$_j$ showed him$_i$ the way and he$_i$ settled into his office just in time for the first candidate. A man$_k$ walked into the room. Alex$_i$ asked him$_k$ to sit down. *The man*$_k$ pulled out a gun. Alex$_i$ recognised that he$_k$ was *the man*$_j$ who had given him$_i$ directions."

2 Reference Resolution

For each utterance input to the system of language understanding, we obtain a logical representation; for example, the following is the representation obtained for "John loved Mary":

$$\exists A : \{ref(\lambda B(speech\_time(B))) > A\}$$
$$\exists C : \{aspect(simple, A, C)\}$$
$$\theta(C, agent, ref(\lambda D(named(D, John))))$$
$$\&love(C)$$
$$\&type(C, event)$$
$$\&\theta(C, object, ref(\lambda E(named(E, Mary))))$$

All referring expressions are represented in this as terms of the form $ref(\lambda x P x)$. Temporal expressions, or tense, also give rise to such terms. We use the inference engine described in (Ramsay, 1995; Cryan and Ramsay, 1997) to anchor the utterance in the common ground and update the common ground with its entailments. Before our representation is input to the theorem prover, it is converted to Skolemised normal form and then to sequent form.

$$[] ==> [ref(\lambda A(speech\_time(A))) > \#72]$$
$$[] ==> [aspect(simple, \#72, \#71)]$$
$$[] ==> [\theta(\#71, agent, ref(\lambda A(named(A, John))))]$$
$$[] ==> [love(\#71)]$$
$$[] ==> [type(\#71, event)]$$
$$[] ==> [\theta(\#71, object, ref(\lambda A(named(A, Mary))))]$$

The order of the terms $ref(\lambda B(named(B, John)))$ and $ref(\lambda B(named(B, Mary)))$ in these retains the ordering of "John" and "Mary" in the utterance, and this is the order in which the terms will be dereferenced. Retaining the order of referring expressions is obviously important for dealing with examples like "John loves himself", but also for determining the relative salience of referents for the dereferencing of subsequent pronominals.

In anchoring referring expressions, we treat each reference term as carrying with it an invitation to carry out an existence proof in the common ground. The nature of the proof is reflected in the properties of the reference term. Names and definite descriptions have no special properties, as they don't bring with them specific instructions as to where the referent is to be located in the discourse structure. We

discuss how we deal with these below. Anaphors are characterised as *salient* and pronouns as *centred*, in the current discourse state (e.g., 4):

"himself" $ref(\lambda X(salient(X,4)\&m(X)))$

"him" $ref(\lambda X(centred(X,4)\&m(X)))$

In order to dereference these terms, we need to consider the discourse structure. In the case of anaphors, we consider the referents already mentioned in the new discourse state being constructed for the utterance currently being processed. In the case of pronouns, we consider all *open nodes* in the discourse tree we have constructed so far, i.e. those on its right-hand frontier:

```
  0
  |
  1      Mentioned: (1, #72, #71, #46, #47)
 / \
2   3    Mentioned: (3, #72, #76, #75, #47, #46)
```

(1) John loved Mary.
(2) He had hated her.
(3) She had hated him.

Each of these nodes represents a discourse state corresponding to an utterance and has associated with it the list of referents mentioned in that utterance. These include, for example, in the case of "She had hated him": 3, the speech time; #72, the reference time; #76, the event time; #75, the event itself; #47, Mary; and #46, John. As well as pronouns, temporal references are dereferenced using the open nodes in the discourse structure, as detailed below.

2.1 Anaphors

As stated above, we characterise the referents of anaphors as salient. This means that we must find their referents in the current discourse state. The other constraint we place on anaphors is that their referents play the role of *arguments*. This rules out cases like:

* "Mary$_i$ took John with *herself*$_i$."

While anaphor resolution is not dependent on discourse structure, we discuss it briefly here mainly for the purposes of comparison with pronoun resolution, which is heavily dependent on discourse structure. Consider the following example:

(1) John$_i$ loves Bill$_j$.
(2) He$_i$ adores himself$_i$.

In dereferencing "himself" in (2), the reference term to be considered is $ref(\lambda B(salient(B,2)\&m(B))))$. In dereferencing this we consider, in the order in which they were mentioned and dealt with by the inference engine, the entities already added to the current discourse state:

[2,#78,#46]

These are speech time (2), the event itself (#78), and John (#46). Of these, we can prove that #46 satisfies the property m, which we use to represent male gender. It also meets the constraint of being an argument, so our reference term is dereferenced to it.

2.2 Pronouns

In the case of pronouns we consider two cases, that where the referent functions as an argument and that where it doesn't. In the following case, "him" plays the role of argument:

(1) John$_i$ loves Bill$_j$.
(2) He$_i$ adores him$_j$.

What this means is that, in order to ensure that the constraint of *disjoint reference* (i.e., between pronouns and the antecedents of anaphors) is met, we must start our search for the *centred* referent of $ref(\lambda B(centred(B,2)\&m(B)))$ not in the current discourse state, but in the previous discourse state, represented by the lowermost open node in the discourse tree. Eliminating from consideration any of its referents which are also current entities, i.e., #46, this leaves us with:

[1,#80,#38]

These entities are, respectively, speech time, the event, and Bill. All of these are considered, in order, and we find that #38, i.e., Bill, matches the referring expression.

In the following example, "her" does not refer to an argument, and we must begin the search for the referent of the corresponding term $ref(\lambda B(centred(B,1)\&f(B))))$ not in the previous, but the current discourse state:

(1) Mary$_i$ took John with her$_i$.

Here, the referents mentioned so far include:

[1, #82, #81, #47, #46]

Of these, we can prove that #47, i.e., Mary, satisfies the property f, for female gender, and so our term is successfully dereferenced to this. We deal with the following cataphoric example in exactly the same way:

"With *her*$_i$ Mary$_i$ took John."

Because we treat "With her" as an instance of left extraposition, we get exactly the same logical form for this utterance as for that above. It is this which determines the order in which referring expressions

are dereferenced and so, in this case, as in that above, "Mary" is dereferenced before, and so provides the antecedent for, "her"[3].

Where we fail to find an antecedent for a pronoun in the previous discourse state, our search considers the next open node in the discourse tree. This is not necessarily the immediately previous discourse state, as illustrated by the following example[4]:

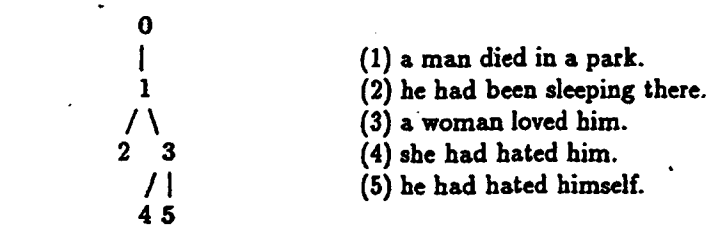

Here (3) is the next open node to be considered after (5), since (4) was closed by the attachment of the latter. It is in the list of referents mentioned in (3) that we are able to find the referent of "her" given the following continuation:

(6) he had loved her.

2.3 Reference Times

For the purposes of reference resolution, we treat referential tenses similarly to pronouns. As an example, consider utterance (5) above. The logical form we obtain corresponding to this is as follows:

$$\exists A :\{ref(\lambda B(centred(B,5)\& ref(\lambda C speech\_time(C))) > B) > A\}$$
$$\exists D :\{aspect(perfective, A, D)\}$$
$$\theta(D, agent, ref(\lambda E(centred(E,5)\&m(E))))$$
$$\&hate(D)$$
$$\&type(D, event)$$
$$\&\theta(D, object, ref(\lambda F(salient(F,5)\&m(F))))$$

This contains a reference to a centred time before speech time:

$$ref(\lambda B(centred(B,5)\& ref(\lambda C(speech\_time(C))) > B))$$

When we are anchoring (5), the open nodes in the discourse tree are (4), (3), and (1). The entities mentioned in discourse state (4), the first open node to be considered, are as follows:

[4, #478, #480, #479, #476, #477]

These correspond to speech time, the reference time (i.e., the event time of (3)), the event time, the event itself, the woman, and the man. Our temporal reference term is dereferenced to the first *centred* time which we can prove is before speech time[5]. This is not #480, the time of (4)'s event, but #478, the time of (3)'s event, which precedes it in the list of entities mentioned in (4). As will be apparent below, this choice of reference time has important implications for the discourse structure we build.

2.4 Names and Definite Descriptions

There are no reliable surface cues for distinguishing short, non-unique NPs used to refer to familiar referents and longer NPs which may denote simply by virtue of their meaning. We therefore need a unified approach to dereferencing the reference terms we obtain for both. As mentioned above, these are all of the form, $ref(\lambda x P x)$, but, unlike in the case of anaphors and pronouns, incorporate no properties characterising where in the discourse structure their referents are to be found:

"the man" $ref(\lambda X(man(X)))$

In finding our referent we have no recourse to *dynamic* properties of entities like *salient* and *centred*, but only *static* properties like *man*. However, we still need to distinguish more and less prominent potential referents of a short NP like "the man". The way in which we do this is to consider entities mentioned in the discourse in order of recency. Given a referring expression containing a property such as $\lambda X(man(X))$, we prove that there is an entity satisfying the property. In this way we obtain either an entity which has been explicitly mentioned or, in the case of bridging descriptions, one associated with such an entity but which has not itself been explicitly mentioned. In the case of a bridging description, we obtain the proof via a meaning postulate, such as the following:

$$\forall X(house(X) => \exists Y(of(Y, \lambda Z(door(Z)), X)))$$

Due to the ordering of facts and meaning postulates in the database, we tend to obtain the most recently mentioned entity or, in the case of bridging descriptions, that with the most recent antecedent, first. The following example serves to illustrate:

[3] There are further cases of cataphora which we don't attempt to deal with here. These differ from the above in that there is no left extraposition of pronoun-containing constituents. Instead, the pronominalised items are *marked* as *new* by being stressed, which is what distinguishes them from unmarked, anaphoric uses (Halliday and Hasan, 1976); for example:

"THIS is what worries me:
I can't get any reliable INFORMATION."

Dealing with such examples seems to require a theory of focus, and so is beyond the scope of this paper.

[4] This was built using the cue of pronominalisation, as described in the following section.

[5] We ignore previous speech times from the point of view of subsequent dereferencing.

(1) Shirley arrived at the house$_i$.
(2) John opened the door$_{f(i)}$.

In resolving the reference to "the door", we obtain a proof that the entity #88(#123) satisfies this property. Although this entity has not previously been mentioned, the house[6], #123, has been recently mentioned, and it is this, together with the meaning postulate above, which forms the basis of our proof.

In order to recognise cases of ambiguity, rather than simply find the single most recent entity satisfying the referring expression, we simultaneously find all entities satisfying the property which are equally recent.

(1) The car$_i$ arrived at the house$_j$.
(2) John opened the door$_{f(i/j)}$.

Our default criterion of equal recency is having been mentioned, or having an antecedent, within the same discourse state. This means that we recognise the ambiguity in the example above. However, given this criterion, we do not detect any ambiguity in the following example[7]:

(1) Shirley got in the car$_i$.
(1) She arrived at the house$_j$.
(2) John opened the door$_{f(j)}$.

Where no entity of the kind has been mentioned, our reference resolution procedure is exactly the same, with all potential referents being considered equally recent. We will obtain a unique referent in this case given a definite description, like "the *first* man on the moon", which denotes such a referent by virtue of its meaning.

Because the referents of NPs needn't have been mentioned and don't have to satisfy the property *centred*, we are able to obtain referents for names and definite descriptions in cases where a pronominal reference would have failed. However, the reverse situation is also possible. In the following example, "it" clearly picks out the red car, whereas "the car" doesn't seem an appropriate way of referring to it.

(1) Mary drives a red car$_i$.
(2) She hates it$_i$/the car$_{i?}$.

Our explanation of this is that, in the same way that we must use an anaphor rather than a pronoun when an anaphor is available (i.e., for arguments within the same clause), there is a convention whereby we normally use a pronoun rather than a definite description when this option is open to us[8]. It is possible to use a full definite description in such circumstances, but there should be a reason for such a choice of referring expression, of which discourse structure is the obvious candidate, since this may well be affected in such a case, as we will show below. We define normal use as follows. A pronoun should be used to realise any entity mentioned in the previous discourse state, except in those cases where this entity is being promoted over a more salient referent with the same pronominal properties. According to this definition, "Bill" in (2) below is a normal use of a full NP, since "He" would have been dereferenced to John, over whom Bill is being promoted.

(1) John loves Bill.	(1) John loves Mary.
(2) Bill loves him.	? (2) Mary loves him.

However, the use of "Mary" is regarded as exceptional. As such, it is likely to be interpreted as conveying some implicature (Grice, 1975). Given some associated discourse move (i.e., a further convention associated with the exception), the use of "Mary" can be regarded as *felicitous* (Austin, 1962). However, in the absence of any reason for the exceptional choice of referring expression, it may be regarded as *infelicitous*. While we have not associated any discourse moves with the use of definite descriptions, we do find such moves *emerging* in the work described below. This happens because we treat pronouns, but not names and definite descriptions, as anchored to some particular discourse state, and this is one kind of information used in building the subsequent discourse structure.

3 Building the Discourse Structure

Once we have anchored an utterance and created a new discourse state corresponding to it, we use information about continuity of reference to attach it to the existing discourse tree. Since there are a variety of referential cues and ways of using (and combining) them to be considered, this attachment information is represented in a variety of interchangeable rules. The referential cues explored so far include:

[6]This was also a first mention. In this case, the referent was accommodated.

[7]This seems appropriate in this case:

(1) Shirley got in the car$_i$.
(1) She arrived at the house$_j$.
(2a) John opened the door$_{f(j)}$ to let her in.
? (2b) John opened the door$_{f(i)}$ to let her out.

However, in other cases, it does seem that we need to consider entities other than those within the same discourse state to be equally recent. Discourse structure may be a factor, and this is something we may be able to capture. However, other factors which are harder to track, like discourse topic, also seem to be involved.

[8]This is not merely an *arbitrary* convention. The use of a pronoun assures us that we are continuing to talk about the same referent in circumstances where a name or definite description might signal a change of referent.

- the theme
- the reference time (where applicable)
- all pronouns (including reference time)

In choosing a discourse state to which to attach the current discourse state, we may specify, for instance:

- one that shares its theme
- one to which any pronouns that it contains can be dereferenced
- the lowest suitable utterance in the tree
- the highest suitable utterance

Furthermore, attachment may be:

- as a daughter, thereby keeping the attachment utterance open
- as a sister, resulting in its closure

Here, we briefly consider how useful the different referential cues seem to be, by comparing the results of some of the more successful attachment rules.

3.1 Theme

The tree below was constructed using a very simple attachment rule whereby a discourse state which continued the existing theme was attached to its predecessor as a sister, while one which changed the theme was instead attached as a daughter[9].

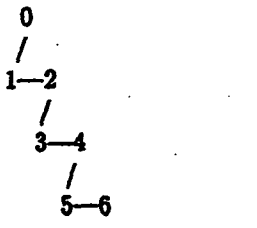

(1) a man died in a park.
(2) he had been sleeping there.
(3) a woman loved him.
(4) she had hated him.
(5) he had hated himself.
(6) he had loved her.

Of the rules we considered which utilised theme, this one was relatively successful, supporting the long-distance reference to "her" in (6). It can be seen that this is because the open nodes in the tree serve to keep track of any previous themes. However, theme, taken by itself, does not seem to provide a good guide to discourse structure. The discourse structure we obtain here is, in effect, linear. Alternative theme-based attachment rules did result in more interesting structures. For instance, (5) could be treated as resuming the theme of (2), and so attached there; however, in this case the reference to "her" in (6) would fail due to the closure of nodes (3) and (4). A further alternative involves treating smooth transitions from mention in the rheme to mention as theme as cases of thematic continuity (Hahn and Strube, 1997). We found this approach to result in some trees which were useful for pronoun resolution but which, again, didn't seem to reflect the discourse structure.

[9]We distinguish between directly attaching n_2 as a sister of n_1 and attaching it as a daughter of n_1's mother. In both cases n_1 would no longer be open either for dereferencing pronouns or as an attachment site for future nodes. The difference becomes apparent when we start labelling the tree with discourse relations, a matter which is not discussed in the current paper.

3.2 Reference Time

We illustrate the cue of reference time using the same example as above. In this case, attachment of a discourse state with a referential tense was to the discourse state which first introduced the time referred to.

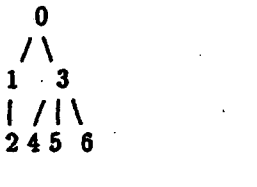

(1) a man died in a park.
(2) he had been sleeping there.
(3) a woman loved him.
(4) she had hated him.
(5) he had hated himself.
(6) he had loved her.

Again, as in the case above, we are able to manage the reference to "her" in (6). There is a difference in this case, however, in that we anchor the referring expression to (3) in this case but to (4) above. Arguably this tree, which has fewer open nodes, does a better job of retaining just that information which is needed for pronoun resolution. What it clearly does better, however, is represent aspects of the discourse structure like the dependency of (2) on (1) and the parallelism of (4), (5), and (6). It can be seen to do this insofar as the temporal structure reflects the discourse structure. Where it is unsuccessful is in capturing that (3) is a continuation of (1), but then this is unsurprising since the simple past tense of (3) is existential rather than referential[10].

3.3 Pronominalization

The results we obtain using the cue of reference time alone can be enhanced either by considering theme together with reference time or by considering all pronominalised items, which we take to include reference time. The problematic attachment of (3) above is not helped by considering its theme, since this is, like its tense, existential rather than referential. However, if we use a pronominalization-based rule, whereby we attach a new discourse state as a daughter of the highest open node at which all its

[10]It refers only to the current speech time.

pronouns can be dereferenced, then we obtain the following tree.

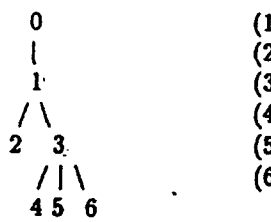

(1) a man died in a park.
(2) he had been sleeping there.
(3) a woman loved him.
(4) she had hated him.
(5) he had hated himself.
(6) he had loved her.

By considering all pronouns, we realise that (3) should not be attached higher than (1), where the referent of "him" was introduced. It is arguable that the attachment of (3) to (1) should have been as a sister rather than a daughter, thereby closing (1) to future pronoun resolution and attachment, but this is a distinction which is hard to capture.

4 Discussion

By experimenting with a variety of interchangeable attachment rules we have found that, at least in those cases where it is available, temporal reference provides a reliable constraint on discourse structure. Where it is not available, tree-building can be guided by other cues such as continuity of theme and pronominalisation. However, unlike reference time, these are heuristics rather than constraints and in some cases can be unreliable. In the first example below, the theme, "She", of (3) seems to cue attachment to (1), where this referring expression is anchored, but (3) is, rather, an elaboration of (2). In the second example, (3) is a continuation of (1) even though its theme, "He", refers to John who was not even mentioned until discourse state (2).

(1) Sam arrived at Hannah$_i$'s house.
(2) He had lost the key.
(3) She$_i$ would be annoyed with him.

(1) Mary ran.
(2) John$_i$ had seen her.
(3) He$_i$ chased her.

While reference time is of limited applicability, we do not regard this as a problem since we can recognise precisely those cases in which it is applicable. There is one other kind of cue which we have found to be useful, and again this only applies to a particular set of cases. This is the use of a full NP to refer to an already familiar referent, particularly the current theme, when, for the purposes of reference resolution, a pronominal reference would be perfectly adequate. While this is not a cue we explicitly set out to investigate, interestingly different discourse structures emerged from certain of the attachment rules specified, depending upon whether or not references were made pronominally. Compare the following discourse structures, both constructed using the cue of pronominalization, as outlined above. The first, with its pronominal references, can be regarded as an *elaboration*, but, if continuity is regarded as a prerequisite of this relation, then the second cannot[11].

(1) Mary$_i$ loves John.
(2) She$_i$ adores him.

(1) Mary loves John.
(2) Mary adores John.

This is not a result we set out to achieve. It emerged as a result of our different approaches to reference resolution for pronouns and full NPs. Recall that the former but not the latter are regarded as anchored within the discourse structure, thus providing us with potential attachment cues. In constructing the first of the trees above, the anchoring of the pronouns in (2) to discourse state (1) provides the basis for the attachment made. There are no alternative rules for full NPs in the second example, however; rather, attachment simply fails to be made to (1) because there is no evidence of a discourse connection.

The different discourse structures we obtain given pronominal and non-pronominal references to the same thematic entity suggest the importance of the choice of referring expression to discourse structure. One of our goals is to use referring expressions to recognise the discourse moves associated with them. The following examples serve to illustrate:

(1) Sam$_i$ arrived at the house.
(2) He$_i$ had lost the key.
(3) He$_i$/Sam$_i$ rang the bell.

(1) Sam$_i$ arrived at the house.
(2) He$_i$ had lost the key.
(3) He$_i$/?Sam$_?$ dropped it.

In the first of these, "Sam" in (3) seems to be a felicitous choice of referring expression. This makes sense, since, in this case, (3) is not a continuation

[11] If there were a change of referent, then the relation would seem to be a parallel or contrastive one, as in:

(1) Mary loves John.
(2) Jane adores him.

The discourse structure we get when the full NP "Mary" is used is one associated with such a discourse relation. However, these relations are inappropriate in this case. Rather than attempt to label the tree resulting from the use of "Mary", we would prefer to regard this choice of referring expression as infelicitous.

of (2). In the second example, however, "Sam" does not seem to be available as a choice of referring expression. Again this makes sense insofar as (3) is here a continuation of (2). The cue of pronominalization seems to indicate continuity. Where it could have been used but isn't, the deliberate choice of an alternative referring expression serves to *implicate* (Grice, 1975) the absence of continuity. However, our rules as currently implemented don't really capture this explanation. Given a non-pronominal reference to the current theme, there may well be a lack of evidence of continuity, so we avoid the attachment of (3) to (2), but we may also fail to attach (3) to (1)! One way of dealing with such cases would be to introduce a new kind of attachment rule. All the rules explored so far rely on choosing a preferred node for attachment. These cases might be better dealt with by negative constraints on attachment. Given a default attachment of (3) to (2)[12], this attachment will be prevented in those cases where (3)'s theme is not pronominalised (but could have been), thereby resulting in the alternative attachment of (3) to (1).

Having identified two kinds of referential cues which seem to provide useful guides to discourse structure in particular cases, it is worth considering what problems remain with these. Firstly, we consider temporal reference as a cue to discourse structure. We have argued that this provides a relatively reliable constraint. However, this assumes that we resolve our temporal references appropriately in the first place.

```
  3
 / \
4   5
```

(3) a woman loved him.
(4) she had hated him.
(5) he had hated himself.

In this fragment from the discourse used above, we attach (5) to (3) only because we first dereferenced its reference time to the event time of (3). Now consider the alternative continuation:

(5) he had beaten her for years.

In this case, we would also attach (5) to (3) because again its reference time would be preferentially dereferenced to the event time of (3)[13], rather than that of (4), but in this case we get it wrong. The mistake is not necessarily one which is irrecoverable. If there were no sensible discourse relation to be inferred between (5) and (3), as in this case, then the mistake could be recognised and both the temporal reference and the attachment revised. However, this example does serve to illustrate the problems with in effect using temporal structure as a cue to discourse structure, when perhaps the two are better treated as *mutually constraining*.

In exploiting the distinction between pronominal and non-pronominal references to the current theme we find a related, albeit less serious, problem to that above. In order to exploit the choice of referring expression in building discourse structure, we need to know that it is felicitous. However, whether it is felicitous itself depends on the discourse structure! Note that, in the examples discussed above, discourse states (1) and (2) taken together provide *identical* contexts for the dereferencing of the referring expressions "He/Sam" in (3). The only difference is the discourse thread, which joins (3) to (1) in the first case, but to (2) in the second case. Since this is not the kind of information we have available to us when we are doing our reference resolution, we simply have to *assume* that the choice of referring expression is felicitous. However, this assumption could be retracted were we to fail to label the discourse structure constructed. In this example, "Sam arrived at the house ...Sam dropped it [the key]." can be rejected as a *narrative* sequence, because Sam can't drop the key he has lost[14].

5 Further Work

Referential continuity is a complex phenomenon, and in considering the impact of reference on discourse structure there are many more examples of reference to be considered than those discussed here. These include referential adverbials which, as an additional source of temporal reference to tense, can be expected to further constrain the discourse structure. Some examples of these are discussed in (Seville, 1999). Another important category of referential expressions are bridging descriptions. While pronominalisation is a key indicator of referential continuity where it is available, where it is not, bridging descriptions seem to provide a similar degree of referential continuity:

(1) The house$_i$ was grand.
(2a) It$_i$ was in the baroque style.
(2b) The door$_{f(i)}$ was carved with gilded cherubs.

[12] By this we mean only that consideration of (2) as a potential attachment node precedes consideration of (1).

[13] A similar preference would be made by other approaches such as Temporal Centering (Kameyama, Passoneau, and Poesio, 1993).

[14] Similarly, "he had beaten her for years" would be rejected as a potential *explanation* of "a woman loved him" in the example above, and, to return to a previous example:

(1) Mary loves John.
(2) Mary adores him.

This could not be labelled as a *contrast*, as no element of contrast can be identified.

Integrating bridging descriptions into our treatment of reference resolution was straightforward, because of the way in which the terms they refer to are dervived, via meaning postulates, from terms representing existing referents. For the same reason, recognising and exploiting this kind of referential continuity to build discourse structure should be relatively straightforward.

6 Conclusion

We have described a unified approach to reference resolution and discourse structure which is implemented in a system of language understanding. As well as exploiting discourse structure for reference resolution, we explore the extent to which referential information can in turn be exploited to provide the discourse structure to be used by subsequent reference resolution. The strengths and weaknesses of the different referential cues investigated have already been discussed in the previous section. Here we highlight some of the other issues raised.

What, if any, is the role of discourse structure in the dereferencing of short, non-unique NPs referring to familiar referents? We have argued that the differences between the felicity conditions for the use of pronouns and non-pronominalised NPs suggest that the former are dependent on discourse structure but the latter are not. This has further implications for our account in that it means we in turn exploit pronominalised references, but not non-pronominalised references, in building discourse structure. The different uses of pronouns and non-pronominalised NPs have also been noted by (Hitzeman and Poesio, 1998), although with a view to reference generation.

While space precludes a full discussion of their approach, it is worth summarising the difference between their approach and ours here. They assume that discourse structure is used for the resolution of all referring expressions, but that, while pronouns can only refer to an entity which was at one point the Most Salient Entity[15] in the discourse, definite descriptions can refer to any entity on the discourse stack. We assume that discourse structure is used only for the resolution of pronominalised referring expressions. Related to this, we exploit the fact that items which *can* be pronominalised normally *should* be, in order to impose further structure on the discourse. If a short NP is used where a pronoun would have been interpretable, we assume that the current utterance should not be attached, either as a daughter or a sister, to the most recent node. Failing to take advantage of an attachment clue, such as the use of pronominalisation of theme, is seen a positive instruction not to attach at a certain point, rather than just as a lack of information.

Our observation of the effects of the choice of referring expression, whether pronominalised or not, on discourse structure raised a number of issues which need to be explored further. The first of these is the need to make a distinction between the effects of *referents* and *referring expressions* on discourse structure, since the choice of a different referring expression for the same referent may convey a different discourse move. Closely related to this, there is the need to consider how referential continuity may be indicated where pronominalisation is not an option, as, for instance, by the use of bridging descriptions linking non-identical but closely related referents. Finally there is the need to consider whether attachment choices may be better explained in terms of negative constraints on attachment rather than (or in addition to) the attachment preferences we have assumed here.

A further observation we made was that discourse structure constrains temporal structure, and so temporal reference, in the same way that temporal structure constrains discourse structure. While reference time may provide us with a reliable cue to discourse structure, if our preferred reference time turns out to be wrong then we also get the discourse structure wrong. Similarly, if we had operated attachment preferences before reference resolution, the converse would have held. Either way, we are faced with the dilemma that it is not until we have *both* a potential set of referents and a potential attachment site that inference can be used to determine whether any plausible discourse relations hold given the choices we have made.

[15] This notion is closely related to that of the *center*.

References

Austin, J. L. 1962. *How to do things with words.* Oxford University Press, Oxford.

Cryan, M and A M Ramsay. 1997. A normal form for Property Theory. In *Fourteenth Conference on Automated Deduction*, Townsville, Australia.

Grice, H. P. 1975. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Speech Acts.* Academic Press, London, pages 41–58.

Gross, B. J. and C. L. Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.

Hahn, U. and M. Strube. 1997. Centering-in-the-large: Computing referential discourse segments. In *ACL/EACL*, pages 104–111.

Halliday, M. A. K. 1985. *An Introduction to Functional Grammar*. Edward Arnold, London.

Halliday, M. A. K. and R. Hasan. 1976. *Cohesion in English*. Longman, London.

Hitzeman, J. and M. Poesio. 1998. Long distance pronominalisation and global focus. In *ACL/COLING*, Montreal.

Kameyama, M., R. Passoneau, and M. Poesio. 1993. Temporal centering. In *ACL*, pages 70–77.

Mann, W. C. and S. A. Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organisation. *Text*, 8(3):243–281.

Ramsay, A. 1999. Does it make any sense? updating = consistency checking. In K. Turner, editor, *The Semantics/Pragmatics Interface from Different Points of View*. Elsevier Science B.V.

Ramsay, A M. 1995. A theorem prover for an intensional logic. *Journal of Automated Reasoning*, 14:237–255.

Seville, H. 1999. Experiments with discourse structure. In *Third International Workshop on Computational Semantics*, pages 233–246, Tilburg.

Seville, H. and A. Ramsay. 1999. Reference: the long and short of it. Draft.