LANGUAGE AND COGNITIVE PROCESSES, 1993, *8* (2), 163–196

Direct Association and Sentence Processing: A Reply to Gorrell and to Gibson and Hickok

Martin Pickering

Human Communication Research Centre, Department of Psychology, University of Glasgow, Glasgow, UK

Papers by Gorrell and by Gibson and Hickok in this issue question Pickering and Barry's (1991) arguments against empty categories in sentence processing. This reply disputes Gorrell's claims that Pickering and Barry's (PB's) interpretation of the data is inadequate, and, in agreement with Gibson and Hickok, reinforces the arguments that the gap location is irrelevant to the formation of an unbounded dependency. However, it accepts that it is possible to provide "predictive" accounts where a gap can be postulated before the gap location is reached. But any such proposal is likely to be more complicated and less natural than PB's, since there is no *processing* reason to assume the reality of gaps. In order to demonstrate the naturalness of the gap-free account, this reply next shows how PB's proposals can be modelled in an approach to incremental processing based on a version of flexible categorial grammar, and how the differential complexity of various recursive constructions can be accounted for. Finally, the paper considers the implications of PB's proposals for various processing strategies, and suggests that PB's conclusions are reinforced by Bock, Loebell and Morey's (1992) evidence from language production. It concludes that processing evidence gives us no reason to postulate empty categories.

INTRODUCTION

The papers by Gorrell and by Gibson and Hickok in this issue argue against Pickering and Barry's (1991) claim that empty categories are not employed by the sentence processor. This paper provides a reply to these

---

Requests for reprints should be addressed to Martin Pickering, Department of Psychology, University of Glasgow, 56 Hillhead Street, Glasgow G12 9YR, UK.

This paper obviously owes a great deal to collaborative work conducted with Guy Barry over the last 4 years. In addition, I would like to thank Gerry Altmann, Nick Chater, Matt Crocker, Elisabet Engdahl and David Milward for their comments and suggestions. Part of this research was carried out at the Centre for Cognitive Science, University of Edinburgh. I acknowledge the support of SERC Postdoctoral Fellowship B/90/ITF/293 and a British Academy Postdoctoral Fellowship.

criticisms. Let us first summarise the most important points in Pickering and Barry (PB), introduce some terminology and outline the structure of this reply.

Transformational grammar (e.g. Chomsky, 1965) assumes that many sentences are generated as a result of a process of movement, so that a question like (1) below is derived from an underlying "canonical" level of representation known as *deep structure*:

1. Which man do you think Mary loves?

At deep structure, *Which man* is located after *loves*, but it is then moved to the front of the sentence. In Government-Binding Theory (GB: Chomsky, 1981), an empty category known as *wh*-trace is left after *loves*, and is co-indexed with the moved element:

2. [Which man]$_a$ do you think Mary loves $Ø_a$?

There is also a need for some way of associating the trace and *loves*, so that *Which man* can be associated with *loves*. A consequence of this approach is that the utterances of a language are divided into those which have basic word orders and those whose word orders are derived from more basic forms. A sharp wedge is placed between two types of construction, one which is derived via movement, the other which is generated directly by a phrase structure component.

In contrast, if transformations are abandoned entirely, all constructions can be treated as equally basic.[1] The similarities between sentence pairs such as actives and passives, or topicalisations and their untopicalised counterparts, can be captured by rules which assign representations to the pairs which are related in systematic ways; for example, the sentence pairs can be given very similar meanings. There is no need to assume that one form is basic and the other derived. This approach is adopted in flexible categorial grammar (Moortgat, 1988; Steedman, 1987) as well as other frameworks mentioned in PB (Hudson, 1984; Kaplan & Zaenen, 1988). On these theories, the association between *Which man* and *loves* is direct:

3. [Which man]$_1$ do you think Mary [loves]$_1$?

These two approaches imply very different accounts of how language is processed (assuming that these grammars reflect mental representations at some level). In transformational theories, relevant properties of deep

[1]This need not be the case. For instance, GPSG (Gazdar, Klein, Pullum & Sag, 1985) has traces but no transformations. It is not clear which kind of processing account such grammars fit most happily with, so we shall not consider them further.

structure have to be accessed. In the context of GB, the link between the moved element and the trace can be established by *gap-filling* (Fodor, 1978). Only after the gap is filled can *Which man* be interpreted as the object of *loves* in (2). Gap-filling is generally thought to involve a specialised component of the sentence processor, such as a HOLD cell and associated subroutine in an ATN (Wanner & Maratsos, 1978) or a module detached from simple syntactic analysis (e.g. Clifton & Frazier, 1989; Nicol & Swinney, 1989). Hence a relatively rare construction type causes a considerable architectural complication in the language processor. The same wedge is found in the processor as in the grammar. But in non-transformational theories, there is no need to assume a gap-filling mechanism, and unbounded dependencies can be treated in the same basic way as other constructions. All that the processor has to do is to establish direct associations such as the one shown in (3) above.

PB argue for this position on the basis of evidence about the processing complexity of various constructions. Most of the relevant experimental evidence (e.g. Crain & Fodor, 1985; Stowe, 1986; Swinney, Ford, Frauenfelder & Bresnan, unpublished) does not distinguish the two accounts, because the trace is adjacent to the verb which subcategorises for the trace [as in (2)]. In contrast, PB concentrate on cases where the verb and the trace are separated by an intervening post-verbal argument of arbitrary length and complexity, as found in (PB 15), repeated below:

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

In such examples, PB claim that gap-filling implies that the filler cannot be interpreted as an argument to the verb until the post-verbal argument has been processed in its entirety. PB provide a good deal of evidence that this does not happen, and therefore argue that the filler is associated with the verb directly. The paper proposes that empty categories play no role in processing, and therefore that a theory of linguistics within cognitive science should avoid empty categories. It is these conclusions which Gibson and Hickok (GH) and Gorrell dispute.

Let us now introduce some terminology which should help clear up some issues. Following Gorrell, we define the Direct Association Hypothesis (DAH) as the claim that the filler is associated with the subcategoriser directly, without going via an empty category. (Note that this makes no claim about processing complexity.) This is contrasted with gap-filling, where the filler is associated with the gap, and the gap with the subcategoriser. Let us also distinguish *standard* gap-filling, which can occur when the processor reaches the gap location, from *predictive* gap-filling, which can occur earlier. As we shall see, Gorrell's and GH's proposals constitute different versions of predictive gap-filling.[2] We define the Immediate Association Hypothesis (IAH) as the claim that the filler can be associated with the subcategoriser immediately the subcategoriser is reached. The IAH is consistent with both the DAH, as proposed by PB, and GH's predictive gap-filling account. Finally, it is important to note that these distinctions are unrelated to the issue of whether the processor employs a *first resort* strategy, where the unbounded dependency is formed immediately it can be (Fodor, 1978). Hence we can have first resort standard gap-filling, first resort IAH, and so on. For most of the following discussion, we shall assume that a first resort strategy is employed, but we return to this important issue later.

The arguments of the two papers concentrate on two claims. The first is that PB have not presented an adequate analysis of the presented data, while the second is that PB are wrong to assume that the evidence against standard gap-filling is necessarily evidence against gaps, and for direct association. This paper will first provide a defence against the first claim, dealing with Gorrell and then with GH. Gorrell disputes much of our data against standard gap-filling and provides alternative explanations of our findings, but then appears to accept some of our data, and eventually proposes an account which does not fit with standard gap-filling after all. On the other hand, GH accept that PB's data disprove standard gap-filling. However, they claim that we fail to provide an adequate account of processing complexity. Much of what they say appears to be correct, but it should be pointed out that PB did not attempt to provide such an account. In considering the criticisms found in both papers, it is important to stress that our account is only concerned with the effect of nesting on processing complexity. It does not claim that all complexity is due to nesting.

The next section discusses the other objection to PB, which is that it is wrong to assume that abandoning standard gap-filling requires abandoning gaps. Both Gorrell and GH claim that it is not necessary for the processor to reach the gap location for a gap to be posited. This reply accepts that it is possible to provide such accounts, and therefore that PB's argument against standard gap-filling is stronger than the argument against gaps themselves. However, the use of gaps remains an avoidable complication if the gap location itself plays no role in processing (as PB's data imply). Both papers present predictive accounts; GH's is indistinguishable from PB's in terms of the processing data in question, but Gorrell's account assumes that unbounded dependencies can sometimes be formed even before the subcategoriser is reached. This paper disputes Gorrell's model in this respect, and criticises his interpretation of Boland's (1991) experimental data.

[2]In fact, predictive accounts may on occasion require that gap-filling occurs at some point after the gap location. For instance, GH's model would predict this behaviour if the subcategoriser came after the gap (as in *who do you think left*?). Such cases are not important for present purposes.

Since it is possible to model PB's data with or without gaps, it is now necessary to show that the DAH is compatible with an elegant and empirically adequate approach to sentence processing. The next section outlines such an account: a gap-free incremental processing model based on flexible categorial grammar (as introduced in PB). The account is based on a grammatical framework proposed in Barry and Pickering (1990), implemented in terms of a shift-reduce parser. It deals with unbounded dependencies with the same mechanisms that are employed for "canonical" constructions. This proposal represents direct association in a straightforward way, and accounts for the complexity of various kinds of nested constructions. This implies that there is no need to employ gaps in processing, and therefore that they can be regarded as an unnecessary complication.

The paper then briefly considers the implications of PB's arguments for various possible strategies of ambiguity resolution when the formation of an unbounded dependency is locally doubtful, and explores the relations between immediate association and first resort strategies. The final section links PB's conclusions with the proposals of Bock, Loebell and Morey (1992). They argue that language production involves binding properties of arguments of a verb to a syntactic representation which is based on surface structure rather than on underlying structure, and hence suggest that no processes akin to transformations occur during generation. The conclusions of this paper and of PB appear to be complementary and suggest that both production and comprehension are "direct". The conclusion is that sentence processing has no need for empty categories.

GORRELL'S DISCUSSION OF PB'S DATA

PB's evidence against standard gap-filling is based on examples where the gap location is a long way after the subcategoriser, as in (4) above. Apart from the intuition that people can interpret *in which box* as an argument of *put* well before the end of the sentence, and the fact that it is technically possible to provide linguistic accounts (such as flexible CG) which are consistent with the IAH, PB provide two arguments against standard gap-filling. The first is based on memory load, and the second on properties of nested constructions. It should be clear that PB regard the second argument as considerably the more important. Gorrell takes issue with both arguments; his points are discussed and replied to below. First, however, it must be stressed that Gorrell's criticisms do not stop him accepting that some of our arguments demonstrate problems for standard gap-filling, and that he eventually presents an account which is definitely not compatible with standard gap-filling. So he does not appear to dispute PB's basic point, even if he disagrees with most of the arguments.

Let us begin by summarising his arguments and our responses. In his section on memory load, Gorrell's criticisms are based on the assumption that filler–subcategoriser distance is an important determinant of memory load. In fact, this is not our claim; rather, we argue that the association between a filler and subcategoriser via a gap should involve at least the same complexity as an association between a subcategoriser and an unmoved argument in the same location as the gap, and that this is not in fact the case; indeed, he appears to accept this argument. In addition, Gorrell objects to most of our arguments regarding nested constructions, although again he accepts one argument (essentially the same as in the memory load section). While Gorrell identifies some interesting perspectives on our data, his arguments do not bear on our fundamental point that nested dependencies must present severe processing difficulties when there are enough levels of nesting, but the multiple pied-piping examples (like PB 42, repeated below) are not hard:

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

As a result, most of his arguments do not affect our conclusions. For instance, he demonstrates that complexity occurs in sentences without nested unbounded dependencies. He also claims that (PB 42) is in fact somewhat hard to process, but explains this in terms of garden-pathing. If (PB 42) involved nested dependencies, then it should be extremely hard to process, like classically nested constructions. Although Gorrell's points have merit, they do not invalidate our conclusions, as we shall see below.

Memory Load

Gorrell criticises PB's argument that gap-filling makes incorrect predictions about memory load. He claims that we fail to take into account a number of important factors, and concludes that filler–subcategoriser distance is not a better predictor of processing difficulty than filler–trace distance. But this claim is never made nor intended. In fact, we are concerned with the relation between the subcategoriser and the trace, not the filler and the trace.

Let us first review our argument concerning (PB 32) and (PB 33) (Gorrell's 15 and 16):

(PB 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.

(PB 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 (PB 32) causes some processing difficulty, which appears to be because *gave* and *a prize* are separated by an argument of considerable length and complexity. Now consider processing (PB 33) under standard gap-filling. First, *which* has to be associated with the empty category at the end of the sentence, and then this "filler-in-the-gap" has to be associated with *gave*. PB claim that this process should cause at least as much difficulty as it does in (PB 32), simply because the same association has to be made. A processing operation in a transformed sentence ought to involve at least the same complexity as the same operation in its canonical counterpart. In addition, (PB 33) involves gap-filling; this might lead to some additional processing complexity, but the argument does not depend on this. The only important point is that the gap-filling account predicts at least the same difficulty for (PB 33) as (PB 32), and this is clearly not the case. On p. 137, Gorrell admits this point, and hence accepts that standard gap-filling is inadequate. This acceptance is enough to justify our position.[3]

In fact, a similar argument can be made with (PB 15) (4 above) and (PB 16):

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

In standard gap-filling, the same association is made in both (PB 15) and (PB 16) after gap-filling: *in which box* is associated with *put* over the long intervening object. Hence we would expect the same processing difficulty in the two cases [any incidental complexity of pied-piping or preposition stranding should be controlled for by the short sentences (PB 9) and (PB 12), where the long object is replaced by *the cake*], but this is not in fact found. The IAH can explain the ease of (PB 16) compared with (PB 15) for

[3]It is in fact possible to object to our argument by claiming that the properties of the "filler-in-the-gap" are different from those of the filler itself. On this account, the complexity effects would be due to some property of the surface form of the sentence rather than the pattern of associations between the elements. It is difficult to see how this could be incorporated into an account of sentence processing employing standard gap-filling, but this is one reason why the argument from nested dependencies is stronger. (Note that this point plays no part in Gorrell's critique.)

two reasons: either because the filler–subcategoriser distance is smaller, or because there is no need to make an association between the verb and an element after the long object in (PB 16).

Gorrell provides us with a number of arguments which suggest that filler–subcategoriser distance is not an important determinant of processing complexity; on the basis of these arguments, we would expect that the second explanation for the difference between (PB 15) and (PB 16) is the right (or more important) one. But these arguments do not provide evidence against the IAH. PB claim that remembering a filler must cause some processing load (though it is certainly possible that the load may be extremely small), but never claim, as Gorrell states, that filler–subcategoriser distance is an accurate indicator of processing difficulty. Other issues may be much more important, as Gorrell points out with his examples (8)–(14), and his suggestion that preposition stranding is in general awkward (after a long object) is probably correct. But standard gap-filling predicts that pied-piping constructions like (PB 15) will be at least as difficult, and this is not in fact so.

Gorrell argues that PB predict that constructions with a long distance between filler and subcategoriser will be harder than ones with a short distance. PB certainly do not argue that this distance is the sole contributor to processing difficulty, nor that remembering a single filler is usually likely to be hard. Hence it is not difficult to accommodate the lack of a marked difference between Gorrell's (13) and (7). In addition, PB do not predict a marked difference between Gorrell's (17) and (16) (PB 33). The DAH is only concerned with the fact that (PB 32) and (PB 33) are different in complexity. If the filler is remembered until after the object in (PB 33), it should be at least as hard to process as (PB 32), irrespective of what causes this processing difficulty. We can conclude that there is no reason for PB to retract the arguments against standard gap-filling.

Nested Constructions

Gorrell disagrees with most of PB's arguments regarding nested constructions, but appears to accept one argument which shows that the processor does not wait until the end of sentences like (PB 42) and (PB 55), repeated below, before associating each filler with its empty category [exactly as in (PB 15) and (PB 33)]:

(PB 55) Jane opened the cupboard in which Bill left the box from which Sue took the tray upon which John placed the saucer on which Mary put the cup into which I poured the tea.

In fact, he uses this conclusion to motivate his own predictive account. There is obviously no reason to dispute this conclusion. However, it is

important to consider Gorrell's arguments, since he claims that PB's main line of reasoning is invalid.

Let us first discuss the point that Gorrell accepts. This is essentially the same argument as the one he has already accepted, extended to cases with more than one filler. In (PB 42), standard gap-filling claims that both fillers are associated with the empty categories after long post-verbal arguments; in the case of the first filler, this intervening argument is very long and complex. Gorrell accepts that the processing difficulty of (PB 42) is not as great as would be found if the empty categories were replaced by lexical phrases; presumably, the contrast is between (PB 42) and the multiple preposition stranding example (PB 60), repeated below:

(PB 60) John found the saucer which Mary put the cup which I poured the tea into on.

However, he does not appear to believe that the fact that (PB 42) involves nested filler–gap dependencies is relevant to this argument.

Gorrell accepts that (PB 42) does not present the processing difficulties normally associated with nested constructions. It is unclear whether he agrees that there is always some processing load associated with nested constructions, or whether he accepts that enough levels of nesting will cause processing breakdown. This is certainly the standard assumption, as found in Chomsky (1965). Gorrell instead makes the point that complexity can be found in sentences without nested unbounded dependencies. This is true, but does not run counter to our claims. As mentioned above, we make no claim that all complexity is a result of nesting. Indeed, there are many English examples of nested constructions with no unbounded dependencies which are hard to process:

5. I bought the man who showed the girl who gave the boy a kiss a painting a cake.
6. John said that Fred mentioned that Sue announced that she was leaving to her mother to his father to everyone in the room.

In both examples, the three sentence-final phrases are arguments of the three verbs in a nested pattern. Notice also that they contain "coherent substructures" in Gorrell's sense, yet this does not make them easy to process. We shall later present a general processing account which represents the fact that such sentences involve nested dependencies.

Gorrell's Korean examples show that nested constructions of a similar kind to those found in (5) and (6) are harder to process than equivalent sentences in which the words are rearranged into a non-tested order with preposed phrases (24b). Analyses may well differ about whether (24b) involves an unbounded dependency or not, but if it does, then this does not seem to make processing difficult. The point is that unbounded dependencies need not be especially hard compared with other constructions involving long-distance dependencies like (5) and (6) (indeed, the proposal is that unbounded dependencies do not have any special status in processing at all). However, PB do assume that all kinds of nested constructions are difficult to process (so long as there are enough levels of embedding). In sum, (PB 42) should be hard under standard gap-filling because it involves nested filler–gap and gap–verb dependencies, but nested unbounded dependencies need not be the only constructions which exhibit processing complexity.

Gorrell next argues that multiple pied-piping sentences are in fact hard to process, though he has previously acknowledged that they are not as hard as standard nested constructions. He suggests that this difficulty is caused by attachment ambiguity; there is a tendency for the processor to attempt to attach the second relative high (to *the saucer* rather than *the cup*). This would appear to serve as a reasonable explanation for any mild processing difficulty associated with this sentence. The obvious conclusion would be that the lack of the expected severe processing difficulty associated with nested dependencies is not explained, and hence that there is no reason to reinstate standard gap-filling.

However, Gorrell's explanation for this tendency to attach high is surely unwarranted. He claims that the parser attaches high in order to try to avoid the nested pattern of dependencies which would be caused by low attachment under standard gap-filling. As corroborative evidence, he cites Blumenthal's (1966) finding that subjects use a similar strategy in interpreting classically nested constructions. There are two problems with this. First, the tendency to attach high can be overridden, in which case the nested pattern of filler–gap dependencies remains. However, Gorrell has already accepted (p. 138) the fact that the processing load associated with (PB 42) is much less than standard nested constructions *when it is interpreted correctly, i.e. with low attachment*. Secondly, there is also a tendency to attach high in recursive constructions which have no nested dependencies, for example multiple subject relatives, or cases where the first post-verbal argument is extracted rather than the second:

7. That's the dog which chased the cat which ran away.
8. John poured the tea which Mary poured into the cup which Bill placed on the tray.

In these examples, the processor may attempt to attach *which ran away* to *the dog* and *which Bill placed on the tray* to *the tea*. This obviously cannot be explained in terms of the processor attempting to avoid nested constructions.

Gorrell claims that we predict (PB 42) should involve no processing complexity because one relative can be completely processed before the next is encountered. But this is not a prediction of our approach. Moderate complexity is likely to come from a number of sources, just as in (7) and (8) above. One reason is that there the processor could be garden-pathed onto the high attachment reading, as Gorrell himself has pointed out. Another is that pied-piping is a somewhat marked construction for most people, so multiple pied-pipings ought to increase complexity. A more interesting reason is that the process of hunting down the appropriate referent is quite complex in (PB 42). Assuming an account where the computation of reference is performed incrementally (e.g. Haddock, 1987), the process of determining which saucer is being referred to will be done in two stages. First, the processor must pick out the set of saucers which have had cups placed on them; it is assumed that the appropriate saucer is one of this set. Then the processor can home in on the cup which I poured the tea into, allowing it to disregard the other cups. From this it is possible to determine which saucer is being referred to (see, e.g. Haddock, 1987). This means that the relevant saucer cannot be discovered until the end of the sentence. This is clearly a harder process than is involved in Gorrell's (25). Certainly PB do not predict a lack of interaction between the processing of the relative clauses in (PB 42).

Gorrell then criticises PB's "extended" example (PB 55), and argues that parsing it is in fact complex. Certainly, there is no reason to be surprised at some complexity, because the problems just discussed ought to be more extreme. Hence PB's proposals are entirely consistent with obtaining the Blumenthal effect. As examples (5) and (6) demonstrate, having coherent substructures does not prevent the "classic" effect of nested constructions.

The next argument is that we have misinterpreted the term *non-null* in Chomsky (1965). PB's discussion is in fact somewhat anachronistic, since empty categories had not been invented at the time of Standard Theory. But it is not clear why Gorrell regards multiple pied-piping sentences as nested in Chomsky's sense, since Chomsky's definitions (quoted in PB) are not made in the context of a discussion about deep structure. However, the terminology clearly does not affect the argument either way.

Fodor (1978) argues for a nested dependency constraint, under which the processor prefers nested patterns of dependencies to *crossing* ones. She does not compare them with disjoint dependencies. There is very strong evidence that the processor prefers disjoint dependencies to nested ones, as we would expect [see PB's discussion of Bach, Brown and Marslen-Wilson (1986), for instance]. Thus the nested dependency constraint provides no reason for assuming that nesting does not contribute to processing difficulty.

Gorrell's criticisms of our arguments concerning nesting have been shown not to bear directly on our basic claim that the processor does not