

Whose Context?

Hans Kamp

University of Stuttgart
University of Texas, Austin

August 6, 2015

Three perspectives on Natural Language Semantics

- Perspective 1: The semantics of a language is a property (or set of properties) of the language as an autonomous, user-independent system.

It can be described without reference to use.

- Perspective 2: Meaning cannot be separated from use.

More specifically, it can't be separated from interpretation.

- Perspective 3: The communication-theoretic perspective:

Language is a vehicle of communication.

The meanings of linguistic expressions manifest themselves in their capacity to transfer information from the speaker or author to the hearer or reader.

A problem about the communication-theoretic perspective

- To make a theory in the spirit of Perspective 3 work, we need a way to identify the thought representations of speaker and hearer independently of the linguistic expressions they use as go-betweens.
- But where are such identifications of thought representations supposed to come from if not from language?
- So how could such an account of meaning escape circularity?

A problem about the communication-theoretic perspective

- A (somewhat tentative, and partial) answer:

In representational approaches (like DRT and SDRT) to discourse anaphora and other dynamic aspects of language, semantic representations play a double role. They

(i) identify, via a formal semantics in the spirit of Perspective 1, the content of what has been interpreted;

(ii) serve as discourse contexts for the interpretation of the next utterance, or next part of the text.

In order that representations can play these two roles systematically and effectively, they must satisfy fairly strict constraints.

Successful representations are thus likely to tell us something about the semantic representations constructed by humans.

An old-fashioned DRT example

(1) a. Joseph has a donkey. It carries his fiancee.

b.

$s_1 \quad j \quad d$
Joseph'(j) donkey'(d)
$n \subseteq s_1$
$s_1: \text{have}'(j,d)$

c.

$s_1 \quad j \quad d \quad s_2 \quad u \quad v \quad f$
Joseph'(j) donkey'(d)
$n \subseteq s_1$
$s_1: \text{have}'(j,d)$
fiancee-of'(f,v) $u = d \quad v = j$
$n \subseteq s_2$
$s_2: \text{carry}'(u,f)$

An issue for all three perspectives: Context Dependence

- Syntactic form is not the only factor that determines semantic value or content.

Context plays an important part as well.

It impinges on content in all sorts of ways.

- Two fairly well understood types of context:
 - (i) Utterance Contexts: determine the content of occurrences of words like *now*, *I*, *you*.
 - (ii) Discourse Contexts. These are the representations of DRT, SDRT and similar approaches mentioned above.
- But even taken together, Utterance Context and Discourse Context are not enough.

Referent identification for Definite Noun Phrases

- A sub-problem of the general problem of context dependence:
What contextual information may be needed for the interpretation of definite noun phrases?
- The definite noun phrases of English:
 - (i) proper names
 - (ii) 1st and 2nd person pronouns
 - (iii) 3rd person pronouns
 - (iv) definite descriptions
 - (v) demonstrative NPs (NPs beginning with *this* or *that*)

Referent identification for Definite Noun Phrases

- Contextual information needed for the interpretation of definite NPs:

1st and 2nd person pronouns (indexical NPs): the Utterance Context

- Anaphoric uses of definite NPs (3rd person pronouns, definite descriptions, demonstrative NPs): the Discourse Context:

Proper names: ??

Non-anaphoric uses of 3rd person pronouns, definite descriptions, demonstrative NPs: ??

Referent identification for Definite Noun Phrases

- Among the NP uses for which Utterance Context and Discourse Context are not enough:

‘Discourse-new’ uses of definite descriptions and proper names in texts.

- Examples of discourse-new descriptions and names:

Mary, Paris, ESSLLI, my cousin from Heidelberg, the independence of Catalonia, the Spanish Succession War, the highest mountain in the Americas, the nearest star.

Referent identification for Definite Noun Phrases

- Why are discourse-new NP uses unproblematic in practice?

The intuitive answer is obvious:

Interpreters can recognize discourse-new NPs as describing entities with which they are already familiar.

(Compare the Familiarity Theory of Definites (Jespersen, Christophersen, Heim):

Definite NPs signal Familiarity; indefinite NPs signal Novelty.)

Entity Representations in MSDRT

- But how can we make this notion of familiarity more precise?
- Here is a proposal. It is rooted in MSDRT (for ‘Mental State DRT’).
- MSDRT is several things at once:
 - (i) a DRT-based theory of the structure of mental states;
 - (ii) a (DRT-based) formalism for the description of mental states;
 - (iii) a Logical Form formalism to be used in an account of the semantics of attitude-ascribing sentences and discourses.

Entity Representations in MSDRT

- According to MSDRT, a mental state consists of
 - (i) Propositional Attitudes
 - (ii) Entity Representations
- Propositional Attitudes have the form $\langle \text{MOD}, K \rangle$, where
 - (a) K is a representation of the content of the attitude (for us: a DRS);
 - (b) MOD is ‘Mode Indicator’, such as BEL for belief, DES for desire, INT for intention, etc.

Entity Representations in MSDRT

- The form of Entity Representations (ERs): $\langle [ENT, x], K, \mathcal{K} \rangle$, where
 - (i) x is a discourse referent;
 - (ii) K is a DRS;
 - (iii) \mathcal{K} is a (possibly empty) set of *internal anchors*.
- An internal anchor is a testimony, at the level of mental representation, of a causal link between (i) ER, (ii) referent and (iii) agent.
- Types of Internal Anchors:
 - (a) perceptual anchors;
 - (b) vicarious anchors;
 - (c) memory-shifted anchors

An example of a perceptual anchor

- Suppose the agent is currently seeing an object (or takes herself to be seeing an object), which she thinks is a gold coin.

This will take the form of her mental state being enriched with an ER $\langle [ENT, x], K, \{K_{anch}\} \rangle$,

where K_{anch} would be:

s
$n \subseteq s$
s: i see x

.

K might be the empty DRS in this case.

An example of a vicarious anchor

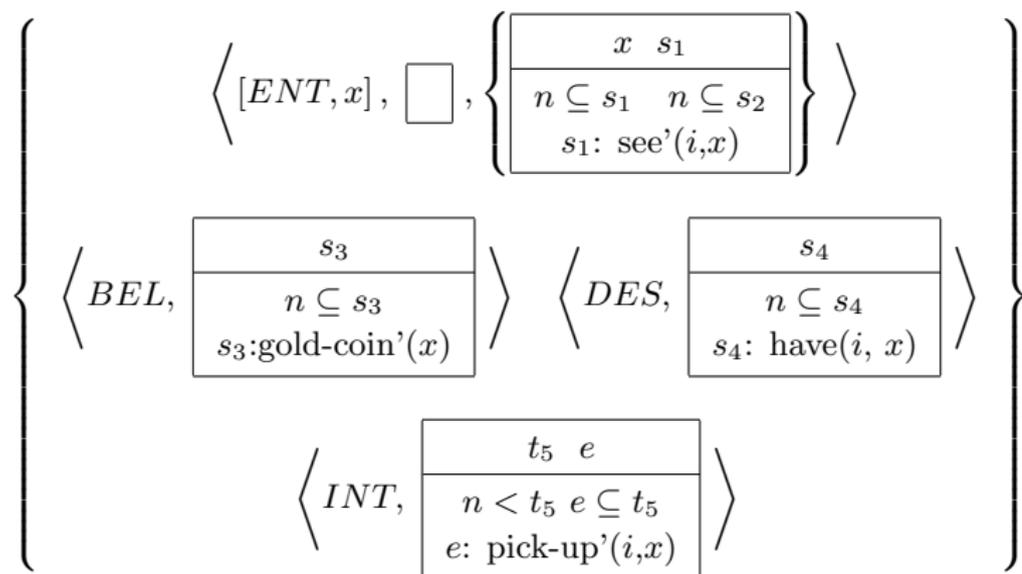
- Vicarious anchor of an ER formed by an interpreter B in response to the use by a speaker A of a noun phrase α :

$$\left\langle [\text{ENT}, x_b], K_b, \left\{ \begin{array}{l} \begin{array}{cc} s & e \\ \hline n \subseteq s & e < n & e \subseteq s \end{array} \\ s: \text{Att}(a_b, \{l : \langle [\text{ENT}, x_a], \emptyset, \emptyset \rangle\}, \{\langle x_a, x_b \rangle\}) \\ e: \text{REAL}(l, \alpha) \end{array} \right\} \right\rangle$$

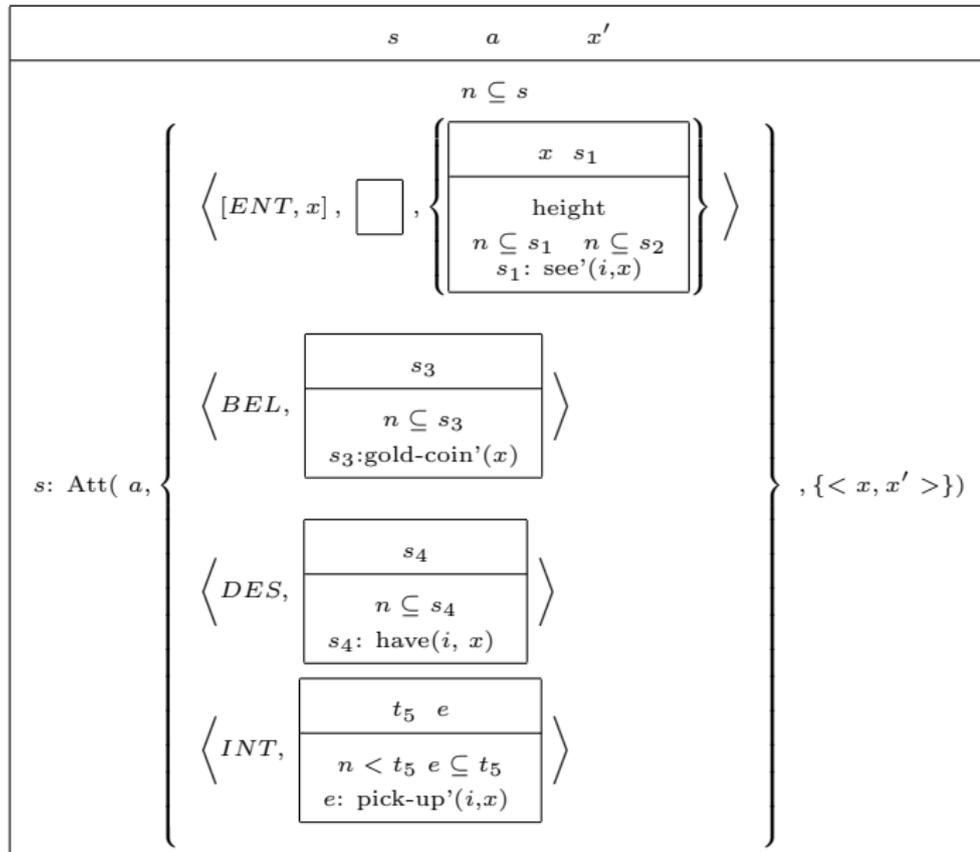
An Example

- A mental state: propositional attitudes + entity representations.

(2)



Attitude Attributions



Unification of Utterance Context and Discourse Context

- We can incorporate the information carried by the Utterance Context into the Discourse Context by making use of *indexical discourse referents*:

n for the utterance time, sp for the speaker/author, ad for the addressee, and perhaps others.

- DRSs with indexical discourse referents have a special significance:

They are representations of sentences or discourses **uttered in particular utterance contexts**.

The Utterance Context of the utterance represented by K determines the values (= embedding targets) of the indexical discourse referents in K .

Articulated Contexts

- Contextual information needed for the interpretation of definite NPs:
 - (i) the ‘Discourse+Utterance Context’: K_{dis} ;
 - (ii) Entity Representations from the interpreter’s ‘entity library’ (the stock of ERs from his mental state): K_{enc} ;
 - (iii) ‘Generic’ information about how the world does and can hang together: K_{gen} ;
(K_{gen} won’t play a role in what follows and we say no more about it)
 - (iv) In case of face-to-face communication about perceptually accessible entities in the environment:
A set of ERs with current or recent perceptual anchors: K_{env} .

Articulated Contexts

- Thus an Articulated Context has the following general form:

$\langle K_{dis}, K_{enc}, K_{gen}, K_{env} \rangle$

- A little more about these components:
 - (i) K_{dis} is the discourse context, built in the course of interpretation of the utterance, discourse or text along the lines of DRT;
 - (ii) K_{enc} is identified with the interpreter's entity library;
 - (iii) K_{env} is the set of those of the interpreter's ERs that have a current or recent perceptual anchor.
 - (iv) K_{gen} is some collection of propositional content representations. (Not our concern here.)

Articulated Contexts as mental constructs

- With this we have reached the crucial point of the argument:
Two of the four components of an AC (viz. K_{enc} and K_{env}) are constituents of mental states.

So NP interpretations involving these can only be described in psychological terms.

But what then of aspects of utterance interpretation that co-occur with such NP interpretations?

How can we tell a coherent story about utterance interpretation if we do not describe those other aspects in psychological terms as well?

Articulated Contexts as mental constructs

- This question arises in particular for the interpretation of other NPs that co-occur with NPs which need K_{enc} or K_{env} .

But if the interpretation of those other NPs must be described in psychological terms, then the contextual information used in those NP interpretations should be part of the interpreter's mental resources as well.

So we conclude:

All components of Articulated Contexts should be treated as parts or aspects of the mental states of language users.

Using Articulated Contexts in discourse interpretation

- Here is an example of the kind of case I have in mind:

Suppose that a speaker A says (3) to a hearer B.

(3) That shrub is a hibiscus. John gave it to me in 2001.

In (3) a deictically used demonstrative NP (*that shrub*) and a name (*John*) occur together with indexicals (*I* and the tenses of the verbs) and an anaphoric NP (the pronoun *it*).

When this communication works as it is meant to, B's interpretation of (3) will lead to a change in the mental state of B that can be described as follows in MSDRT:

Using Articulated Contexts in discourse interpretation

- Presuppositions:

- (i) B must have an ER for the person referred to by A as ‘John’;

Assume that this ER has the form

$$\left\langle [\text{ENT}, j], \begin{array}{|c|} \hline \\ \hline \text{Named}(j, \text{'John'}) \\ \hline \end{array}, \emptyset \right\rangle$$

- (ii) B must either already have perceptually anchored ER for the shrub that A is talking about, or be able to put such an ER in place in response to A’s use of *that shrub*.

Assume that this second ER has the form

$$\left\langle [\text{ENT}, x], \begin{array}{|c|} \hline \\ \hline \text{shrub}'(x) \\ \hline \end{array}, \left\{ \begin{array}{|c|} \hline s \\ \hline n \subseteq x \\ s: \text{see}(i, x) \\ \hline \end{array} \right\} \right\rangle$$

Using Articulated Contexts in discourse interpretation

- B's K_{dis} after processing of the first sentence of (3):

(4)

hibiscus'(x)

- B's K_{dis} after processing the second sentence of (3), building on the discourse context in (4):

(5)

e t v
hibiscus'(x)
$e < n$ $t = 2001AD$ $e \subseteq t$ $v = x$
$e: \text{give}'(j, v, sp)$

The double function of deictic uses of demonstratives

- A could have used the NP *that shrub* assuming that B has already noted the shrub, and has a perceptually anchored ER for it.

But she need not have.

For her use of the NP, perhaps accompanied by a pointing or some other deictic gesture, has the power to draw B's attention to the shrub;

and that may in fact have been part of A's intention.

If this is the case, then A's utterance of *that shrub* will *prompt* B to form a perceptually anchored ER.

And having done that, B can then use this ER for the interpretation of the NP, just as he would have in case he already had an ER for the shrub.

The double function of deictic uses of demonstratives

- I assume that in either of these two cases B's ER also gets a vicarious anchor, as a witness that according to him the NP refers to the entity represented by his ER.

So in either case the ER will end up with (at least) two anchors, a perceptual anchor and a vicarious anchor.

- For an ER to have more than one anchor is quite common.

It also arises with recognition:

Suppose you encounter an entity for which you have an ER anew and recognize it as the referent of your ER.

Then you will add to the anchor set of the ER a new anchor testifying to the current encounter.

Thus anchor sets will grow.

Using Articulated Contexts in discourse interpretation

- Note:

(i) The DRS (5) represents the truth conditions of the discourse as a proposition that is ‘singular’ both about John (the John A is talking about) and to the shrub A has referred to with *that shrub*.

(ii) (5) is a user-independent representation so long as embeddings of it in models (in the sense of the model theory for DRSs) are restricted to functions that map j to John and x to the shrub.

(iii) Whenever an utterance or text does not involve K_{enc} , K_{env} or K_{gen} , then its processing could be given a non-psychological interpretation (just as for older versions of DRT).

Interpretation and Production

- So far our focus has been on interpretation.

But from a communication-theoretic perspective this is only half of the story:

From this perspective a producer encodes a thought content into language and the recipient(s) try to reconstruct that content by decoding the linguistic input he receives.

But production is not just the inverse of interpretation.

The principal difficulty: What are the thought representations from which producers ‘derive’ the utterances that express them?

Interpretation and Production

- Here is a hypothesis about the inputs to production (but a VERY tentative one):

Input representations to production are of the same form as the outputs of interpretation.

In any case, it seems reasonable to assume that a producer will end up (as a kind of production spin-off) with a representation that can be obtained as an interpretation of the linguistic encoding she has herself produced.

(This representation can serve the producer as a record of what she has said so far, and as a guide when phrasing her next utterance or sentences.)

This representation will typically be very close in form to the interpretations constructed by her audience.

Interpretation and Production

- In the normal case producer and interpreter(s) have the sense that they share these representations with each other, in the sense of Common Ground (or David Lewis' *scoreboard*).

This kind of sharing is the *result* of successful communication.

It must be distinguished from the much weaker notion of 'de facto sharing' that typically is all that is needed to guarantee communicational success.

Again, we take the production and interpretation of definite NPs as an illustration.

In order to correctly interpret a discourse-new definite description or name, the interpreter should have an ER matching the ER identifies the NP's referent for the speaker.

But the interpreter doesn't have to have any advance knowledge that the speaker has such an ER.

Interpretation and Production

- It is true that the speaker will have to make certain assumptions about the entity library of the hearer, so that she can choose an NP that the hearer can be expected to interpret correctly.

But in the normal case of discourse-new uses of definite NPs these are the only assumptions that one participant needs to make about the other.

There is no need for an infinite hierarchy of ever more complex mutual attributions (as in fixed point definitions of Common Ground).

- On the other hand, when a discourse context has been constructed through processing of part of the discourse, then that will normally be Common Ground.

Thus interpretation moves that only make use of the discourse context, as in anaphora resolution, will be part of the Common Ground as well.

Interpretation and Production

- The by now familiar versions of formal semantics – both static and dynamic – can be seen as dealing only with that part of interpretation that is covered by the Common Ground (i.e. by the Discourse Context+Utterance Context).
- This brings us to our final assessment.

The appeal of the first, traditional approach to natural language semantics is that it captures what must be shared by all competent speakers of the language.

For without such sharing verbal communication could not work and language would be useless.

One might worry that this crucial aspect of language gets lost when we switch from the first perspective to the second or the third.

Interpretation and Production

- But all that semantic theories developed in the spirit of the first Perspective can make visible about what competent speakers share can be captured by theories designed according to the second or third Perspective as well.

The difference is that theories of the second kind can also tell us something about how the same piece of information can have a different significance for you and for me.

This is because the way in which that piece is embedded within your mental state will as a rule differ from the way in which it is embedded in mine.

- But that is the beginning of a new and much longer story.

● THANK YOU

for sticking it out with me through this one.