When Something Is Missing: Ellipsis, Coordination and the Chart

Alberto Lavelli and Oliviero Stock

Istituto per la Ricerca Scientifica e Tecnologica

I - 38050 Povo TN, Italy

lavelli@irst.it / stock@irst.it

Abstract

This paper deals with two linguistic phenomena which are usually considered cases of ill-formedness by the computational linguistics community: intersentential ellipsis and coordination (possibly with *gaps*). We present an original solution, if compared to those already known for the two phenomena. This solution is conceived within a relevant approach to parsing, i.e. chart parsing, and is coherent with the basic ideas of this approach.

1. Introduction

The ability to face and resolve problems associated with ill-formedness is fundamental in order to make natural language interfaces usable (see [Carbonell and Hayes, 1983] for a review of the problems of ill-formedness and of techniques used to resolve them). We shall focus on two phenomena: intersentential ellipsis and coordination (possibly with gaps).

Ellipsis is a very common phenomenon and is frequently encountered in dialogues between persons. Up to the present, studies on natural language interaction with computers generally highlight the frequency of this phenomenon, (see, for example, [Eastman and McLean, 1981]).

For this reason, ellipsis has received much attention and different solutions have been proposed based on the mechanism used to analyze the sentence: semantic grammars (the LIFER/LADDER system [Hendrix, 1977]), ATN [Kwasny and Sondheimer, 1981, Weischedel and Sondheimer, 1982], or case-frame instantiation (the XCALIBUR system [Carbonell *et al.*, 1983]).

As far as coordination is concerned, it is also frequently considered a phenomenon of ill-formedness for the following reasons:

- since every pair of constituents of the same syntactic category may be coordinated, if the grammar specifically included all these possibilities it would greatly increase the size of the grammar itself;
- a constituent inside a coordination may have gaps (that is, missing elements) that, in general, are not allowed in constituents of the same type.

Even in the most purely linguistic area, coordination has not received in-depth treatment, if not recently (for example, see [Sag *et al.*, 1984] and [Kaplan and Maxwell, 1988]).

Until the beginning of the 80's almost no system (the most relevant exception being the SYSCONJ module, present in the LUNAR system [Woods, 1973]) has confronted coordination in a generalized

manner. The 80's have seen renewed computational interest in coordination that has brought new efforts (see [Kwasny and Sondheimer, 1981], [Dahl and McCord, 1983], [Fong and Berwick, 1985], [Lesmo and Torasso, 1985], [Kosy, 1986] and [Proudian and Goddeau, 1987]).

In this paper we present a solution to both the problems outlined above (that is, ellipsis and coordination). The solution is original with respect to those presented in the literature in that it is based on chart parsing [Kaplan, 1973, Kay, 1980], an approach little used to treat ill-formedness until now.¹ For both problems the solution is based on a strategy that uses the information contained in the structures produced by the parser (as is the case in almost all the work mentioned in this introduction; for a pragmatics-based approach to ellipsis, see [Carberry, 1989]).

Both the solutions proposed have been inserted into WEDNESDAY 2 [Stock, 1989], a chart-based parser used in a dialogue system for Italian; attention has been paid to distinctive aspects of this language (for example, the relative liberty of order inside a single constituent). The process of building the analysis is based on a sort of unification. The parser is built so that it may be connected to a semantic (and possibly pragmatic) discrimination component; in this way, the discrimination is carried out incrementally as the syntactic analysis proceeds. All examples in this paper are in Italian.

2. Ellipsis

First, consider an example of ellipsis (upon which an informal description of this solution will rely):

User: Giotto ritrasse S.Francesco in
un'opera di Assisi?
[Did Giotto portray St.Francis in a
masterpiece located in Assisi?]

System: Sì, ad esempio nella Predica agli
uccelli.
[Yes, for example in the Sermon to the
birds.]

User: Gioacchino in un affresco di Padova?
[Gioacchino in a fresco located in
Padova?]

In the elliptical sentence two constituents may be distinguished: the noun phrase *Gioacchino* and the prepositional phrase *in un affresco di Padova*. Substitution of some constituents of the preceding sentence is foreseen with the recognized constituents

¹ Always in the area of chart-based parsing see [Mellish, 1989], who, however, deals with other problems of ill-formedness, i.e. the presence of unknown words and of spurious constituents.

in the elliptical sentence, respecting, in essence, the ordering of the two sentences (this means that the elliptical sentence *In un affresco di Padova Gioacchino?* [lit. *In a fresco located in Padova Gioacchino?*] would not be recognizable and the associated ellipsis would not be resolvable). In this case it is necessary to find the points in the preceding sentence in which to insert the fragments *Gioacchino* and *in un affresco di Padova*. Once the points of insertion are found, the parser restarts. The semantic correctness of the proposed hypothesis is checked through the usual semantic control during unification. In this example, the first possible solution, starting from the left and moving right, is:

- *Gioacchino* instead of *Giotto*
- *in un affresco di Padova* instead of *in un'opera di Assisi*

which gives the sentence *Gioacchino ritrasse S. Francesco in un affresco di Padova?* [*Did Gioacchino portray St. Francis in a fresco located in Padova?*]. Actually, already during an initial unification this possibility is rejected for semantic and contextual reasons (given that there is no painter named *Gioacchino*). The following substitution is then tried:

- *Gioacchino* instead of *S. Francesco*
- *in un affresco di Padova* instead of *in un'opera di Assisi*

which produces the sentence *Giotto ritrasse Gioacchino in un affresco di Padova?* [*Did Giotto portray Gioacchino in a fresco located in Padova?*]. This represents the correct interpretation of the ellipsis.

It must be emphasized the fact that work is never duplicated; this positive characteristic of the chart is respected by our solution since (as will be seen in the algorithm) the base mechanism of the general case is preserved.

2.1. The algorithm

Now the algorithm for resolving ellipsis will be described more formally. That will be done by considering the most complex case, substitution (even multiple), in the general case. Firstly, provision is made for saving the chart, called *model chart*, built for the preceding sentence (in which all the edges, active and inactive, that have contributed to building the correct analysis for that sentence are saved). Then, an attempt is made to analyze the next elliptical sentence, partially building another structure, called *fragments chart* (the construction of this second structure can be carried on and concluded during resolution of the ellipsis).

The algorithm for treatment of ellipsis consists of alternate scanning of the *model chart* and *fragments chart*. When a vertex is found in the *model chart* from which a cycling edge exits that is of the same category as the inactive edge *I'* exiting the vertex of the *fragments chart* upon which it is positioned, then the inactive edge *I* (generated from the given cycling edge) of the *model chart* is substituted by the inactive edge *I'* of the *fragments chart*. Thereafter the parser is restarted by the usual mechanism (the substitution concerns not only edge *I*, but all the inactive edges that that edge

subsumes and all the active edges that subsume edge *I* or some of the edges subsumed by *I*. This aspect of the algorithm is consistent with some of the techniques used in [Wirén, 1989]). In essence the algorithm is expectation-driven, given that the fragment is inserted where there is a cycling edge that wants it. In this way, the expectations created by the preceding sentence are taken into account to guide the search for a solution to the ellipsis.

Because an agenda is used to manage the chart, the algorithm is sufficiently simple and causes only limited changes in normal functioning of the parser. There are two different phases:

- i) activation of the mechanism (that is, the search for the vertex where to insert the first fragment that restarts the usual mechanism);
- ii) the next phase brings only effects of the modification of the function that selects the tasks contained in the Agenda.

The algorithm uses pointers to the vertices of the two charts: *Remaining-ModelChart* for the *model chart* and *Remaining-FragmentsChart* for the *fragments chart*, which are initialized at the initial vertex of the respective chart. First of all, consider the activation phase:

1. for all cycling edges *C* exiting from the vertex of *Remaining-ModelChart*, a check is done to determine whether among the inactive edges leaving the vertex of *Remaining-FragmentsChart* there is one, *I'*, of the same syntactic category as *C*; if yes, go to point 2, otherwise move *Remaining-ModelChart* to the next vertex in the *model chart* and return to the beginning of point 1;
2. save the preceding context (what this means will be made clearer at the end of the algorithm formulation);
3. remove the edge *I*, generated from the cycling edge *C*, from the *model chart* (together with those specified above) and put a task into the Agenda that inserts the edge *I'* into the *model chart*;
4. move the pointers *Remaining-ModelChart* and *Remaining-FragmentsChart* to the arrival vertex of the inactive edges selected in the *model chart* and *fragments chart* respectively;
5. start the normal mechanism (with the only change being that to the selection function indicated below) on the *model chart* thus modified;
6. if the process in point 5 does not succeed, it is necessary to backtrack, which means return to point 1 after having reestablished the preceding context and having moved *Remaining-ModelChart* to the next vertex of the *model chart*.

Now, we analyze the way in which the function that selects the tasks contained in the Agenda is modified:

1. when a task that extends an active edge *A* with an inactive edge *I* is to be executed, check whether among the inactive edges leaving the vertex of *Remaining-FragmentsChart* there is one, *I'*, of the same syntactic category as *I*; if yes, go to point 3, otherwise go to point 2;
2. move *Remaining-ModelChart* to the arrival vertex of edge *I* and exit from the function normally

executing the task selected;

3. save the preceding context (what this means will be made clearer at the end of the algorithm formulation);
4. the task selected is removed from the Agenda;
5. remove edge I from the *model chart* (together with those specified above) and put a task into the Agenda that inserts the edge I into the *model chart*;
6. move the pointers Remaining-ModelChart and Remaining-FragmentsChart to the arrival vertex of the inactive edges selected in the *model chart* and *fragments chart* respectively;
7. start the normal mechanism on the *model chart* thus modified;
8. if the process in point 7 does not succeed, backtrack, which means go to point 2 after having reestablished the preceding context.

Some aspects of the backtracking mechanism are now specified (and with them the reason for saving and reestablishing the contexts): backtracking is here intended exclusively as a nondeterministic strategy relative to the insertion of elliptical fragments. This does not impinge in any way upon the chart mechanism, which continues to provide flexible and efficient nondeterministic management. Furthermore, intuitively, the ellipsis resolution algorithm will only have to work on a very limited number of fragments, as they occur in man-machine interaction. This considered, the complexity of the algorithm remains, in practice, within acceptable limits (as well-known for the chart).

How this is applied to the example given at the beginning of this section (and repeated here) will now be considered:

User: Giotto ritrasse S.Francesco in un'opera di Assisi?

[Did Giotto portray St.Francis in a masterpiece located in Assisi?]

System: Sì, ad esempio nella Predica agli uccelli.

[Yes, for example in the Sermon to the birds.]

User: Gioacchino in un affresco di Padova?

[Gioacchino in a fresco located in Padova?]

Figure 1 shows the model sentence with the relevant cycling edges; Figure 2 shows the inactive edges for Gioacchino and in un affresco di Padova. The activation phase immediately brings into operation the substitution of the inactive edge that goes from vertex 1 to vertex 2 (Giotto) with the inactive edge of the *fragments chart* that represents Gioacchino; after this substitution and opportune changes to Remaining-ModelChart and to Remaining-FragmentsChart, the parser is restarted, but it fails because of a violation of semantic restrictions. This failure causes backtracking to begin and therefore to seek another vertex in which to make the substitution. This vertex is number 3, from which an inactive edge terminating in 4 that has recognized S.Francesco exists. This edge is then substituted with the inactive edge of the fragments chart that represents Gioacchino. The parser works normally until Remaining-ModelChart points to vertex 4 and the task selection function must extend the active S edge with an inactive PP edge (which extends from vertex 4 to vertex 9). In this case it can substitute the inactive PP edge [in un'opera di Assisi] with in un affresco di Padova and therefore restart the chart, which reaches the conclusion of the analysis.

Returning to a more general consideration on the working of the algorithm, note that sluicing (for example, Giotto dipinse affreschi in Veneto? [Did Giotto paint frescoes in Venetia?] Sai quando? [Do you know when?]) since it includes wh-words, causes the bottom-up introduction of a cycling edge for a sentence with long distance dependencies.

Expansion is much simpler - it is convenient to allow adjunctive adverbs to be inserted only at a fixed position in the sentence (e.g., extreme left), with obvious advantages of efficiency for the parser.

Figure 1. Model sentence with the cycling edges relevant for the example.

Figure 2. The inactive edges of the fragments Gioacchino and in un affresco di Padova.

Figure 3. Coordination rule.

We conclude this section with a couple of remarks. The first one concerns preferences for the insertion of fragments: for many kinds of dialogues it seems reasonable to use a heuristics that favors the insertion of the elements of the *fragments chart* all at the same level inside one constituent. The second one concerns unification (or any other mechanism that one would use for functional control in connection with the chart); this would prevent the analysis of: *Giotto nacque a Padova?* [*Was Giotto born in Padova?*] *I Lorenzetti a Siena?* [*The Lorenzetti brothers in Siena?*] (in which the subjects of the two sentences differ in some features, in this case the number). This aspect emphasizes the need to employ relaxation techniques in the unification mechanism in order to be able to accept this kind of elliptical sentence.

3. Coordination

If the parser uses a top-down strategy (as is usual in ATN and logic grammars), it must hypothesize a structure for the second conjunct without any knowledge of its actual structure. Since this may be any structure that parallels that of a constituent that ends immediately before the conjunction, the parser must build and check all these possibilities to find the right ones; all this leads to a combinatorial explosion of possibilities.

A chart-based approach allows a natural factorization in constituent construction, thereby limiting the exponential explosion of possible analyses, a negative characteristic of SYSCONJ and other systems. Moreover, a bottom-up strategy provides further information to guide the parser in the analysis of successive coordination constituents.

In the following, only coordination of constituents belonging to the same syntactic category and adjacent is considered.

Three changes to the normal chart mechanism are necessary to treat coordination (plus the introduction of metarules for the coordination of constituents containing gaps, see § 3.1) and concern:

- 1) the configuration of an inactive edge I_X followed by an inactive edge I_{And} (*conjunction type*), that causes the insertion of new active edges;
- 2) features that must be associated with coordinated edges;
- 3) for a system that incorporates a mechanism based on unification (see, for example, [Shieber, 1986]), the way in which the mechanism is activated when an active edge incorporates an inactive one.

The first change (shown in Figure 3) causes the insertion of an active edge A_X (of the same category as I_X and marked as destined to recognize a coordinated

structure) that *covers* the two inactive edges I_X and I_{And} , and of a cycling edge A'_X of the same category as I_X (if not there already). The cycling edge A'_X is meant to recognize the coordinated constituent.

As for the features associated with the edge, in case of coordination it may occur that a noun phrase has some features (in particular gender and number) different from those of the coordinated elements: for example, the noun phrase *il ragazzo e la ragazza* (*the boy and the girl*) has the features "Gender Masculine, Number Plural", in part different from those of the noun phrases that compose it. This modification is necessary for running control of agreement (as, for example, between subject and verb).

Finally, it is necessary to modify the unification mechanism so as to permit manipulation of sets of elements (as it occurs in coordination): such modifications allow correct management of the features of sets (as illustrated in the preceding paragraph) and must keep track of the semantic representation with which to label the sets.

An example of coordination of complete constituents is now considered, to illustrate what happens inside the chart:

Giotto e Orcagna dipinsero un affresco a Padova?
 [Did Giotto and Orcagna paint a fresco in Padova?]

Here coordination concerns the two noun phrases *Giotto and Orcagna*; obviously, the verb is conjugated in the third person plural because it agrees with the two coordinated noun phrases that must therefore have the features "Gender Masculine, Number Plural" (confirming the necessity of the change previously presented on the features of the sets formed because of coordination).

3.1. Coordination of constituents with gaps

Our approach to the problem of constituents containing gaps consists of introducing metarules associated with some configurations of the rules. The metarules allow an active edge to be inactivated under conditions for which this normally is not permitted (for example, when the *head* of a constituent has not yet been found). These metarules must be applicable only to active edges contiguous with a *conjunction type* inactive edge (thereby limiting the growth of inactive edges introduced by metarules). Introduction of metarules carries out only a part of the work necessary to treat coordination of constituents containing gaps; the remaining part must be carried out during unification of the edges that have recognized the two conjuncts. At this point it is necessary to fill the gaps, using information brought

by the complete constituent. Several proposals have been made for carrying out this role (among them [Sag *et al.*, 1984] and [Kaplan, 1987]). For example, *priority union* (proposed in [Kaplan, 1987] as a means of assigning correct interpretations to constructions containing gaps in the framework of Lexical-Functional Grammar) which, in Kaplan's original formulation, is an operator that transforms a pair of f-structures into a new f-structure, so that the values of one of the two f-structures (that with priority) are conserved and the other f-structure furnishes default values, if any. The suitability of this method for confronting the phenomenon needs further study, as do many aspects of metarules.

The following sentence will be used as an example:

Giotto dipinse un affresco e Orcagna un polittico?
[Did Giotto paint a fresco and Orcagna a polyptych?]

The parser works normally until an active S edge (that covers the sentence fragment *Orcagna un polittico*) is inserted into the chart to the immediate right of the conjunction. Such an edge may be made inactive by a metarule that establishes that a type S (coordinated) constituent may be accepted even without its head. The parser then continues working regularly until unification of the edges that recognize the conjunction is attempted and an effort is made to fill the gap present in the second conjunct using the head of the first conjunct.

In applying metarules it is possible to use heuristics that put restrictions on the ordering of the constituents contained in the second conjunct. For example, it is more likely that the correct interpretation is that in which the last constituent present (linearly) in the input part of the sentence recognized by the right conjunct corresponds to the constituent that is found furthest to the right in the first conjunct. Another possibility is that of imposing that the order of the constituents inside the two conjuncts must be parallel; but, in this case, the sentence *Giotto dipinse un affresco e un polittico Orcagna?* [lit. *Did Giotto portray a fresco and a polyptych Orcagna?*] would not be correctly interpretable. Certainly, sentences such as this last are to be considered correct, even if unusual in spoken language. On the other hand, these restrictions serve to limit the proliferation of interpretations that afflicts languages such as Italian that have relatively free ordering of the elements inside single constituents.

4. Conclusions

We have presented a solution for two phenomena of ill-formedness (that is, ellipsis and coordination), a solution that fits coherently into a chart-based approach.

As for intersentential ellipsis, it has been shown that no changes are needed for either the grammar or the basic parser: the algorithm requires only a resettlement of the chart (that is, the working memory) and the introduction of a new selecting function. Evidently, this is a great advantage in terms

of clarity and modularity that is combined with the efficiency of the entire approach.

For coordination it has been shown how changes that are brought to the apparatus are modularized so as not to fall into the intractability of other approaches.

Both the algorithms have been implemented as enhancements of the WEDNESDAY 2 parser. The parser is used in the ALFresco interactive system, a multimodal dialogue prototype for the exploration of art history.² The examples in this paper refer to that prototype. Note, however, that the solutions proposed for ill-formed input in this paper are generally valid for other chart-based approaches.

References

- [Carberry, 1989] Sandra Carberry. A Pragmatics-Based Approach to Ellipsis Resolution. *Computational Linguistics*, 15(2):75-96, 1989.
- [Carbonell *et al.*, 1983] Jaime G. Carbonell, W. Mark Boggs, Michael L. Mauldin, Peter G. Anick. XCALIBUR Project Report #1: First Steps Towards an Integrated Natural Language Interface. Technical Report CMU-CS-83-143, Carnegie-Mellon University, Pittsburgh, Pennsylvania, 1983.
- [Carbonell and Hayes, 1983] Jaime G. Carbonell and Philip J. Hayes. Recovery Strategies for Parsing Extragrammatical Language. *American Journal of Computational Linguistics*, 9(3-4):123-146, 1983.
- [Dahl and McCord, 1983] Veronica Dahl and Michael C. McCord. Treating Coordination in Logic Grammars. *American Journal of Computational Linguistics*, 9(2):69-91, 1983.
- [Eastman and McLean, 1981] C.M. Eastman and D.S. McLean. On the Need for Parsing Ill-Formed Input. *American Journal of Computational Linguistics*, 7(4):257, 1981.
- [Fong and Berwick, 1985] Sandiway Fong and Robert C. Berwick. New Approaches to Parsing Conjunctions Using Prolog. In *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics*, Chicago, Illinois, 1985.
- [Hendrix, 1977] Gary G. Hendrix. Human Engineering for Applied Natural Language Processing. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, Massachusetts, 1977.
- [Kaplan, 1973] Ronald M. Kaplan. A General Syntactic Processor. In: R. Rustin (ed). *Natural Language Processing*, Prentice-Hall, Englewood Cliffs, New Jersey, 1973.
- [Kaplan, 1987] Ronald M. Kaplan. Three Seductions of Computational Psycholinguistics. In: P. Whitelock *et al.* (eds). *Linguistic Theory and Computer Applications*. Academic Press, Englewood Cliffs, New Jersey, 1987.

²The ALFresco (Automatic Language-Fresco) interactive system has been developed in InterLisp and CommonLisp on Xerox 1186 and Sun4.

- [Kaplan and Maxwell, 1988] Ronald M. Kaplan and John T. Maxwell. Constituent Coordination in Lexical-Functional Grammar. In *Proceedings of the 12th International Conference on Computational Linguistics*, Budapest, Hungary, 1988.
- [Kay, 1980] Martin Kay. Algorithm Schemata and Data Structures in Syntactic Processing. Technical Report CSL-80, Xerox Palo Alto Research Centers. Palo Alto, California, 1980.
- [Kosy, 1986] Donald W. Kosy. Parsing Conjunctions Deterministically. In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, New York, New York, 1986.
- [Kwasny and Sondheimer, 1981] Stan C. Kwasny and Norman K. Sondheimer. Relaxation Techniques for Parsing Grammatically Ill-Formed Input in Natural Language Understanding Systems. *American Journal of Computational Linguistics*, 7(2):99-108, 1981.
- [Lesmo and Torasso, 1985] Leonardo Lesmo and Pietro Torasso. Analysis of Conjunctions in a Rule-Based Parser. In *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics*, Chicago, Illinois, 1985.
- [Mellish, 1989] Chris S. Mellish. Some Chart-Based Techniques for Parsing Ill-Formed Input. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, Vancouver, British Columbia, 1989.
- [Proudian and Goddeau, 1987] Derek Proudian and David Goddeau. Constituent Coordination in HPSG. Technical Report CSLI-87-97. Center for the Study of Language and Information, Stanford, California, 1987.
- [Sag *et al.*, 1984] Ivan A. Sag, Gerald Gazdar, Thomas Wasow and Steven Weisler. Coordination and How To Distinguish Categories. Technical Report CSLI-84-3. Center for the Study of Language and Information, Stanford, California, 1984.
- [Shieber, 1986] Stuart M. Shieber. An Introduction to Unification-Based Approaches to Grammar. CSLI Lecture Notes Series No.4, University of Chicago Press, Chicago, Illinois, 1986.
- [Stock, 1989] Oliviero Stock. Parsing with Flexibility, Dynamic Strategies and Idioms in Mind. *Computational Linguistics*, 15(1):1-18, 1989.
- [Weischedel and Sondheimer, 1982] Ralph M. Weischedel and Norman K. Sondheimer. An Improved Heuristics for Ellipsis Processing. In *Proceedings of the 20th Annual Meeting of the Association for Computational Linguistics*, Toronto, Ontario, 1982.
- [Weischedel and Sondheimer, 1983] Ralph M. Weischedel and Norman K. Sondheimer. Meta-rules as a Basis for Processing Ill-Formed Input. *American Journal of Computational Linguistics*, 9(3-4):161-177, 1983.
- [Wirén, 1989] Mats Wirén. Interactive Incremental Chart Parsing. In *Proceedings of the 4th Conference of the European Chapter of the Association for Computational Linguistics*, Manchester, England, 1989.
- [Woods, 1973] William A. Woods. An Experimental Parsing System for Transition Network Grammars. In: R. Rustin (ed), *Natural Language Processing*. Prentice-Hall, Englewood Cliffs, New Jersey. 1973.