

Sentence Processing without Empty Categories

Martin Pickering and Guy Barry

Centre for Cognitive Science, University of Edinburgh, Edinburgh, U.K.

Theories of sentence processing have standardly made use of grammatical theories with empty categories, and have therefore postulated a process known as "gap-filling". In contrast, this paper provides evidence that the processing of unbounded dependencies does not make use of empty categories. We propose instead that there is a direct association between the extracted element and its subcategoriser. To show that gap-filling cannot take place, we consider a number of examples where there is material separating the assumed empty category and the subcategoriser, and then present a formal argument, based on patterns of dependencies in sentences involving multiple cases of extraction. We then sketch a linguistic account of unbounded dependencies that does not use empty categories, and which can serve as the basis of a processing model. We conclude that empty categories are not psychologically real.

INTRODUCTION

Linguistic analyses of sentences such as (1) below, which we shall call *unbounded dependencies*, have traditionally assumed that they are derived from an underlying or "canonical" form:

1. Which man do you think Mary loves?

Which man is taken to be an argument of *loves*, and because in most other constructions arguments have to be in close proximity to their subcategoriser, it is assumed that there is a level of representation where this is in fact the case. Here, a declarative sentence would have the object argument

Requests for reprints should be addressed to Martin Pickering, Centre for Cognitive Science, University of Edinburgh, 2 Buccleuch Place, Edinburgh EH8 9LW, U.K.

We would like to thank Kathrin Cooper, Robin Cooper, Elisabet Engdahl, Mark Hepple, Dick Hudson, Pauline Jacobson, Glyn Morrill, Antonio Sanfilippo, Mark Steedman, Michael Tanenhaus, Max Volino and two reviewers for their comments and suggestions. M.P. was supported by ESRC studentship C00428722002 and SERC Postdoctoral Fellowship B/90/ITF/293; G.B. was supported by SERC studentship 88306971.

directly after *loves*, so it is assumed that this is the case at this underlying level. In transformational grammars such as Standard Theory (Chomsky, 1965), the element *which man* is found after *loves* in Deep Structure, but then moves to its actual location at Surface Structure. In Government-Binding Theory (Chomsky, 1981), it is further assumed that an empty category known as *wh*-trace is left after *loves* as a result of movement, and that this is co-indexed with the moved element. In this paper, we shall represent this co-indexation with alphabetic subscripts. Therefore (1) gets a representation like (2) below (with all irrelevant details suppressed), where \emptyset represents the empty category:

2. [Which man]_a do you think Mary loves \emptyset_a ?

By this method it is possible to show that *which man* is an argument of *loves*, and also that it serves as the object rather than the subject.

Psycholinguistic accounts of the processing of unbounded dependencies have generally assumed the existence of *wh*-trace, and assume that there is a mechanism that forms an association between the extracted element, known as the *filler*, and the empty category, referred to as a *gap*. Hence this process is called *gap-filling* (Fodor, 1978) (this is directly related to the HOLD hypothesis of Wanner and Maratsos, 1978). The extraction site therefore plays a key role in mediating the association between the extracted element and its subcategoriser. Notice that there is a need not only for gap-filling, but also for an association between the gap and the subcategoriser, in this example *loves*, so that we have to assume a two-stage process. In this case, the gap is adjacent to its subcategoriser.

However, only a subset of generative theories assume the existence of empty categories in the analysis of unbounded dependencies. For instance, there are no empty categories in most categorial grammars (e.g. Ades & Steedman, 1982), in Word Grammar (Hudson, 1984) or in Kaplan and Zaenen's (1988) reformulation of Lexical-Functional Grammar. None of these theories assumes that unbounded dependencies are derived from "canonical" sentences or representations, and they do not require the existence of extraction sites where empty categories are located. The details of their treatments of unbounded dependencies differ, but the important feature they appear to have in common is that an association is made between the extracted element and the subcategoriser. Clearly, there has to be some way of indicating the grammatical role of the extracted element, but this is not a serious problem, because the other roles will be filled by elements which have not been extracted. We shall outline one possible formulation of this mechanism later, within the general framework of categorial grammar. We shall indicate the association between the extracted element and its subcategoriser by numeric co-indexation as follows:

3. [Which man]₁ do you think Mary [loves]₁?

From a psycholinguistic perspective, an association has to be made between the filler (as we shall continue to call the extracted element) and the subcategoriser directly, and this corresponds to the formation of the unbounded dependency. It does not assume the existence of an extraction site, so there is no sense in which the processing of this type of construction requires the formation of an underlying, "canonical" or "untransformed", representation. This is in some sense more parsimonious, because there is no need to assume the existence of a phonologically unrealised intermediary, and it frees us from having to assume that a relatively rare construction type requires the use of special mechanisms. But for a sentence like (1), it is difficult to see any empirical difference between the two classes of theories. The reason is that the trace in (2) is adjacent to the subcategoriser *loves*.

In this context, let us consider the experimental literature on the processing of unbounded dependencies. It turns out to be the case that there is excellent evidence for the psychological reality of the formation of unbounded dependencies, which indicates the sensitivity of the experimental techniques used. However, there is no need to interpret the data by making use of empty categories. For instance, Swinney, Ford, Frauenfelder and Bresnan (1988) argue that a gap causes reactivation of properties of the filler. They found associative priming effects for the filler at the offset of the subcategoriser of the filler in sentences like (4):

4. The policeman saw the boy that the crowd at the party accused of the crime.

Associates of *boy* (but not of the other nounphrases) were primed after *accused*, but not just prior to *accused*. Hence there must be a reactivation of properties of *boy* at *accused*. They interpret this result in terms of an empty category analysis as in (5), but it could be interpreted equally well by direct association as in (6):

5. The policeman saw the boy [that]_a the crowd at the party accused \emptyset_a of the crime.
6. The policeman saw the boy [that]₁ the crowd at the party [accused]₁ of the crime.

Similar reinterpretations can be made of other experimental results. For instance, Crain and Fodor (1985) and Stowe (1986) discuss the so-called "filled-gap" effect, where in their terms the processor has encountered a filler and is looking for a gap. Stowe uses (7) and (8):

7. My brother wanted to know if Ruth will bring us home to Mom at Christmas.

8. My brother wanted to know who Ruth will bring us home to at Christmas.

Using a word-by-word self-paced reading time technique, she showed that subjects take longer to process *us* in (8) than in (7). The gap, she argues, is sited by the processor after *bring* in (8). Hence when the word *us* is encountered, we realise that the gap cannot be located directly after *bring*, and so we are forced to backtrack. This is in contrast to *us* in (7) where no gap is possible. However, if we avoid appealing to empty categories, we can claim instead that the filler is immediately associated with the verb directly, given that this is possible at this stage in the sentence. When we encounter *us*, we are forced to undo this association. (Note that these results can equally well be interpreted in parallel terms.)

All the examples that we have considered involve constructions where the "canonical" location for the extracted element, and hence the purported empty category, is adjacent to the subcategoriser. But this is not necessarily the case. In the next section, we shall consider a number of examples where there is material separating the assumed gap and the subcategoriser. We shall see that it appears highly unlikely that fillers cannot in general be interpreted until their extraction sites. Next, we shall present a formal argument, based on patterns of dependencies in sentences involving multiple cases of extraction, to show that this is in fact impossible. Then, we shall outline one possible linguistic account of unbounded dependencies without empty categories, which makes use of categorial grammar. Finally, we shall address empty category accounts of passives, and briefly discuss incremental interpretation. We conclude that empty categories cannot have any psychological reality.

PROCESSING WITHOUT GAPS

Multiple Object Verbs

The psycholinguistic literature on unbounded dependencies has concentrated on constructions where the subcategoriser is adjacent to the gap location. This need not, however, be the case. Consider a sentence like (9), where the second post-verbal argument is extracted:

9. In which box did you put the cake?

Because *put* "canonically" takes two post-verbal arguments in the order NP PP, as in *put the cake in the box*, and the PP-argument has been extracted, the empty category would be located after *the cake*, and not after *put*:

10. [In which box]_a did you put the cake \emptyset_a ?

The processing realisation of this in gap-filling terms is that the filler must be remembered until the end of the sentence, because it is impossible to conclude that it is to be associated with *put* until the empty category is located. This means that there can be no immediate interpretation of the fragment *in which box did you put*. Intuitively, this seems wrong, because *in which box* is an argument of *put*.

In contrast, the account outlined in the previous section, which we shall henceforth call the *gap-free* account, associates *in which box* with *put* directly:

11. [In which box]₁ did you [put]₁ the cake?

Assuming that it is possible to describe a parsing mechanism that makes this association, it would be strangely inefficient if we did not do so, but instead hunted down phonologically null elements at later points in the sentence. Hence the psychological reality of empty categories is at least questionable.

Contrast (9) with (12):

12. Which box did you put the cake in?

In this case, the extracted element, *which box*, is not itself an argument of *put*, but instead is subcategorised for by the preposition *in*. Therefore, the empty category would be after *in*:

13. [Which box]_a did you put the cake in \emptyset_a ?

The gap-free account assumes an association between *which box* and *in*, as in (14):

14. [Which box]₁ did you put the cake [in]₁?

In this case, the assumed gap location is adjacent to the subcategoriser, so the accounts do not exhibit a clear difference. In processing terms, both approaches are consistent with there being no direct association between *which box* and *put*.

In (9), the distance between the gap location and the verb *put* is not very great. However, if we extend the distance between the verb and the gap by replacing *the cake* with a longer nounphrase, we can make a clearer distinction between the two accounts and show that the use of an empty category becomes more unlikely:

15. In which box did you put the very large and beautifully decorated wedding cake bought from the expensive bakery?

The gap-filling account now has to assume that the filler is not associated with the gap, and hence the verb *put*, until the end of the complex NP, whereas the gap-free account associates the filler directly with *put* as in (9).

This is now much more efficient and intuitively sensible. In contrast, if we make a similar replacement in (12), the resulting sentence is rather awkward:

16. Which box did you put the very large and beautifully decorated wedding cake bought from the expensive bakery in?

The obvious explanation for this is that the filler has to be remembered while the complex NP is being processed, and, if this is correct, the fact that (15) is not so awkward suggests that in this case the filler *in which box* is associated with *put* directly, rather than via an empty category.

We can apply similar principles to constructions involving adjuncts. In English, post-verbal arguments standardly come before adjuncts, so extracting an adjunct should leave a gap that is not adjacent to the verb that is associated with it. Again, it would allow more rapid interpretation if the adjunct were associated with the verb directly rather than via a gap. Consider:

17. When do you think John ate the very large and beautifully decorated wedding cake bought from the expensive bakery?

The empty category would once more be at the end of the sentence. In addition, most phrase structure accounts associate the adjunct with the embedded verbphrase or sentence rather than simply the verb *ate*, which does not seem to be reflected in processing.

Empty Categories Somewhere Else?

We have questioned the plausibility of associating the filler with a gap sited at a particular location, rather than that of associating the filler with a gap anywhere. It is conceivable that gap-filling does occur, but that the gap is not located at the "canonical" site for an argument. If we can show that the gap could not be located elsewhere, we have the basis of an argument against using empty categories in sentence processing at all. There are two possibilities that we should consider: heavy NP shift and extraposition from NP. We shall also briefly address how word-order freedom affects gap-filling.

Heavy NP Shift. In many cases it is possible to reverse the order of post-verbal arguments, or to place arguments after adjuncts, so long as the argument we want to position last is sufficiently "heavy":

18. I put in the box the very large and beautifully decorated wedding cake bought from the expensive bakery.

Therefore, it is possible that the extracted PP in (9) is co-indexed with an empty category located immediately after *put*:

19. [In which box]_a did you put \emptyset_a the cake?

One problem with this is that the last argument, *the cake*, is not heavy, but because "heaviness" has never been defined well, it may be that the crucial issue is the relative weight of the arguments, as a trace is presumably very light. Also, it is possible that only heavy NPs are shifted past empty categories, and that with light NPs the empty category is located after the NP.

However, this explanation cannot be true in general, because not all post-verbal arguments allow variable order. For instance, the order is fixed for two NP arguments after the verb. Under the reading where the book is given to the woman, only (20) is possible, and the heaviness of the arguments is irrelevant:

20. I gave the woman in the heavy winter coat the book.
21. *I gave the book the woman in the heavy winter coat.

Therefore, a gap-filling account of (22) must assume the empty category to be at the end of the sentence (under the pragmatically sensible reading):

- 22a. [Which book]_a did you give the woman in the heavy winter coat \emptyset_a ?
22b. *[Which book]_a did you give \emptyset_a the woman in the heavy winter coat?

Hence in general the argument from heavy-shifting cannot be used to support the idea that empty categories are used in sentence processing at "non-canonical" locations. But it is also highly probable that extraction is impossible in combination with heavy-shifting (Ross, 1967; Wexler & Culicover, 1980):

23. Which box did you put the very large and beautifully decorated wedding cake bought from the expensive bakery in?
24. *Which box did you put in, the very large and beautifully decorated wedding cake bought from the expensive bakery?

This suggests that heavy-shifting is not possible at all in (9) and (15), and lends support to the argument using double object verbs above.

Extraposition from NP. It is sometimes possible to extrapose part of the first post-verbal argument, beyond the second:

25. I gave a prize to one boy which I knew was worth quite a lot of money.

Therefore, it is possible that the trace of a second post-verbal argument can be located within the first post-verbal argument, immediately after the head noun:

26. [To which boy]_a did you give a prize \emptyset_a which you knew was worth quite a lot of money?

However, it is clear that this cannot serve as a general explanation. Not all arguments can be interrupted in this way. Most obviously, it is impossible to extrapose any part of an argument in which the head noun is final:

27. I gave the very generous and extremely talented young Australian poet the award.

Here there is no alternative word order where *the award* is found before the end of the sentence. Therefore, a gap-filling account of (28) would have to assume an empty category sentence-finally:

28. Which award did you give the very generous and extremely talented young Australian poet?

Also, note that extraposing the modifier *capable of* . . . from one NP past another NP is very bad (see Hudson, 1990):

29. ?*We gave every student a prize capable of answering all the questions.

Hence again the empty category has to be sentence-final in (30):

30. That's the prize which we gave every student capable of answering all the questions.

Gap-filling and Word-order Freedom. Where the gap is located is, in fact, a general problem for gap-filling accounts of the processing of unbounded dependencies, because it is often the case that the "canonical" or underlying structure of a construction involving extraction is not clear. Consider an example with adjuncts:

31. In which park did Bill meet Tom on Tuesday?

The point is that both *meet Tom in the park on Tuesday* and *meet Tom on Tuesday in the park* are acceptable, so there is more than one possible site for the empty category. This phenomenon is fairly infrequent in English because of the relative rigidity of the word order, but it would be more apparent in a language with clear constituents and word-order freedom within each constituent (because it would be clear which constituent an element had been extracted from, but not which position within that constituent). We could deal with this by treating one word order as "basic", or alternatively by assuming that there are multiple canonical

forms and that the parser has to reconstruct any one of them by means of some predefined strategy (e.g. siting the gap as soon as possible). The point of this, however, is that any account making use of gaps has added complications that can only be resolved by making additional assumptions in the processor or the grammar (or both). None of these assumptions are needed if we avoid empty categories, because *in which park* is simply associated with *meet* directly.

Conclusions. We can thus conclude that the empty category cannot always be assumed to be close to the verb. It remains to be shown that the suggestive evidence above that a filler cannot simply be associated with its "canonical" location is in fact generally correct. If we can do so, then we have evidence for the stronger conclusion that empty categories are not used in sentence processing at all. We shall return to this in the next section.

Memory Load

We have argued that processing unbounded dependencies by locating a gap at the point where supposedly related "canonical" constructions would have a phonologically realised phrase is in general inefficient and intuitively unlikely. But we also know that having to remember a filler (just like remembering any other element) must require using some memory resources, and so the longer we have to remember it, the more the processing demands on the system should be. More interestingly, we can test whether the filler is remembered until the gap or just until the verb by considering the memory load between the verb and the gap.

Let us look at a pair of sentences related to (29) and (30) above:

32. We gave every student capable of answering every single tricky question on the details of the new and extremely complicated theory about the causes of political instability in small nations with a history of military rulers a prize.
33. That's the prize which we gave every student capable of answering every single tricky question on the details of the new and extremely complicated theory about the causes of political instability in small nations with a history of military rulers.

Sentence (32) is unwieldy and somewhat difficult to process, presumably because the first post-verbal argument is so complicated that we expend a lot of resources on processing it, and we are confused by having to return to a different point in our analysis when we associate the other argument *a prize* with *gave*. In an account with empty categories, extracting this argument ought to produce at least as strong an effect. Yet clearly this is

not the case. This fact is predicted by the gap-free account, because the filler can associate with the verb directly. In this account, the heavy argument is the final argument of *gave*, which is where heavy arguments prefer to be placed. Hence intuitive evidence from memory load and ease of processing supports not using empty categories.

Adjuncts appear to behave similarly:

34. I don't know where John found the book.

35. I don't know whether John found the book.

In an empty category account, (34) would have an empty category at the end of the sentence, and therefore require the employment of a gap-filling mechanism, whereas (35) would not. If we extend the post-verbal argument, (34) ought to become harder to process in relation to (35), yet this does not seem to be the case:

36. I don't know where John found the very long and extremely controversial book capable of inspiring such strong views both for and against the conclusions of the author.

37. I don't know whether John found the very long and extremely controversial book capable of inspiring such strong views both for and against the conclusions of the author.

Hence there is no evidence for assuming an empty category at the end of (36). We have already argued against locating an empty category elsewhere [(24) suggests that it cannot be located after *found*], so this militates against using empty categories in general.¹

Both memory load considerations and the ease of processing various constructions therefore give evidence against using empty categories in sentence processing. We shall present stronger evidence in the next section, where instead of simply lengthening the post-verbal argument, we consider sentences that involve recursion.

PROCESSING RECURSIVE CONSTRUCTIONS

Recursion, Nesting and Self-embedding

Much discussion of the distinction between linguistic competence and performance has dealt with constructions such as (38–41), in which one grammatically complex element appears within another:

¹Note also that a gap-filling account cannot straightforwardly explain the processing of (i) (cf. Jacobson, 1990):

i. I don't know where or whether John found the book.

38. I believe that Mary met the student who solved the problem.
 39. I know the professor who taught the student who solved the problem.
 40. I believe that the boy who the students recognised is a genius.
 41. The man who the boy who the students recognised pointed out is a friend of mine.

We shall refer to all these and similar constructions as *recursive*. In (38), a relative clause occurs within an embedded sentence; in (39), a relative clause occurs within another relative clause. Sentences (40) and (41) are like (38) and (39) respectively, with the additional property that the lower-level clause occurs in the middle rather than on the periphery of the higher-level clause.

In describing such constructions, it is useful to make reference to the concepts of *nested construction* and *self-embedded construction* as defined by Chomsky (1965). Chomsky's definitions run as follows:

The phrases *A* and *B* form a nested construction if *A* falls totally within *B*, with some nonnull element to its left within *B* and some nonnull element to its right within *B*. . . . The phrase *A* is self-embedded in *B* if *A* is nested in *B*, and, furthermore, *A* is a phrase of the same type as *B*.

For example, in (40), *who the students recognised* is nested in *that the boy who the students recognised is a genius*. Similarly, in (41), *who the students recognised* is both nested and self-embedded in *who the boy who the students recognised pointed out*.²

In a discussion of how performance factors might affect the acceptability of various types of grammatical structure, Chomsky observes that "repeated nesting contributes to unacceptability" and that "self-embedding contributes still more radically to unacceptability", whereas "there are no clear examples of unacceptability involving only left-branching or only right-branching". He suggests that this is "simply a consequence of finiteness of memory", giving the following reasons:

In some measure, these phenomena are easily explained. Thus it is known . . . that an optimal perceptual device, even with a bounded memory, can accept unbounded left-branching and right-branching structure, though nested (hence ultimately self-embedded) structures go beyond its memory capacity.

²The additional term "centre-embedding" has been used more recently to refer both widely to nesting and narrowly to self-embedding, causing some confusion of terminology. Because self-embedded constructions are merely a subset of nested constructions, to save confusion we shall usually refer to sentences like (40) and (41) simply as "nested constructions", although nearly all the examples we give of nesting will also be examples of self-embedding.

However, let us now consider:

42. John found the saucer on which Mary put the cup into which I poured the tea.

By Chomsky's definitions, *into which I poured the tea* is neither nested nor self-embedded in *on which Mary put the cup into which I poured the tea*, because no lexical material from the latter lies to the right of the former. This seems reasonable given that (42) is not hard to process in the way that (41) is.

But now consider the analysis of this sentence in terms of a theory with empty categories. This would require two empty categories at the end of the sentence (because *the cup into which I poured the tea* is all one argument of *put*). Hence the associations that the empty elements make with their antecedents are nested. This seems wrong for a sentence that is not hard to process as the nested constructions are. The sentence processor would be required to remember both fillers right until the end of the sentence when we encounter both of their associated gaps, and would only then be able to begin interpretation. But if we associate the extracted arguments with the verbs *put* and *poured* directly, then we have no need to wait. So the account without empty categories is much more straightforward, requires less memory resources, and ties in with the intuitive lack of difficulty we have with the construction. The problems that (42) raises for the theory of empty categories and filler-gap association are severe.

In the rest of this section, we shall address the processing of such sentences in more detail. We first give a more precise description of the relevant aspects of the empty category account, and the difficulties that it leads to regarding processing complexity. We then contrast it with the gap-free account, and show that this account does not lead to such difficulties. We end with a discussion of various related points.

An Empty Category Analysis of Recursive Constructions

Let us return to our initial example (1), repeated below. Recall that in an empty category account it has the analysis in (2):

1. Which man do you think Mary loves?
2. [Which man]_a do you think Mary loves \emptyset_a ?

The use of such a theory of grammar entails that there are two types of association to be made between the gap and the other parts of the sentence. One type is the co-indexation relation between the filler and the gap, which we shall refer to (naturally) as a *filler-gap* association. The other type is a dependency relation between the gap and whatever ele-

ments the filler would be associated with in the canonical sentence. For example, in (1) the filler *which man* would depend on the verb *loves* in the canonical sentence, and so a dependency must be established between the gap and *loves*. In all our initial examples, the relevant element will be a verb, and we shall refer to this type of association as a *gap-verb* association, although we shall see later (pp. 246–247) that this is a special case of a more general type of association. To distinguish gap-verb associations from filler-gap associations in examples, we shall use superscript Greek letters, so that (1) would be annotated as follows:

43. [Which man]_a do you think Mary [loves]_a \emptyset_a^a ?

If we look at sentences in which more than one deviation from canonical form occurs, we obtain some interesting patterns of association. Consider, for instance, the well-known contrast between pairs of English sentences such as the following:

44. I saw the farmer who owned the dog which chased the cat.
45. The cat which the dog which the farmer owned chased fled.

We shall refer to sentences such as (44) as *multiple subject relative* constructions, and sentences such as (45) as *multiple object relative* constructions. Let us consider the structures of these two sentences in terms of the patterns of association described above. The annotated versions of the sentences run as follows:

46. I saw the farmer [who]_a \emptyset_a^a [owned]_a the dog [which]_b \emptyset_b^b [chased]_b the cat.
47. The cat [which]_a the dog [which]_b the farmer [owned]_a \emptyset_a^a [chased]_b \emptyset_b^b fled.

There is a clear contrast between the patterns of filler-gap association in the two sentences: the pattern in (46) is *aabb*, but the pattern in (47) is *abba*. We shall refer to association patterns of the form *aabb*, *aabbcc*, etc., as *disjoint* association patterns, and patterns of the form *abba*, *abccba*, etc., as *nested* association patterns. (There are, of course, various other possible patterns of association; we discuss cross-serial association patterns at the end of this section.) However, there is no difference between the patterns of gap-verb association in the two sentences; both have the disjoint pattern *aa $\beta\beta$* . [In (46) the gaps precede their verbs and in (47) they follow them, but this does not affect the pattern.] Thus we may characterise sentences such as (44) as exhibiting disjoint filler-gap and disjoint gap-verb associations, and sentences such as (45) as exhibiting nested filler-gap and disjoint gap-verb associations.

In further contrast, the German multiple subject relative construction in (48) has disjoint filler-gap and nested gap-verb associations:

