Some Different Approaches to DRT

Alan W Black
Centre for Cognitive Science, University of Edinburgh
2 Buccleuch Place, Edinburgh EH1 9NW
awb@ed.ac.uk

July 21, 1993

Abstract

This paper discusses some issues in various approaches to Discourse Representation Theory. This is not intended to be detailed descriptions of various systems but general comparisons between both implementations of DRT and descriptions of DRT in some other context (e.g. a situation theoretic framework). The intention is to show how close (or how different) these approaches are and which aspects have been considered important by their developers. The view point is taken from a computational rather than a linguistic slant but these issues are not unrelated.

1 Background

This paper discusses a number of approaches to Discourse Representation Theory, both in implementation and/or in description. DRT was originally described in [Kamp 81] but has been developed in many directions. Although the original work was concerned with tense, DRT to many people has been inextricably linked with "donkey sentences", for which it gives good treatment. In this paper we will look at the following treatments of DRT:

- The "original" DRT [Kamp 81] and its later definitions in [Kamp & Reyle 93].
- [Johnson & Klein 86] (J&K) which offers a simple description of the core aspects of DRT and an implementation in Prolog.
- [Pinkal 91] which concerns itself with a strict "compositional" approach to DRT.
- [Cooper 93] offers a description of DRT within the situation theoretic framework of EKN [Barwise & Cooper 93].
- [Black 92] shows a treatment of DRT within the computational situation theoretic language of ASTL ([Black 93])

- [Groenendijk & Stokhof 91b] describe the dynamic logic DPL which although is not a description of DRT itself is specifically designed to offer a treatment of the same semantic phenomena as DRT and hence within this context seems relevant to our discussions. Also [Black 93] gives DPL a treatment in ASTL and hence allows a closer comparison with DRT through the DRT description in ASTL.

It must be emphasized that although we will give short descriptions of each of the above systems this paper is not designed as a full tutorial on these different approaches. Some knowledge of these systems will be assumed. What we will do is highlight certain aspects of these systems that allow us to give a closer comparison.

Although we will try to make statements about what certain systems can and cannot do, ultimately this cannot always be determined merely from published papers (which typically are old). There may have been developments in later versions of which I am not aware, or I feel that the development really constitutes a new system. Also stating the X does not do Y is usually far too strong a statement as X may be able to do Y with some simple extension or merely that that aspect has been omitted in a discussion of X in order to be compressed into some conference paper page limit.

Of course all systems under discussion are very similar and at least at some level they can be considered the same. It is the purpose of the various researchers involved to address exactly those issues addressed by the original DRT (and file change semantics [Heim 82]).

During the descriptions of each system we will be concentrating mainly on three aspects of Discourse Representation Structures. That is how DRSs are *represented*, how they are *interpreted* and probably most interesting of all how they are *constructed*. Although comparisons are drawn throughout each description a more detailed overall comparison is given in the final section.

2 Kamp's "Original" DRT

The following may not be exactly as Kamp first described DRT [Kamp 81] but the following will be accepted by most as a description of the basic core aspects of DRT. It is these core aspects which will be considered as the basic requirements when looking at the later systems.

DRT offers a representation for natural language discourses. A DRS (Discourse Representation Structure) consists of two parts a set of *discourse markers* which are used to represented objects introduced in the discourse; and *conditions* on these objects. DRSs are characteristically drawn as boxes. The markers are shown in the upper part of the box with the conditions below. A typical DRS for the sentence *"A man walks"* is

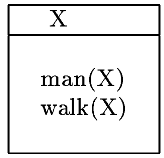

Conditionals and the determiner *"every"* are represented as a relation between two DRSs. Thus a DRS for the sentence *"every man loves a woman"* would be drawn as

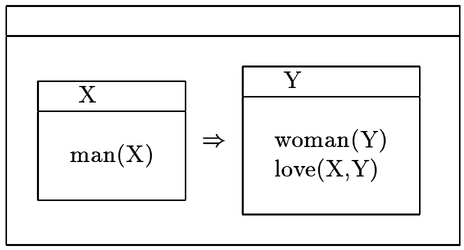

The interpretation of a DRS can be defined as follows. A DRS, D is true in a model M, when

- There exists a binding B with assigns each discourse marker in the top part of D to individuals in M such that each condition in D with respect to B becomes a true proposition in M.
- A $\Rightarrow$ condition is defined to be true if for all bindings B' (which are an extension of B, the binding the $\Rightarrow$ condition is within, for just those markers introduced in the left DRS) that make the left DRS true in M, there exists B'', an extension of B' for just those markers introduced in the right hand DRS, which makes the right hand DRS true in M.

This treatment of the conditional is the basis of DRT's popular claim to fame. That is its treatment of "donkey sentences". Donkey sentences first considered in [Geach 62] offer a problem to standard logical treatments of language. The classic example is

Every man who owns a donkey beats it.

The problem can be seen when we try to give a first order representation of the sentence. In translating indefinite noun phrases in simple sentences like *"a man walks"*, the indefinite introduces an existential quantifier, while in the donkey sentence such a simple translation might give

$$\forall x \exists y[[man(x) \wedge donkey(y) \wedge owns(x,y)] \rightarrow beats(x,y)]$$

But the above is not a meaning of the English utterance. The above expression is true in the following model

$$\begin{array}{ll} farmer(a) & own(a,b) \\ donkey(b) & cat(c) \end{array}$$

where a owns a donkey but does not beat it. Other placements of the existential will also fail to capture the meaning. In order to get a reasonable translations it is necessary to translate the indefinite that is within the scope of the universal as a universal quantifier rather than a simply an existential. Thus

$$\forall x \forall y[[man(x) \wedge donkey(y) \wedge owns(x,y)] \rightarrow beats(x,y)]$$

will give us a correct reading. Although this problem can be solved by translating indefinites as universals in such a context it is considered strange to need to deal with indefinites in different ways depending on their context. DRT however solves this offering a consistent treatment of indefinites irrespective of their context. Indefinites are always simple added to the DRS (there is an implicit existential for all discourse markers). Universal quantification is not over variables but over DRSs, thus existentials in the sub-DRS of a $\Rightarrow$ relation will effectively be treated as universals.

The analysis of anaphora is also an aspect of DRT. DRSs offer a representation which can be accessed to find which objects in a discourse are available for pronominal reference. For example given the simple discourse "*A man walks. He talks.*" After the first sentence we would have[1]

X
man(X) walk(X)

In dealing with the pronoun "*He*" at the start of the following sentence we must find an *accessible* marker that already exists in the DRS which we can related this pronoun too. Here we will introduce a new marker for the pronoun and use the condition *is* to related to to an appropriate existing marker. Thus after the second sentence we would have a DRS of the form

X Y
man(X) walk(X) is(X,Y) talk(X)

[1]The exact form of the DRS at the end of the first sentence depends on the particular treatment of DRT. As we will see the way DRSs are constructed and related to points in the discourse is a major area of difference in the versions presented here.

The definition of *accessibility* states which markers are available as pronoun referents. A marker is accessible in the DRS which introduces it (i.e. where it appears in the top part of the box). It is also accessible in all sub-DRS of a DRS in which it is accessible. The third condition is that any marker introduced in the left sub-DRS of a $\Rightarrow$ relation is accessible in the right sub-DRS. A more formal definition is not possible without giving a better characterisation of which DRSs we are talking about. The above characterisation would allow for cataphora (as in "*He walks. A man talks*") which is not normally allowed in DRT. Accessibility is really part of how we go about constructing DRSs.

The construction of DRSs from utterances (called the *construction algorithm*) is considered part of DRT. It is not sufficient for a semantic theory to merely offer a representation for natural language utterances, a method for constructing such representations from utterances is also required. DRT rightly considers this to be part of the theory and includes the *construction algorithm* to build DRSs from simple parse trees. As we will see the construction algorithm varies quite substantially in different systems of DRT but in the original Kamp DRT, the translation is defined by a number of re-write rules that will take a conventional parse tree (which obviously must be from a grammar suitable for DRT) and recursively re-write the tree building a DRS for the utterance. Thus during conversion objects will exists that are neither parse trees or DRSs but somewhere in between. The conversion can be viewed as happening at points in the syntactic tree thus for the sentence *"John owns a Porche"* the DRS/tree object that will exist after processing the subject noun phrase will look something like

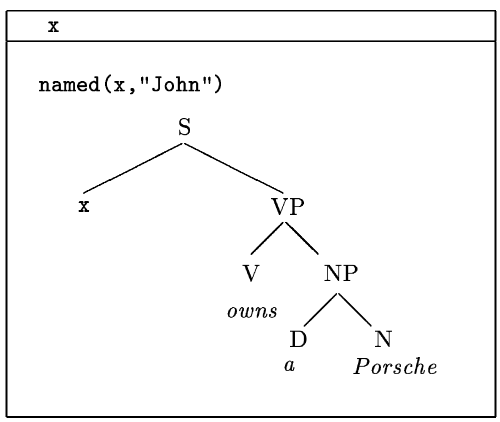

It must also be said that the above is a very simple description of what is considered the core parts of DRT. There have been many extensions by Kamp himself as well as others. [Kamp & Reyle 93] offers treatments of generalised quantifiers and plurals. But for the sake of the descriptions here most of the extensions to the above core DRT that exist could as easily have been added to the other systems as they have been to the Kamp DRT (in fact some extensions have been added to other versions of DRT).

Some aspects which we will consider in more detail that have sometimes been used as arguments against the DRT approach are "*compositionality*" and "*representational*". Unfortunately these terms are vague and not always easy to see when they can be used as criticisms or virtues. We will also look at the dynamic aspects of DRT which were perhaps not emphasized in the early work but could be considered as an important aspect of the theory. That is we can view the meanings of utterances as relating incoming DRSs, which define their context, to outgoing utterances which act as context for succeeding phrases in the discourse.

3 Johnson and Klein's DRT

[Johnson & Klein 86] is a short conference paper which neatly describes the basic aspects of DRT from an abstract view and in the particular context of a Prolog program. Their fragment is just enough to deal with donkey sentences and simple anaphora but is sufficient to show the core aspects of DRT. One specific point that is made explicit in this paper which perhaps was not made before is the dynamic properties of DRT. J&K view that utterances are in the context of an incoming DRS and an outgoing DRS. The outgoing DRS is a function of the Incoming DRS *plus* information from the current utterance. For example consider the discourse, *"A man walks. He talks"* After processing the first sentence we would have a DRS thus

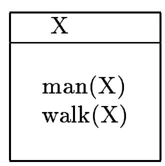

This output DRS from the first sentence would act as the input DRS to the second sentence. The subject noun phrase of the second sentence would have input and output DRSs like

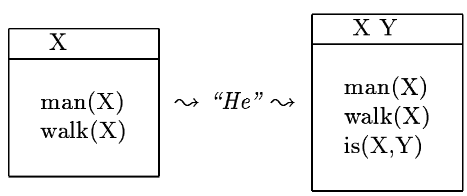

This output DRS would then be fed as the input DRS for the verb phrase *"talks"* producing the output DRS of

X Y
man(X) walk(X) is(X,Y) talk(Y)

This dynamic aspect although it exists in the early DRT is emphasized more in the J&K description.

The J&K description is not quite as conventional as the above example shows. Because they are working in Prolog (actually Pratr, Prolog with attributes) and are actually computing the DRSs J&K have to use a linear form of expression for DRSs rather than actual boxes. Their representation for the DRS

X
man(X) walk(X)

is actually as a Prolog list

```
[m.man(m),walk(m)]
```

Notably discourse markers are not partitioned from conditions (nor need they even appear at the front of the list). This aspect is not very important and we could easily construct a partitioned DRS from this. The reason for holding both markers and conditions as a single list is in order to make searching for accessible markers easier.

A second point that differs from the example above is that J&K do not introduce new discourse markers for pronouns, but simply pick up the referent marker directly. Thus "*A man walks. He talks.* would be represented as

```
[m,man(m),walk(m),talk(m)]
```

This simplification is not unusual in DRT descriptions but is admitted to be a simplification which loses information. Introducing a new marker for a pronoun seems to allow a great distinction which would probably be necessary if we were to deal with examples of a form like VP ellipsis as in "*Mary gave her mother some flowers and so did Jane*".

Because J&K is a conference paper they are pushed for space so there are a number of short cuts. The whole Prolog program is not presented but the idea is put over.

J&K's interpretation of the construction algorithm, as stated earlier, is interesting. They seriously treat the idea of dynamics. Each syntactic constituent is related to an incoming and outgoing DRS. The outgoing DRS is constructed from the incoming one plus information from that constituent. DRSs are *threaded* through the utterance. The following diagram for the utterance *"every man walks"* goes some way to illustrate the flow.

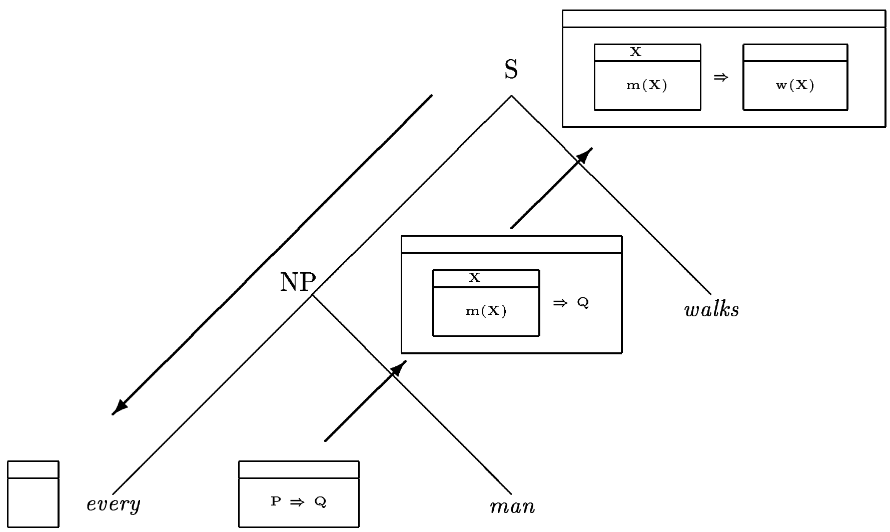

The above diagram shows how a DRS is passed in through each syntactic part of an utterance. The whole story is in fact a little harder to show in a single diagram. There are some crucial aspects to this form of construction algorithm which we need to highlight in order to compare them with later systems. The DRS which is fed into the verb phrase *"walks"* is not actually the DRS that is shown above coming out of the noun phrase. The verb phrase contributes to the right sub-DRS of the ⇒ constraint in the main DRSs. This threading of appropriate DRSs is achieved by the specification of the appropriate variables in the (Prolog) grammar rules. Thus the main DRS actually bubbles up from the determiner node in the subject NP rather than from the verb phrase.

Crucially when we look at the output DRS from the subject noun phrase. At the time of processing that NP (if we can talk about processing in such terms). The DRS would look like

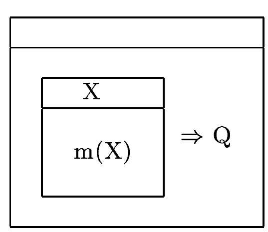

The important aspect is that this structure itself is not a valid DRS but some form of abstraction over DRSs (where Q represents the "missing" sub-DRS). Of course this variable Q is a Prolog variable that will be bound later in the analysis such that on completion of the sentence it will have been given a value. Although to some this point may be disregarded as irrelevant others have considered such a problem as serious and have given explicit solutions (notably [Pinkal 91] and [Cooper 93]). Therefore we will be specifically looking in each system discussed in this paper what the representation of an isolated noun phrase containing the determiner *"every"* is, as this highlights an important aspect of *compositionality* in DRT.

Compositionality (that the meaning is a function of the meaning of the parts) has been waved as a problem for DRT. It has been said (by many) that DRT is not compositional ([Groenendijk & Stokhof 91b] use this a specific criticism of DRT and make it an essential characteristic of their alternative). It is unclear at what level compositionality really is a problem. [Zadronzy 92] shows how compositionality can be achieved for any formalism merely at the cost in either the compose function or in the information content of the structures being composed. Also there seems to be a particular aspect of typical sentences treated by DRT that cannot be compositional in the classic sense. That is discourse anaphora. In a strict sense we cannot give a meaning for the isolated sentence *"He talks"* without a context in which there is a referent for the pronoun. Of course we can consider that context (what becomes before it in the discourse, or the situation in which the utterance occurs, in situation semantic terms) to be a part of the compose function and then we achieve our goal of compositionality, but purists may regard that as cheating. Others, particularly [Pinkal 91] and [Cooper 93], rather than explaining the problem away have specifically addressed this issue to offer strict compositionality in the classic sense (as we will see later).

The J&K description of DRT is small but not too small as to avoid the basic aspects of the theory. The major aspect that they emphasized is that of threading and the dynamic aspect of utterances transforming incoming DRSs to produce outgoing ones. But some aspects were left undefined. Particularly their DRSs may contain variables for which a semantics is not given (though in following the Prolog code their meaning can be found). Also during construction they will allow "DRSs" which may not be regarded as DRSs in the strict sense. After processing the verb in an utterance like *"A man likes a donkey"* the DRS to be fed to the donkey noun phrase would be

X
man(X) likes(X,Y)

Thus a discourse marker is used before it is introduced. The Y may be viewed as some form of abstraction in the same way the Q was above or we could say that the threading is not in the right order (more will be said on this issue in discussion of the treatment of DRT in ASTL).

4 Pinkal's compositional DRT

Although there are not too many details published about Pinkal's system, the only generally available publication is [Pinkal 91], it is obvious that he has addressed the specific issue of compositionality and found a reasonable solution. As mentioned above in building a DRS for a utterance such as *"Every man likes a donkey"* we have a problem in what the DRS representations would be for its syntactic constituents. When we look at the desired DRS for the whole sentence

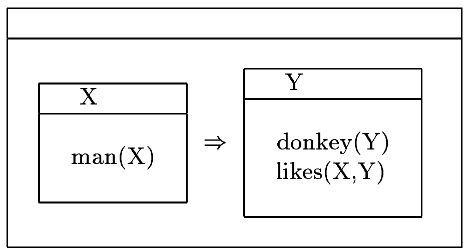

we cannot not find any sub-part of this DRS which corresponds to the noun phrase or verb phrase of the given utterance. While the J&K version solved this by naming variables appropriately during the construction in Prolog, Pinkal has given a much more explicit treatment to this issue. He has a concept of abstractions over DRSs such that he represents the noun phrase *"every man"* as the abstraction

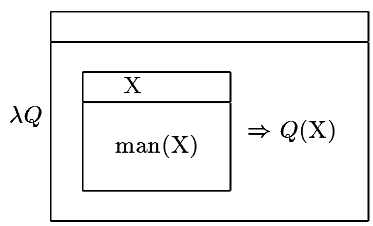

Unlike J&K here we have an explicit identification of the part being abstracted over. But when we look at the representation for the verb phrase we see that we require something more complex than simple lambda application to achieve our desired result. The verb phrase representation is

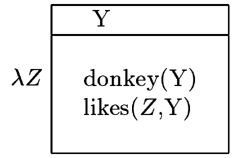

There are two issues here. First, the (X) argument of Q given in the first abstraction identifies the marker introduced by the head noun of the subject noun phrase. It is given here so that the verb phrase may attach it to the appropriate role. J&K use other techniques to attach arguments to roles. But it is not just the argument from the subject noun phrase that needs to be passed. Consider the sentence, *"Every man who owns a donkey likes it"*. This time the subject NP must identify two arguments as the two components must be

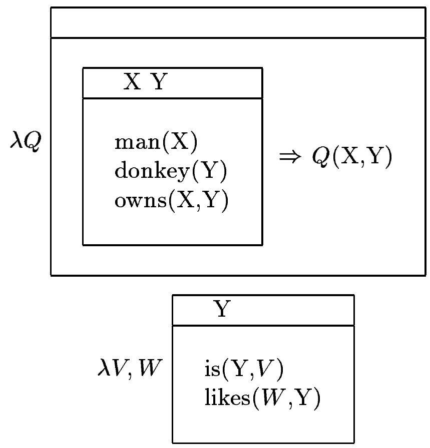

Here we have two arguments but crucially we cannot depend on each part having the same number (or in fact that they are in the correct order). The compose function $\oplus$ has to be defined in such a way as to allow for such matching thus making it significantly more complex than simple lambda application.

[Pinkal 91] does not discuss this issue directly but considers other syntactic constraints on the composition of abstractions over DRSs. Because he works within a general framework of GB various aspects "binding theory" are included in his discussion.

Although we have only given a very brief summary of Pinkal's treatment of DRT (which actually is implemented in a GB based system which considers many other issues including quantifier scope) the most important aspect is highlighted. Pinkal's system defines a $\oplus$ operator which allows the composition of objects representing abstractions over DRSs. The composition is not simple lambda application as there is really a set of variables that are abstracted over (discourse markers). This set does not necessarily need to be fully applied to the argument. Part of the application function must match the markers in the "lambda" expression as appropriate. We will return to this issue in the following description of Cooper's DRT in EKN which attempts to define exactly the same function.

5 Cooper's DRT in EKN

The next treatment of DRT we wish to describe (and the one following this) started from a very different basis to the previous descriptions. Both [Cooper 93] and [Black 92] are working within a situation theoretic framework. In both cases they are interested in DRSs as semantic (structured) objects.

EKN is a graphic notation for the representation of situation theoretic objects, [Barwise & Cooper 93]. The notion of using boxes to represent objects was taken from DRT but EKN offers a much richer ontology.

Like the next description ([Black 92]) Cooper's treatment is part of an overall plan to show how contemporary semantic theories may be recast within a situation theoretic framework. Cooper offers two methods for representing DRSs. First we can view a DRS as a relation, whose arguments would act as "values" for discourse markers. In this case the EKN DRS representation for the discourse *"A man owns a donkey"* would be

$i \rightarrow X,\ j \rightarrow Y$
$\text{man}(X)$ $\text{donkey}(Y)$ $\text{own}(X, Y)$

An important difference in Cooper's DRSs is that the discourse markers are labelled. This makes explicit the fact that the discourse markers are unordered (by simultaneous abstraction) but still allows us to refer to them outside the abstraction.

Alternatively we can view DRSs as abstractions of situation types. The idea being that given some assignment of markers to individuals a DRS can be made a situation type. In this case the representation would be

$i \rightarrow X,\ j \rightarrow Y,\ k \rightarrow S$
S
$\text{man}(X)$ $\text{donkey}(Y)$ $\text{own}(X, Y)$

Cooper states that this second form is ultimately that which is required if his fragment were to be extended to cover DRT treatments of tense and attitudes.

The addition of the situation argument allows a reasonable form for events and states.

Like Pinkal's DRT, Cooper also allows abstractions over DRSs (abstractions are a fundamental construct in EKN). Thus his representation for sub-sentential constituents is as follows. The EKN representation for the isolated noun phrase *"every man"* is

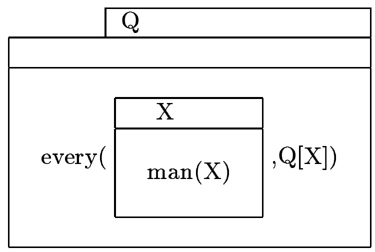

While the verb phrase *"walks"* is represented by

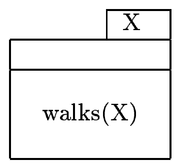

Again, like Pinkal, composition of such objects cannot be achieved by standard lambda application. Cooper describes a composition operator which achieves the required result. This operator is effectively (and intended) to be the same as Pinkal's $\oplus$ operator. Cooper gives the composition function in abstract from DRT and posits it as a general function. Although the mechanism is defined for assigning any argument of an abstraction to an argument of the composition function, the selection of which discourse marker should be selected to resolve which unresolved marker is not. This is a non-trivial problem but probably can be abstracted from the actual composition and left to linguistic (and discourse) constraints. (Pinkal too does not give a full discussion on this issue either).

Thus Cooper's DRT in EKN offers a strict compositional construction of DRSs. It also gives DRSs a specific semantics which opens obvious routes for treatments of tense and attitudes from the use of the situation object. Cooper does not propose his DRT in EKN as an alternative to "original" DRT but is specifically looking at how DRT can be described in a situation theoretic context.

6 DRT in ASTL

ASTL is a computational language based on situation theory which allows contemporary natural language semantic theories to be described in it. Because

ASTL's semantics are given in terms of situation theoretic objects, descriptions in ASTL are effectively given a situation theoretic treatment ([Black 93]). Apart from a treatment of situation semantics and dynamic semantics [Black 93] also presents a description of DRT. ([Black 92] also presents a treatment of DRT in ASTL but that description has been superseded by the later one in [Black 93].)

The ASTL description of DRT is not dissimilar to Cooper's EKN. The representation of DRSs in ASTL is as parametric situation types (which is almost the same as Cooper's alternative representation). In ASTL (which also offers an EKN mode for displaying objects) the utterance *"a man walks"* is given a DRS of the form

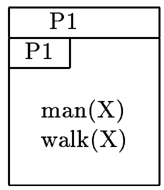

The difference with Cooper's representation is that the ASTL one does not include the the discourse markers as part of the abstraction. However it does abstract over the situation in which the conditions are supported.

The concept of threading is important in the ASTL description. Unlike J&K, threading is defined abstractly from the DRSs. A threading relation is defined over *utterance situations* (effectively syntactic constituents of a discourse). DRSs are then defined in terms of the threading relation. The relationship between an incoming and outgoing DRS is defined by constraints. For example, in the case of proper nouns the following constraint is used.

```
*S:[S ! S != <<DRSOut,S,
              *A :: *AType &
                    [A ! A != <<accessible,*X,*TYPE,1>>],
              *DRSIn &
               [D ! D != <<named,*X,*N,1>>],1>>]
<=
 *S:[S ! S != <<cat,S,ProperNoun,1>>
         S != <<use_of,S,*N,1>>
         S != <<sem,S,*X,1>>
         S != <<type,S,*TYPE,1>>
         S != <<DRSIn,S,
               *A1 :: *AType,
               *DRSIn,1>>].
```

A short explanation of ASTL constraints is necessary in order to understand the above constraint. Variables in ASTL are written as atoms prefixed with *. The above can be read as for all situations which are of the type as specified on the right hand side of the constraint there exists a situation of the type specified on the left hand side (i.e. before the <=). Because *S is used to identify the situation on both sides on the <= the constraint states that the same situation that is of type ProperNoun etc. must also be of the type which the DRSOut relation as specified.

In particular, the outgoing DRS of a proper noun utterance situation is the incoming one (*DRSin) *plus* the new condition named for the marker introduced by the proper noun itself. Also the DRSIn and DRSOut relations include an argument identifying which discourse markers are accessible at that point in the discourse (and the types of those markers). In the proper noun example above, the new marker is added to the outgoing accessibility situation (with the appropriate type) so that it is available a a pronoun referent later in the discourse.

Accessible markers are explicitly held rather than found from a function of the current DRS (and the context that DRS is within) as it was felt that the accessibility relation as defined in the "original" DRT section above depended too much on positional aspects of the boxes.

Another aspect of the ASTL description that differs is the treatment of threading. Threading is treated quite seriously and is defined abstractly (without reference to DRSs). The threading relation is defined as an alternative structure to the syntactic structure defined over the same syntactic entities (called utterance situations). This relation does not necessary follow the syntactic tree from left to right top down. The ASTL description of DRT defines the threading relation more abstractly. Each entity (except the initial discourse situation) appears as the second argument to the threading relation (i.e. all utterance situations have an incoming thread). Determiner utterance situations are also related to two situations via the range and body relations. These situations are not threaded to any other situation (i.e. they are ends of threads).

An annotated tree for the sentence *"Every man likes Hanako."* is shown below. The threading relation is shown with bold arrows

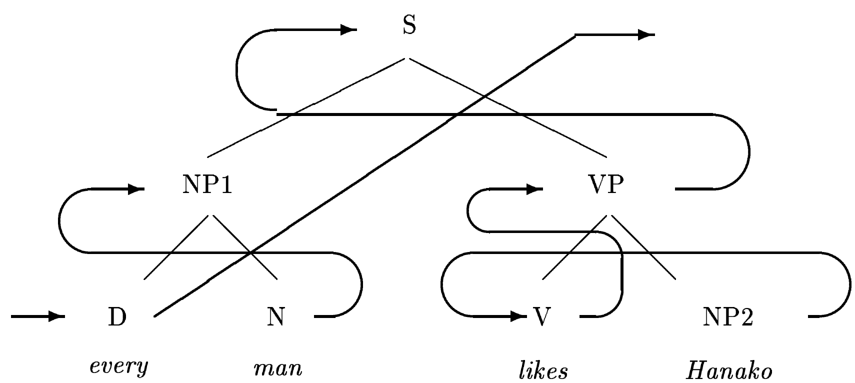

In addition, DS (the discourse start situation) is threaded to D, N and NP2. The main discourse thread will go through `D`. There are two other threads ending at `NP1` and `S`. `D` will be related to `NP1` by the relation `range` and to `S` by the relation `body`.

While in J&K processing in done top-down, here processing is done bottom up. Importantly because of the constraints defining the relation between outgoing and incoming DRSs, an outgoing DRS cannot be fully determined unless all its dependencies already exist. This means that the outgoing DRS of the utterance situation representing *"every"* in *"every man walks"* cannot be assigned until both the range and scope sub-threads have been determined (i.e. until after the words "man" and "walks" have been dealt with). Where J&K (and more explicitly Pinkal and Cooper) would build an abstraction over a DRSs for *"every"* the ASTL description requires that the dependencies have been processed. At some level of abstraction these two treatments are the same. After processing of the full discourse the outgoing DRS of a J&K determiner node will be fully filled in and be the same as that on the ASTL determiner utterance situation. Pinkal and Cooper however would still have the abstraction and only the sentence (or at least a node further up the tree) would have the fully fleshed out DRS.

Arguably the ASTL description of DRT is not compositional in the strict sense. But defining the incoming and outgoing DRS relationship as a constraint does not seem unreasonable.

One problem with the ASTL representation of DRSs is due to the representation of DRSs. An ASTL DRS for the discourse *"A man arrived. Every woman saw him."* is

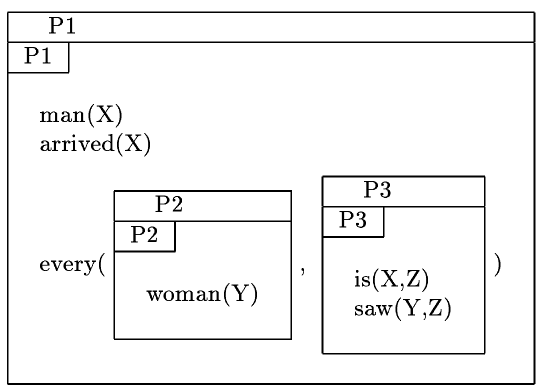

The problem is with the X in the right sub-DRS. The definition given for the treatment of the relation `every` requires that the scope of discourse markers be known. Just looking at the DRS arguments to `every` alone one cannot really determine if X should be taken as fully within the `every` clause or depend on something outside it. In order to give the proper treatment it is necessary to interpret a DRS top-down so that the X will have been assigned before we interpret the `every` relation. Again this may or may not a problem depending on one's viewpoint. Not explicitly stating the point of introduction could introduce confusion in scope so that the treatment given in Cooper's EKN where DRSs are abstractions over situations types and discourse markers, seems more attractive. In the current version of ASTL such a general abstraction object is not possible but extensions that would allow such a form as discussed in [Black 93, Ch. 7].

7 Dynamic semantics and DPL

The work on dynamic semantics, Dynamic Predicate Logic and Dynamic Montague Grammar ([Groenendijk & Stokhof 91b] [Groenendijk & Stokhof 91a]) also deserve some mention here. DPL is devised specifically to offer a more logic treatment of the phenomena that DRT was designed form. Particularly in its treatment of donkey anaphora. DPL itself gives a translation for *"every man who owns a donkey beats it"* as

$$\forall x[[man(x) \wedge \exists y[donkey(y) \wedge owns(x,y)]] \rightarrow beats(x,y)]$$

Although the y argument to *beats* appears to be outside the scope of the existential introducing y within DPL this is not the case. The semantics of DPL expression is a set of pairs of *assignments*. Assignments are functions from DPL variables to objects in the model. The semantics of a DPL expression is given in terms of threading assignments to sub-expressions. Because of the definition, existentials introduced within the range of universal quantifiers are still available as referents in the scope. However these are not available outside the universal quantifier. These conditions are effectively the accessibility conditions of DRT.

DPL itself is not given with respect to an English fragment so it contains no equivalent of DRT's construction algorithm. However [Lewin 92] and [Black 93, Ch. 6] both give such a conversion. Black's description (again in ASTL) shows how closely the threading is when compared with DRT using exactly the same threading relation of the DPL description as for the DRT one. Lewin's treatment too has obvious similarities to a DRT construction algorithm.

In DPL, the equivalent of DRSs seem to be assignments. When a definition states that DPL variable x must be bound to something that is a man, that is capturing the same information as the following DRS

X
man(X)

Looking at DPL from the point of view of the other discussions of DRT given above we can say that the original DPL was present as compositional. (Compositionality being one of the specific points of comparison with DRT in the original paper.) It is interesting that this compositionality is not achieved by the use of any form of explicit abstraction but by a redefinition of logical operators. It is also interesting that DPL only offers compositionality between (natural language) sentences and not within sentences. DPL's representation for the noun phrase *"every man"* would need to be some form of abstraction.

Admittedly Groenendijk and Stokhof offer DPL only to illustrate the use of dynamic semantics in natural language semantics and propose DMG as the actual framework to use. (Unfortunately my knowledge of DMG is not really sufficient to comment on this more.)

8 Comparison

One further aspect of DRT which is often quoted, and sometimes made explicit within treatments can be best exemplified by DRT's treatment of proper names. It is assumed that when proper nouns are introduced into a discourse they are available for pronominalisation irrespective of whither they are introduced within the scope of a universal or not. In the (deliberately over-simplified) version of "original" DRT, the DRS for "Every woman likes John." would be

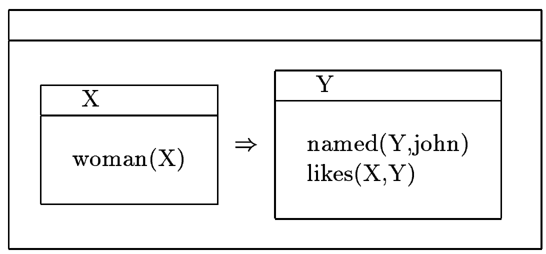

That is the discourse marker introduced by the proper noun is introduced into the right sub-DRS of the ⇒ relation, thus making it inaccessible for later pronoun reference. This is obviously not the case as the above sentence can naturally be followed by *"He starred in a recent film."* where *"He"* refers to John. In order for this to be possible within a DRT framework it is necessary for proper names to add markers to the outermost box, so that they are available as potential referents.

It is not just proper names that require this treatment, definites (and others) may also require the ability to add markers to the outermost context. From a logical perspective there is a direct analogy with quantifier scoping. If proper names are treated as existential quantifier's then it is necessary for them to have widest scope in the sentence so they may be used anywhere with in the logical expression.

In order to achieve this within the threading versions of DRT (i.e. J&K, Cooper's EKN, and ASTL) it is necessary to also thread a "global" DRS through the constituents. Cooper's EKN version does do this while J&K and the ASTL version do not—though such a mechanism could easily be added. The strict compositional treatment of Pinkal's may preclude such a treatment (though I am unsure of enough details of that system to be sure). The DPL treatment also has a problem, though as no construction algorithm is defined it is the responsibility of implementors to find a solution. The ASTL version of DPL does not but the description could be extended in the way the DRT in ASTL description would need to be extended. Lewin's treatment of proper names in DPL is adequate because the quantifiers introduced by proper names are given widest scope.

The following table tries to sum up the various systems discussed above. Of course such tables are arguable but the following does give an interesting comparison.

	Compositional	Threading	ProperNames
Kamp	No	No	Yes
J&K	sort-of	Yes	No
EKN	Yes!	Yes	Yes
ASTL	No	Yes!	No
Pinkal	Yes!	No	??
G&S	Yes!	Sort-of	No

The exclamation mark after "yes" is used to denote that fact that these treatments of DRT have specifically addressed that point.

Cooper's EKN appears to use both the compose function and threading although either of these two operations seem sufficient. Cooper actually uses the compose function for DRS construction but also uses threading for other information through a discourse.

9 Conclusions

What this paper shows is that there are number of systems which (rightly) claim to be treatments of DRT, but they achieve their goal in a quite different way. These different aspects of DRT are relevant to a computational treatment of the theory and it is not yet obvious which are better or worse (or possibly irrelevant). One major choice we have highlighted seems to be in the construction of a DRS from an utterance. We can choose the "compose" method taken by Pinkal and Cooper or the dynamic threading method chosen by Johnson & Klein and Black.

Also there are other aspects of later versions of DRT which would also benefit from a similar comparison. For example, treatments of quantifier scope, plurals, aspect etc. all have slightly different treatments by different people. A close side by side comparison would give a better understanding of the treatments.

Note that it is possible to borrow techniques from one system to another. Cooper does exactly that. He takes Pinkal's compose operator but also takes threading from Johnson & Klein for other information. Variation within the theory is possible and indentifying such variations make it clearer what the computational aspects are.

Although the description of DRT in ASTL makes certain choices about its treatment of DRT, ASTL is really a more general system. It is possible to describe Pinkal's compose function within ASTL and hence implement either Pinkal's version of DRT or Cooper's within ASTL. In fact the use of a computational tool makes it much easier to identify these different techniques as well as offering an environment within which they can be implemented.

References

[Barwise & Cooper 93] J. Barwise and R. Cooper. Extended Kamp Notation: a graphical notation for situation theory. In *Situation Theory and its Applications, III*, CSLI Lecture Notes. Chicago University Press, forthcoming 1993.

[Black 92] A. Black. Embedding DRT in a Situation Theoretic framework. In *Proceedings of COLING-92, the 14th International Conference on Computational Linguistics*, pages 1116–1120, Nantes, France, 1992.

[Black 93] A.W. Black. *A Situation Theoretic approach to computational semantics.* Unpublished PhD thesis, University of Edinburgh, Edinburgh, UK., 1993.

[Cooper 93] R. Cooper. Towards a general semantics framework. Prepared for the Dagstuhl Seminar: Semanitc Formalisms in Natural Language Processing, 1993.

[Geach 62] P. Geach. *Reference and Generality.* Cornell University Press, Ithaca, NY, 1962.

[Groenendijk & Stokhof 91a] J. Groenendijk and M. Stokhof. Dynamic Montague Grammar. In *Quantification and Anaphora I*, DYANA Deliverable R2.2A, pages 1–37. Centre for Cognitive Science, University of Edinburgh, 1991.

[Groenendijk & Stokhof 91b] J. Groenendijk and M. Stokhof. Dynamic Predicate Logic. *Linguistics and Philosophy*, 14:39–100, 1991.

[Heim 82] I. Heim. *The Semantics of Definite and Indefinite Noun Phrases in English.* Unpublished PhD thesis, University of Massachusetts, Amherst, Mass., 1982.

[Johnson & Klein 86] M. Johnson and E. Klein. Discourse, anaphora and parsing. In *Proceedings of the 11th International Conference on Computational Linguistics*, pages 669–675, Bonn, West Germany, 1986.

[Kamp & Reyle 93] H. Kamp and U. Reyle. *From discourse to logic: Introduction to Model Theoretic Semantics of Natural Language, Formal logic and Discourse Representation Theory.* Studies in Linguistics and Philosophy, 42. Kluwer, Dordrecht, forthcoming 1993.

[Kamp 81] H. Kamp. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof, editors, *Formal Methods in the Study of Language.* Mathematical Center, Amsterdam, 1981.

[Lewin 92] I. Lewin. *Dynamic Quantification in Logic and Computational Semantics.* Unpublished PhD thesis, University of Edinburgh, Edinburgh, UK., 1992.

[Pinkal 91] M. Pinkal. On the syntactic-semantic analysis of bound anaphora. In *Proceedings of the 5th conference of the European Chapter of the Association for Computational Linguistics*, pages 45–50, Berlin, Germany, 1991.

[Zadronzy 92] W. Zadronzy. On compositional semantics. In *Proceedings of COLING-92, the 14th International Conference on Computational Linguistics*, pages 260–266, Nantes, France, 1992.