Weak Lexical Semantics and Multiple Views

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

Keywords: lexical semantics, weak lexical items

1. The Role of Inference in Interpretation

Most formal/computational treatments of natural language semantics discuss the production of formal paraphrases in some suitable language. There is some debate as to whether this can be done strictly compositionally, e.g. by function application and composition, as proposed in Montague grammar and its descendants (Dowty et al., 1981; Kohlhase et al., 1996), or by simple structure sharing (Pollard and Sag, 1988; Pollard and Sag, 1994; Fenstad et al., 1987); or whether some more indirect construction process is involved (Dalrymple et al., 1996; van Genabith and Crouch, 1997). In nearly every case, however, the goal of the process is a formal paraphrase that 'means the same as the original'.

This belief that there is a single end-product that constitutes the result of the semantic analysis has a number of undesirable consequences. (i) It requires the use of a formal language with all the interesting properties of natural language. In particular, any such language must be 'dynamic' – it must express some relationship between situations, or states of affairs, or discourse states, or ... (ii) The formal paraphrase must contain all and only the information that is carried by the sentence. This will mean that the results of your analyses will be inflexible and overly precise, and considerable amounts of work will be required in order to choose between alternative 'ambiguous' readings.

I have shown elsewhere that the dynamic aspects of meaning can be dealt with using an orthodox 'static' language if you allow utterances to include statements about the shared and private knowledge of the speaker and hearer and if you further think about the reasoning that the discourse participants will perform when considering a given utterance (Ramsay, 1999). I will argue in the current paper that the same notions will support much subtler and more parsimonious accounts of lexical semantics. The key notions underlying this analysis are as follow:

- comprehending an utterance involves building a model of the world which admits the information it encodes.
- considerable flexibility can be achieved by considering multiple 'views' of a given entity.

The discussion below will concentrate on lexical items with particularly 'light' semantics, i.e. things like tense and aspect markers and prepositions which appear to mean different things when used in (linguistically) different contexts. Such items are notoriously difficult to characterise. They are not exactly 'ambiguous' – you wouldn't want to say that the present participle markers in

(1) He is sleeping.

(2) He is hiccupping.

are different items; but they do produce different effects – (1) denotes a single unfinished sleeping event, (2) denotes an unfinished *sequence* of hiccups. I have argued elsewhere (Ramsay, 1996) that the two effects can be seen as the result of combining a very sparse description of the 'meaning' of the marker *'-ing'* with appropriate descriptions of the temporal properties of *'sleep'* and *'eat'*. This contrasts with (Moëns and Steedman, 1988), where operators which *change* the meaning of the aspect markers when they are combined with verbs with specific temporal properties are invoked. The current paper extends this notion to cover other items with the same general properties.

2. Meaning Postulates

Consider the sentence

(3) A dead dog ordered a drink.

A formal paraphrase of this might look like

$$
\begin{aligned}
&\exists A :: \{dead(A, \lambda B dog(B))\} \\
&\quad \exists C \exists D :: \{D < now\} \\
&\qquad aspect(perfective, D, C) \\
&\qquad \wedge \forall E :: \{C.E\} \\
&\qquad\quad type(E, [entity, abstract, \ldots, order]) \\
&\qquad\quad \wedge\ \theta(E, object, \lambda F \exists G :: \{drink(G)\} F.G) \\
&\qquad\quad \wedge\ \theta(E, agent, A)
\end{aligned}
$$

This paraphrase has several desirable properties. It encodes the temporal properties of the reported event set; it marks the relationships between this event set and its participants; it treats the adjective as a relationship between an individual and a class (thus leaving space for treatments of non-intersective and non-factive adjectives); it uses a notion of type which makes it possible to compare types very directly, by taking a type descriptor to be a path through a predefined type hierarchy, so that two types T_1 and T_2 can be seen to be compatible if one of them is a prefix of the other; and it makes the object of the event some abstraction over situations in which a drink is present, thus avoiding any commitment to the existence of a particular drink.

But what it doesn't do is tell you what (3) actually *means*. Or at least, it doesn't tell you how someone who utters (3) might see the world. It doesn't, for instance, tell you that there is something very odd about the speaker's view of the world, which there must be since dead dogs don't go around performing volitional acts.

Faced with a sentence like (3), any attempt to give a formal account of its significance must involve constructing a formal paraphrase like the one given above. The problem is that unless you back such paraphrases up with some indication of how the various terms are linked to other concepts and with an engine for drawing out the consequences of these links then you are not going to be able to say anything about the speaker's view of the world. The approach I will take in this paper involves acknowledging that I have a responsibility to specify the connections between words and concepts to the best of my ability; and using an adaptation of a standard theorem prover in order to construct concrete models that are compatible with what has been said.

As far as specifying the connections between words and concepts is concerned, you just have to do it. You may be able to extract some of what you want automatically from corpora or other lexical resources, though most attempts to do so only extract rather superficial semantic features. Wherever you get them from, however, you need to write them down in some appropriate form. A reasonably interesting one, relating to the example above, is

$$\begin{aligned}
\forall A \; &:: \; \{type(A, [entity, abstract, situation, event, order])\} \\
&\quad \forall B \; :: \; \{\theta(A, agent, B)\} \\
&\quad\quad \forall C \; :: \; \{\theta(A, object, C)\} want(B, [C.\lambda D have(B, D)])
\end{aligned}$$

This says that if B is the agent of some ordering event where C describes what was ordered, then B wanted a situation in which he or she had something of the specified kind. Suppose we consider a more reasonable sentence, say,

(4) A man ordered a drink.

with the following paraphrase:

$$
\begin{array}{l}
\exists A :: \{man(A)\} \\
\quad \exists B \exists C :: \{C < now\} \\
\quad\quad aspect(perfective, C, B) \\
\quad\quad \wedge\ \forall D :: \{B.D\} \\
\quad\quad\quad type(D, [entity, abstract, \ldots, order]) \\
\quad\quad\quad \wedge\ \theta(D, object, \lambda E \exists F :: \{drink(F)\} E.F) \\
\quad\quad\quad \wedge\ \theta(D, agent, A)
\end{array}
$$

Given this interpretation and the above rule concerning ordering events, we would want to know that this man wanted a situation where he had a drink. We can obtain a picture of the world that reveals this fact as follows.

1. Ask some theorem prover whether it can prove that the formal paraphrase is inconsistent with whatever background knowledge has been provided.

2. a) If it can, report that this sentence makes no sense in the current context.

   b) If, more interestingly, it can't, then inspect the state of the theorem prover at the point when the proof failed. For any tableau-like theorem prover, failure to prove inconsistency amounts to having an open branch: if you inspect such a branch you should be able to extract a set of positive and negative literals which provide a partial model of the world which is consistent with the starting point.

This approach to model construction resembles the processes described by (Johnson-Laird, 1983). Johnson-Laird argues that when people are asked to draw out the consequences of pairs of premises such as *'All bakers are candlestick-makers'* and *'No athletes are bakers'*, they typically explore the models that are consistent with what they have been told in order to see if there are any commonalities between them. Johnson-Laird suggests that people perform less reliably on this task if there are large numbers of potential models, since they are likely to fail to spot them, or to fail to inspect them all properly when checking out hypotheses about what they have in common.

The task that Johnson-Laird sets his subjects is rather unnatural. Ordinary people (i.e. not logicians!) do not typically spend a large amount of time exploring the consequences of pairs of premises. The task I am describing, of trying to find at least one model of a set of premises, seems more common. Given that someone has said something to you, you want to construct a view of the world which is compatible with what they said. Exploring a tableau in the way described above does seem like a reasonable, if idealised, computational analogue of the mechanisms that Johnson-Laird believes his subjects employ when performing the tasks he sets them.

I use a version of Manthey & Bry's (1988) model generation theorem prover, adapted to work with the intensional logic described by (Turner, 1987), to perform this task. For details of this theorem prover, see (Ramsay, 1995; Cryan and Ramsay, 1997). Given the example above, this process produces the model in Fig. 1. Note that the model contains the fact that what #2 wants is for the proposition $\exists F :: \{drink(F)\} have(\#2, F)$ to be true. This emerges from combining the intensional reading of *'a drink'* as $\lambda E \exists F :: \{drink(F)\}\ E.F)$ with the meaning postulate for *'order'*.

$$
\begin{array}{l}
type(\#1, [entity, abstract, situation, event, order]) \\
\theta(\#1, object, \lambda D \exists E :: \{drink(E)\} D.E) \\
\theta(\#1, agent, \#2) \\
man(\#2) \\
human(\#2) \\
male(\#2) \\
alive(\#2) \\
type(\#2, [entity, concrete, solid, living]) \\
intended(\#2, \#1) \\
want(\#2, [\exists F :: \{drink(F)\} have(\#2, F)])
\end{array}
$$

Figure 1. Model for *'A man ordered a drink'*

This technique of backing up your formal paraphrases with appropriate 'meaning postulates'[1], and using these in order to construct concrete models of what has been said, is particularly fruitful for terms such as tense and aspect markers which seem to pick up a great deal of their significance from the linguistic context. In the remainder of this paper I will consider the use of this notion for giving precise semantics

[1] See (Cruse, 1986) for a discussion of the relationship between meaning postulates and sets of necessary and sufficient conditions.

to prepositions like *'in'* which can be used to link items of different kinds with radically different consequences.

3. Prepositions and Multiple-views

Consider the following three sentences:

(5) A man reads a book in a park.

(6) A man reads a story in a book.

(7) A man reads a story in a park.

Each of these has an interpretation under which the prepositional phrase modifies the VP rather than the final NN (so that (5–7) are closely related to the questions *'Where did a man read a book/story?'*).

Under these interpretations, it seems reasonable to say that (5) supports the inference that the book was also in the park, and that (6) supports the inference that the story was in the book, but that no corresponding conclusion follows from (7). Intuitively we feel that the object of a reading event will generally be in the same place as the event itself; but that stories and parks are of such different types that it doesn't make much sense to talk of a story being in a park (that is not to say that a story can't be *set* in a park; but that's a different question).

To cope with this, I introduce the idea that it is possible to have multiple views of the same entity. So a reading event can be seen as a physical entity, involving some physical object with marks on it; as a mental entity, involving some abstract object with a 'propositional content'; as a temporal entity, entering into relationships with other temporal entities; and maybe in other ways too. Thus the interpretation of a simple sentence such as

(8) A man reads a book.

leads to the construction of the model in Fig. 2.

The key item in this model is the event #2. The first entry in the model says that #2 is an event of a particular kind. The next three statements introduce different views of this event. #11, for instance, is the temporal view of the event: as such, it is capable of entering into relationships with other temporal entities, so that the model contains an instant #8 which is the start point of #11 and which precedes the reference instant *now*, and likewise for the other temporal entities in

$type(\#2, [entity, abstract, situation, event, read])$
$view(\#2, \#5)$
$view(\#2, \#6)$
$view(\#2, \#11)$
$\theta(\#2, object, \#3)$
$\theta(\#2, agent, \#4)$
$book(\#3)$
$type(\#3, [entity, abstract, mental])$
$view(\#3, \#7)$
$envelope(\#3, \#10)$
$man(\#4)$
$human(\#4)$
$male(\#4)$
$type(\#4, [entity, concrete, solid, living])$
$type(\#5, [entity, concrete])$
$type(\#6, [entity, abstract, mental])$
$type(\#7, [entity, concrete, solid])$
$type(\#8, [entity, abstract, time, instant])$
$\#8 < now$
$start(\#8, \#11)$
$type(\#9, [entity, abstract, time, instant])$
$end(\#9, \#11)$
$now < \#9$
$type(\#10, [entity, abstract, mental])$
$type(\#11, [entity, abstract, time])$

Figure 2. Model for *'A man reads a book'*

the model. #5, on the other hand, is the view of the event as a physical entity, and #6 is its view as an abstract mental activity.

These different items can each enter into appropriate relationships. In particular, we can capture the fact that the object of a reading event can generally be found at the location where the event took place with a restricted rule like the following:

$$\begin{aligned}\forall A \forall B \forall C \forall D \forall E\, sharedobjloc(A) \wedge view(A, D) \\ \wedge\ location(D, C) \wedge\ \theta(A, object, B) \\ \wedge\ view(B, E) \wedge E \approx C \\ \rightarrow\ location(E, C)\end{aligned}$$

$E \approx C$ here means that E and C should have compatible types in the sense noted above. So this rule says that if C is the location, from some point of view, of an event A that shares its location with its object, and E is a view of A's object, then C will also be E's location if they have compatible types.

We also need to characterise the meaning of *'in'*. The following rule says that if A is in B then if A and B have views with compatible types, say C and D, and B has an envelope E then C can be found by finding E:

$$\begin{aligned}\forall A \forall B\ &::\ \{in(A, B)\} \\ &\forall C\ ::\ \{view(A, C)\} \\ &\quad \forall D\ ::\ \{view(B, D)\} \\ &\quad\quad \forall E\ ::\ \{envelope(D, E)\} \\ &\quad\quad\quad E \approx C \rightarrow location(C, E)\end{aligned}$$

This rule just seems to replace one undefined term, *in*, by another, *envelope*. We seem to be no further forward. Note, however, that the envelope is associated with a particular view of the ground B. This enables us to see what *'in'* means when it is used with any kind of entity that can be seen as having an envelope – physical spaces, which have simple physical envelopes, temporal entities such as intervals whose envelopes are something like open subsets of the real line, sets of entities which are presumably their own envelopes, and so on.

This rather object-oriented approach makes it possible to give a word like *'in'* a very flexible interpretation, without making it ambiguous and without appealing to metaphorical relations. The word *'in'* picks out the envelope of the ground as the place where you will find the figure. If the figure and the ground have views with compatible types and you know what the envelope of the relevant view of the ground is like, you know what follows from saying that the figure is in the ground.

Of course the notion of envelope itself may have historical or developmental roots in the spatial notion, as suggested by e.g. (Jackendoff, 1983). Nonetheless, it is not self-evident that the spatial notion is particularly central to everyday use of the word *'in'*. In, for instance, a collection of 492 sentences from the 8000 or so sentences containing this word from the British National Corpus, just 91 involved items that could be seen as physical locations. The others included many phrasal verbs and special cases such as *'in particular'*, but even ignoring these the proportion of cases where *'in'* introduces a physical relation between an object and physical space is not all that impressive (*'in'* a period of time, or *'in'* a collection, have comparable numbers of occurrences).

A man reads a book in a park	A man reads a story in a book	A man reads a book in January
$type(\#2,[\ldots,read])$ $in(\#2,\#3)$ $view(\#2,\#6)$ $\theta(\#2,object,\#4)$ $\theta(\#2,agent,\#5)$ $park(\#3)$ $type(\#3,[\ldots,concrete])$ $envelope(\#3,\#12)$ $book(\#4)$ $view(\#4,\#8)$ $man(\#5)$ $human(\#5)$ $male(\#5)$ $type(\#5,[\ldots,living])$ $type(\#6,[\ldots,concrete])$ $location(\#6,\#12)$ $type(\#8,[\ldots,concrete])$ $location(\#8,\#12)$ $type(\#12,[\ldots,concrete])$	$type(\#2,[\ldots,read])$ $in(\#2,\#3)$ $view(\#2,\#7)$ $\theta(\#2,object,\#4)$ $\theta(\#2,agent,\#5)$ $book(\#3)$ $type(\#3,[\ldots,mental])$ $envelope(\#3,\#11)$ $story(\#4)$ $type(\#4,[\ldots,mental])$ $location(\#4,\#11)$ $man(\#5)$ $human(\#5)$ $male(\#5)$ $type(\#5,[\ldots,living])$ $type(\#7,[\ldots,mental])$ $location(\#7,\#11)$ $type(\#11,[\ldots,mental])$	$type(\#2,[\ldots,read])$ $view(\#2,\#13)$ $\theta(\#2,object,\#4)$ $\theta(\#2,agent,\#5)$ $in(\#2,\#3)$ $month(\#3)$ $name(\#3,January)$ $type(\#3,[\ldots,interval])$ $envelope(\#3,\#11)$ $book(\#4)$ $man(\#5)$ $type(\#9,[\ldots,instant])$ $\#9 < now$ $member(\#9,\#3)$ $member(\#9,\#13)$ $start(\#9,\#13)$ $type(\#10,[\ldots,instant])$ $now < \#10$ $end(\#10,\#13)$ $type(\#11,[\ldots,interval])$ $type(\#13,[\ldots,time])$ $location(\#13,\#11)$

Figure 3. Models for (5), (6) and (9)

The rules given above will produce the models in Fig. 3 for (5) and (6). What matters is that the model of *'A man reads a book in a story'* locates the physical (*concrete*) view of the book in the physical envelope of the park, whereas the model of *'A man reads a story in a book'* locates the story, which is a mental object, in the envelope of the mental view of the book. Various different kinds of entity will have their own envelopes. If you see the word *'in'*, you just have to decide what kind of envelope(s) the ground has and which of them are compatible with some view of the figure. As a final illustration, consider (9), which exemplifies the second commonest use of *'in'* in the extract of the BNC which I inspected:

(9) A man reads a book in January.

The third column in Fig. 3 shows the model obtained for this sentence. The relevant view of the event is the temporal one, which is located with respect to the envelope of the month called January. The fact that the envelope of some temporal entity is the location of the temporal view of the event will have the consequence that the start and end points of the event are included in January. This arises from the properties of temporal envelopes: it is at this point that the 'object-oriented' character of such notions pays off.

We can treat other items with light semantics in much the same way. Consider, for instance, the prepositions *'from'* and *'to'*. (Dowty, 1991) (p. 569) suggests that you can deal with (10) by viewing a driving event as a path, and then equating the locations of London and Paris with the start and end points of that path, by requiring the prepositions to pick out these two points. If we separate out (i) the view of the event as a path, (ii) the fact that paths have start and end points, and (iii) the function of *'from'* and *'to'* as picking out start and end points, then we can get a quite general analysis which deals with (10) exactly as suggested by Dowty, but which also deals with (11) by drawing on the fact that the temporal view of an event also has a start and end point.

(10) He drove from London to Paris.

(11) He drove from dawn to dusk.

(10) and (11) contain referring expressions. For the purposes of this paper I am simply treating these as existential quantifiers – there are, of course, numerous serious treatments of referring expressions, but using one of these (even the one I usually use) would simply obscure the issues under consideration. The formal paraphrase of (10) is therefore as follows, with the paraphrase of (11) too similar to be worth repeating.

$$\begin{aligned}
&\exists B :: \{named(B, Paris)\}\\
&\quad \exists C :: \{named(C, London)\}\\
&\quad\quad \exists D :: \{male(D)\}\\
&\quad\quad\quad \exists E \exists F :: \{F < now\}\\
&\quad\quad\quad\quad aspect(simple, F, E)\\
&\quad\quad\quad\quad \wedge\ \forall A :: \{E.A\}\\
&\quad\quad\quad\quad\quad to(A, B) \wedge from(A, C)\\
&\quad\quad\quad\quad\quad \wedge\ type(A, [event, \ldots, drive])\\
&\quad\quad\quad\quad\quad \wedge\ \theta(A, agent, D)
\end{aligned}$$

This paraphrase of (10), and the corresponding one for (11), produce the models in Fig. 4.

These models arise from the interactions between the paraphrases of (10) and (11) and the following rule about *'from'* (and a similar one for *'to'*):

He drove from London to Paris	He drove from dawn to dusk	It copies a disk to a file
$type(\#2,[\ldots,drive])$	$type(\#2,[\ldots,drive])$	$type(\#2,[\ldots,copy])$
$to(\#2,\#4)$	$to(\#2,\#4)$	$to(\#2,\#4)$
$from(\#2,\#5)$	$from(\#2,\#5)$	$view(\#2,\#11)$
$view(\#2,\#9)$	$view(\#2,\#9)$	$view(\#2,\#10)$
$view(\#2,\#10)$	$view(\#2,\#10)$	$\theta(\#2,object,\#5)$
$\theta(\#2,agent,\#6)$	$\theta(\#2,agent,\#6)$	$\theta(\#2,agent,\#3)$
$\#3 < now$	$\#3 < now$	$unsexed(\#3)$
$\#3 < \#8$	$\#3 < \#8$	$disk(\#4)$
$named(\#4,Paris)$	$dusk(\#4)$	$view(\#4,\#8)$
$type(\#4,[\ldots,place])$	$type(\#4,[\ldots,instant])$	$file(\#5)$
$named(\#5,London)$	$dawn(\#5)$	$view(\#5,\#9)$
$type(\#5,[\ldots,place])$	$type(\#5,[\ldots,instant])$	$type(\#6,[\ldots,instant])$
$male(\#6)$	$male(\#6)$	$start(\#6,\#10)$
$\#7 < \#3$	$\#7 < \#3$	$\#6 < now$
$type(\#7,[\ldots,instant])$	$type(\#7,[\ldots,instant])$	$member(\#6,\#10)$
$member(\#7,\#10)$	$start(\#7,\#10)$	$type(\#7,[\ldots,instant])$
$start(\#7,\#10)$	$member(\#7,\#10)$	$end(\#7,\#10)$
$type(\#8,[\ldots,instant])$	$location(\#7,\#5)$	$type(\#8,[\ldots,information])$
$end(\#8,\#10)$	$type(\#8,[\ldots,instant])$	$type(\#9,[\ldots,information])$
$path(\#9)$	$end(\#8,\#10)$	$type(\#10,[\ldots,time])$
$type(\#9,[\ldots,concrete])$	$location(\#8,\#4)$	$transfer(\#11)$
$type(\#10,[\ldots,time])$	$path(\#9)$	$type(\#11,[entity])$
$end(\#11,\#9)$	$type(\#9,[\ldots,concrete])$	$end(\#12,\#11)$
$location(\#11,\#4)$	$type(\#10,[\ldots,time])$	$location(\#12,\#8)$
$member(\#12,\#9)$	$end(\#11,\#9)$	$similar(\#13,\#9)$
$start(\#12,\#9)$	$start(\#12,\#9)$	$location(\#13,\#8)$
$location(\#12,\#5)$	$member(\#12,\#9)$	$start(\#14,\#11)$
		$member(\#14,\#11)$

Figure 4. Models for (10), (11) and (12)

$$\begin{array}{l}\forall A \forall B :: \{from(A,B)\}\\ \quad \forall C :: \{view(A,C)\}\\ \quad\quad \forall D :: \{view(B,D)\}\\ \quad\quad\quad \forall E :: \{start(E,C)\}\\ \quad\quad\quad\quad C \approx D \rightarrow location(E,D)\end{array}$$

This rule says that if the relationship *'from'* holds between the figure A and the ground B then if C and D are views of A and B and E is the start of C and C and D have compatible types then B and E are in the same place. In the model for (11) this 'place' is in fact an instant, but that's OK.

The final example is also drawn from (Dowty, 1991), where it is suggested that (12) 'asserts that the information in the file, *viewed abstractly*, 'moves' from one place to another' (p. 597, my italics).

(12) It copies a file to a disk.

The third model in Fig. 4 says that #13, which is something which resembles #9, the view of the file as a piece of information, is located at #8, which is the view of the disk as a piece of information. This comes about as a result of combining the formal paraphrase of (12) with the following facts about copying events:

$$\begin{aligned}
&\forall A :: \{type(A, [entity, abstract, situation, event, copy])\}\\
&\quad \forall B :: \{\theta(A, object, B)\}\\
&\quad\quad \forall C :: \{view(B, C)\}\\
&\quad\quad\quad \forall D :: \{view(A, D)\}\\
&\quad\quad\quad\quad \forall E :: \{end(E, D)\}\\
&\quad\quad\quad\quad\quad D \approx C\\
&\quad\quad\quad\quad\quad \rightarrow \exists F similar(F, C)\\
&\quad\quad\quad\quad\quad\quad \wedge \forall G :: \{location(E, G)\}\\
&\quad\quad\quad\quad\quad\quad\quad location(F, G)
\end{aligned}$$

$$\begin{aligned}
&\forall A :: \{type(A, [entity, abstract, situation, event, copy])\}\\
&\quad \exists B view(A, B) \wedge transfer(B)
\end{aligned}$$

$$\forall A(transfer(A) \rightarrow \exists! C end(C, A))$$

The predicate *similar* which records the resemblance between #13 and #9 is yet another 'weak' term which picks up a great deal of its significance from the properties of its arguments. To say that two pieces of information are similar to one another is quite different from saying that two pictures are, or two people are, or ...

4. Zeugma

Consider the following sentences:

(13) This book is on the shelf and on the history of the Tudors.

(14) Get me the book on the shelf on the history of the Tudors.

(13) sounds very odd – much odder than (14), which I find quite acceptable. Why?

Recall that we have multiple views of books, including at least one view as a physical object and another as a mental object. Being on the shelf is a property that would locate the physical view of the book with respect to the surface of the physical view of the shelf. Why? Because the meaning postulates associated with prepositions pick out compatible views of the figure and ground and link them in some way. The meaning postulate for *'in'*, for instance, said that if C and D and are compatible views of A and B and E is the envelope of D then there is some connection between C and E. Likewise, the MP for *'from'* picked out compatible views of the figure and ground and then located one of these with respect to the start of the other. The MP for *'on'* is just like the one for *'in'*, except that it looks for the surface of the ground rather than its envelope.

So in (14) we have two properties being used in an attempt to isolate the book in question. It's the one for which (a) the physical view is located with respect to the surface of the physical view of the shelf, and (b) the mental view is located with respect to the surface of the mental view of 'the history of the Tudors'. Leaving aside for the moment the question of what the surface of a mental view is, it seems reasonable to suppose that you can use a variety of properties in order to pick out an individual.

In (13), however, we are predicating something of an entity – saying something like

$$holds\_of(book36, \lambda X(on(X, shelf19) \wedge on(X, `the\ \ Tudors')))$$

At the very least, coping with the first conjunct of the property would emphasise the physical view of *book36*, so that we might be surprised when we have to invoke the mental view as well. At worst we might have a meaning postulate for *holds_of* which explicitly invoked a view of the entity, e.g.

$$\forall X \forall P holds\_of(X, P) \rightarrow \exists Z\ ::\ \{view(Z, X)\} P.Z)$$

This MP would force you to choose a view of the target entity which was going to make sense with the predication. The physical view would make sense with the first part of the predication, but not with the second, thus leading to the sense of disorientation evoked by (13).

5. Conclusions

My aim in this paper was to show that you can obtain fine-grained descriptions of the *consequences* of uttering a sentence by constructing a fairly sparse logical form and then using an inference engine to induce descriptions of the world which are compatible with this logical form and with any meaning postulates and general world knowledge. What I am doing here is constructing, purely compositionally, a formal object which represents the 'linguistic meaning' (Barwise and Perry, 1983) of the words in the sentence, and then letting a full-scale inference engine loose in order to see what follows from this in the context in which it was produced. This contrasts with approaches such as (Pustejovsky, 1991; Copestake and Briscoe, 1996) which attempt to produce a richer initial analysis by allowing restricted amounts of real-world inference to contribute to the analysis. It is clear that the approach described here will produce a more accurate final analysis, at the cost of performing considerably more inference. I believe that with modern theorem proving techniques the burden of carrying out this inference is quite supportable, and that the results are well worth having. The proof of the pudding is, of course, in the eating: readers who are sceptical about just what this inferential effort costs are invited to try the demo at `http:\\ubatuba.ccl.umist.ac.uk`.

The models given in Figs. 1–6 are all obtained in this way: construct the logical form, try to prove that this is incompatible with the meaning postulates, and when this fails read the model off the current open branch of the tableau.

The analyses given above depend heavily on the use of multiple views of an object. This idea extends suggestions by various people. Cruse (1986) introduces a very similar notion by talking about the 'facets' of an objects, though it should be noted that Cruse (personal communication) explicitly disavows a close connection. Dowty (1979) refers to ' ... van Fraassen's notion of *logical space*. There will be as many axes of logical space as there are kinds of measurement. [...] *Each axis might have a different mathematical structure according to the discriminations that can be appropriately made in each case*' (p. 126, my italics). There are clear parallels between this notion and the idea of multiple views, with the 'different mathematical structures' corresponding to the different characterisations of each such view. The fact that a single concept such as *envelope* or *path* can be fleshed out differently depending on the type of the object it is being applied to – what I referred to as the object-oriented character of such concepts – provides further flexibility, so that the significance of of a preposition such as *'in'* or *'to'* depends in rather delicate ways on the properties of the entities it links. This is not to say that such prepositions are ambiguous: *'in'*, for instance, always says that the location of some view of the figure is the envelope of some view of the ground, and *'from'* always says that location of the start of some view of the figure is some view of the ground. The context sensitivity arises because the various views have to be of compatible types, and because the characterisation of *envelope* and *from* depends on the type of the relevant entity.

References

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

Copestake, A. and E. Briscoe: 1996, 'Semi-productive polysemy and sense extension'. In: J. Pustejovsky and B. Boguraev (eds.): *Lexical Semantics.* Oxford: Clarendon Press, pp. 15–67.

Cruse, D. A.: 1986, *Lexical Semantics.* Cambridge: Cambridge University Press.

Cryan, M. and A. M. Ramsay: 1997, 'A Normal Form for Property Theory'. In: *Proc. 14th Conference on Automated Deduction*, Vol. 1249 of *Lecture Notes in Artificial Intelligence.* Berlin, pp. 237–251, Springer-Verlag.

Dalrymple, M., J. Lamping, F. C. N. Pereira, and V. Saraswat: 1996, 'A deductive account of quantification in LFG'. In: M. Kanazawa, C. Piñón, and H. de Swart (eds.): *Quantifiers, deduction and context.* pp. 33–58.

Dowty, D. R.: 1979, *Word Meaning and Montague Grammar.* Dordrecht: D. Reidel.

Dowty, D. R.: 1991, 'Thematic Proto-roles and Argument Selection'. *Language* **67**, 547–619.

Dowty, D. R., R. E. Wall, and S. Peters: 1981, *Introduction to Montague Semantics.* Dordrecht: D. Reidel.

Fenstad, J. E., P.-K. Halvorsen, T. Langholm, and J. van Bentham: 1987, *Situations, Language and Logic.* Amsterdam: Kluwer Academic Publishers (SLAP-34).

Jackendoff, R. S.: 1983, *Semantics and Cognition.* Cambridge, MA: MIT Press.

Johnson-Laird, P. N.: 1983, *Mental Models: towards a cognitive science of language, inference and consciousness.* Cambridge, Mass.: Harvard University Press.

Kohlhase, M., S. Kuschert, and M. Pinkal: 1996, 'A type-theoretic semantics for λ-DRT'. In: *Proceedings of the Tenth Amsterdam Colloquium.* ILLC, Amsterdam.

Manthey, R. and F. Bry: 1988, 'SATCHMO: a Theorem Prover in Prolog'. In: *Proc. 9th Inter. Conf. on Automated Deduction*, Vol. 310 of *Lecture Notes in Artificial Intelligence.* Berlin, pp. 415–434, Springer-Verlag.

Moëns, M. and M. Steedman: 1988, 'Temporal Ontology and Temporal Reference'. *Computational Linguistics* **14(2)**, 15–28.

Pollard, C. J. and I. A. Sag: 1988, *An Information Based Approach to Syntax and Semantics: Vol 1 Fundamentals.* Chicago: CSLI lecture notes 13, Chicago University Press.

Pollard, C. J. and I. A. Sag: 1994, *Head-driven Phrase Structure Grammar.* Chicago: Chicago University Press.

Pustejovsky, J.: 1991, 'The Generative Lexicon'. *Computational Linguistics* **17(4)**, 409–441.

Ramsay, A. M.: 1995, 'A Theorem Prover for an Intensional Logic'. *Journal of Automated Reasoning* **14**, 237–255.

Ramsay, A. M.: 1996, 'Aspect and Aktionsart: fighting or cooperating?'. In: *Proceedings of the 16th International Conference on Computational Linguistics (COLING-96).* Copenhagen, pp. 889–894.

Ramsay, A. M.: 1999, 'Dynamic and Underspecified Semantics without Dynamic and Underspecified Logic'. In: H. Bunt, L. Kievit, R. Muskens, and M. Verlinden (eds.): *Computing Meaning*, Vol. 1. Dordrecht, pp. 208–220, Kluwer Academic Publishers (SLAP 73).

Turner, R.: 1987, 'A Theory of Properties'. *Journal of Symbolic Logic* **52(2)**, 455–472.

van Genabith, J. and R. Crouch: 1997, 'How to glue a donkey to an f-structure'. In: H. Bunt, L. Kievit, R. Muskens, and M. Verlinden (eds.): *2nd International Workshop on Computational Semantics*. Tilburg, pp. 52–65, Univ. of Tilburg.