

- Pereira, F., & Warren, D. (1983). Parsing as deduction. In *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*, MIT, Cambridge, MA.
- Tamaki, H., & Sato, T. (1984). Unfold/fold transformation of logic programs. In *Proceedings of the Second International Logic Programming Conference*, p. 127-138, Uppsala University, Uppsala, Sweden.
- Van Riemsdijk, H., & Williams, E. (1986). *Introduction to the theory of grammar*. Cambridge, MA: M.I.T. Press.

A Computational Model of Human Parsing

Steven P. Abney¹

This report describes licensing-structure parsing (LS-parsing), a computational model of human parsing. LS-parsing corresponds to human parsing at three points: (1) the order in which the LS parser builds nodes is psychologically more accurate than the orders in which either LL or LR parsers build nodes, (2) the LS parser's preferences in resolving local ambiguities are preferable to Frazier's strategies on both empirical and theoretical grounds, and (3) the backtracking strategies the parser uses when it has made an error at an earlier choice point model the distinction between weak and strong garden paths—strong garden paths being irrecoverable, weak garden paths causing psychological difficulty, but not preventing recovery of the correct structure.

INTRODUCTION

This report describes a computer program that is intended as a computational model of human parsing. Two aspects of human parsing receive particular attention: (1) the role of attachment preferences and (2) the distinction between *weak* and *strong* garden paths—strong garden paths being irrecoverable, weak garden paths causing psychological difficulty, but not preventing recovery of the correct structure. The program I describe is intended as an explicit hypothesis about the mechanisms underlying human syntactic analysis, though of course with the caveat that it characterizes human cognitive processes only at a certain level of abstraction.


¹Bell Communications Research, Morristown, New Jersey 07960-1910.

I will assume without argument that the human parser has the following properties: It is deterministic on most input it encounters; when it does parse nondeterministically, it backtracks, as opposed to pursuing multiple parses in parallel; and the parser's backtracking is experimentally detectable. For argumentation on these points, I refer the reader to the work of Lyn Frazier and her colleagues (Frazier, 1978; Frazier & Fodor, 1978; Frazier & Rayner, 1982).

Standard Deterministic Parsers

Most natural language parsers are thoroughly nondeterministic because of the high degree of local ambiguity in natural language. There have been preliminary attempts at adapting deterministic parsers to natural language by providing them with humanlike preferences for resolving local ambiguities (Shieber, 1983; Pereira, 1985). But standard deterministic parsers are not well suited to psycholinguistic modeling. Shieber and Pereira proposed supplementing a standard bottom-up deterministic parser (i.e., an LR parser; see Aho & Johnson, 1974, 1975) with conflict-resolution strategies to model human parsing preferences. However, LR parsers cannot model the incrementality of human parsing. Namely, in right-branching structures, LR parsers build no structure until all the input has been read. LR parsers build such structures literally bottom-up. The lowest node is the first node built, and it is not built until all the input has been read. By contrast, people clearly build structure before the end of the sentence and pass it incrementally to the semantic processor.

Standard top-down deterministic parsers (e.g., LL parsers) are capable of producing structures incrementally. Unfortunately, though, they introduce spurious local ambiguities that force them to backtrack where people do not. Consider, for example, a grammar that includes the rules $VP \rightarrow V NP PP$ and $VP \rightarrow V NP$. On a sentence beginning "John found," the parser builds the following structure and attempts to expand VP:



However, the parser has insufficient information to determine which expansion to use. Therefore, it must guess, and backtrack if it guesses

wrong. This means it will be forced to backtrack on one of "John found the book" or "John found the book on the desk"; this is unsatisfactory since there is no indication that people backtrack on either of these sentences.

Licensing-Structure Parsing

The parser I will consider here, a *licensing-structure* (LS) parser, does not suffer from either of these drawbacks. Intuitively, the common problem with LR and LL parsing is that they use global characterizations of phrasal expansions. In LR parsing, for example, a node can be built only when *all* daughters of the node have been found; and in LL parsing, when a node is expanded, the parser must decide on the identity of *all* the daughters of the node at once. In licensing-structure parsing, by contrast, node expansions are collections of smaller pieces of structure, i.e., binary relations between head and sister—a sister appearing in the expansion of XP if and only if it is assigned a *licensing relation* by the head of XP. Unlike in LR parsing, we can build a VP as soon as we have seen the V, without waiting to see the remaining daughters of VP; but unlike in LL parsing, we do not have to specify what those daughters are going to be at the point of VP expansion. We can incorporate them as they appear, whether only an NP follows or both an NP and a PP.

A licensing "grammar" consists of a set of specifications of the licensing relations various words (or classes of words) can assign (cf. Abney, 1987). The following are examples of phrase structure rules and the corresponding licensing-relation sets:

- (2) $VP \rightarrow V_{kiss} NP (PP)$
 $kiss: [\rightarrow NP \theta],$
 $[\rightarrow PP M]$
- (3) $S \rightarrow NP VP$
 $[I\epsilon]: [\leftarrow NP S],$
 $[\rightarrow VP F]$

The licensing relation $[\rightarrow NP \theta]$, for example, specifies that *kiss* can license a noun phrase to its right by θ -assignment (i.e., as an argument). In (3), we are forced to assume an empty head for the exocentric category S. ("I" or "Infl" stands for "Inflection"; see, e.g., Chomsky, 1986.) All four relation types are illustrated in this small example: θ -assignment (θ), modification (M), subjecthood (S), and functional selection (F). The first three are self-explanatory. *Functional selection*, for want of a standard term, is what I call the relation between a functional ("non-

- (4) a. Prefer θ -attachment
- i. I saw with a telescope
a man
 - ii. I thought in the Volvo
about his interest
 - iii. the destruction with a 20M-ton warhead
of the city
- b. Else, prefer V
- i. I sang in the kitchen
to the cat
 - ii. I wrote to a friend
a letter
- c. Else, attach low
- a gift in a box
to a boy

(The last example, by the way, underlines the fact that attachment preferences are genuinely *syntactic* preferences, not semantic preferences: the preferred analysis in (c) is semantically anomalous.)

This set of preferences is intended as an alternative to Frazier's Minimal Attachment and Late Closure. I have serious misgivings about Minimal Attachment on both empirical and theoretical grounds. The section on minimal attachment provides a discussion of the issues.

As for Late Closure, because I am approaching parsing preferences in terms of next-action conflict resolution, rather than in terms of structural properties of competing analyses, it is necessary to revert to something similar to Kimball's original Right Association and Closure principles (Kimball, 1973). P3 (Low Attachment) corresponds roughly to Kimball's Right Association principle. Kimball's Closure principle corresponds in the present model to P4 below.

Shift/Attach, Shift-Attach-L Conflicts

The general principle determining the resolution of Shift/Attach and Shift/Attach-L conflicts is "if you can build structure now, do so." In other words, Shift is an action of last resort: P4. Prefer Attach over Shift. P5. Prefer Attach-L over Shift.

Attach-L/Attach-L Conflicts

A consequence of resolving Shift/Attach-L conflicts in favor of Attach-L is that Attach-L/Attach-L conflicts never arise. In an Attach-L/Attach-L conflict, the Left node (L) can be licensed by two heads in the near edge of the current tree, call them A and B. At an earlier point in the

parse, only L and *one* of A and B were on the stack. At that point, we had a Shift/Attach-L conflict. Resolving Shift/Attach-L conflicts in favor of Attach-L effectively resolves potential Attach-L/Attach-L conflicts in favor of low attachment.

This appears to make the right predictions, empirically. Consider examples like the following:

- (5) a. [quite surprisingly] big (low)
b. [quite surprisingly big] (high)
- (6) a. those [hundred-pound] bags (low)
b. those [hundred-pound bags] (high)

(5a) means "big to a quite surprising degree"; (5b) means "quite big, and surprisingly big." (5a) seems to be the preferred interpretation. (6a) and (6b) are adapted from Marcus (1980). In (6a), the bags each weigh 100 pounds; in (6b) the bags weigh a pound, and there are 100 of them. Again, the low-attachment structure (6a) is preferred.

Shift/Shift Conflicts

A Shift-Shift conflict arises exactly when the next word is categorially ambiguous, the conflict being which category is to be shifted onto the stack. I will not hazard a proposal about the resolution of lexical ambiguity beyond suggesting that the simple expedient of ranking alternatives in the lexicon will take us a good way. There are quite complex issues involved, though, that I cannot hope to do justice to here.

Attach/Attach-L Conflicts

A final type of conflict that could arise in theory is Attach/Attach-L. For an Attach/Attach-L conflict to arise, both L and C must be able to license the other. I have been unable to construct any examples where that occurs.

Conclusion

To sum up, we require basically three preferences: (1) Prefer θ -attachment, or failing that, attachment to a canonical θ -assigner. (2) As a default, prefer low attachment. (3) Shift only as a last resort. These resolve all conceivable conflicts, with the exception of Shift/Shift conflicts, about which I have little insightful to say, and Attach/Attach-L conflicts, which, though conceivable, appear not to arise in fact.

ON MINIMAL ATTACHMENT

In this section, I would like to consider Frazier's Minimal Attachment principle (Frazier, 1978): "Minimal Attachment (MA): Choose the analysis which involves building the least number of nodes." Minimal Attachment (and its companion, Late Closure) have been widely adopted in the literature. However, I have felt it necessary, for a number of reasons, to make the alternative proposal presented in the previous section.

Empirical Predictions

First, MA makes empirical predictions that disagree with the attachments assumed in the sentences of (4). For example, MA predicts that PP-attachment to NP will never be preferred over attachment to VP—Frazier assumes that PPs are uniformly Chomsky-adjoined to NP but simply attached to VP, hence that it takes one more node to associate a PP with NP than to associate it with VP. I claim that attachment to NP will be preferred when the noun θ -marks the PP but the verb does not. The crucial example is sentence (4.a.ii.), "I thought about his interest in the Volvo." My intuitions are strong that attachment to NP is preferred—though, of course, final resolution of the issue awaits experimentation.

Consider also cases where two NPs are involved. MA makes no predictions; hence, Late Closure takes over, predicting uniformly low attachment. On the other hand, my strategies predict high attachment when the lower attachment is a non- θ -attachment and the higher attachment is a θ -attachment. The crucial examples, then, are ones like (4.a.iii.), *the destruction of the city with a 20M-ton warhead*. My strategies predict high attachment ("destruction by means of a warhead"); Frazier's strategies predict low attachment ("the city that possesses a warhead"). Again, my intuitions are strong that high attachment is empirically the correct preference.

Structural Assumptions

My second reservation concerning MA is that it depends essentially on certain structural assumptions that I find linguistically unjustifiable. Frazier assumes that PPs uniformly Chomsky-adjoin to NPs, requiring the creation of a new NP node, whereas they simply attach to VPs,

requiring no extra node. This assumption has not been defended on linguistic grounds, to my knowledge, and I consider it highly implausible.

A position that *would* be defensible is that adjunct PPs are Chomsky-adjoined, while argument PPs are simply attached. In conjunction with this position, Minimal Attachment would make roughly the predictions of P1 above, without P2. Hence, it would make no prediction in cases where both sites or neither site involves argument-attachment (examples (4.b.)), with the result that Late Closure takes over, predicting low attachment—exactly the wrong prediction.

If we make what is perhaps the standard structural assumption—that base-generated PPs are never Chomsky-adjoined—then there is no distinction between attachment to VP and attachment to NP in terms of number of nodes built; hence, MA makes no predictions at all. Late Closure takes over, predicting low attachment, and the result is that PPs are uniformly attached to NPs—a hopelessly wrong prediction.

Parsing Assumptions

My third qualm with MA is that it depends crucially on contestable assumptions about the order in which nodes are built. Consider the fact that utterances are frequently not complete sentences, but PPs or NPs standing alone. It would be reasonable, then, if the parser did not build an S node until it had evidence that it was in fact dealing with a complete sentence—i.e., until it encountered a word that could not be incorporated into the NP it was building. Under this view, in the main clause analysis of *the horse raced past the barn fell*, the parser builds both VP and S when it encounters *raced*. If we assume that the "reduced relative" reading involves building only a PtcP over *raced* (rather than a true relative clause), and if we assume that this PtcP attaches directly under NP, rather than Chomsky-adjoining, then PtcP is the only node built under the reduced relative analysis—versus two nodes, VP and S, under the main clause reading. Given these assumptions, MA falsely predicts that the reduced relative reading is preferred.

In short, MA works only if one makes a number of auxiliary assumptions that have not been defended, to my knowledge. Certain of these assumptions I cannot accept—in particular, Chomsky-adjunction to NP but not to VP, and the assumption that S is built as soon as the first word is encountered. Hence, I am unable to adopt MA.

STATE

Any parser that recovers a unique initial parse will necessarily misparse certain sentences, initially. We desire that those sentences be precisely those that people initially misparse; it is the role of the conflict-resolution preferences we have discussed to guarantee that this is the case. But we also desire that the parser be able to recover from misparses in precisely those cases where people are able to recover. In the remainder of the paper, I discuss when and how the parser backtracks from an initial misparse.

LR-States

First, the parser must be able to recognize that it has misparsed. A simple LS-parser is not able to recognize that it has made an error, at least not until it reaches the end of the input and discovers that it has multiple trees on the stack. This is so because Shift is always a legal action. We must permit the parser to shift when it cannot attach because, in left-branching structures, the licenser of the current tree is still waiting in the input. What we need, then, is an encoding of local state, such that the parser knows when it is conceivable that a licenser is yet to come, and when no following word could possibly rescue the current partial structure.

I will use standard techniques to accomplish this. A *state* in LR parsing is a characterization of how much of a phrasal expansion one has parsed. For example, the state

(7) [NP → Det.AP N]

encodes the following information: (a) The parser is building an NP, (b) it has seen a Det, and (c) it is expecting an AP next, followed by an N. On any word that can start an AP, state (7) has what I will call a *continuation*. For example, if the next word is A, the continuation is the state [AP → A .]; i.e., we are now building an AP, and we have seen an A. As long as there is a continuation on the next word in the input, the parser may Shift. If there is no continuation on the next word—if the next word is a verb, for example—the parser recognizes that it has gone down a garden path: Either it made an error at some choice point, or the sentence is ill-formed.

When we reach the end of an expansion encoded in an LR-state, we perform a “reduction,” which consists of a series of Attach-L actions. For example, when we reach the state [NP → Det AP N .], N is on the top of the stack. We project it to NP, and do two Attach-L’s to incorporate the AP and Det preceding it.

In brief, then, LR-states keep track of unassembled pieces of a




Fig. 3. LR states in an LS parse.

phrasal expansion, and encode information about which continuations will lead to a well-formed phrase. They are necessary only until the licensing head has been found—i.e., on left-branching portions of the structure. On right-branching structures, words can be immediately assimilated as they are shifted onto the stack; hence, there is no need for local state to keep track of the pieces of an unassembled phrase.

Figure 3 illustrates the use of LR-states in an LS-parse. For reasons that will become clear as we proceed, we show the stack explicitly, by connecting members of the stack with dotted lines. We include LR-states on the stack for ease of recognizing which is the current state (i.e., the one on top of the stack), but they are otherwise “invisible.” The next action in Figure 3 will be a reduction—i.e., a sequence of two Attach-L actions—which will assemble the NP.


An exception to the rule that LR-states are used only on left-branching pieces of structure arises with functional categories like Infl. Infl is frequently phonologically empty, presenting a problem for recognizing when it has been encountered. Let a licensing grammar’s *LR-relations* be the class of relations used to construct LR states. A solution to the empty-head problem is to include left-directional relations *plus* functional-selection relations among the LR-relations.

For example, instead of using the LR-state [IP → . NP (I)], we will use the state [IP → . NP (I) VP]. We will know to build an IP when we have seen NP and VP, even if the Infl is phonologically null. (If the Infl is null, we will insert it when we do the reduction.)


This contrasts with the LR-state used in parsing VP, which will not include any complements of V, inasmuch as complements of V are not functionally selected but are either θ -complements or modifiers. Correspondingly, V must be overt; we cannot permit empty V’s.²

²Taking this approach to the problem of empty heads requires us to put strong constraints on functional selection. If we allowed V to functionally select IP, for example, we would introduce into LS-parsing the problem that caused us to reject LR-parsing—namely, that in certain right-branching structures, no structure is built until the end of the input. Basically, we cannot permit functional-selection cycles of the form

$$[xP X -F \rightarrow [yP Y \dots -F \rightarrow [xP \dots$$



(a) At the point of conflict



(b) At the point of backtracking

Fig. 6. A strong garden path.

the verb *fell*. Hence, the parser recognizes that it has misparsed and must backtrack. (The state given in parentheses is the abandoned continuation (1) of Figure 5a.)

At this point, we require a heuristic for finding the relevant abandoned continuation. A candidate is the following: Search the right edge of the Left tree for an abandoned state which has a continuation on the Current node. This heuristic will find the abandoned state [IP → CP NP . VP] in Figure 5b. Making this the current state will ultimately yield the correct analysis, in which *the sock* is the subject of the main clause.

Let us contrast this with the parser's actions on a strong garden path, *the boat floated sank*. Let us assume a very simple structural analysis, in which the conflict is between a postnominal participle phrase and a matrix verb. At the point of conflict, the structure is that given in Figure 6a. We have a Shift/Shift conflict. Assuming that the verb reading is preferred over the participle reading, continuation (1) becomes the new LR-state, and (2) is abandoned.

At the point of backtracking, we have the situation in Figure 6b. The current state, [→ . PP], has no continuation on V. If we search the right edge of the left tree, we find no relevant abandoned states. Even if we look at lexical alternatives and find the abandoned state (2), state (2) has no continuation on V either. The chain of deductions that would be

necessary to determine that state (2) is relevant would go something like this: If we build the PtcP and attach it to N, we can get *floated* "out of the way," and the current state will become [IP → NP . VP], which *does* have a continuation on V. In effect, to determine if state (2) is relevant, the parser has to blindly resurrect it and see what happens. My claim is that weak garden paths arise where an abandoned state is both accessible and clearly relevant; the parser fails to find the correct alternative parse otherwise.

CONCLUSION

In sum, I have proposed a particular way of modeling human parsing. First, we require a basic parsing algorithm that provides us with the right "traces" in the default case—i.e., that builds nodes at the right time. (Thus, for example, we rejected standard LR parsing because nodes were built much too late in right-branching structures.) Second, there is a set of preferences for local ambiguity resolution. I believe the preferences I have proposed are empirically more adequate than Frazier's proposals and are not dependent on questionable assumptions about syntactic structure. Finally, there are heuristics for deciding which parse to pursue next at the point of backtracking. If we suppose that the parser has only a restricted set of "simple" heuristics—on an intuitive notion of "simplicity" that is at least adequate for certain clear cases—then we can give an account of the difference between strong and weak garden path sentences.

REFERENCES

- Abney, S. (1987). Licensing and parsing. *Proceedings of NELS 17*, University of Massachusetts, Amherst.
- Aho, A.V., & Johnson, S.C. (1974). LR parsing. *Computing Surveys*, 6(2), 99-124.
- Aho, A.V., & Johnson, S.C. (1975). Deterministic parsing of ambiguous grammars. *Communications of the Association for Computing Machinery*, 18(8), 441-452.
- Chomsky, N. (1986). *Knowledge of language*. New York: Praeger.
- Frazier, L. (1978). *On comprehending sentences: Syntactic parsing strategies*. Unpublished doctoral dissertation, University of Connecticut.
- Frazier, L., & Fodor, J. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, 6, 291-325.
- Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, 14, 178-210.

- Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition*, 2, 15-47.
- Marcus, M. (1980). *A theory of syntactic recognition for natural language*. Cambridge, MA: M.I.T. Press.
- Pereira, F.C.N. (1985). A new characterization of attachment preferences. In D. Dowty et al. (Eds.), *Natural language parsing* (pp. 307-319). New York: Cambridge University Press.
- Shieber, S. (1983). Sentence disambiguation by a shift-reduce parsing technique. In *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*, pp. 113-118.

Movement Errors and Levels of Processing in Sentence Production

José E. García-Albea,^{1,4} Susana del Viso,² and José M. Igoa³

From a corpus of 3,530 slips of the tongue in Spanish, a sample of 753 cases of movement errors was analyzed, comprising those tokens that could be unambiguously assigned to the major categories of anticipations, perseverations, exchanges, and shifts. The analysis was performed according to two main criteria: (a) the degree of correspondence between the linguistic elements interacting in an error, and (b) the distance between such elements in terms of the type and number of the intervening linguistic boundaries. The results of this analysis converge with those obtained in English, supporting a model of sentence planning with different levels of representation and processing. Furthermore, Spanish provides a clear case to attest the role of syllabic structure in production processes, the constraints set by word boundaries in sublexical errors, and the contribution of inflectional suffixes to the assignment of grammatical category to the root morphemes.

INTRODUCTION

The analysis of spontaneous speech errors has become a classical research method in the study of the processes underlying language production (Boomer & Laver, 1968; Dell & Reich, 1981; Fay & Cutler, 1977; Fromkin, 1971; Garrett, 1975, 1980; Shattuck-Hufnagel, 1975). Despite the technical difficulties such analysis must endure (Butterworth, 1980; Cutler, 1981, 1988), its major justification seems unquestionable:

¹Universidad Complutense de Madrid.

²Universidad de Oviedo.

³Universidad Autónoma de Madrid.

⁴Address all correspondence to José E. García-Albea, Facultad de Psicología, U.C.M., Campus de Somosaguas, 28023-Madrid, Spain.