Introduction to Discourse Representation Theory (DRT)

Alistair Knott
Dept of Computer Science
University of Otago

What DRT is about

DRT is a theory of natural language semantics. The aim: to asssociate sentences with expressions in a logical language which represent their meaning.

What language should we use to express the meaning of a sentence like

A dog barked ?

This question was originally explored by philosophers (logicians), then by linguists, and now it's relevant to researchers in LT.

- I'll begin by looking at predicate logic as a candidate representation language, and motivate DRT as an extension of this.

Some history: Russell's theory of NL semantics

Bertrand Russell proposed that **first-order predicate logic** can be used to represent sentence meanings.

For instance:

Sentence:	*Every child plays.*
Means:	$\forall x[child(x) \supset plays(x)]$

A key part of his analysis is that indefinite NPs (e.g. *a dog*) are translated using an existential quantifier ($\exists(x)$).

For instance:

Sentence:	*A dog barked.*
Means:	$\exists x[dog(x) \wedge barked(x)]$

Russell's treatment of indefinite NPs

Russell noted that the NP *a dog* doesn't always introduce a particular dog, especially around **scope elements** like *not* and *every*.

For instance:

Sentence:	*A dog didn't bark.*
Means:	$\neg\exists x[dog(x) \wedge barked(x)]$

Sentence:	*Every child owns a dog.*
Means:	$\forall x[child(x) \supset \exists y[dog(y) \wedge owns(x, y)]]$

Using $\exists(x)$ to translate an indefinite NP works nicely in these cases.

- *A dog* always translates as $\exists x$.
- This expression *combines* with other logical expressions in the right ways.

Some problems for Russell

Russell's account doesn't scale up very well to **discourse**—i.e. to a (coherent) *sequence* of sentences.

A key problem: how to interpret definite NPs?

A dog arrived.	The dog barked.
$\exists x[dog(x) \wedge arrived(x)]$	?

We really want to keep asserting predicates about the variable x. But any subsequent occurences of x will be outside the scope of the existential quantifier.

In this context, *a dog* does seem to introduce a particular dog, which the definite NP then refers back to.

Towards a more general conception of *a*

Russell already has a neat way of making *a dog* behave in different ways in different contexts.

We want to extend this, so that

- *a dog* always contributes the same expression to the representation of a sentence, but
- this expression can *combine* in different ways with other parts of the sentence's representation, so that in some contexts it introduces a new entity, and in other contexts it acts like an existential quantifier.

Discourse representation theory (DRT) does this.

DRT: Representing discourse contexts

We need a representation language in which we can talk about operations like ‘introducing a new entity into the discourse context’, and ‘referring back to the current discourse context’.

DRT represents ‘the discourse context’ as a **discourse representation structure** (or **DRS**). A DRS is:

- A set of **referents**: the entities which have been introduced into the context; and
- A set of **conditions**: predicates which are known to hold of these entities.

Here’s an example:

A, B
primeminister(A) securityguard(B) frisked(A,B)

DRT: Sentences as context update operations

In DRT, a sentence's meaning is taken to be an *update operation on a context.*

- Each sentence is interpreted in a context. The result of interpretation is a new context.

A sentence is also represented as a DRS. For instance, here's the DRS for *A dog arrived*:

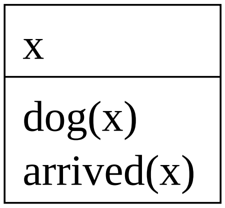

x

dog(x)
arrived(x)

The current context is **merged** with the sentence DRS to yield the new context.

For instance:

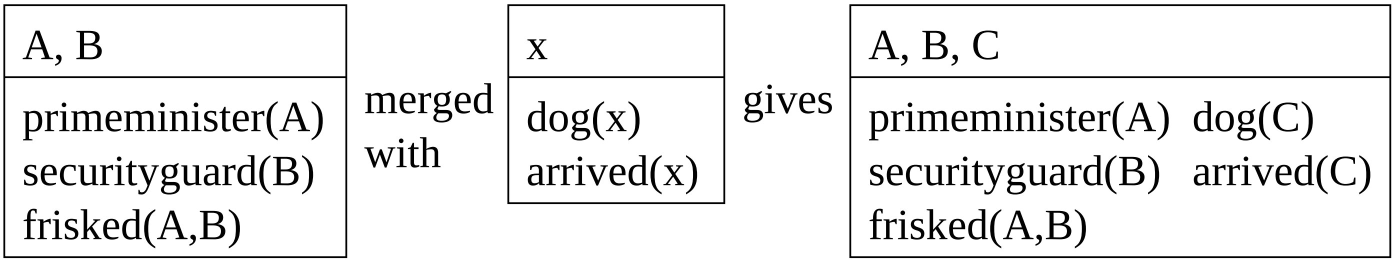

A, B

primeminister(A)
securityguard(B)
frisked(A,B)

merged with

x

dog(x)
arrived(x)

gives

A, B, C

primeminister(A) dog(C)
securityguard(B) arrived(C)
frisked(A,B)

DRT: Presuppositions

A sentence can also make **presuppositions** about the kinds of context in which it can be interpreted. For instance:

- *The dog barked* presupposes that there is a dog in the discourse context.

Each presupposition can also be represented as a DRS. So the general representation of a sentence is:

- A single **assertion** DRS;
- A *set* of presupposition DRSs.

Here's the DRS for *The dog chased a cat*:

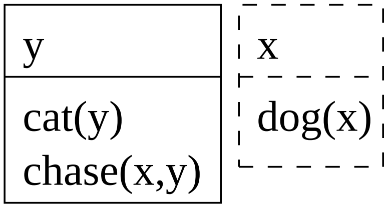

y

cat(y)
chase(x,y)

x

dog(x)

DRT: Presupposition resolution

A presupposition is basically a simple *query* to execute on the discourse context.

A presupposition is **resolved** if the query is successful:

- Any variable bindings returned by the query are carried over to the assertion DRS.
- The assertion DRS is then merged with the context DRS.

If context is $\begin{array}{|l|}\hline \text{A, B} \\ \hline \text{dog(A)} \\ \text{man(B)} \\ \hline\end{array}$ & sentence DRS is $\begin{array}{|l|}\hline \text{y} \\ \hline \text{cat(y)} \\ \text{chase(x,y)} \\ \hline\end{array}$ $\begin{array}{:l:}\hdashline \text{x} \\ \hdashline \text{dog(x)} \\ \hdashline\end{array}$

the assertion DRS after resolution is $\begin{array}{|l|}\hline \text{y} \\ \hline \text{cat(y)} \\ \text{chase(A,y)} \\ \hline\end{array}$,

and the context after merging is $\begin{array}{|ll|}\hline \text{A, B, C} & \\ \hline \text{dog(A)} & \text{cat(C)} \\ \text{man(B)} & \text{chase(A,C)} \\ \hline\end{array}$.

Translating DRSs to predicate logic

There's a simple translation from a DRS to an expression in first-order predicate logic.

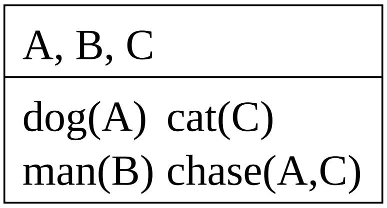

A, B, C

dog(A) cat(C)
man(B) chase(A,C)

- For each referent in the DRS, create an existential quantifier. E.g. for the above DRS, $\exists A \exists B \exists C$
- Join all the conditions together with the connective $\wedge$. E.g. for the above DRS, $[dog(A) \wedge man(B) \wedge cat(C) \wedge chase(A,C)]$.

A context DRS is really just a notational variant of a predicate calculus formula. But crucially, it's one which supports a various context-update operations.

Indefinite NPs in DRT

Recall: we're trying to find a representation of indefinite NPs which shows how they sometimes introduce new discourse referents, and sometimes behave like quantifiers.

In DRT, an indefinite NP like *a dog* always contributes a DRS which looks like this:

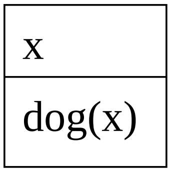

x

dog(x)

To make this structure behave in different ways, DRT introduces the idea of **sub-DRSs**.

- Sub-DRSs are created by scope elements such as quantifiers and negation.
- When we process a phrase inside a scope element, we add material to the sub-DRS.

A sub-DRS for negation

Consider the sentence *A dog did not bark.*

- First, create an empty sentence DRS.
- Then process the negation.
- Then process the sentence *A dog barked* inside the sub-DRS.

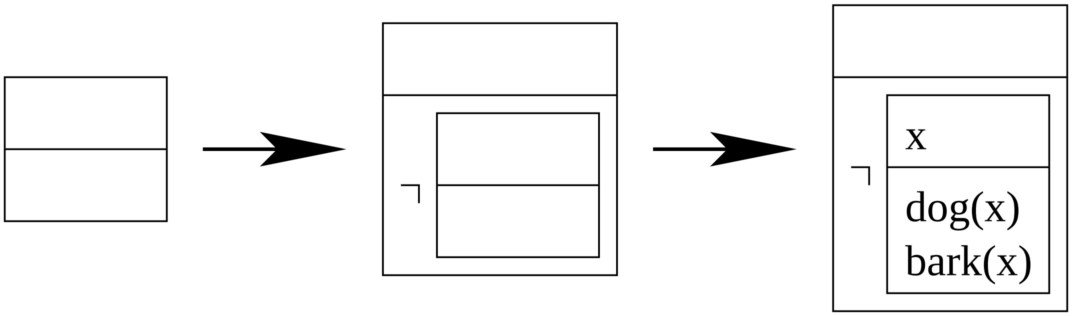

Translating back to predicate logic is almost transparent:

$$\neg\exists x[dog(x) \wedge bark(x)]$$

The concept of accessibility

DRSs now have two separate functions:

- Modelling context updates & presupposition;
- Modelling indefinites under negation.

These two functions come together to describe one other discourse phenomenon: **indefinites in the scope of negation can't be used to resolve presuppositions.**

A dog did not bark. ?The dog was big.

DRT can handle this very easily, by specifying that material within sub-DRSs is **inaccessible** to presuppositions.

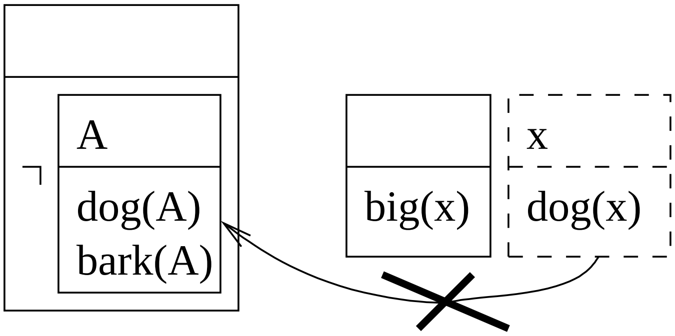

A sub-DRS for quantification

Russell's account of indefinites as existentials was also good for dealing with indefinites in conjunction with quantifiers.

Every child owns a dog

In DRT, quantifiers also introduce a sub-DRS structure. (A slightly more complex one.)

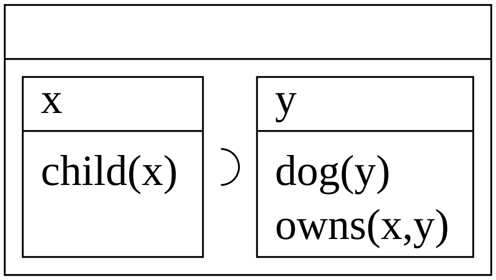

x

child(x)

⊃

y

dog(y)
owns(x,y)

- The material in the noun introduced by the quantifying determiner goes in the left-hand DRS. The rest goes in the right-hand DRS.

This is also easy to translate to predicate logic:

$$\forall \ \underline{\hspace{2cm}} \ [\ \underline{\hspace{2cm}} \ \supset \ \underline{\hspace{2cm}} \ ]$$

Donkey sentences

DRT is famously associated with sentences like this one:

> *Every farmer who owns a donkey beats it.*

Geach noted that these pose problems for Russell's theory of indefinites.

- Geach traces these sentences back to medieval philosophers.
- Heim traces them back to Chrysippos (5th century BC).
- I think I have discovered an even earlier antecedent.

Donkey sentences in DRT

Here's another donkey sentence.

Every child who owns a dog loves the dog.

Here's what the sentence needs to mean:
$\forall x \forall y [child(x) \wedge dog(y) \wedge owns(x,y)] \supset loves(x,y)$

- *A dog* needs to introduce a <u>universal</u> quantifier. (This is a case Russell couldn't cover.)
- The variable it introduces must be available for presuppositions elsewhere in the sentence.

Both requirements fall out naturally from the DRS analysis.

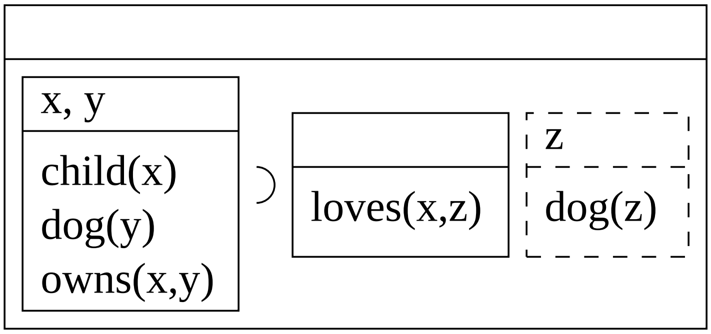

Some conclusions about DRT

DRT achieves three main things.

1. It provides a framework for expressing the meaning of a sentence as an operation of updating contexts. This opens up elegant treatments of 'ordinary' indefinites, and of presuppositional constructions.
2. It provides a way of giving a single denotation of the indefinite determiner *a*, which works in different contexts
    - to introduce a new individual into the discourse;
    - to introduce quantified variables.
3. The combination of 1 and 2 allow the formulation of a very nice theory of how pronouns can refer back to quantified variables.

Who cares about donkey sentences?

Achievements 1 and 2 are the main ones.

Donkey sentences are like Eddington's measurements of light bending when it passes the sun.

- We're trying to choose between several alternative theories.
- One theory predicts something that happens in certain rare circumstances. The other theories have nothing to say about these circumstances.
- No-one cares much about the circumstances themselves, because they're rare.
- But we might as well use the theory which makes the right predictions.