

Syntactic disambiguation using presupposition resolution

Alistair Knott and Peter Vlugter

Department of Computer Science

University of Otago

alik/pvlugter@cs.otago.ac.nz

Abstract

If a sentence is ambiguous, it often happens that the correct reading is the one which can most easily be incorporated into the discourse context. In this paper we present a simple method for implementing this intuition using the mechanism of presupposition resolution. The basic idea is that we can choose between the alternative readings of an ambiguous sentence by picking the reading which has the greatest number of satisfied presuppositions. We present two uses of the disambiguation algorithm in our bilingual human-machine dialogue system.

1 Introduction

Syntactic ambiguity is a well-known problem for natural language interpretation systems, and it is one which becomes increasingly serious as the syntactic coverage of a system increases. A great deal of research in NLP has focussed on how to avoid or alleviate the problem, and several different (and reasonably complementary) approaches have been devised. Four broad classes of approach can be distinguished. The oldest approach involves devising structural heuristics for disambiguation; these are often psycholinguistically-inspired, such as the minimal attachment strategy of Ferreira and Clifton (1986). Another class of approaches rely on the deployment of world or situation knowledge to assess alternative interpretations of a sentence. These approaches re-

quire interpretation systems which generate fairly rich semantic representations for sentences, along with large knowledge bases of facts about the domain of interpretation, along with sophisticated inference mechanisms to operate on them. Limitations on the size of knowledge bases and on the speed of theorem-provers mean that this approach is often hard to scale to real applications. (However, there are a number of promising approaches for tackling the scalability problem; see for instance Bos (2001), Hobbs (1993), Stone (1998).) But by far the dominant approach to syntactical ambiguity resolution at the moment involves the use of statistical techniques in one form or another. In particular, using statistical grammars which take into consideration lexical dependencies (see e.g. Collins (1996), Goodman (1998), Magerman (1995)) has proved a very sensitive way of capturing the kind of grammatical contexts that particular words or word combinations typically appear in; this information can be effectively used to decide between alternative parses.

However, there remain some species of ambiguity which are hard to resolve using statistical parsers. It frequently happens that at least two of the alternative readings of a sentence are relatively frequently attested in the corpus used to train the parser. Some of the classical illustrations of syntactic ambiguity can be used to make the point. For instance, there are two readings of the following well-known sentence:

- (1) The man saw the girl with the telescope.

In one reading, the PP *with the telescope* attaches to the NP *the girl*, while in the other, it attaches

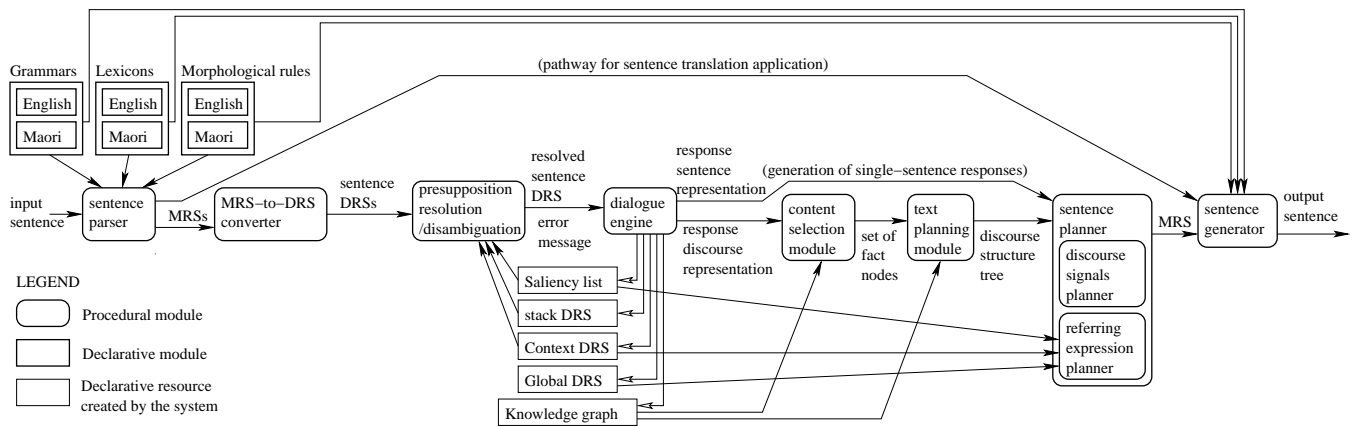


Figure 1: Architecture of the Te Kaitito system

to the VP *saw the girl*. The ambiguity in this sentence is clear, because both the alternative readings describe situations which are consistent with our world knowledge: telescopes can be used for seeing, and people can possess telescopes. In a representative corpus of English, we can expect *with [telescope]* (or perhaps, backing off, *with [instrument]*) to occur quite frequently modifying NPs as well as VPs. For the same reason, we cannot expect world knowledge to be very useful in distinguishing between the two alternative parses. A more useful source of information for disambiguation is the immediate discourse context. Simply put, if the reader knows that the man has a telescope *in the current discourse context*, one of the readings is preferred; if the reader knows that the girl has a telescope in the current context, the other reading is preferred.

In this paper, we present a simple method for making use of this kind of contextual information in syntactic disambiguation. The key idea is to allow the process of **presupposition resolution** to provide information about the contextual appropriateness of each alternative reading of a sentence (or perhaps each of the readings which a syntactic parser considers reasonably likely). In Section 2, we introduce the general framework we will be working in, by describing our human-machine dialogue system and the semantic representations it uses. In Section 3, we outline the (fairly simple) approach to disambiguation we have in mind, and give an example. In Section 4, we describe one particularly useful application of


the approach, for dealing with sentences which have ambiguous information structure.

2 Presuppositional DRT in the Te Kaitito system

Our sentence interpretation system is embedded in a larger human-machine dialogue system called Te Kaitito¹ (see Knott *et al* (2002) for a general overview). The architecture of Te Kaitito is given in Figure 1. When it is the user’s turn to contribute to the dialogue, (s)he enters a sentence, in written form. (Te Kaitito is bilingual between English and Māori, so either language can be used.) The sentence is first parsed, using the LKB system (Copestake, 2000), and a set of syntactic analyses are computed. Each analysis is associated with a semantic interpretation in the Minimal Recursion Semantics (MRS) formalism of Copestake *et al.* (1999). Each MRS representation is then converted to a representation in Discourse Representation Theory (DRT) (Kamp and Reyle, 1993), as extended by van der Sandt (1992) to incorporate a treatment of presuppositions. Each of the sentence’s DRSs then has its presuppositions resolved using information about the current discourse context, and one of the DRSs is selected as the preferred interpretation of the sentence—this operation is the focus of the current paper. The preferred interpretation is then passed to the dialogue engine, and a response is generated.

¹Online demos of Te Kaitito can be found at <http://tutoko.otago.ac.nz:8080/teKaitito/>.

Here is a brief introduction to the relevant concepts from DRT. The dialogue context and incoming sentences are modelled with Discourse Representation Structures (DRSs). A DRS is a structure with two fields, one for representing **discourse referents**, and one for representing **conditions** or predications over these referents. DRSs are typically drawn as split boxes, where referents appear at the top, and conditions below. For example, here is the DRS for the sentence *A cat walked*:



The discourse referent x is created by the indefinite NP *a cat*. This shows that *a cat* has introduced a new discourse referent. The conditions $cat(x)$ and $walk(x)$ were placed in the bottom part by the NP *a cat* and the VP *walked*.

A sentence's **presuppositions** are elements of its content which the speaker assumes are already part of the common ground; they are constraints on the kinds of context in which the sentence can be uttered. Here are two examples.

- (2) The dog chased a cat.
- (3) John's cat slept.

Sentence 2 presupposes that there is a dog in the discourse context (or more precisely, that there is exactly one *salient* dog in the context). Sentence 3 presupposes that there is someone called John, and also that this person has a cat. Presuppositions are triggered by lexical items such as the definite article, proper names, and possessive forms. These triggers determine what is asserted information, and what is presupposed in a given sentence.

As already mentioned, we use a DRT-based treatment of presuppositions as proposed by van der Sandt (1992). A sentence is modelled as an **assertion DRS** and a set of **presupposition DRSs**. The DRSs for Examples 2 and 3 are shown in Figures 2 and 3. Notice that the presupposition DRSs are distinguished by dashed lines.²

²One difference between our model and that of van der Sandt's is that there is no hierarchy of presuppositions in




Figure 2: *The dog chased a cat*




Figure 3: *John's cat slept*

The presuppositions of a sentence need to be **resolved** or **satisfied** in the current discourse context before its assertional content can be processed. In van der Sandt's DRS-based treatment, this is modelled as a binding operation: the referents in each of the sentence's presupposition DRSs need to be bound to referents in the context DRS which have the properties identified in the presupposition DRS. Once this binding has been done, if the presuppositions of a sentence are not satisfied, referents with suitable properties can be (charitably) assumed to exist, and added to the context DRS, in an operation called **accommodation**.

3 Presuppositional weight for sentence disambiguation

Our basic idea for disambiguating between alternative readings of an ambiguous sentence is to choose the reading that can most easily be incorporated into the discourse context. To implement this we use the following general strategy: first, generate the DRS for each alternative interpretation of an ambiguous sentence; then perform presupposition resolution for each interpretation within the current discourse context; finally, choose the most contextually appropriate interpretation. We use the following two principles to determine preference between the alterna-

our model: they all appear at the same level. In van der Sandt's model, a presupposition DRS can itself have presuppositions.

tives:

1. **Accommodation principle:** the most contextually appropriate interpretations of a sentence are those whose presuppositions can be resolved with the minimum amount of accommodated material.
2. **Presuppositional weight principle:** if two interpretations of a sentence can both have their presuppositions resolved *without* accommodation, the most contextually appropriate interpretation is the one with the greatest amount of presuppositional material.

One point to clarify in relation to these principles is how to determine the ‘amount’ of accommodated or presupposed material in the interpretation of a sentence. We will assume that this relates simply to a count of DRS conditions. For instance, for the accommodation principle, an interpretation which requires only one DRS condition to be accommodated into the discourse context is preferable over an interpretation which requires two conditions to be accommodated. For the presuppositional weight principle, given two interpretations whose presuppositions are both satisfied without recourse to accommodation, we simply count the total number of presupposed DRS conditions in each interpretation and choose the one with the highest total. It may be that there are alternative definitions for this notion of ‘amount of semantic material’, for instance including a count of the number of referents introduced, or the number of separate presupposition DRSs. But a simple account will suffice for the examples we deal with in this paper.

An example

Let us return to Example (1), repeated below:

- (4) The man saw the girl with the telescope.

There are two possible readings for this sentence. Both readings presuppose a man, a girl, and a telescope, and both assert that the man saw the girl. The DRS that represents the reading in which the man used the telescope to see the girl is given in Figure 4. The other possible reading, in which the girl possesses the telescope, is given in Figure

5. Notice that the second reading has more presuppositional content, or greater presuppositional weight, than the first.




Figure 4: ... [saw [the girl] [with the telescope]]




Figure 5: ... [saw [the girl [with the telescope]]]

If we know of a girl who has a telescope (see e.g. Figure 6(a)), then the reading in which the girl possesses the telescope would be preferred. In a context like this all the presuppositions of both readings could be resolved. However, we can choose the second reading, by the presuppositional weight principle, since it is ‘heavier’ in resolvable presuppositions. In a discourse context where there is a man, a girl, and a telescope, but the girl does not possess the telescope (see e.g. Figure 6(b)), only the presuppositions of the first reading can be satisfied. In this case, this reading is preferred, by the accommodation principle.

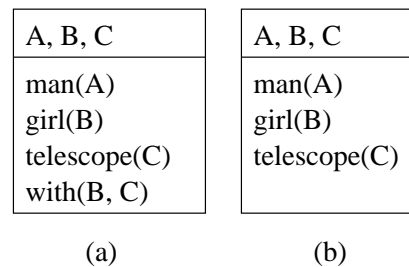


Figure 6: Two alternative contexts for Sentence 1

4 Application: sentences with ambiguous information structure

In this section, we describe a use for the two principles just outlined in the interpretation of sentences with ambiguous information structure. To begin with, we will illustrate this kind of ambiguity. Consider the following sentence:

(5) The cat chased the dog.

The most obvious reading of the sentence is that the speaker is simply asserting a new fact a propos of nothing, namely that the cat chased the dog. However, there are alternative possible readings in which the speaker is answering a wide variety of *questions*: for instance, ‘What did the cat chase?’, ‘What chased the dog?’, ‘What did the cat do to the dog?’ and so on. If the answers were spoken, they would have different prosodic and intonational structures: *THE CAT chased the dog*, *The cat chased THE DOG*, *The cat CHASED the dog*, and so on. In a grammar of English, it is important to distinguish these alternative readings even when there are no prosodic cues; any declarative sentence can genuinely be interpreted in these different ways, and the grammar should reflect this. However, computational grammars typically do not pay much attention to information structure ambiguity,³ principally because once the ambiguity is created, it is hard to resolve. Information structure ambiguity is prototypically a kind of ambiguity which cannot be resolved using statistical techniques or knowledge about the domain; it requires an analysis of the relation between the sentence and its immediate discourse context.

In this section, we describe how information structure ambiguity of this kind can be resolved using a presuppositional framework. We begin by providing some background about the way questions and answers are represented using presuppositions in our system.

4.1 Background: presuppositions in questions and answers

In the Te Kaitito system we model questions and answers using presuppositions (for details, see de

³Categorial grammar is a notable exception; see Steedman (2000).

Jager *et al.* (2002)). A question can be thought of as asserting that something is ‘unknown’,⁴ and presupposing the content of the query to exist within the current discourse context. For example, consider the question:

(6) What chased the dog?

This question presupposes that something did chase the dog. The only information it asserts is that the something which chased the dog is unknown to the speaker. This is represented as a DRS in Figure 7.




Figure 7: *What chased the dog?*

An answer to a question also makes use of presuppositions. An answer presupposes a question containing an unknown element, and asserts that this unknown element is identical to some referent within the discourse context. An answer to the question in Example 6 could be given as *It was the cat* or just *The cat*. This answer is represented as a DRS in Figure 8. For further details




Figure 8: *The cat / It was the cat*

about this presuppositional treatment of questions

⁴Strictly speaking, we are here asserting something about the speaker’s *knowledge* about the object, rather than about the object itself. But in our system, we draw no distinction between this kind of epistemic predicate and ordinary first-order predicates.

and answers, and a discussion of some of its additional benefits in a dialogue management system, see de Jager *et al.* (2002).

4.2 Using presuppositional weight for disambiguation

One could also answer the question in Example 6 using the full sentence *The cat chased the dog* (c.f. Example 5), with understood emphasis *THE CAT chased the dog*. How do we distinguish this reading from the other possible readings? The DRSs in Figures 9-11 show three ways of interpreting *The cat chased the dog*. Using this rep-

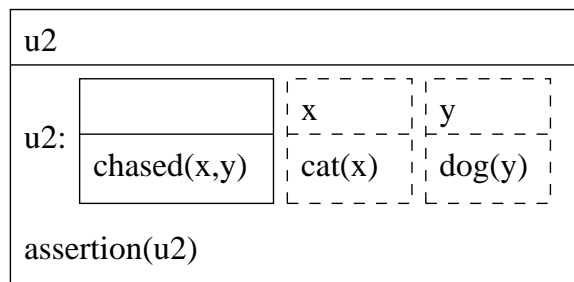


Figure 9: *The cat chased the dog*

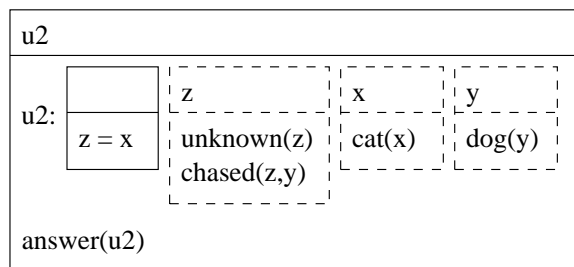


Figure 10: *THE CAT chased the dog*

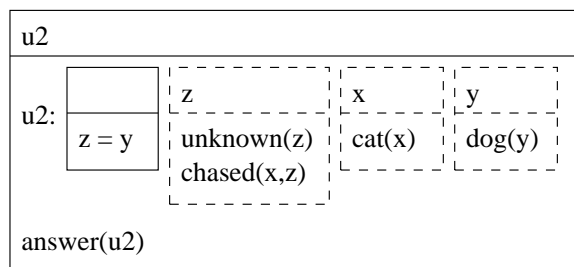


Figure 11: *The cat chased THE DOG*

resentation, the principles presented in Section 3

can be used directly to derive the right interpretation of the sentence, namely that given in Figure 10. Firstly, note that both the interpretation in Figure 9 and that in Figure 10 can have all their presuppositions resolved without accommodation in the context created by the question (Figure 7), while the interpretation given in Figure 11 contains a presupposition which needs to be accommodated in this context (an unknown object which is chased). By the accommodation principle, we can reject the interpretation in Figure 11. Now note that the interpretation in Figure 10 has a larger number of presupposed DRS conditions than the interpretation in Figure 9 (4 versus 2). By the presuppositional weight principle, we can therefore choose the interpretation in Figure 10, which is the correct interpretation for the sentence.

5 Summary and conclusions

In this paper, we have presented a simple approach to sentence disambiguation, which makes use of presupposition resolution, and we have described two example contexts where the approach is useful. We have argued that this approach is likely to be helpful in cases where other methods for disambiguation have difficulties, in particular in cases where statistical methods cannot deliver a clear verdict because two (or more) interpretations of a sentence contain constructions which are frequently attested in training corpora, and where methods based on determining consistency with world or domain knowledge have a similar problem; information structure ambiguity is a good case in point.

Naturally, the information about contextual appropriateness delivered by the presupposition resolution mechanism should just be considered as one source of information about the appropriate reading of a sentence. Other sources of information should also be taken into account. A statistical parser is still likely to be of considerable use in weeding out very unsuitable readings at the very start, because presupposition resolution is computationally quite expensive. And information from world or domain knowledge about the absolute likelihood of alternative readings is also bound to be important. Our main conclusion is simply that information about presuppositions should have a

useful role to play in a complete framework for sentence disambiguation.

Acknowledgements

This work was supported by University of Otago Research Grant MFHB10, and by the NZ Foundation for Research in Science & Technology (FRST) grant UOOX02.

References

- J Bos. 2001. DORIS 2001: Underspecification, resolution and inference for discourse representation structures. In *Proceedings of ICoS-3: Inference in Computational Semantics*.
- M Collins. 1996. A new statistical parser based on bigram lexical dependencies. In *Proceedings of the 34th Meeting of the ACL*, Santa Cruz.
- A Copestake, D Flickinger, I Sag, and C Pollard. 1999. Minimal Recursion Semantics: An introduction. Manuscript, CSLI, Stanford University.
- A Copestake. 2000. The (new) LKB system. CSLI, Stanford University.
- S de Jager, A Knott, and I Bayard. 2002. A DRT-based framework for presuppositions in dialogue management. In *Proceedings of the 6th workshop on the semantics and pragmatics of dialogue (EDILOG 2002)*, Edinburgh.
- F Ferreira and C Clifton. 1986. The independence of syntactic processing. *Journal of Memory and Language*, 25:348–368.
- Joshua Goodman. 1998. *Parsing Inside-Out*. Ph.D. thesis, Harvard University.
- J Hobbs, M Stickel, D Appelt, and P Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63.
- H Kamp and U Reyle. 1993. *From discourse to logic*. Kluwer Academic Publishers, Dordrecht.
- A Knott, I Bayard, S de Jager, and N Wright. 2002. An architecture for bilingual and bidirectional nlp. In *Proceedings of the 2nd Australasian Natural Language Processing Workshop (ANLP 2002)*.
- David M. Magerman. 1995. Statistical decision-tree models for parsing. In *Proc. of the 33rd Annual Meeting of the Association for Computational Linguistics*. Cambridge, MA, 26–30 Jun 1995.
- M Steedman. 2000. *The syntactic process*. MIT Press/Bradford Books.

M Stone. 1998. *Modality in dialogue: planning, pragmatics and computation*. Ph.D. thesis, Institute for Research in Cognitive Science, University of Pennsylvania.

R Van der Sandt. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics*, 9:333–377.