Comparative Computational Semantics: Capturing Meaning by Boxes at the Berlin Workshop

Johan Bos
University of Groningen

November 10, 2014

Contents

Preface

At the WeSearch project meeting in Berlin 2014, various semantic formalisms will be compared to one another. This document gives some background and comments on the semantic analysis of selected examples in the framework of Discourse Representation Theory, in the variant produced by the semantic parser Boxer.

1 Introduction

One of the central aims of computational semantics is to associate meaning representations with natural language expressions. But how such meaning representations should look like, ideally (i.e. causing to predict everything we want our formalism to predict, and nothing more), is a question without a simple answer. Perhaps it is wrong question to ask in the first place, and we shouldn't judge meaning representations by their appearance but rather on how they perform in typical semantic tasks. I can think of three ways of evaluating meaning representations in a formalism-neutral way:

1. Predicting entailments between and/or contradictions within texts;
2. Predicting the truth conditions of a sentence (text) given some sensory input;
3. Predicting acceptability of translations.

The first method has long been practised by formal semanticists (let's call it the proof-theoretic appoach), and has recently become a standard shared task challenge, known as Recognizing Textual Entailment (RTE). One way to view this approach is to consider it as the *study of the relationship between meaning representations*. The proof-theoreric evaluation approach is illustrated in Table 1.

Table 1: Predicting informativeness and consistency

Given text	**New sentence**	**Informativeness**	**Consistency**
Abrams barked and was old.	Abrams was old.	UNINFORMATIVE	CONSISTENT
Abrams barked and was old.	Abrams was tired.	INFORMATIVE	CONSISTENT
Abrams barked and was old.	Abrams didn't bark.	INFORMATIVE	INCONSISTENT

The second approach is known as the model-theoretic approach to semantics, but hasn't matured to a practical exercise yet (but I think this will change soon, given the advances in image analysis recently and the current interdisciplinary efforts between language and vision engineers). This approach can be characterized as the *study of the relationship between meaning representations and situations*. To illustrate the model-theoretic evalation method, consider the situations and predictions in Table 2.

Table 2: Predicting truth values

Sentence	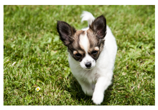	
There is a dog outside.	YES	YES
A dog is lifting his left paw.	YES	NO
A dog is carrying a stick in his mouth.	NO	YES
A dog is chasing a cat.	NO	NO

The third approach is, admittedly, not worked out in great detail, but the idea is here that if a translation of a sentence is meaning-preserving, the meaning representation of the target sentence ought to be identical to that of its source sentence. (There is one catch: translations aren't always meaning-preserving for various reasons.)

If a semantic formalism, when given many of such unseen examples, and does make the right predictions most of the time for any of the three evaluation approaches I have just sketched, then it is very hard to deny that it knows something about meaning, in whatever way it does it. In computational semantics it is more or less taken for granted that one needs an explicit level of semantic representation to compute with meanings. Meaning representations used nowadays come in in various formats and have a lot of common, but they also differ on many points. Most of them resemble first-order representations and have some notion of scope, variables ranging over entities, properties of entities and relations between entities. But subtle differences exist. One issue is the degree of underspecification (of scope, word senses, and relations). The other issue is the level of granularity of analysis (perhaps this is even depending on the semantic task). The fact that there are many differences shows that semantic analysis is a hard task. It therefore makes sense to compare semantic analyses performed in different formalisms (developed by different schools of thought).

In these notes we produce gold-standard meaning representations following a variant of Discourse Representation Theory for ten sets of selected linguistic phenomena. I first introduce the method used for constructing these representations. Then (quite informally) I present the box language used for the meaning representations. In the final part I discuss the selected phenomena.

2 Method

2.1 Followed procedure

We manually tokenized the example sentences. Then parsed it with the C&C parser, and produced meaning representations with Boxer [3]. Boxer's output was manually corrected where needed. The objective is to show fully resolved meaning representation (except for word senses). For readability we left out information on tense and aspect (although Boxer supports this) because it didn't seem relevant for the examples under discussion. We show presuppositions by putting them all together in a separate box. We don't distinguish between singular and plural entities for simplicity's sake. We assume an inventory of thematic roles similar to VerbNet [6] in a neo-Davidsonian event semantics (unlike Kamp & Reyle, [5]).

2.2 The box language for capturing meanings

The box language that we use is a variant of the Discourse Representation Structures used in Discourse Representation Theory [4], and has the following ingredients:

- constants and variables (ranging over entities or contexts)
- logical symbols for negation (and disjunction and implication)
- non-logical symbols for properties of entities and relations between entities or contexts
- contexts (represented by boxes)

A box comprises a set of discourse referents (variables) and a set of conditions. Conditions are either basic conditions or complex conditions. Basic conditions are formed by applying properties to entities or forming relations between entities (or contexts). Complex conditions are constructed by combining the logical symbols with boxes. The box language therefore describes a recursive data structure. It is not hard to show that this box-language is a fragment of first-order logic. Yet it has enough expressive power to represent the meanings of a wide range of linguistic expressions.

2.3 A note on sorts

All non-logical symbols are given meaning by arranging them in an ontology of some kind. For readability some concepts within this ontology are not explicit in the representation using non-logical symbols, but by using different letters for discourse referents: x for entities, e for events, s for states, and p for propositional contexts.

3 Studied Phenomena

3.1 Clausal Complements

In 0101 the discourse referent p1 is introduced that associates with the context "they barked". This context is the theme of the thinking event e1. Both x2 ("they") and x1 ("we") are presupposed.

```
0101 We thought that they barked .
  __________    __________________
 |x2 x1     |  |e1 p1             |
 |..........|  |..................|
(|thing(x2) |+|theme(e1,p1)      |)
 |person(x1)|  |agent(e1,x1)      |
 |__________|  |think(e1)         |
               |      ____________ |
               |     |e2          ||
               |  p1:|............||
               |     |agent(e2,x2)||
               |     |bark(e2)    ||
               |     |____________||
               |__________________|
```

3.2 Coordination

In 0201 we assume a collective interpretation, resulting in a dancing event performed by x1, a group that has x2 and x3 as members. Both x2 (an entity with the name "Abrams") and x3 (an entity named "Browne" are presupposed.

```
0201 Abrams and Browne danced .
  ____________________    ____________
 |x3 x2               |  |e1 x1       |
 |....................|  |............|
(|named(x3,browne,per)|+|agent(e1,x1)|)
 |named(x2,abrams,per)|  |dance(e1)   |
 |____________________|  |x3 C x1     |
                         |x2 C x1     |
                         |____________|
```

3.3 Ellipsis

It is impossible to give an (interpretable) semantic analysis for elliptical sentences such as “Abrams is” without context (it could be Abrams is ten years of age, Abrams is going to the party, Abrams is in the house, etc). Of course I could show one of the possible lambda-expressions for expressions such as “Abrams is”, but I don’t think that makes much sense.

3.4 Identity Copula

Intuitively the copula connects two discourse referents with an equality condition. But it is a bit more complicated than this, because just doing so would not be able to distinguish between “Abrams is Browne” and “Abrams was Browne”. So what I do here is to make identity effectively a three-place relation between two entities (x1 and x2) in a context (p1). Any tense information can be attached to the p1 variable.

```
Abrams is Browne .

  ____________________   _____________
 |x2 x1               | |p1           |
 |....................| |.............|
(|named(x2,browne,per)|+|     _______ |)
 |named(x1,abrams,per)| |    |       ||
 |____________________| | p1:|.......||
                        |    |x1 = x2||
                        |    |_______||
                        |_____________|
```

3.5 Nominalization

I more or less follow Asher’s proposal formulated in DRT, where nominals trigger their own neo-Davidsonian event structure [1].

```
0504 The voters support the government 's repeal of the law .

  ______________   ______________
 |x4 x3 x2 x1   | |e1            |
 |..............| |..............|
(|law(x4)       |+|theme(e1,x3)  |)
 |repeal(x3)    | |agent(e1,x1)  |
 |by(x3,x2)     | |support(e1)   |
 |government(x2)| |______________|
 |voter(x1)     |
 |of(x3,x4)     |
 |______________|
```

3.6 Comparatives

There is a lot of syntactic variation in comparatives. This makes it challenging for compositional semantics. For the example below I have chosen for a semantic analysis that is quite close to the surface syntax, but this requires additional background knowledge (shown below as well).

```
0601 The dog is older than the cat .

  _________   _______________
 |x2 x1    | |s1             |
 |.........| |...............|
(|cat(x2)  |+|than(s1,x2)    |)
 |dog(x1)  | |topic(s1,x1)   |
 |_________| |be-older(s1)   |
             |_______________|
```

$\forall s \forall x \forall y$ [be-older(s) $\land$ topic(s,x) $\land$ than(s,y) $\rightarrow$ $\forall m$ [age(y,m) $\rightarrow$ age(x,m)]]

Note that 0601 does not entail that the dog is old. The meaning representation supports this fact. Further note that if the cat is old, then 0601 entails that the dog is old. This is harder to formalize!

3.7 Control Relations

Not much to say here. Control relations are simply dealt with by assigning the right discourse referents to the thematic roles of the events in question. This is powered by lexical semantics. A rich resource for syntactic derivations including lexical semantics is the Groningen Meaning Bank [2].

```
0701 They forgot to vote.

  ___________   _______________________
 |x1         | |e1 p1                  |
 |...........| |.......................|
(|thing(x1)  |+|theme(e1,p1)           |)
 |___________| |agent(e1,x1)           |
               |forget(e1)             |
               |       ______________  |
               |      |e2            | |
               |  p1: |..............| |
               |      |agent(e2,x1)  | |
               |      |vote(e2)      | |
               |      |______________| |
               |_______________________|
```

3.8 Measure Phrases

I analysed measure phrases as playing a thematic role within a state that measures something. So in 0801 we have a state of being old, with the dog being the topic, and the unit of measurement being years. The number of units is spelled out as a relation between the type of unit (x2) and the number itself (n1).

```
0801 The dog was ten-and-a-half years old .

  _______   __________________________
 |x1     | |n1 s1 x2                  |
 |.......| |..........................|
(|dog(x1)|+|unit-of-measurement(s2,x2)|)
 |_______| |topic(s2,x1)              |
           |be-old(s2)                |
           |year(x2)                  |
           |ten-and-a-half(n1)        |
           |number_of_units(x2,n1)    |
           |__________________________|
```

3.9 Parentheticals

Some of the parentheticals seem to project, perhaps always when they function as appositive.

3.10 Relative Clauses

If 1002 were “a dog that barked disappeared” its meaning representation would have been indistinguishable from “a dog that disappeared barked”. But since we have a definite article and a restrictive relative clause, the barking becomes part of the presupposition that there is a barking dog.

```
1002 The dog that barked disappeared .

  _____________   _________________
 |x1 e1        | |e2               |
 |.............| |.................|
(|dog(x1)      |+|disappear(e2)    |)
 |bark(e1)     | |theme(e1,x1)     |
 |agent(e1,x1) | |_________________|
 |_____________|
```

References

[1] Nicholas Asher. *Reference to Abstract Objects in Discourse*. Kluwer Academic Publishers, 1993.

[2] Valerio Basile, Johan Bos, Kilian Evang, and Noortje Joost Venhuizen. Developing a large semantically annotated corpus. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey, 2012. European Language Resources Association (ELRA).

[3] James Curran, Stephen Clark, and Johan Bos. Linguistically Motivated Large-Scale NLP with C&C and Boxer. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 33–36, Prague, Czech Republic, 2007.

[4] Hans Kamp. A Theory of Truth and Semantic Representation. In Jeroen Groenendijk, Theo M.V. Janssen, and Martin Stokhof, editors, *Truth, Interpretation and Information*, pages 1–41. FORIS, Dordrecht – Holland/Cinnaminson – U.S.A., 1984.

[5] Hans Kamp and Uwe Reyle. *From Discourse to Logic; An Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and DRT*. Kluwer, Dordrecht, 1993.

[6] Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. A large-scale classification of English verbs. *Language Resources and Evaluation*, 42(1):21–40, 2008.