A Constructive View of Discourse Operators *

Allan Ramsay
Department of Computation
UMIST, PO Box 88,
Manchester M60 1QD, UK
Allan.Ramsay@umist.ac.uk

Helen Gaylard
Department of Computer Science
University of Exeter
Exeter, EX4 4PT, UK
H.L.Gaylard@exeter.ac.uk

Abstract

Dialogue systems have to consider not just the propositional content of the user's utterances, but also the user's attitude to that content. It is common practice to treat these issues separately: we will argue that they can, and should, be dealt with at the same time and using the same mechanisms.

1 Interpretations, Proofs and Meaning Postulates

Language is used to convey ideas – to produce in the hearer's mind a picture of the world that corresponds to the picture which is already in the speaker's. More than that, however, it is used to convey *attitudes* – this bit of what I am saying is interesting, this bit undermines what you just said, this bit follows on from what you said, ...

Any system that is to process *and respond to* user utterances will have to be sensitive to the way that utterances encode attitudes. We suggest that the best way to capture this information is to build it into the basic meaning representation, and then to elaborate the consequences of using one construction rather than another in exactly the same way that you elaborate the consequences of using a specific lexical item.

Consider for instance (1) and (2):

(1) a. Mary and John got divorced in March.
b. In March Mary and John got divorced.

(2) a. Mary and John got divorced in *March*.
b. Mary and John got *divorced* in March.

These sentences all report the same event – they all have the same PROPOSITIONAL CONTENT. Nonetheless, they differ in the way they express the speaker's attitude to that content, or to parts of it.

We can capture the propositional content itself by constructing a LOGICAL FORM in the usual way. The LF in Fig. 1 is fairly orthodox, and we will use it to illustrate what you have to do in order to include facts about the speaker's attitude. The details of how many θ-roles there should be[1], the decision to deal with definite NPs by including REFERENCE TERMS inside the LF rather than treating them separately as constraints (Barwise and Perry, 1983; Kamp, 1984), the specific treatment of aspect are all open to debate, but nothing much hangs on these issues in the remainder of the paper and we will simply assume that they are at least defensible.

$$\begin{array}{l}\exists A : \{aspect(simple, ref(\lambda B past(P)), A)\} \\ \quad \theta(A, object, ref(\lambda E(name(E, Mary))))^2 \\ \quad \&\theta(A, object, ref(\lambda F(name(F, John)))) \\ \quad \&divorce(A) \\ \quad \&A\ is\ event \\ \quad \&in(A, ref(\lambda G(name(G, March))))\end{array}$$

Figure 1: Logical form for (1a)

This LF was constructed COMPOSITIONALLY, i.e. on the basis of the meanings of the parts and their mode of combination. There is no alternative: when you hear an utterance, or a read a text, you have to base your understanding on what you hear or see. The key claim in this paper is that when you are trying to participate in a dialogue it is even more important to extract all the information that is encoded by the choice and arrangement of words than when you are merely trying to extract the propositional content.

There is no point, however, in constructing an LF unless you link it to other things that you know. When people say things to one another, they as-

[0]This work was partially supported by EU grant IST-2000-29452 "Dynamic Universal Mobility for Adaptive Speech Interfaces"

[1]see (Dowty, 1989; Dowty, 1991) for a detailed discussion of how many thematic roles there are.

sume that the other person will construct a rich model of what was said, incorporating a combination of general background knowledge and a model of what has already been said (by any participant in the conversation). This model goes by various names – the discourse model, the common ground, ... We follow fairly common practice in using the term SLATE for this object.

The curious thing about the slate is that although both the speaker and hearer rely on having the same view of the slate, neither of them can directly inspect the other's version. Dialogues only flow coherently if both parties have very similar slates, but since neither of them can see directly into the other person's head they cannot be sure that they do.

We take it, then, that the basic steps in assimilating an utterance are as follow:

1. construct a logical form that captures all and only the information explicitly encoded by the form of the utterance.
2. check that any claims that this logical form makes about the current slate are indeed supportable
3. update the slate to include the new information in the utterance

If you get step 1 wrong, you've got a problem. If you fail to capture any of the information encoded by the utterance, no amount of subsequent inference can retrieve it for you; and if you include anything erroneous you will have great trouble spotting it and removing it. You may need to note that there are alternative analyses, and either follow up one choice, with the possibility of backtracking when it turn out to be wrong, or somehow delay making a decision until it really matters (at which point you may have the information required for making an informed choice), but you do have to look closely at the utterance and get what you can from it.

Step 2 involves verifying any presuppositions that are encoded in the utterance. (van der Sandt, 1992) argues that referential expressions are, in fact, a form of presupposition, and that determining the item 'pointed to' by such an expression is just a by-product of the process of verifying the presupposition. We argue that if you are the hearer then this process involves checking that your version of the slate Σ_H supports the same proofs as the speaker's Σ_S. This is the closest you can get to inspecting each other's views of the slate – if they support the same proofs then they can't be all that different. The uniqueness element of definite NPs is particularly critical at this point – if Σ_S supports a proof of $man(t)$ for exactly one t and Σ_H does likewise then S and H are probably thinking about the same man, who can therefore safely be referred to as *'the man'* (Ramsay, 1999; Gaylard and Ramsay, 2002).

Step 3 involves adding the new information included in the current utterance to the logical form, producing Σ_H^{i+1} from Σ_H^i. This step may just involve adding the new elements of the logical form, or it may involve adding other things that can be inferred from the logical form in the current context. We have argued that you can combine steps 2 and 3 by trying to build a model of the current utterance (Ramsay and Seville, 2000a). Other authors have taken a similar view (Wedekind, 1996; Blackburn et al., 1997; Baumgartner and Kühn, 2000; Gardent and Konrad, 2000).

No matter how you approach steps 2 and 3, there is no doubt that you need substantial amounts of background knowledge. To find out what someone who utters (1a) has in mind, you have to know that divorce is an action that involves terminating an earlier contract, and that usually when it happens the two parties concerned no longer love each other, and ... In other words, you need to have access to the kind of information contained in Fig. 2, and you need to be able to reason with this information.

$$
\begin{array}{l}
\forall A : \{divorce(A)\} \\
\quad \forall B : \{\theta(A, object, B)\} \\
\quad\quad \forall C : \{\theta(A, object, C)\ \&B \neq C\} \\
\quad\quad\quad \exists D marry(D) \\
\quad\quad\quad\quad \theta(D, object, B) \\
\quad\quad\quad\quad \&\ \theta(D, object, C) \\
\quad\quad\quad\quad \&termination(A, D)
\end{array}
$$

Figure 2: You can only get divorced if you were married

We are thus working in a framework where we construct logical forms and use these to build models that support them, taking the model constructed at each stage to be the current version of the slate. For a variety of reasons we assume that it is convenient to use a fine-grained intensional logic as the formal substrate of this activity. The logic we use is a constructive version of (Turner, 1987)'s PROPERTY THEORY. Other intensional logics are available – (Bealer, 1989)'s fine-grained intensional logic, non-well-founded set theory (Aczel, 1988), higher-order unification(Pulman, 1993). We choose property theory because it lends itself to a simple extension of a standard first-order theorem prover, which is what we use for constructing mod-

els (Ramsay, 2001).

2 Discourse Operators as Attitude Reports

Can we extend the update process described above so that it deals with the differences between the various versions of the report of John and Mary's divorce in (1) and (2)? If we can, then we can make the critical information about how these sentence fit into the discourse as a whole available to our dialogue manager.

2.1 Theme & Rheme

The first move is to see just what the differences are. We reconsider (1):

(1) a. Mary and John got divorced in March.
b. In March Mary and John got divorced.

The difference between them lies in the fact that by putting the temporal modifier first in (1b) we are drawing attention to what haappened in March, rather than saying what John and Mary did.

Informally we can say that (1a) is 'about' John and Mary, and (1b) is about what happened in March. Formally we can say the same thing:

$$\begin{array}{l} about(\lambda B(B.ref(\lambda C(name(C, Mary))) \\ \quad \& \; B.ref(\lambda D(name(D, John)))), \\ \quad \lambda E \exists A : \{aspect(simple, ref(\lambda F past(F)), A)\} \\ \quad\quad E \; . \; \lambda I(\theta(A, object, I)) \\ \quad\quad\quad \& divorce(A) \& A \; is \; event \\ \quad\quad\quad \& in(A, ref(\lambda J(name(J, March))))) \end{array}$$

Figure 3: (1a) is 'about' what happened to John and Mary

$$\begin{array}{l} about(\lambda C(\lambda D(C \; . \; D \; \& \; in(D, ref(\lambda E(name(E, March)))))), \\ \quad \lambda B \exists A : \{aspect(simple, ref(\lambda F past(F)), A)\} \\ \quad\quad (B.\lambda I \theta(I, object, ref(\lambda J(name(J, John)))) \\ \quad\quad\quad \& \theta(I, object, ref(\lambda K(name(K, Mary)))) \\ \quad\quad\quad \& divorce(I) \& I \; is \; event).A) \end{array}$$

Figure 4: (1b) is 'about' what happened in March

We do this by noting, following (Halliday, 1985), that the leftmost phrasal daughter of an English sentence – its THEME – seems to be particularly significant. It is easy enough to spot which item is the leftmost phrasal daughter, so we can mark this item carefully when we are constructing our logical form. In fact all we have to do is to delay combining the meaning of the theme and the meaning of the rheme (everything else).

Having obtained such a logical form, however, we have to specify the meaning of 'about'. This comes in two parts: (i) we have to recover the normal propositional content – if what I tell you about John and Mary is that they got divorced then they did indeed get divorced – and (ii) we have to make use of the fact that one part of the utterance is marked as being more important.

The first part of this is easy, given that we are using a highly intensional logic. Fig. 5 just says that the rheme holds of the theme, which recovers the standard meaning (in fact it just carries out the λ-reduction step that we delayed when we constructed the logical form).

$$\forall X \forall Y about(X, Y) \rightarrow Y.X$$

Figure 5: If I say Y about X then Y is true of X

What do we do with the theme:rheme division itself?

One obvious function is to help link the current utterence to an earlier one. If, for instance, the current utterance is an answer to an earlier question then they should have the same rheme:

(3) Who got divorced in March? John and Mary got divorced in March.

(4) Who got divorced in March? Peter and Susan got divorced in April.

(4) doesn't make a sensible question:answer pair because the rheme of the answer is different from the rheme of the question. We can capture this constraint with Fig. 6[3].

$$\begin{array}{l} \forall U_i \forall U_{i+j} answer(U_i, U_{i+j}) \\ \quad\quad \rightarrow \forall R(rheme(U_i, R) \rightarrow rheme(U_{i+j}, R)) \end{array}$$

Figure 6: An answer should address the *rheme* of the question

Similarly, if the current utterance is a narrative continuation of some previous utterance then the two are likely to share the same theme:

$$\begin{array}{l} \forall U_i \forall U_{i+j} narrative(U_i, U_{i+j}) \\ \quad\quad \rightarrow \forall T(theme(U_i, T) \rightarrow theme(U_{i+j}, T)) \end{array}$$

Figure 7: The theme stays the same as a story unfolds

Thus the theme:rheme distinction can be used to look for connections between utterances of the kind described in RST (Mann and Thompson, 1988; Mann, 1999). It is also worth noting that objects

[3]Note that this rule can also be used to *reconstruct* the rheme of an elliptical answer.

referred to within the theme are likely to be particularly salient in the discourse, so that they are good candidates for when you are trying to dereference pronouns[4].

2.2 Focus

The difference between (2a) and (2b) is in some ways similar. Again we have individual items picked out, the difference being that here we use intonation or typography rather than dislocation. Again the marked items are in some way interesting, and again we have to be able to reconstruct the basic propositional content.

(2) a. Mary and John got divorced in *March*.
b. Mary and John got *divorced* in March.

Isolating the semantics of the focussed item is a slightly more complicated process this time, but the resulting analyses are fairly similar:

$$
\begin{array}{l}
focus(\lambda B(name(B, March)),\\
\quad \lambda C \exists A : \{aspect(simple, ref(\lambda D past(D)), A)\}\\
\qquad \theta(A, object, ref(\lambda G(name(G, Mary))))\\
\qquad \&\theta(A, object, ref(\lambda H(name(H, John))))\\
\qquad \&divorce(A) \& A \ is \ event\\
\qquad \&in(A, ref(\lambda I(C \ . \ I))))
\end{array}
$$

Figure 8: They got divorced in March, not in April

$$
\begin{array}{l}
focus(\lambda B(divorce(B) \& \ B \ is \ event\\
\qquad \& \ in(B, ref(\lambda C(name(C, March))))),\\
\quad \lambda D \exists A : \{aspect(simple, ref(\lambda E past(E)), A)\}\\
\qquad \theta(A, object, ref(\lambda H(name(H, Mary))))\\
\qquad \&\theta(A, object, ref(\lambda I(name(I, John))))\\
\qquad \& D \ . \ A)
\end{array}
$$

Figure 9: They got divorced, not married, in March

As before, we have a rule for reconstructing the standard propositional content by combining the two parts that have been left unreduced, so that each of (2a) and (2b) entails that they did get divorced in March:

$$\forall X \forall Y focus(X, Y) \rightarrow Y.X$$

Figure 10: If I say Y about X then Y is true of X

The extra information carried by the fact that one item is put in focus has to be consumed by some other operator. In cases like (2a) and (2b) where there is no overt discourse operator, we assume that the given sentence is being contrasted with some proposition which is already present in the discourse, and which is now being denied. We therefore have the rule in Fig. 11, which picks out, via the referential term, some item which is currently believed to satisfy Q and claims that this is not in fact the case. Note the explicit reference to the slate Σ in this rule: the current utterance is being contrasted with some other salient proposition which is entailed by the slate.

$$
\begin{array}{l}
\forall P \forall Q(focus(P, Q)\\
\qquad \rightarrow \neg(Q.ref(\lambda P'(\Sigma \rightarrow (Q.P')))))
\end{array}
$$

Figure 11: Q does not hold of P'

The act of focussing one element of the sentence, however, does not always indicate a direct contrast with some proposition in the discourse. Its task is to split the propositional content into two parts. In many cases, such as (5) and (6), there is an explicit lexical item which requires two such arguments. The assumption that there is a contrast should only be made when there is no such lexical item.

(5) a. I only borrowed your *bike*
b. I only *borrowed* it.

(6) a. Susan even kissed *Peter*.
b. She even *hugged* him.

The logical forms for (5a) and (6b) are given in Fig. 12 and Fig. 13: the corresponding pair are similar.

$$
\begin{array}{l}
only(\lambda F(bike(F)),\\
\quad \lambda G \exists H : \{H \ is \ interval \ \& \ past(H)\}\\
\qquad \exists E : \{aspect(simple, H, E)\}\\
\qquad\quad \theta(E, agent, ref(\lambda J speaker(J)))\\
\qquad\quad \&borrow(E)\\
\qquad\quad \& E \ is \ event\\
\qquad\quad \&\theta(E,\\
\qquad\qquad object,\\
\qquad\qquad ref(\lambda L(G.L\\
\qquad\qquad\qquad \&of(L, ref(\lambda M(hearer(M))))))))
\end{array}
$$

Figure 12: (5a): what I borrowed was your bike, not something else of yours

$$
\begin{array}{l}
even(\lambda E(hug(E) \ \& \ E \ is \ event),\\
\quad \lambda F \exists G : \{G \ is \ interval \ \& \ past(G)\}\\
\qquad \exists D : \{aspect(simple, G, D)\}\\
\qquad\quad \theta(D, agent, ref(\lambda I f(I)))\\
\qquad\quad \&F.D\\
\qquad\quad \&\theta(D,\\
\qquad\qquad object,\\
\qquad\qquad ref(\lambda K m(K))))
\end{array}
$$

Figure 13: (6b): she hugged him, which was the last thing anyone expected her to do

As usual, we have to say what follows from these operators. They are, as before, truth preserving, so

[4]The relationship between the theme:rheme distinction and the phenomena investigated in CENTERING THEORY (Joshi and Weinstein, 1998) is rather unclear: what is clear is that devices for indicating which elements of an utterance are particularly prominent carries a great deal of information that can be used for organising an extended dialogue.

we need the rule in Fig. 14. *'only'* again contrasts the current proposition with something which is already present in the discourse, with the extra constraint that this proposition should be in some sense 'stronger' [5].

$$\forall A \forall B(only(A,B) \rightarrow B.A)$$

$$\forall A \forall B(even(A,B) \rightarrow B.A)$$

Figure 14: *'even'* and *'only'* are truth preserving

$$\forall A \forall B(only(A,B) \rightarrow \neg B.ref(\lambda A'(A' \gg A \& B.A'))$$

Figure 15: B doesn't hold of the 'stronger' item A' which you thought it did

The effect of *'even'* is to emphasise how unlikely the reported event is: we capture this with Fig. 16, which notes that there is some similar but more probable proposition:

$$\forall A \forall B(even(A,B) \rightarrow \exists A'(A' \approx A \& prob(B.A') > prob(B.A))$$

Figure 16: There is some similar A' for which $B.A'$ is more likely than $B.A$

Note that *'even'* and *'only'* are components of the actual utterance, and hence would naturally be dealt with in its logical form, whereas the implied contrastive stress in (2) is a relationship *between* utterances. Since the device used to mark the focussed elements is the same in all these cases, it is highly desirable to provide a uniform treatment by including the constrastive stress in the logical form as well.

2.3 Mood

The theme:rheme distinction and the use of marked stress, then, enrich the propositional content of an utterance by linking it to the surrounding discourse and by expressing complex relations to various components. By capturing these facets of the meaning inside the logical form, we obtain a smooth connection between the two.

[5] You can reasonably say *'You think I stole it, but actually I only* ***borrowed*** *it'*, whereas *'You think I borrowed it, but actually I only* ***stole*** *it'* sounds strange: the notion of relative strength here is similar to that in (Gazdar, 1979)'s use of 'expression alternatives' for dealing with implicature, and to (Kruijff-Korbayová and Webber, 2001)'s notion of 'alternative sets'

It is also clear that choosing the way the words are arranged allows you to express different attitudes to the truth of the proposition as a whole:

(7) Did John and Mary get divorced in March?

(8) Get divorced in March.

The basic event type depicted by (7) is the same as the one depicted by (2) and (1), and the one depicted by (8) is clearly closely related. It is therefore worth seeing whether we can capture the attitudes underlying these two sentences within the logical form as well.

As ever, we can easily include a term in the logical form which corresponds to the specified mood, just on the basis of the surface form:

$$\begin{array}{l} claim(\exists A : \{aspect(simple, ref(\lambda B past(B)), A)\} \\ \quad \theta(A, object, ref(\lambda E(name(E, Mary)))) \\ \quad \& \theta(A, object, ref(\lambda F(name(F, John)))) \\ \quad \& divorce(A) \& A \text{ is event} \\ \quad \& in(A, ref(\lambda G(name(G, March))))) \end{array}$$

Figure 17: (1a): I'm telling you it's true

$$\begin{array}{l} query(\exists A : \{A \text{ is interval} \\ \quad \& past(A)\} \\ \quad \exists E : \{aspect(simple, A, E)\} \\ \quad \quad \theta(E, object, ref(\lambda F(name(F, Mary)))) \\ \quad \quad \& \theta(E, object, ref(\lambda G(name(G, John)))) \\ \quad \quad \& divorce(E) \& E \text{ is event} \\ \quad \quad \& in(E, ref(\lambda H(name(H, March))))) \end{array}$$

Figure 18: (7): I don't know whether it's true

$$\begin{array}{l} useful(\exists C \theta(C, object, ref(\lambda Dhearer(D))) \\ \quad \& divorce(C) \\ \quad \& C \text{ is event} \\ \quad \& in(C, ref(\lambda F(name(F, March))))) \end{array}$$

Figure 19: (8): I'd like you to make it true

And as ever, having included such terms we have to devise rules to account for their effects.

Most treatments of mood follow (Austin, 1962; Searle, 1969) in assuming that there is a connection between the overt 'force' of the utterance and the speaker's underlying goals, with a substantial tradition of AI work linking this to AI planning theory (Cohen and Perrault, 1979; Cohen and Levesque, 1980; Cohen et al., 1990; Allen and Perrault, 1980; Bunt and Black, 2000). We will argue that this connection has to be made explicit in the logical form, rather than being dealt with as a separate phenomenon.

Consider (9):

(9) a. John fancies every woman he sees.

b. John fancies any woman he sees.

There doesn't seem to be much difference between (9a) and (9b). They both seem to claim that if X is a woman and John sees X then John fancies X. In other words, *'any'* and *'every'* both seem to be universal quantifiers.

In (10), however, the effects of *'every'* and *'any'* are very different:

(10) a. You can't invite John. He'll drink everything.

b. You can invite John. He'll drink anything.

The difference between (10a) and (10b) is that (10a) says that there will be a future state of affairs where *'John drinks everything'* is true, whereas (10b) says that for any drink D there is a future state of affairs where *'John drinks D'* is true.

The fact that *'any'* has very wide scope is also reflected in the interpretation of negative sentences like *'I didn't get anything I wanted for Christmas'*, which can be paraphrased as saying that if X is something I wanted then I didn't get X for Christmas. Just how wide its scope is, however, becomes clear when we consider (11) and (12):

(11) a. Do you know everyone at the party?

b. Do you know anyone at the party?

(12) a. Take any lane for Macclesfield.

b. Take every lane for Macclesfield.

Fig. 20 says that there is something I would like to do, but that I can't do it unless I have evidence that you know someone at the party. What kind of thing might I be able to do under those circumstances? If I'm the host then I might go and look for them so that you have someone to talk to, if I'm a fellow guest then I might ask you to introduce them to me, ...In general, however, when you ask a question it's because there is something that you could do if you knew the answer. The key point about Fig. 20 is that it says that it doesn't matter which person you know: for every one who is at the party, if you know them then I can do whatever it is I have in mind.

We thus encode the notion that *'any'* gives you a 'free choice' (Vendler, 1967) within a perfectly standard logic. It is still just a universal quantifier, but it has *very* wide scope – wide enough to be quantifying over questions that I would like the answer to, rather than as part of the queried proposition as would happen with (11a), as in Fig. 21.

Similar considerations apply to (12a) and (12b). (12a) says that whichever lane you choose will be appropriate for getting to Macclesfield. (12b)

$$
\begin{array}{l}
\exists C : \{intend(speaker, do(speaker, C))\} \\
\quad \forall D : \{one(D) \\
\qquad \& \; at(D, ref(\lambda E(party(E))))\} \\
\qquad \exists F : \{F \; is \; interval \; \& \; past(F)\} \\
\qquad\quad \exists H : \{aspect(simple, F, H)\} \\
\qquad\qquad \theta(H, agent, ref(\lambda I(hearer(I)))) \\
\qquad\qquad \& know(H) \\
\qquad\qquad \& H \; is \; event \\
\qquad\qquad \& \theta(H, object, D) \\
\qquad \rightarrow \; do(speaker, C)
\end{array}
$$

Figure 20: There's something I could do if I knew that you knew someone at the party

$$
\begin{array}{l}
\exists C : \{intend(speaker, do(speaker, C))\} \\
\quad \exists D : \{D \; is \; interval \; \& \; past(D)\} \\
\qquad \forall F : \{one(F) \; \& \; at(F, ref(\lambda G(party(G))))\} \\
\qquad\quad \exists H : \{aspect(simple, D, H)\} \\
\qquad\qquad \theta(H, agent, ref(\lambda I(hearer(I)))) \\
\qquad\qquad \& know(H) \& H \; is \; event \\
\qquad\qquad \& \theta(H, object, F) \\
\quad \rightarrow \; do(speaker, C)
\end{array}
$$

Figure 21: There's something I could do if I knew that you knew everyone at the party

doesn't really make any sense, since it seems to require you to drive in several lanes at once. We therefore paraphrase (12a) as in Fig. 22.

$$
\begin{array}{l}
\forall B : \{lane(B)\} \\
\quad useful(\exists A take(A) \\
\qquad\qquad \& A \; is \; event \\
\qquad\qquad \& \theta(A, object, B) \\
\qquad\qquad \& for(A, \\
\qquad\qquad\quad ref(\lambda C(name(C, Macc)))) \\
\qquad\qquad \& \theta(A, \\
\qquad\qquad\quad agent, \\
\qquad\qquad\quad ref(\lambda D speaker(D))))
\end{array}
$$

Figure 22: For any lane, it makes sense to take it if you want to go to Macclesfield

The wide scope of *'any'* allows us to interpret (12a) as saying that there are several things that would be useful in the current situation, namely all the propositions that result from choosing a lane. We have weakened the force of the imperative to *useful* because, as is clear from the example, not all commands actually relate to things the speaker wants. As always, however, having introduced the term *useful* into our logical form we are under an obligation to provide an account of what it signifies.

$$
\begin{array}{l}
\forall P(useful(P) \\
\quad \rightarrow \exists G : \{\forall S : \{speaker(S)\} want(S, G) \\
\qquad\qquad \vee \forall H : \{hearer(H)\} want(H, G)\} \\
\qquad\qquad P \rightarrow G
\end{array}
$$

Figure 23: Something is useful if it will help either the hearer or speaker achieve some goal

G in Fig. 23 will usually be something that the speaker wants, but in examples like (12a) it may be

the hearer's underlying goal that will be satisfied by the specified action. Note that the logical form for (12a) said that certain kinds of situations would be useful. Fig. 23 then says that if a situation that is described by the proposition P is useful then it must be the case that G will be achievable in any situations of this general type.

It is very hard indeed to see how the interactions between *'any'* and the various moods can be captured other than by including the mood in the logical form. (11b) involves quantifying over questions whose answer would satisfy my needs, (12a) involves quantifying over useful/desirable actions. If you try to separate the logical form and the mood you simply cannot retrieve this interaction. By including the mood inside the logical form, we make it possible to account for the discourse effects of *'any'*, and we continue to work with a unified two-stage framework – anchor the logical form and then think about its consequenes. This is particularly significant when considering the interactions between negation, mood and quantification. Consider (13):

(13) Don't touch anything.

We suggest the logical form in Fig. 24 for this.

$$\begin{array}{l}\forall C : \{thing(C)\} \\ \quad useful(\neg(\exists D\theta(D, agent, ref(\lambda Espeaker(E))) \\ \qquad\qquad \& touch(D) \\ \qquad\qquad \& D\ is\ event \\ \qquad\qquad \& \theta(D, object, C)))\end{array}$$

Figure 24: For each thing C it would be helpful if you didn't touch C

Ensuring the right relations between the universal quantifier, the negation and *useful* would be extemely difficult if they were not dealt with at the same time and in the same place.

3 Conclusions

We have show that it is possible to capture various aspects of the speaker's attitude to what he is saying inside the logical form. We believe that it is also extremely convenient to do so. If you do not deal with information structure and intonation in the logical form, you will find it extremely difficult to localise their effects and coordinate them with the propositional content, since they dfeal with local elements of the propositional content (and in the case of focus, they deal with *arbitrary* elements, so that you cannot rely on them to pick out discourse referents or other simple entities). If you do not deal with mood in the logical form you will simply find it impossible to cope with the interactions between illocutionary force and the quantifier *'any'* (so that you will not, for instance, be able to react appropriately to a question such as *'Do I have any messages about Viagra?'* oe a command like *'Don't delete any emails from John'*). We deal with these cases by constructing models which reveal the critical issues. The logical form contains all the information encoded by the surface form. The inference engine unpacks this information in a way that makes it possible to plan appropriate responses.

Computing logical forms of the kind shown above can be done using the standard compositional techniques. All the logical forms in this paper were obtained in this way, using the parser described in (Ramsay and Seville, 2000b). You have to make use of standard rescoping algorithms (van Eijck and Alshawi, 1992; Milward and Cooper, 1994) to ensure that the relations between the various discourse operators are handled correctly, but you need this anyway for handling scopes.

Carrying out the required inference is harder. In (Ramsay and Seville, 2000a) we showed how to construct models by combining the literal content of an utterance, as encoded in the logical form, with a rich collection of background knowledge. To do this with the meaning postulates given above for discourse operators we need to be able to carry out this activity with meaning postulates couched in a higher-order logic. The theorem prover described in (?) allows us to do exactly that.

References

P Aczel. 1988. *Non-Well-Founded-Sets.* CSLI Publications, Stanford.

J F Allen and C R Perrault. 1980. Analysing intention in utterances. *Artificial Intelligence*, 15:148–178.

J Austin. 1962. *How to Do Things with Words.* Oxford University Press, Oxford.

J Barwise and J Perry. 1983. *Situations and Attitudes.* Bradford Books, Cambridge, MA.

P Baumgartner and M Kühn. 2000. Abducing coreference by model construction. *Journal of Language and Computation*, 1(2).

G Bealer. 1989. Fine-grained type-free intensionality. In G Chierchia, B H Partee, and R Turner, editors, *Properties, types and meaning: vol I, foundational issues.* Kluwer Academic Publishers, Dordrecht/Boston/London.

P Blackburn, J Bos, M Kohlhase, and H de Nivelle. 1997. Inference and computational semantics. In

H Bunt and E Thijsse, editors, *3rd International Workshop on Computational Semantics*, pages 5–19, University of Tilburg.

H Bunt and W J Black, editors. 2000. *Abduction, Beliefs and Context: Studies in Computational Pragmatics*. John Benjamins, Amsterdam/Philadelphia.

P R Cohen and H Levesque. 1980. Speech acts and the recognition of shared plans. In *Proceedings, Canadian Society for Computational Studies of Intelligence*, pages 263–270.

P R Cohen and C R Perrault. 1979. Elements of a plan-based theory of speech acts. *Cognitive Science*, 7(2):171–190.

P R Cohen, J Morgan, and M E Pollack. 1990. *Intentions in Communication*. Bradford Books, Cambridge, Mass.

D R Dowty. 1989. On the semantic content of the notion of 'thematic role'. In G Chierchia, B H Partee, and R Turner, editors, *Properties, Types and Meaning II: Semantic Issues*, pages 69–130, Dordrecht. Kluwer Academic Press.

D R Dowty. 1991. Thematic proto-roles and argument selection. *Language*, 67:547–619.

C Gardent and K Konrad. 2000. Interpreting definites using model generation. *Journal of Language and Computation*, 1(2):215–230.

H Gaylard and A M Ramsay. 2002. A unified theory of reference resolution. In *4th Discourse and Dialogue Colloquium*, Lisbon.

G Gazdar. 1979. *Pragmatics: Implicature, Presupposition and Logical Form*. Academic Press, New York.

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

A Joshi and S Weinstein. 1998. Formal systems for complexity and control of inference: A reprise and some hints. In M. A. Walker, A. K. Joshi, and E. F. Prince, editors, *Centering Theory in Discourse*, pages 31–38. Oxford University Press, Oxford.

H Kamp. 1984. A theory of truth and semantic representation. In J A G Groenendijk, T M V Janssen, and M B J Stokhof, editors, *Formal Methods in the Study of Language*, pages 277–322, Dordrecht. Foris Publications.

I Kruijff-Korbayová and B L Webber. 2001. Concession, implicature and alternative sets. In H Bunt, I van der Sluis, and E Thijsse, editors, *4th International Workshop on Computational Semantics*, pages 227–248, University of Tilburg.

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

W C Mann. 1999. An introduction to rhetorical structure theory.

D Milward and R Cooper. 1994. Incremental interpretation: Applications, theory and relationship to dynamic semantics. In *Proceedings of the 15th International Conference on Computational Linguistics (COLING-94)*, pages 88748–754, Kyoto.

S G Pulman. 1993. Higher order unification and the interpretation of focus. *Linguistics & Philosophy*, 20(1):73–115.

A M Ramsay and H Seville. 2000a. Models and discourse models. *Journal of Language and Computation*, 1(2):167–181.

A M Ramsay and H Seville. 2000b. Unscrambling English word order. In M Kay, editor, *Proceedings of the 18th International Conference on Computational Linguistics (COLING-2000)*, pages 656–662, Universität des Saarlandes, July.

A M Ramsay. 1999. Does it make any sense? updating = consistency checking. In K Turner, editor, *The Semantics-Pragmatics Interface from Different Points of View*, London and Amsterdam. Elsevier.

A M Ramsay. 2001. Theorem proving for untyped constructive λ-calculus: implementation and application. *Logic Journal of the Interest Group in Pure and Applied Logics*, 9(1):89–106.

J R Searle. 1969. *Speech Acts: an Essay in the Philosophy of Language*. Cambridge University Press, Cambridge.

R Turner. 1987. A theory of properties. *Journal of Symbolic Logic*, 52(2):455–472.

R van der Sandt. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics*, 9:333–377.

J van Eijck and H Alshawi. 1992. Logical forms. In H Alshawi, editor, *The Core Language Engine*, pages 11–40, Cambridge, Mass. Bradford Books/MIT Press.

Z. Vendler. 1967. *Philosophy in Linguistics*. Cornell University Press, Ithaca.

J Wedekind. 1996. On inference-based procedures for lexical disambiguation. In *Proceedings of the 16th International Conference on Computational Linguistics (COLING-96)*, pages 980–985, Copenhagen.