

Interpretation of Definite Reference with a Time-Constrained Memory

Jean-Pierre Corriveau

Department of Computer Science

University of Toronto,

Toronto (Ontario) CANADA M5S 1A4

(819) 682-1279

jeanpier@bnr.ca

Abstract

In this paper, I demonstrate how complex cases of definite reference resolution can be processed within the independently-motivated framework of *time-constrained memory*, and, most importantly, without having to resort to the complex mechanisms assumed by Haddock (1987). The key idea of the solution is to initiate a referent search for the complex definite NP formed by the attachment of a prepositional phrase to a definite noun phrase.

Haddock's Problem

Introduction

The goal of this paper is to solve a problem in the interpretation of definite reference that was first put forth by Haddock (1987). The linguistic framework provided by *time-constrained memory* (Corriveau, 1991) offers an approach that avoids the complexity and drawbacks of the mechanisms that Haddock assumed to be necessary for solving this problem.

The problem is this: Suppose we have a context in which there are three rabbits, R1, R2, and R3, two hats, H1 and H2, and one box, B1. One of the rabbits, R2, is in one of the hats and another one is in the box, as illustrated in figure 1. Rabbit R2 can be referred to by means of the following complex NP:

0.1 *the rabbit in the hat*

What is interesting about this expression is that even though as a whole it is perfectly natural, the NP *the hat* uses the definite determiner *the*, despite the fact that there are not one but two hats in context. Compositional accounts of NP semantics would judge (1.1) to be infelicitous. Haddock uses this fact to argue for an incremental approach




Figure 1: The context for Haddock's example

that “evaluates a semantic representation—after each word, say” and will never consider the empty hat as a “viable candidate for the inner NP” (1987, p.661).

When the word *rabbit* is reached, a hearer can collect together in his mind the set of rabbits in the context. After the preposition, this set is refined to contain only rabbits which are *in* something and, most importantly, the hearer can start thinking about another set of objects, those which have rabbits in them. There is only one hat in this new set and so by the time the inner NP is processed a definite determiner sounds natural. (*ibid.*)

In more explicit computational terms, Haddock suggests that, given the predicates and variables corresponding to the input discourse, the algorithm first determines, for each variable, all entities in the discourse that satisfy the predicates, and second, collects these into candidate sets for each variable. From the fragment *the rabbit in*, we get the predicates *rabbit*(e_1) and *in*(e_1, e_2).

Given the context of (1.1), the algorithm will label the variable e_1 with the set $\{R2, R3\}$ and e_2 with $\{H1, B1\}$. Haddock emphasizes that the advantage of this scheme

resides in that satisfying the constraints of the two predicates determine a candidate set for e_2 as well as

e1. Because the phrase is evaluated strictly word-by-word, in left-to-right order, this scheme is also supported by psycholinguistic evidence (Steedman, 1987).

Haddock assumes that it is syntax that controls the assembling of semantic representations, and thus, in order to provide incremental semantics, adopts the framework of Combinatory Categorical Grammar (Ades and Steedman, 1982). My intent in this paper is to present an alternative approach to the interpretation of (1.1), one that does not require all the complexities, such as *type-raised* categories, of Combinatory Categorical Grammar.

Difficulties with Haddock's Solution

Unfortunately, Haddock's approach exhibits drawbacks. First, because his approach is syntax-driven, it depends heavily on the unambiguous attachment of *in* to *the rabbit*. But it is commonly accepted that structural disambiguation requires more than syntax and often involves semantics and inference (Hirst, 1987). Thus, because later inferences, for example, may influence the attachment of the PP, it seems incorrect to have the semantic processing of the preposition *immediately* restrict the set of rabbits to those that are *in* something. Furthermore, it may not even be the case that the preposition heads a locative proposition. In other words, further input text might be required for correct attachment.

Second, a solution based on the attrition of arbitrary large sets is not cognitively appealing. If, for example, there were a multitude of rabbits in various objects, the real-time processing constraint of human comprehension would not allow the user to mentally assemble the set of all rabbits nor the set of all containing objects. Instead, the attentional mechanisms of the comprehender's memory would in practice restrict the set of possible referents.

Third, the comprehender does not necessarily start thinking about objects that have rabbits in them after processing *in*. Not only could this approach lead to a large, unmanageable set of rabbit containers, but it also seems to depend on the *inference strategy* of the comprehender (Granger and Holbrook, 1983). Haddock's approach corresponds to the particular case in which a comprehender, on processing the preposition *in*, sets up an expectation about its attachment to the NP *the rabbit* and, consequently, comes to expect that what follows the preposition must refer to one of the objects that contain a rabbit. This is just one of several possible reference resolution strategies. Another tactic may be to wait for further input and for correct attachment, and only then look for a referent for the resulting complex NP. In other words,

another valid resolution strategy is to wait for the complex NP *the rabbit in the hat* to be constructed before looking for a referent. The precise reference resolution scheme can vary across comprehenders, and even, for a single comprehender, can change depending

on his/her state of mind.

Finally, though I will not elaborate on this topic, I have serious concerns about any linguistic approach that assumes a more or less direct correspondence between words and simple predicates (see Corriveau, 1991, chapter 2).

The Proposal

The approach I suggest circumvents the identified drawbacks of Haddock's. In essence, I assume that each definite NP is given a short amount of *time* to find a referent. Finding a referent is taken to consist in searching the cognitive structures that have been built from the processing of previous input for an entity that 'matches' the definite NP. The set of possible referents is dynamically dependent on the amount of time allocated to the search, as well as the traversal order of cognitive structures in memory. If, at the end of the time-constrained search, a single referent is found, the definite NP is resolved. If no referent is found, then the definite NP is left unresolved. And, most importantly, if ever a second possible referent is found during the search, the search is immediately abandoned and the definite NP is taken to be ambiguous.

So, assuming all possible referents are in context, the processing for (1.1) would be the following:

- The NP *the rabbit* is assembled and the search for a referent is initiated. This search finds a first possible referent, then a second one, at which point it is abandoned.
- The preposition *in* is input, and initiates a process that will eventually attach the NP it governs, namely *the hat*, to the NP *the rabbit*. Such an attachment would be prevented in sentences such as *Put the rabbit in the hat*.
- The words *the* and *hat* are input and assembled into an NP. The search for a referent for this NP is initiated. This search fails because there are two hats in context, that is, two possible referents.
- Eventually the PP *in the hat* is attached to the NP *the rabbit*. This forms a new definite NP for which a new search for a referent is initiated. This search reaches R2, which is considered a possible referent because it is a rabbit in a hat. Because no other possible referent is found before the end of the search, it succeeds and the complex NP *the rabbit in the hat* is resolved.

Overview of Time-Constrained Memory

Let me now express this solution in more explicit computational terms.

The solution I propose for the interpretation of (1.1) is rooted in the notion of a strictly quantitative *time-constrained memory*. In Corriveau (1991), I motivate and develop a linguistic framework specified in terms of a computational model of memory that, although

following a simple, knowledge-independent algorithm, is able to account for many problems of understanding written language. Here, I can provide only a brief review.

Of the multitude of possible factors that constrain linguistic comprehension, I focus primarily in this work on the role and importance of quantitative *time*, *i.e.*, time as it pertains to memory management and memory processes such as retrievals (see Corriveau, 1987). More specifically, the fundamental and most pervasive hypothesis of this work is that, generally, linguistic comprehension is a *time-constrained process*—a *race*.

I abandon the notion of a correct interpretation of a text and, rather, adopt a *reader-based* approach to text understanding, one in which ‘meaning’ is constituted by the interaction between text and reader (Holub, 1984). The reader is not seen as an idealized competent entity, but rather as an individual (with all the idiosyncrasies implied by this term) that comes to a personal interpretation of a text with respect to a private idiolect. Therefore, the concept of linguistic *competence* (which is inaccessible by definition) is abandoned. I focus only on actual *performance*, which is taken to be idiosyncratic.

Thus, this approach shifts the focus from rules and algorithms to the underlying cognitive architecture which is taken to be strictly quantitative: “With the abandonment of algorithmic control, the brain can be hypothesized to follow a strictly quantitative ‘trivial algorithm’ that fixes only the general form of operations; all qualitative information (*e.g.*, rules, concepts, etc.) is treated as data.”

I choose to root the cognitive architecture that I propose in the basic metaphor of human memory. The system is partitioned into static and dynamic components of memory. Static memory consists of a large network of simple computing units that operate in parallel and exchange simple signals. These units are called *knowledge units* (KUs) and have the ability to modify the current context, which is simply taken to be the set of *clusters* in dynamic memory. More specifically, a KU has its *expansion procedure* execute upon it becoming activated. Expansion procedures consist of ordered sequences of primitive cluster operations, which are built-in instructions of the model and are independent of any linguistic consideration. These operations provide the basic functionality required for generic data structures: access, addition, deletion, comparison, traversal, etc. Among these operations, one that is particularly relevant to a solution to the problem of definite reference: *findInclusiveReference*. This operation is intended for general matching. It takes as an argument a cluster, and searches dynamic memory for another cluster whose set of features includes all those of the argument. The operation is given a limited amount of time for its search, and it fails either if multiple referents can be found during its time-span or

if no referent at all can be found within the time-span. A straightforward approach to reference resolution can be developed in a single KU using this instruction.

A cluster is defined as a non-empty set of knowledge units, each

governing a (possibly empty) set of clusters. In other words, a cluster is a hierarchical structure whose ‘leaves’ are knowledge units. Clusters are general enough to allow the construction of almost any kind of representation (*e.g.*, parse trees, scripts, etc.) that might be hypothesized for comprehension.

Among the six kinds of signals that KUs can exchange, the *expectation signal* allows the specification of exactly the sort of expectation that Haddock assumes of the hearer who starts thinking about objects that contain rabbits.

Within the framework of my basic metaphor:

- All memory processes are taken to be strictly quantitative, *i.e.*, mechanical and deprived of any linguistic and semantic knowledge.
- All ‘knowledge’, that is, all qualitative information, is assumed to be specified separately and act as data for the underlying system.
- The processes of linguistic comprehension must be defined only in terms of memory processes.

A prototype of time-constrained memory has been implemented and a knowledge base of more than 300 knowledge units has been specified to address the problems of syntax, reference resolution, word-sense and structural disambiguation, and bridging inferences required for written text comprehension.

Examples

I now show a number of working examples of definite reference resolution in this system.

- **0.2** *R1 is a rabbit. R2 is a rabbit. The rabbit eats.*
Multiple referents are found, and thus there is no resolution.
- **0.3** *R1 is a rabbit in a hat. The rabbit in the hat eats.*
In the second sentence, the PP is attached to the NP *the rabbit* and then a unique referent is found.
- **0.4** *R1 is a rabbit. R2 is in a box. R1 is in a hat. R2 is a rabbit. The rabbit in the box eats.*
In the last sentence, the disambiguation of *the rabbit* fails because of multiple referents. Then a unique referent is found for *the rabbit in the box*.
- **0.5** *The young rabbit eats. The rabbit sleeps.*
A unique referent is found, although the second NP is not identical to the first.
- **0.6** *Put the rabbit in the hat.*
Simple VP-PP attachment occurs, which blocks, in this case, the NP-PP attachment.

- 0.7 *R1 is a rabbit in a hat. Put the rabbit in the hat.*

In the second sentence, a referent is found for *the rabbit* and the PP is attached to the verb, not to *the rabbit*.

- 0.8 *R1 is a rabbit. R2 is a rabbit. Put the rabbit in the box.*

In the third sentence, no referent is found for *the rabbit*, and there is only simple VP-PP attachment.

- 0.9 *R1 is a rabbit in a hat. R2 is a rabbit. Put the rabbit in the hat in the box.*

The third sentence is interpreted as: *Put in the box the rabbit that is in the hat.*

- 0.10 *Put the rabbit.*

A syntactic conflict is detected, which prevents any attempt at reference resolution.

Conclusion

In this paper, I have demonstrated that complex cases of definite reference resolution can be processed in a straightforward way within the independently motivated framework of *time-constrained memory*, and, most importantly, without having to resort to the complex mechanisms assumed by Haddock (1987). Moreover, the proposed strategy hinges on a strictly quantitative memory operation, which is to say it is not a special-purpose linguistic rule. The key idea of the solution is to initiate a referent search for the complex definite NP formed by the attachment of a PP to a definite NP. This solution is very efficient because it immediately abandons a search and leaves a definite NP unresolved if more than one possible referent can be found.

References

Ades, A. and Steedman, Mark J. (1982). "On the Order of Words". *Linguistics and Philosophy*, 4:517-518.

Corriveau, Jean-Pierre (1987). "On the Role of Time in Reader-Based Comprehension", *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, Seattle, WA, pp. 794-801.

Corriveau, Jean-Pierre (1991). *Time-Constrained Memory for Reader-Based Text Comprehension*, Doctoral Dissertation, Department of Computer Science, University of Toronto, Toronto, Canada.

Granger, Richard and Holbrook, Jennifer (1983). "Perseverers, Recencies and Deferrers: New Experimental Evidence for Multiple Inference Strategies in Understanding", *Proceedings of the Fifth Annual Conference of the Cognitive Science Society*, Rochester, NY.

Haddock, Nicholas (1987). "Incremental Interpretation and Combinatory Categorical Grammar", *Proceedings of the Tenth Joint Conference on Artificial Intelligence*, Milan, Italy, pp. 661-663.

Hirst, Graeme (1987). *Semantic Interpretation and the Resolution of Ambiguity*, Cambridge University Press, Cambridge, England.

Holub, Robert (1984). *Reception Theory—A Critical Introduction*, Methuen, New York.

Malsburg, Charles von der (1985). "Algorithms, Brain, and Organization", In: Demongeot, J., Golès, E. and Tchente, M. (eds.) *Dynamical Systems and Cellular Automata*, Academic Press, London.

Steedman, Mark J. (1987). "Combinatory Grammars and Human Language Processing", In: Garfield, J. (ed.) *Modularity in Knowledge Representation and Natural Language Processing*, MIT Press, Cambridge, MA.