

*Suzanne Stevenson & Paul Smolensky*

In this chapter we explore the possibility that within Optimality Theory, a single syntactic grammar directly yields not only competence theoretic results on the grammatical distributions of elements, but also performance theoretic results on relative preferences when processing sentences with various syntactic ambiguities. Whereas the competence theory applies the grammar at the level of an entire sentence, the performance theory incrementally optimizes interpretation one word at a time. Rankings of syntactically motivated constraints yield correct predictions for a number of ambiguities in English sentence processing, and accounts for a range of cross-linguistic variation in parsing preferences for a widely studied ambiguity.

---

### Contents

1. The theory .....	829
2. The English constraint ranking.....	836
2.1. Two ambiguities that fix the primary constraint ranking .....	837
2.2. Predictions of the account.....	839
2.3. LOCALITY is a power hierarchy .....	843
2.4. *CASE rankings .....	845
3. Relative Clause attachment typology .....	847
3.1. Three attested attachment preference patterns.....	849
3.2. Typologically excluded systems.....	851
3.3. Typologically predicted systems .....	852
4. Reanalysis .....	853
5. Conclusion.....	856

---

A primary focus of research in sentence processing is to determine how, in the face of pervasive ambiguity, the human parser chooses the best interpretation of a sentence. The idea that such decisions result from a process of simultaneously satisfying multiple violable constraints of differing strengths is a familiar one. In the past, this conception of on-line parsing – “linguistic performance” – has been at odds with conceptions of grammar in

theoretical linguistics – “linguistic competence.” This has created a fundamental problem: connecting the competence grammar to the on-line parser.

Several approaches (e.g., Gibson 1991; Stevenson 1994b) reinterpreted the discrete, inviolable conditions from syntactic theory (Chomsky 1981; Rizzi 1990) as weighted, violable constraints in parsing, with the result that grammar and parser, while sharing underlying constraint knowledge, were based in different computational frameworks. Phillips 1996 took an alternative approach of redefining grammar itself as a process of incremental structure-building, but this view did not connect with the notion of multiple constraint satisfaction in parsing. Thus, the linkage between grammar and parser has remained unclear, as grammatical influences on parsing have simply been de-emphasized in the prevailing constraint-based models of sentence understanding (MacDonald, Pearlmuter and Seidenberg 1994, Trueswell and Tanenhaus 1994, among many others).

The advent of Optimality Theory (Prince and Smolensky 1993/2004; Legendre, Vikner and Grimshaw 2001) as a framework for competence grammars changes this situation profoundly: both the parser and the grammar are in the business of simultaneously satisfying conflicting violable constraints. The possibility arises that the parser uses both the same underlying knowledge and the same computational mechanism of constraint optimization as the grammar. Specifically, in our work, we explore the idea that the disambiguation mechanism of the human parser – the component that chooses the best interpretation of a sequence of words – is in fact the grammar itself.

How exactly does the nature of OT enter into this proposal? Within the prevailing view of sentence processing, words are integrated one at a time as they are received. During this process, various ambiguities arise as the sentence unfolds over time. The incremental nature of on-line interpretation entails that preferences in resolving those ambiguities are determined on the basis of sentence fragments. Yet processing a sentence fragment is fundamentally impossible within a grammatical framework of inviolable constraints, or serial rule derivation: fragments simply do not meet inviolable constraints and simply have no derivation. This is what has made it difficult in the past for a model to directly use grammatical principles to guide the ambiguity resolution process. By contrast, in OT, *any* input receives an analysis, by a uniform mechanism. Thus an OT grammar can be asked to determine the optimal structural description of a sentence fragment, just as it can an entire sentence. In the former case, the grammar is functioning as an on-line parser.<sup>1</sup> In the latter case, it is functioning like a traditional competence grammar. The only difference is the completeness of the input.

In this chapter we examine this possibility by asking whether an OT grammar that is well-motivated from the perspective of theoretical syntax can

---

<sup>1</sup> Again, more specifically, the grammar is functioning as the disambiguation component of an on-line parser.

---

explain on-line parsing preferences of comprehenders, as evidenced by empirical data on the processing of sentences which, at intermediate positions, have various structural ambiguities. We focus on English but also consider some cross-linguistic evidence.

The first work exploring this possibility was Gibson 1995 (see Gibson and Broihier 1998). Gibson examined several plausible constraint systems, and concluded that none of them was satisfactory when implemented in an OT grammar. Using a different set of constraints, our conclusion will be that an OT grammar can explain a wide range of sentence processing facts, including those examined by Gibson. Despite our different conclusions, we wish to clearly acknowledge the debt our work owes to Gibson's original idea and its thoughtful development.

In addition to forming a clearer connection between linguistic competence and performance, we argue that the OT approach has the further advantage of providing a relatively restrictive framework within which to interpret the empirical data. Specifically, we show that, with regard to an important set of structural ambiguities, the restrictions entailed by the OT framework do not permit unattested preference patterns, at least some of which a system of numerical constraints would allow.

## 1. THE THEORY

The key idea is expressed in (1) as a hypothesis linking grammar and performance; we will test it in the empirical domain of structural ambiguity resolution.

(1) Linking hypothesis:

- a. Processing word  $w_i$  means building the tree structure that has  $w_1 w_2 \dots w_i$  as its terminal string, and that is optimal according to the OT grammar.
- b. Processing difficulties occur when the optimal parse at word  $w_i$  is not consistent with the optimal parse at word  $w_{i-1}$ .

The goal here is not to offer a general algorithm for OT parsing – i.e., a procedure, composed of primitive computational operations, for constructing candidate representations for an input and selecting the optimal one (see [Section 12.2](#)). Rather, as noted above, we are proposing to use the grammar to specify the disambiguation decisions of the parser – i.e., to identify among a set of possible candidate structures that which is preferred. We focus here on modeling initial preferences in parsing, as indicated in (1a). We will not address in detail the problem of predicting the degree of difficulty induced if a preference is proven wrong by later input, forcing the parser to *reanalyze* the originally preferred structure. For now, we simply assume the informal notion suggested by (1b), that processing difficulty is proportional to the degree to which an earlier preferred structure diverges from what is required by subsequent input. We

return to this question in Section 4, laying out an initial proposal for determining the cost of reanalysis in our framework.

Given hypothesis (1), we should find that, at each step of the incremental parsing process, the *preferred* parse is the *optimal* parse. Our technique for testing whether this goal can be achieved is as follows:

- (2) General method
  - a. Empirically determine the preferred structure from the set of candidate structures for each of a set of example ambiguities.
  - b. Rank the constraints so that the preferred parse is optimal at the point of ambiguity in each case.
  - c. Verify that the ranking holds across additional example ambiguities.

This approach is dependent on several assumptions about preferences, candidate structures, and constraints, which we motivate here.

The first assumption concerns what constitutes the preferred parse. We follow standard practice in adopting as the preferred structure for an ambiguous input the structure that corresponds to the preferred interpretation as determined through psycholinguistic experimentation or accepted intuitions discussed in the sentence processing literature. Typically, these are revealed by inferring back from processing difficulty to preferences in ambiguity resolution: if a particular sentence induces processing difficulty at some disambiguating word  $w_i$ , it must be because the parse needed at  $w_i$  is not consistent with the previously established preferred structure for that ambiguity (1b). For example, in processing the sentence *John put the candy on the table into his mouth*, people experience difficulty at the PP *into his mouth*. Since this PP must be the locative argument of *put* ('put  $x$  into his mouth'), the perceived difficulty indicates an earlier preference to interpret the PP *on the table* as the locative argument of *put* ('put  $x$  on the table').

The second assumption concerns the candidate parse structures. Our theory employs the extension of OT presented in [Chapter 12](#) which incorporates not just production- (or generation-) directed optimization, but also comprehension- (or interpretation-) directed optimization. For interpretation-directed optimization, candidate parses are provided by a function *Int*, the counterpart of *Gen* in generation-directed optimization. Both *Gen* and *Int* establish a relationship between an underlying interpretation and an overt string, which is mediated by a structured representation. In our domain, an interpretation is a predicate-argument structure which incorporates and interrelates the lexical specifications of the words employed. As spelled out in [Chapter 12](#), for *Gen*, an 'input' is such a predicate-argument structure, and an 'output' is a structured syntactic representation (a parse tree) and its corresponding (pronounced) string. In contrast, for *Int*, the 'input' is a string of overtly pronounced words, and the

---

‘output’ is a parse tree and its corresponding predicate/argument structure, including all relevant lexical specifications.<sup>2</sup>

The candidate parses provided by *Int* are, we assume, syntactic trees each of which has as its (overt) yield the ‘input’ thus far (the sequence of words  $w_1$  through  $w_i$ ). We assume that all such trees are candidates subject to evaluation by our grammar. As discussed in Section 12:2, a parsing algorithm would never actually generate more than a small set of these trees. The incremental nature of human parsing especially would be naturally modeled by algorithms which generate a small number of potentially-optimal candidates as they proceed left-to-right through the input string. But here we are not proposing a parsing algorithm per se; rather we are examining the proposal that the grammar can distinguish the best candidate from the entire logically possible set. We leave it to future work to determine how, in the human sentence processor, the grammar and parsing algorithm interact at an earlier stage of structure-building to exclude most suboptimal parses.

We assume furthermore that each candidate tree provided by *Int* satisfies  $X'$  theory (Box 12:1) and thus each candidate structure is a fully connected tree with a single root node, like all trees. However, we do not assume that the trees *Int* produces are all rooted in S (i.e., a sentence node; also labeled as IP or CP in  $X'$  representation). If the initial words of the sentence are parsed as an NP, for example, the parse tree is rooted in NP; later, this NP may be embedded as, say, the subject of a matrix clause, at which point the tree will be rooted in S.

Faithfulness concerns the relation between a paired interpretation and structured representation (Chapter 12). In our model, a candidate output for an overt word string  $w_1w_2\dots w_i$  contains a parse tree  $p$  and a corresponding predicate-argument structure. The tree may be *unfaithful* to the predicate-argument structure, but only in ways that do not prevent  $p$  from having the string  $w_1w_2\dots w_i$  as its yield. (Recall that the predicate argument structure incorporates the lexical specifications of  $w_1w_2\dots w_i$ .)

As in generation-directed OT, there are two general ways in which a tree may be unfaithful. The first is *overparsing*, in which the parse tree  $p$  contains “extra” structure not specified by the lexical properties of the input thus far – specifically, phrases which have no head. (By comparison, in generation-directed phonology, overparsing amounts to the addition or epenthesis of material not present in an underlying form). The second type of unfaithful output involves *underparsing*. Here, the parse tree lacks structure needed to meet the lexical specifications of the words in the predicate-argument structure; for example, the predicate-argument structure may include a verb  $V$ , but the parse tree may fail to include all of  $V$ ’s required arguments. (In generation-directed phonology, underparsing is deletion of underlying material). We return to examples of over- and underparsing in our framework after we present the precise constraints we use.

---

<sup>2</sup> In this chapter, we present the parse tree part of the output of *Int*, and assume that the corresponding predicate/argument structure can be read off from the tree.

Our final, and perhaps most crucial, assumption concerns what constitutes the set of constraints that we apply to these candidate structures in determining the optimal form. Our goal is to bring together the grammar and parser, and so as far as possible we want to adopt existing grammatical constraints from OT work on syntax. But numerous constraints have been proposed, and we risk having a large number of constraints, many of which may be irrelevant to the structural ambiguity phenomena under investigation. On the other hand, the particular ambiguities we look at may necessitate constraints on aspects of syntax that have not yet been explored in OT research. Due to this potential mismatch, we look to work in sentence processing to guide us in focusing on a subset of grammatical constraints relevant to the phenomena of interest.

Our approach is to work from OT grammar and from an established computational model of sentence processing, the competitive attachment (CA) model (Stevenson 1994b), itself formulated as a grammar-based parser in which inviolable grammatical constraints are translated into weighted influences on activation levels in a connectionist framework. We proceed as follows: First, we analyze the CA parser to reveal the primary influences on activation (or strength) levels, which are the determinants of preferences in the model. Then we translate those numeric influences into discrete but violable OT constraints. In each case, we find a direct correlate to either an existing OT constraint, or to a well-established condition of grammar proposed in another grammatical framework; in the latter case, we adopt the constraint in violable form.

We motivate the constraints we adopt by presenting in (3) – (6) below: (a) a numeric influence on preferences in the CA model, and its motivation, and (b) the corresponding grammatical (violable) constraint that we adopt. Note that the numeric influences in (a) are not absolute; rather they should be read as ‘All else being equal, configuration X gains more activation (is stronger or more preferred) than configuration Y.’

- (3) a. Hypothesized phrases (phrases projected in the parse without explicit input to serve as the head of the phrase) have less activation than structure headed by overt input. Hypothesized phrases are created when an overt phrase cannot be directly connected to the rest of the parse tree, as for example in English when the parser processes the subject of a complement clause but has not yet processed the embedded verb. Essentially, input words are the source of activation in the model, so hypothesized structure entails that the same amount of activation must be shared among a greater number of phrases. Specifically, a hypothesized phrase must share activation with the overt phrase that triggers it.
- b. **OB-HD** (Obligatory Heads): The head of a phrase must be filled (Grimshaw 1997).

Note that the head of the phrase may be filled by a word in the input, or by a coindexation relationship between the phrase and

---

another element in the sentence, as with a trace of syntactic movement (Box 12:1).

- (4) a. Argument attachments have higher activation than adjunct attachments.

Arguments are phrases that have an essential relation to a predicate, while adjuncts provide more peripheral modification information (Box 15:1). To capture their different properties, the CA model must encode these relationships differentially. Argument and adjunct attachment sites exhibit differing competitive behavior, in which argument sites compete more strongly for attachment to a potential argument phrase (see Stevenson 1995).

- b. **ASSIGN- $\theta$** : A predicate must assign all of its thematic roles (Chomsky 1981; Box 15:1).

We adopt this “half” of the theta criterion as a violable constraint.

- (5) a. More recent attachment sites have higher activation than less recent sites.

This is due to decay of activation of attachment sites over time, a typical behavior of connectionist models (see Stevenson 1994a).

The result of this, in terms of preferred attachment configurations, is that low right attachment for a phrase XP is preferred over a higher attachment in which XP would c-command (Box 16:3) material intervening between it and its attachment site.

- b. **LOCALITY**: If XP asymmetrically c-commands YP, then XP precedes YP (the Linear Correspondence Axiom, or LCA, Kayne 1994; see also Phillips 1996 for a similar principle within grammar and processing).

By definition, X asymmetrically c-commands Y iff X c-commands Y but Y does not c-command X. (That is, X does not dominate Y, the parent node of X dominates Y, and the parent node of Y does not dominate X. This employs a standard variant of the c-command definition of Box 16:3).

Again, we adapt an inviolable constraint from another grammatical framework as a violable constraint within our OT approach.

- (6) a. Attachments which satisfy grammatical constraints have higher activation than those which don't.

This derives from reinterpretation in the CA model of inviolable grammatical constraints on a complete sentence as weighted violable constraints within an incremental parse.

The constraints of interest here are thematic and agreement requirements. Assignment of thematic requirements is covered by (4b) above, so here we are concerned with agreement, and we focus

on Case (Box 15.1) due to its observed role in ambiguity resolution (Bader 1996; Meng and Bader 1997).

- b. **\*CASE Hierarchy:** \*GEN  $\gg$  \*DAT  $\gg$  \*ACC  $\gg$  \*NOM.

This is a universal markedness hierarchy which ranks possible Case assignments to NPs, genitive being the most marked and nominative the least marked Case (Woolford 2001: 513).<sup>3</sup>

We assume in our work that Nom is assigned to the subject position, Acc to the object of a verb, Dat to the object of a preposition (or to the indirect object of a verb), and Gen to the possessive (i.e., to *Sara* in either *Sara's servant* or *the servant of Sara*).

**Note:** 'Gen' or '\*GEN' refers to genitive Case; *Gen* refers to the candidate generator (not truly relevant to the interpretation-directed optimization studied here; its role is played by *Int*, as discussed above).

- c. **AGR-CASE:** A relative pronoun must agree in Case with the NP that its relative clause modifies. (Sauerland and Gibson 1998; Artstein 2000; Fanselow, Cavar, Kliegl and Schlesewsky 1999).

This constraint, which has grammatical motivation as observed by Fanselow et al. 1999, has been proposed as a violable constraint in each of the frameworks cited above.<sup>4</sup>

We assume that **AGR-CASE** applies whatever the overt expression of the relative pronoun (as in *the woman who/whom/that I saw*), or even when it is not overtly expressed (as in *the woman I saw*).

The five constraints we have adopted fall into the two broad categories seen throughout OT work, FAITHFULNESS and structural MARKEDNESS (Chapter 12, (14)). OB-HD and ASSIGN- $\theta$  enforce faithfulness between a candidate's predicate-argument structure and its parse tree. LOCALITY, \*CASE, and AGR-CASE are markedness constraints on the morpho-syntactic properties of the parse tree. We discuss them briefly in turn.

OB-HD penalizes overparsing – the parse tree created by *Int* contains structure that is unspecified by the lexical properties of the observed sequence of words. For example, consider processing the input *James donated his books to the*

<sup>3</sup> Woolford 2001 advocates \*DAT  $\gg$  \*ACC  $\gg$  \*NOM; Grimshaw 2001: 226 employs \*DAT  $\gg$  \*ACC; Aissen 2001: 65 proposes thematic and relational hierarchies aligned with \*ACC  $\gg$  \*NOM. Tom Wasow (p. c.) calls this last ranking into question, citing evidence that accusative is the default case in English. Woolford 2001: 538, fn 9 explicitly considers this concern, although an actual account of the relevant English facts does not seem available at this time.


<sup>4</sup> This constraint is motivated by sentence processing phenomena involving the attachment of relative clauses, which we discuss in Section 3 below, as well as by grammatical and processing properties of relative clauses observed previously. In the CA model, this idea has been developed as a general agreement process between a relative pronoun and the modified noun (Stevenson in preparation), but here we focus on the more widely studied contributions of Case.



~~library~~, at the word *his*. (Strike-out marks words not yet received in the input.) The parser must posit an XP phrase in which *his* fills the specifier position. At this point in the parse, the head  $X^0$  of this XP projection is unfilled since the input word *books*, which will head the phrase, has not yet been reached. Thus, at the word *his*, there is a mismatch between the lexical specifications of the input string and the parse tree, which constitutes the OB-HD violation.

In our examples below, we will see an OB-HD violation in candidates in which a noun phrase can be interpreted as the specifier of an embedded clause whose verb has not yet been processed, as in *Mary knows Jane left* (14b); the relevant portion of the associated parse tree is shown in (7).

(7) OB-HD violation



ASSIGN- $\theta$ , on the other hand, penalizes underparsing – the parse tree created by *Int* lacks structure that is specified by the lexical properties of the input words. ASSIGN- $\theta$  is violated any time the thematic requirements of a word, as specified in its lexical entry, are not realized in the parse tree. As a violable constraint, it applies equally to obligatory and optional thematic roles. The example above, *James donated his books to the library*, has an instance of underparsing as well: the second argument of *donate* has not been processed yet, leading to an ASSIGN- $\theta$  violation. Note that we assume that a lexical head can assign a thematic role to a phrase as soon as that phrase is projected, even if the head of the recipient phrase is as yet unfilled – i.e., even in the case of an OB-HD violation, discussed in the previous paragraph. Thus, in this example, there is a violation of ASSIGN- $\theta$  for the goal argument of *donate* (~~to the library~~) but not for the theme argument (*his books*), even though ~~books~~ has not been processed yet.

The remaining constraints assess structural markedness. As is standard in OT, a structural constraint may be elaborated as a *universal markedness hierarchy* of related but more finely distinguished constraints. For example, the \*CASE constraint we adopt from morpho-syntactic OT research is specified above as such a markedness hierarchy in (6b), which asserts, for example, that it is worse for an NP to have genitive Case than to have dative Case, with nominative Case being the least marked. AGR-CASE could similarly be further discriminated according to the particular Case features being matched, but for our purposes it is sufficient to consider it as a single constraint that is violated whenever a relative pronoun does not agree with the Case of the modified NP, whatever that Case is.

LOCALITY, another structural constraint, admits of a slightly different type of hierarchy. In LOCALITY (our violable version of the LCA), we assume violations

apply to an XP that asymmetrically c-commands but *follows* (in linear precedence) a phrase YP. We have no evidence, at least as yet, for a difference in markedness according to the syntactic properties of the phrase YP which causes the LOCALITY violation, however, LOCALITY is a *gradiently* violable constraint. That is, if a phrase XP asymmetrically c-commands and follows multiple phrases,  $YP_1 \dots YP_n$ , in a particular candidate structure, then LOCALITY will be violated  $n$  times by XP in that candidate.

Following Legendre, Wilson, Smolensky, Homer and Raymond 1995 and Legendre, Smolensky and Wilson 1998 (see [Chapter 16](#)), we assume that gradiently violable constraints may also lead to a hierarchical family of constraints, a *power hierarchy* generated by multiple local conjunction of a gradiently violable constraint with itself (see [12:\(39c\)](#) and [Section 14:6.1](#)). Such a hierarchy is ordered by number of violations; abbreviating LOCALITY by LOC, the power hierarchy is:

$$(8) \quad \dots \gg \text{LOC}^i \gg \text{LOC}^{i-1} \gg \dots \gg \text{LOC}^2 \gg \text{LOC}$$

where  $\text{LOC}^i$  is violated if a single XP suffers  $i$  violations of LOCALITY. If no other constraints are ranked between the members of this hierarchy, the family of constraints can be collapsed into a single gradiently violable constraint. However, if there is evidence that some constraint  $\div$  dominates  $i$  violations of LOCALITY but is dominated by  $i+1$  violations, then the LOCALITY constraint must be formulated as a power hierarchy with  $\text{LOC}^{i+1} \gg \div \gg \text{LOC}^i$ . As in all power hierarchies, here there is the universal ranking  $\text{LOC}^{i+1} \gg \text{LOC}^i$  for all values of  $i$ . Initially, we will assume the simpler formulation in which LOCALITY is a single constraint, although later we will encounter evidence that LOCALITY indeed must be formulated as a power hierarchy.

## 2. THE ENGLISH CONSTRAINT RANKING

Recall that our technique is to observe the empirically preferred structure of a set of candidate structures at the point of an ambiguity, and then rank the constraints so that the preferred parse is in fact the optimal candidate. We first do this for two types of ambiguity, deriving a hierarchy of the above constraints. We then verify that this ranking makes correct predictions in six other types of ambiguities. (In these examples, we follow the standard approach of ignoring the influence of prosody, not because prosody plays no role in ambiguity resolution, but simply to try to isolate the structural factors which also play a role. Since most sentence processing experiments use textual stimuli, we can be assured of robust preference effects in the absence of prosody.)

Here we focus on the ranking of OB-HD, ASSIGN- $\theta$ , LOCALITY, and the \*CASE constraints, and return to AGR-CASE in [Section 3](#).<sup>5</sup>

---

<sup>5</sup> Some example sentences in this section, (17) – (20), contain relative clauses, but in no situation does AGR-CASE play a role in determining the preference, since the attachment of

---

## 2.1. Two ambiguities that fix the primary constraint ranking

We begin with data on the relative ranking of OB-HD, ASSIGN- $\theta$ , and LOCALITY. Consider the example in (9), which illustrates the general preference to attach a PP as an argument of a preceding verb, rather than as a modifying adjunct of the object NP (Frazier 1978; Gibson 1991). (In this and all remaining examples, the less and more difficult continuations are respectively marked 'T' and '#'.)

(9) John put the candy on

a. T the table.

b. # the table into his mouth.

PP Argument/NP Adjunct		OB-HD	ASSIGN- $\theta$	LOCALITY
John put the candy on [T the table.]				*NP
[# the table into[ his mouth.]			* $\epsilon$	

Tableau (9) shows the two structures possible when processing the word *on*. In the first (preferred) tree, the PP headed by *on* is attached as an argument of the verb *put*. In the second (dispreferred) tree, the PP headed by *on* is attached as a modifier (adjunct) of the preceding NP, *the candy*. (Attachments formed to parse the most recent word are shown with dashed lines; attachments to previous words, with solid lines.) The preferred parse violates LOCALITY because the PP (in bold) asymmetrically c-commands the direct object NP (underlined) but does not linearly precede it. (The LOCALITY violation bears the subscript of the c-commanded phrase.) The dispreferred interpretation violates ASSIGN- $\theta$  because the locative thematic role of *put* has not been assigned: no phrase has been projected in the position of the missing locative argument. The position of the absent locative argument is marked by  $\epsilon$  in our trees: this is no actual phrase at this location in the parse tree. (The violation arises even if the locative argument is optional, as in *Sara brought the letter to Mary*, since ASSIGN- $\theta$  applies equally to obligatory and optional roles.)

Given the indicated violations, it must be that

(10) ASSIGN- $\theta$   $\gg$  LOCALITY

---

the relative clause is not at issue. AGR-CASE can only influence the choice between two or more relative clause attachments.

---

for the preferred structure to be optimal. Note that the constraint OB-HD is satisfied by both candidates (the NP and PP both have their heads filled overtly); thus these competitors provide no information concerning the ranking of OB-HD.<sup>6</sup>

Next consider example (11), which illustrates the preference to group a second noun following a verb with a preceding noun, if possible, rather than to interpret the second noun as the subject of a forthcoming embedded clause. (This example is similar to ones in Pritchett 1992, showing an analogous structural ambiguity with the verb *convince*.)

- (11) I told the department committees
- a. T that budgets were cut.
- b. # would be formed.

Compound Noun/Sentential Complement		OB-HD	ASSIGN- $\theta$	LOCALITY
I told the department <b>committees</b> [T that budgets were cut.]			* <sub>e</sub>	
[# would be formed.]		* <sub>X</sub>		* <sub>NP1</sub>

At the word *committees*, an NP may be formed from *department* and *committees*, of which *committees* is the head, as in the preferred interpretation. This structure violates ASSIGN- $\theta$  because *told* is (at this point) missing its sentential complement (again, the missing argument is indicated by *e*, in the first tree). The sentential argument is provided in the alternative (dispreferred) parse: a phrasal complement is hypothesized, of which *committees* is the specifier (subject). This phrase is denoted XP because it does not (yet) have a head X; this produces an OB-HD violation. There is also a LOCALITY violation because XP asymmetrically c-commands NP<sub>1</sub> (*the department*) but does not precede it. We already know that

<sup>6</sup> Information concerning OB-HD can be obtained by considering another candidate, one just like the second candidate in (9) except that in the argument position marked by *e*, a phrase XP has been projected. Since the head of this phrase has not (yet) appeared in the input, OB-HD is violated. The conclusion is that, like ASSIGN- $\theta$ , OB-HD must dominate LOCALITY. This follows from the stronger result (12) obtained in the subsequent example.

ASSIGN- $\theta$   $\gg$  LOCALITY (10); thus, in order for the dispreferred candidate in (11) to lose, it must be that:

(12) OB-HD  $\gg$  ASSIGN- $\theta$

Combining this with the earlier ranking (10) yields the hierarchy:

(13) OB-HD  $\gg$  ASSIGN- $\theta$   $\gg$  LOCALITY

## 2.2. Predictions of the account

The hierarchy (13) has been derived by considering two candidate parses of each of two strings. The following examples illustrate that this hierarchy correctly accounts for a number of other structural preferences, i.e., that the grammar correctly generalizes well beyond the two examples from which it was determined.

It will be seen that some examples depend only on the *existence* of the constraints we propose, rather than the specific ranking. This arises when the preferred structure violates a subset (possibly empty) of the constraints violated by the dispreferred structure; the preferred candidate must therefore win under any ranking of the constraints. Also, some dispreferred continuations are easy for people to process (as in (14b)), and others are quite difficult (as in (17b) or (18b)). Here we consider only initial preferences, and return briefly to the issue of reanalysis difficulty – how hard a dispreferred continuation is – in Section 4.

The annotated tableaux are self-explanatory; the preference judgments are based on experimental or intuitive evidence from, among others, Frazier 1978 (example (14)); Frazier and Rayner 1982 (15); Traxler, Pickering and Clifton 1996 (16); Gibson 1991 (17); Pritchett 1992 (18), (19). In relative clauses, there is an understood empty operator in the Specifier position of its highest phrase 'XP', which may have an unfilled head position (it is labeled 'CP' when headed by *that*, as in (19) – (20)).

(14) Mary knows Jane

a. T well.

b. # left.

<i>NP/S Complement (Minimal Attachment)</i>		OB-HD	ASSIGN-θ	LOCALITY
Mary knows <b>Jane</b>	VP			
[T well.]	L V' --- V NP knows <b>Jane</b>			
[# left.]	VP V' --- V XP NP X' X knows <b>Jane</b>	*X		

(15) While Mary was mending the socks

- a. T she sang.
- b. # fell.

<i>Direct Object/Main Subject (Late Closure)</i>		OB-HD	ASSIGN- $\theta$	LOCALITY
While Mary was mending the socks				
[T she sang]				
[# fell]		*X	*e	*IP *VP

- (16) The hippie from the tenement on  
 a. T Haight died.  
 b. # drugs died.

<i>Modifier attachment</i>		OB-HD	ASSIGN-θ	LOCALITY
The hippie from the tenement on  [T Haight died.]	NP <sub>1</sub> N' N PP L hippie P' P NP from NP <sub>2</sub> PP the tenement on			
[# drugs died.]	NP NP <sub>1</sub> PP N' N PP hippie P' P NP <sub>2</sub> from the tenement			*PP *NP <sub>2</sub>

- (17) The computer companies  
 a. T failed.  
 b. # buy stinks. [*The computer (that) companies buy...*]

<i>Compound Noun/Relative Clause</i>		OB-HD	ASSIGN-θ	LOCALITY
The computer companies [T failed.]	NP <sub>1</sub> NP <sub>2</sub> N' N' N' N N companies computer			
[# buy stinks.]	NP NP <sub>1</sub> XP N' X' N X XP computer NP <sub>2</sub> X' X companies		*X <sub>1</sub> *X <sub>2</sub>	



(18) John gave the child the dog

- a. T for Christmas.
- b. # bit medicine. [*the child (that) the dog bit ...*]

<i>Double Object/Relative Clause</i>		OB-HD	ASSIGN-θ	LOCALITY
John gave the child <b>the dog</b> [T for Christmas.]				*NP <sub>1</sub>
[# bit medicine.]		*X *X	*e	

(19) Mary told the doctor that

- a. T she had fainted.
- b. # examined her that she felt fine.

<i>Sentential Complement/Relative Clause</i>		OB-HD	ASSIGN-θ	LOCALITY
Mary told the doctor <b>that</b> [T she had fainted.]				*NP
[# examined her that she felt fine.]			*e	

### 2.3. LOCALITY is a power hierarchy

Consider the example in (20), in which the PP headed by *on* can be attached as the second argument of the verb *put* or as an adjunct (modifier) of the verb *seen*.

- (20) I put the biscuits that Jane had seen on  
 a. T the table back in the tin.  
 b. # the table.

<i>PP Argument/NP Adjunct</i>	LOC <sup>2</sup>	ASSIGN-θ	LOC
<p>John put the biscuits that Jane had seen on</p> <p>[T the table back in the tin.]</p>		*e	*NP <sub>2</sub>
<p>[# the table.]</p>	*		*NP <sub>1</sub> *CP *IP *VP *NP <sub>2</sub>

In the account developed so far, even with the numerous LOCALITY violations engendered by the argument attachment, it should be preferred here exactly as in (9), *John put the candy on ...*, because ASSIGN-θ ≫ LOCALITY. However, evidence indicates that with a greater distance between the argument site and the phrase to be attached, a more recent adjunct site may prevail (Kamide 1998). We suggest that this is a case of a gradiently violable constraint for which multiple violations (by a single entity) have a stronger effect than merely more marks at a single level in the hierarchy. As mentioned above, the device of a power hierarchy has been employed in formally similar situations in other OT analyses (e.g., Chapters 12 and 14), and we propose to use it here. Multiple violations of LOCALITY by a single attachment cause a violation of higher-ranked constraints in the universal power hierarchy (8):

- (21) ... ≫ LOC<sup>i</sup> ≫ LOC<sup>i-1</sup> ≫ ... ≫ LOC<sup>2</sup> ≫ LOC

We make the minimal assumption permitted by these data, that it is LOC<sup>2</sup> that is at work in example (20). (The losing candidate violates not only LOC<sup>2</sup>, but also

LOC<sup>3</sup> – LOC<sup>5</sup>. Until we are forced to deploy higher constraints, we will employ LOC<sup>2</sup>.) With LOC<sup>2</sup> dominating ASSIGN- $\theta$  as indicated in the tableau (20), the correct preference is obtained. Thus while a slightly-non-local attachment (with a single XP violating LOC) is preferable to a violation of ASSIGN- $\theta$ , a more highly non-local attachment (with two or more XP's violating LOC, yielding a violation of LOC<sup>2</sup>) is not preferable to a violation of ASSIGN- $\theta$ .

#### 2.4. \*CASE rankings

In the examples thus far, the constraints from the \*CASE subhierarchy play no role in determining preferences. Interestingly, one area where they appear to play a visible role is in filler-gap ambiguities – i.e., in determining the preferred association of a filler (*wh*-phrase) with its *gap* or trace *t*, an empty element with which it is co-indexed (Box 12:1) (Artstein and Stevenson 1999; Artstein 2000).

For example, at the point of processing the verb *give* in (22), the *wh*-phrase *which dog* may be: (a) interpreted as the object of the verb (receiving Acc Case, compatible with the continuation in (22a)); (b) interpreted as the indirect object of the verb (receiving Dat Case, compatible with the continuation in (22b)); or (c) unassociated with either complement position (compatible with the continuation in (22c)):

(22) Which dog<sub>*i*</sub> did John give ...

- a. *t<sub>i</sub>* to Mary? [The trace receives the theme role and Acc Case.]
- b. *t<sub>i</sub>* a bone? [The trace receives the recipient role and Dat Case.]
- c. a bone to *t<sub>i</sub>*? [There is no trace when the verb is processed, so the thematic role is unassigned.]

Artstein and Stevenson 1999 find that the preferred continuation is (a): it is better to associate a filler with an argument position that is assigned Acc Case (a), rather than to leave it unassociated (c). (The preference for (a) over (b) is given by the \*CASE hierarchy: \*DAT  $\gg$  \*ACC.) If the filler is associated with an argument position, as in (a) (or (b)), then the verb can assign a thematic role to that position and avoid an ASSIGN- $\theta$  violation. If the filler is left unassociated, as in (c), then it cannot. This indicates that:

(23) ASSIGN- $\theta$   $\gg$  \*ACC

Continuing this example, however, reveals that it is better to leave a filler unassociated (and a thematic role unassigned) than to associate it with an argument position that is assigned Dat Case. Consider the situation arising if the partial input in (22) is continued with *a bone*, compatible with one of the dispreferred continuations in ((22b) or ((22c) above: which of those possibilities will then be preferred? The filler may be associated with a trace in the Dat-receiving indirect object position (as in (24a)), or may be left unassociated (as in (24b)):

- (24) Which dog<sub>*i*</sub> did John give a bone ...
- a. Which dog<sub>*i*</sub> did John give *t<sub>i</sub>* a bone?
  - b. Which dog<sub>*i*</sub> did John give a bone *e*? [second argument not filled]
  - c. Which dog<sub>*i*</sub> did John give a bone to *t<sub>i</sub>*? [continuation of (b): second argument filled with subsequent PP]

The preferred continuation is of (24) is (c), indicating the preference to avoid a Dat-receiving trace in favor of leaving the filler unassociated – that is, (b) must have been preferred over (a) given the input in (24). Here, the \*CASE violation is worse than the ASSIGN-θ violation; that is:

- (25) \*DAT ≫ ASSIGN-θ

Based on these kinds of preferences in the incremental interpretation of a filler, Artstein and Stevenson 1999 derive the following sub-ranking:

- (26) \*DAT ≫ ASSIGN-θ ≫ \*ACC

The general idea behind the analysis is that there is a tension in binding a gap by a filler: on the one hand, the binding eliminates a violation of ASSIGN-θ (because the filler/gap chain serves as the expected argument of a verb); on the other hand, it induces a \*CASE violation (because the bound gap – the co-indexed empty element *t<sub>i</sub>* – receives Case from the verb). The upshot of the above ranking is that if the gap is in a position that receives accusative Case then it is better to bind it by the filler, but if the gap is in a dative-receiving position, it is better to leave it unbound.

When we incorporate the \*CASE subhierarchy into our earlier analysis, we also determine the following from example (18), in which the preferred interpretation has an additional \*DAT violation (not shown in the earlier tableau):

- (27) OB-HD ≫ \*DAT

That is, the high rank of OB-HD entails that it is not possible to avoid a \*DAT or \*ACC violation by positing an empty phrase to allow restructuring of NPs to achieve lower ranked Case violations.

These observations lead to the following hierarchy (28), merging our earlier rankings with the new \*CASE rankings. Note that the examples thus far provide no evidence concerning the relative ranking of the \*CASE and LOCALITY constraints.

- (28) OB-HD ≫ {LOC<sup>2</sup>, \*DAT} ≫ ASSIGN-θ ≫ {LOC<sup>1</sup>, \*ACC} ≫ \*NOM
-

### 3. RELATIVE CLAUSE ATTACHMENT TYPOLOGY

A relatively well-studied ambiguity is the attachment of a relative clause when there is more than one NP that it could modify, as in (29) (Cuetos and Mitchell 1988).

(29) The servant of the actress who ...

The relative clause introduced by the relative pronoun *who* could modify either *the servant* ('high attachment') or *the actress* ('low attachment'). Further embedding of NPs within PP modifiers can occur, making even more than two NPs possible attachment sites (as in *the servant of the actress with the director who...*). Experimental studies have investigated the preferences in two- and three-NP attachment site cases, revealing that high or low attachment preference varies with differences in the number of attachment sites, the type of prepositions, and the language.

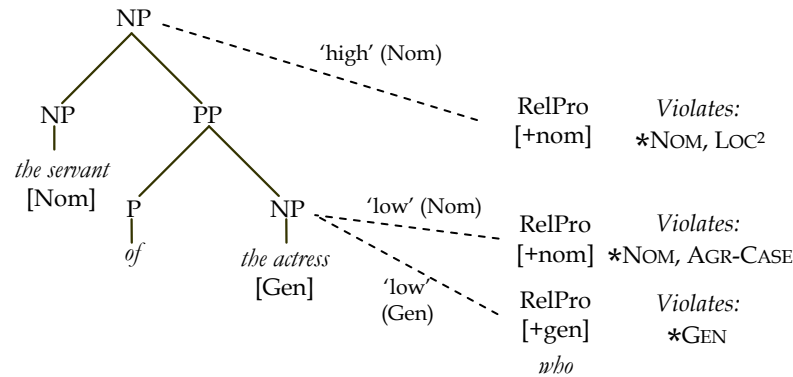
Our goal here is to determine, under the varying conditions, both the possible and impossible patterns of attachment preferences (i.e., high vs. middle vs. low attachment). We propose an account of the syntactic contribution to these patterns (as distinct from the contributions of prosody and discourse, as in Fodor 1998, Hemforth, Konieczny and Scheepers 2000). Our proposal is like those of Gibson, Pearlmutter, Canseco-Gonzalez and Hickok 1996 and Hemforth et al. 2000 in its reliance on two factors, but in contrast, our account involves no 'processing' principles or constraints that are distinct from the syntactic constraints proposed as elements of the competence grammar.

Specifically, we propose to explain relative clause attachment preferences as the result of the interaction of two classes of constraint introduced above: the LOCALITY constraints (5b), (8), and the Case constraints (6), the latter comprised of AGR-CASE (6c) and the universal Case markedness hierarchy \*GEN >> \*DAT >> \*ACC >> \*NOM (6b). Clearly, the LOCALITY constraints are sensitive to the relative height of the alternative NPs to which the relative clause can attach. Because the relative pronoun (e.g., *who*, *which*, *that*) must agree in Case with the noun the relative clause modifies (according to AGR-CASE (6c)), the Case constraints are sensitive to the type of prepositions, which vary in the Case they assign to their complement NPs.

Conflict arises among the Case and LOCALITY constraints when the head of the complex NP (the 'highest' NP, such as *the servant* in (29)) is assigned Nom or Acc Case, and the prepositions assign the more marked Cases Dat or Gen. In such configurations, the Case markedness hierarchy favors high attachment; LOCALITY always favors low attachment. This is illustrated in (30), which shows three candidate attachments for *who* in (29). Two 'low' attachment candidates are shown, one with Nom Case, the other Gen Case. The latter satisfies AGR-CASE because it agrees with the Gen Case of the NP to which it attaches (*the actress*). But this candidate has the most marked Case, violating \*GEN. The low-attached candidate with the least marked Case, Nom, improves the \*CASE violation to \*NOM, but incurs a violation of AGR-CASE. The high attachment can satisfy AGR-CASE and still have the least-marked Case, since the highest NP bears Nom Case.


But this violates LOCALITY, specifically,  $LOC^2$ : with the high attachment, *who* asymmetrically c-commands a PP (*of*) and an NP (*actress*) that precede it.

(30) The servant of the actress who ...



Which attachment is preferred – that is, optimal? If LOCALITY is dominated by the Case constraints –  $\{*\text{GEN}, \text{AGR-CASE}\} \gg \text{LOC}^2$  – the high attachment wins; otherwise – if  $\text{LOC}^2 \gg **\text{GEN}$  or  $\text{LOC}^2 \gg \text{AGR-CASE}$  – one of the low attachments is optimal. The possibilities are richer when we consider prepositions that assign other Cases, and when we consider three nested NPs, where attachment to the highest NP will violate  $LOC^4$ , and to the middle NP,  $LOC^2$ , as shown in (31).

(31) LOCALITY power hierarchy violations



We now explore the typology of attachment preferences that results from reranking the proposed constraints. In this analysis, we consider the portion of the full typology in which the following conditions hold:

(32) Conditions of analyzed structures

- a. The highest NP has Nom or Acc Case (that is, is the subject or object of a verb, and not itself the object of a preposition)

- b. The prepositions employed assign Gen Case (*of*) or Dat Case (all other prepositions)
- c. The relative pronoun is Case-ambiguous (i.e., *who*, *which*, or *that*, each of which is compatible with any Case)

Under these conditions, the following holds.

- (33) Attachment of a relative clause to the complement of a preposition requires violation of at least one of the following constraints:

AGR-CASE, \*GEN, or \*DAT.

Agreeing in Case violates either \*GEN or \*DAT because those are the Cases assigned by a preposition; disagreement in Case (presumably to avoid a higher-ranked Case constraint) violates AGR-CASE.

### 3.1. Three attested attachment preference patterns

As mentioned above, our initial goal is to derive from the competence grammar a typology of possible (and impossible) patterns of relative clause attachment preferences. Several types of patterns have been proposed to describe the observed preferences in various languages. For example, it's been suggested that English and Brazilian Portuguese consistently show low attachment preference (Gibson et al. 1996; Miyamoto 1998). A number of languages, Spanish and German among them, appear to exhibit high attachment in two-site cases and low attachment in three-site cases (Cuetos and Mitchell 1988; Gibson et al. 1996; Hemforth et al. 2000). Finally, it's been proposed that, at least for some two-site cases in some languages (English and Spanish), high attachment is preferred after some prepositions (*of*) but not others (Gilboy, Sopena, Clifton and Frazier 1995; Frazier and Clifton 1996). In this section, we explain how each of these attested patterns results from a relative ranking of Case and LOCALITY constraints in our framework.

A simple system is described in (34): (a) gives the ranking conditions, and (b) the resulting preference pattern.

- (34) LOCALITY strictly dominates Case
- a.  $LOC^2 \gg AGR-CASE$  or  $LOC^2 \gg *GEN$
  - b. Low attachment is optimal in every configuration.

In this type of language, it is more important to attach to the most recent (lowest) NP than to satisfy a Case-based condition, i.e., than to ensure that the Case of the relative pronoun matches the Case of the NP it attaches to, or that it avoids the marked Cases Gen and Dat.

More precisely, if  $LOC^2 \gg AGR-CASE$  then any high attachment is dispreferred to low attachment with Nom Case. If the relative pronoun in the high attachment candidate has Nom Case, then high attachment violates \*NOM and at least  $LOC^2$ , while low attachment violates \*NOM and AGR-CASE. Since  $LOC^2 \gg AGR-CASE$ , low attachment is preferred. A high attachment with a non-

Nom Case can only make the high attachment even worse, so any high attachment is worse than low attachment with Nom Case. (There may be an even better low attachment candidate – that is, one with agreeing Case C, if  $\text{AGR-CASE} \gg *C$  – but in any event, a low rather than high attachment will be optimal.)

Similarly, if  $\text{LOC}^2 \gg *GEN$  then any high attachment is dispreferred to low attachment with Case agreeing with that NP. Whatever this Case may be, the worst violation it can incur is  $*GEN$ , universally top-ranked on the  $*CASE$  hierarchy; and even a  $*GEN$  violation is preferred to the LOCALITY violation resulting from attaching to any higher NP, which is at least  $\text{LOC}^2$ . Since  $\text{AGR-CASE}$  is satisfied by this low attachment, its ranking is irrelevant. Thus no high attachment can be optimal. (Again, there may be a better low attachment candidate: low attachment with Nom Case will be optimal if  $*C \gg \text{AGR-CASE}$ , where C is the agreeing Case. But either way, similarly to the situation above, a low attachment is best.)

Thus under either sub-case of the rankings (34a), low attachment is optimal.

A more complex pattern is described in (35).

(35) Severe, but not mild, LOCALITY violations are worse than Case violations

a.  $\text{LOC}^4 \gg \text{AGR-CASE}$  or  $*\text{LOC}^4 \gg *GEN$

**AND**  $\{\text{AGR-CASE}, *GEN, *DAT\} \gg \text{LOC}^2$

b. High attachment is preferred in two-site cases; low attachment in three-site cases where both prepositions assign the same Case, or the higher preposition assigns a “worse” Case (Gen) than the lower preposition (Dat). (That is, low attachment obtains with three sites when the Case assigned to the middle NP is the same or worse than that assigned to the lowest NP. The alternative situation – in which the middle NP is assigned a less marked Case than the lowest NP – is described in (40) below.)

In this class of language, a Case-based condition is more important than a low-ranked LOCALITY constraint, but less important than a high-ranked LOCALITY constraint.

More precisely, with two attachment sites, high attachment with agreeing Case (Nom or Acc) will violate  $\text{LOC}^2$ , but satisfy the constraints that out-rank it,  $\text{AGR-CASE}$ ,  $*GEN$ , and  $*DAT$ ; low attachment will violate at least one of those higher-ranked constraints (33).

Recall that, with three attachment sites, high attachment violates  $\text{LOC}^4$ , intermediate attachment violates  $\text{LOC}^2$ , and low attachment incurs no LOCALITY violations (31). Given the rankings in (35a), high attachment with three sites cannot be optimal because of the  $\text{LOC}^4$  violation: if  $\text{LOC}^4 \gg \text{AGR-CASE}$ , a lower attachment with Nom Case has higher Harmony, since it will at worst violate  $\text{AGR-CASE}$ ; and if  $\text{LOC}^4 \gg *GEN$ , a lower attachment that agrees in Case will have

---



higher Harmony because it will at worst violate \*GEN. This is parallel to the previous case (34), where LOC<sup>4</sup> here plays the role of LOC<sup>2</sup> there.

Also, with three sites, intermediate attachment cannot be optimal when an intermediate NP is assigned the same or a more marked Case than the lower NP. By (33), both violate one of AGR-CASE, \*GEN, or \*DAT, with the Case violation of the middle NP no better than that of the lowest NP. The lowest NP is then preferred because it does not incur the additional \*LOC<sup>2</sup> violation arising with intermediate attachment.

With three sites, then, low attachment is optimal, except in the single configuration discussed below in (40). This establishes (35).

A third still more complex pattern is given in (36).

- (36) The relative importance of Locality and Case depends on the particular Case being assigned.
- a. LOC<sup>4</sup> » AGR-CASE or LOC<sup>4</sup> » \*GEN  
AND {AGR-CASE, \*GEN} » LOC<sup>2</sup> » \*DAT
  - b. High attachment in two-site cases with a genitive-assigning preposition (*of*), low attachment in two-site cases with a dative-assigning preposition, and low attachment in three-site cases where both prepositions assign the same Case, or the higher preposition assigns a “worse” Case (Gen) than the lower preposition (Dat). (Compare (35b) above.) In this type of language, the importance of LOCALITY relative to a Case-based condition varies according to the identity of the \*CASE constraint involved.

The ranking here is the same as in (35a), except for the domination of \*DAT by LOC<sup>2</sup>. The resulting pattern is thus the same as in (35b), except when the preposition in a two-site structure assigns dative Case. In that situation, it is, under the ranking (36), optimal to attach low with Dat Case and avoid the now higher-ranked LOC<sup>2</sup> violation. Thus, for example, the ranking (36a) yields a distinction in preference between *the sweater with sleeves that...* and *the sweaters of wool that...*, the former yielding a low-attachment preference (*the sweater with sleeves that are too short*), and the latter a high-attachment preference (*the sweaters of wool that are too expensive*). (See Gilboy et al. 1995; Frazier and Clifton 1996.)

### 3.2. Typologically excluded systems

The proposed constraints predict that certain patterns are impossible, due to properties of the grammatical system: the universal rankings of the LOCALITY power hierarchy and of the \*CASE markedness hierarchy, and their limited interactions in constraint re-ranking. We discuss two such examples in (37) and (38).

- (37) Low attachment in all two-site configurations and high attachment in a three-site configuration.

For low attachment to be preferred in all two-site configurations,  $LOC^2$  must dominate either AGR-CASE, or both \*GEN and \*DAT (see (30), which illustrates the situation for a genitive-assigning preposition). Since universally  $LOC^4 \gg LOC^2$ , this entails that  $LOC^4$  also dominates either AGR-CASE or both \*GEN and \*DAT, and by the same logic low attachment is preferred with three sites, contrary to (37).

- (38) With two sites, high attachment after a dative-assigning preposition and low attachment after a genitive-assigning preposition.

To favor high attachment after a dative-assigning preposition, both AGR-CASE and \*DAT must dominate  $LOC^2$ . Since \*GEN  $\gg$  \*DAT universally, this entails that \*GEN also dominates  $LOC^2$ . Thus, low attachment with Gen Case (or with disagreeing Nom Case) is less harmonic than high attachment with Nom Case, so high attachment is required for genitive- as well as dative-assigning prepositions.

The impossibility of these two patterns in our OT system is interesting because, while the “reverse” of each has been proposed to correspond to behavior in some language (i.e., high attachment in two-sites and low in three sites, or high attachment after a genitive-assigner and low after a dative assigner), the patterns in (37) and (38) have not, to our knowledge, been empirically observed.

### 3.3. Typologically predicted systems

Two patterns are predicted possible by the proposed analysis but do not, to our knowledge, correspond to patterns that have been documented in the literature: (39) and (40).

- (39) Case extremely high-ranked relative to LOCALITY
- a. \*AGR-CASE, \*GEN, \*DAT  $\gg$   $LOC^4$ .
  - b. High attachment with three- (and two-) sites

In this type of language, Case constraints are so much stronger than LOCALITY that attachment to the highest NP to have agreeing Nom or Acc Case is compelled even when this incurs four simultaneous violations of LOCALITY. Thus NPs embedded in PPs are simply unavailable for modification by a relative clause (at least in two- and three-site cases). Languages uniformly showing high attachment in three-site cases have not, to the best of our knowledge, been attested. If in fact such languages do not exist, this might suggest some grammatical or functional factor limiting the constraint interaction involving the LOCALITY and \*CASE hierarchies, to prevent a ranking as in (39a). This factor may correspond to a limitation on the Case constraints, prohibiting a situation in which both (i) \*GEN and \*DAT are very highly ranked, and (ii) AGR-CASE is very highly ranked. Condition (i) makes the use of such Cases difficult in general. Condition (ii) is in direct tension with (i): by forcing Case agreement between a relative pronoun and the modified NP, it strongly restricts the use of less marked

---

Cases for relative pronouns. Alternatively, the factor in question may correspond to limitations on the LOCALITY power hierarchy, prohibiting its domination by certain classes of constraints. We leave further investigation of such restrictions on constraint rankings to future work.

- (40) Intermediate attachment with three sites
- a. Rankings of (35) or (36)
  - b. Intermediate attachment with three sites when the intermediate-level PP assigns dative Case, and the lower-level PP, genitive. Otherwise, low attachment with three sites.

The relevant situations here are those, alluded to in (35) and (36) above, in which the lowest NP is assigned a worse Case than the middle NP in the three-site configuration. The Harmony benefit of \*DAT over \*GEN arising from intermediate rather than low attachment, with agreeing Case, exceeds the Harmony cost of incurring LOC<sup>2</sup>. This is similar to the two-site situation described for (36), in that a worse Case violation (\*GEN) on the lowest NP can force higher attachment, where a lesser Case violation (\*DAT) cannot. But in the three-site instance, attachment is forced to the middle NP, because the additional benefit of \*NOM over \*DAT which derives from highest attachment is exceeded by the cost of LOC<sup>4</sup>.

Thus, the typology derived from the above constraints predicts the existence of languages for which intermediate attachment is possible in a phrase such as *the thread from the sweater of wool that was hand-dyed* .... Although to our knowledge this situation has not been experimentally investigated, it seems an intuitively plausible result for English, in which it has been proposed that the genitive-assigning *of* does not block higher attachment.

Obviously, at this time, the extent of psycholinguistic investigation of such questions is sufficiently limited that the empirical status of this prediction, and that of (39), remain to be determined.

#### 4. REANALYSIS

Thus far we have developed a theory of initial preferences – that is, we explain the preferred interpretation for an ambiguity as the candidate structure that is optimal according to our OT constraints (1a). The other class of data that a theory of ambiguity resolution must account for is the behavior when an ambiguity is (later) disambiguated with a result inconsistent with the initial preference (1b). The parser at this point must abandon its initial preference and adopt a different structure as its newly preferred interpretation – a phenomenon known as reanalysis.<sup>7</sup>

---

<sup>7</sup> We use the term *reanalysis* to mean any change in the parser's representation of what is the preferred interpretation of the input. The actual operation may be one of reparsing, modifying the current parse, selective backtracking, or (in a parallel framework) adopting

Two issues must be addressed in any theory of reanalysis:

(41) What qualifies as reanalysis?

Since simply the act of integrating each new word into the developing parse tree constitutes a change to the preferred parse, we must define what it means to be a change that qualifies as a change in the preferred interpretation.

(42) What determines the degree of difficulty induced by any particular reanalysis?

Difficulty in changing a preferred interpretation lies on a continuum from mild influences on processing (e.g., if the sentence fragment *Mary knows Jane* is continued with the verb *left*), to severe garden path effects (as in *Jane gave the child the dog* continued with *bit a band-aid*). A theory of reanalysis must therefore specify the relation between the parse tree changes identified in (41) with a measure of processing difficulty.

Regarding our first question (41), we assume, following Stevenson 1998, that *any* change to the parse constitutes a reanalysis. That is, every incremental step in parsing is a process of revising the prior interpretation, either by “filling in” empty structure (the elements of the parse indicated by ‘*e*’ or ‘XP’ in our tree diagrams), or by actually changing earlier attachment relations. This is the simplest assumption that can be made about reanalysis — that is, there is no special set of reanalysis operations, rather *Int* uniformly provides candidates for each and every sentence fragment.

Not only is this the most parsimonious approach, it also is most consistent with our OT framework, which treats all sentence fragments (partial or complete sentences) uniformly: candidates for each incremental step in the input string are constructed uniformly, without singling out particular operations in the working of *Int* that qualify as reanalysis operations. Similarly, there are no constraints in the evaluation mechanism that compare types of changes from the previously preferred parse to the candidate being evaluated. This contrasts with other approaches to reanalysis in which certain incremental parsing (reanalyzing) operations are distinguished as more costly (such as Lewis 1998; Sturt and Crocker 1996; Weinberg 1995). Here, each candidate is produced with the same set of operations, and evaluated solely according to the constraints provided by the grammar.

Given that every incremental incorporation of an input word is “reanalysis”, and therefore there are no special reanalysis operations or constraints, how do we determine the *difficulty* of reanalysis as required by (42)? Again, we are guided by the simplest assumption under our OT framework: the degree of difficulty at each stage in processing arises from the pattern of constraint violations that the optimal candidate induces — that is, its Harmony.

---

an alternative structure that was previously computed but simply not the preferred choice.

---

Specifically, processing difficulty at word  $w_i$  is measured by the *change in Harmony* (or ‘unmarkedness’) from the preferred parse at word  $w_{i-1}$  to the preferred (optimal) candidate at word  $w_i$ . Note that here we depart from standard practice in OT by comparing two structures which correspond to two different inputs: the optimal structure for words  $w_1$  to  $w_{i-1}$ , vs. the optimal structure for words  $w_1$  to  $w_i$ . If the pattern of constraint violations worsens (from one preferred interpretation to the next) due to additional violations of low ranked constraints, there will be only mild effects, while additional violations of high ranked constraints can cause significant processing difficulty.

As an illustration of this range of processing difficulty, we consider the dispreferred continuations of three of our earlier preference examples. In these examples, the following subhierarchy of our English ranking will play a role in the account of relative difficulty:

- (43) OB-HD  $\gg$  LOC<sup>2</sup>  $\gg$  ASSIGN- $\theta$   $\gg$  {LOC<sup>1</sup>, \*ACC}  $\gg$  \*NOM

In example (14), the preferred interpretation of the input:

*Mary knows Jane*

violates only \*CASE constraints (\*NOM and \*ACC), with *Jane* attached as the argument of *knows*. If this fragment is continued with the verb *left*, as in *Mary knows Jane left*, people experience little processing difficulty. (This difficulty is generally below the level of conscious awareness, but has been attested experimentally, Frazier and Rayner 1982). In our framework, the new parse will have a single violation of ASSIGN- $\theta$  (due to the missing optional object argument of the verb *left*) and two violations of \*NOM. Thus the Harmony decreases, but only by the additional violation of a constraint intermediate in our hierarchy (and in fact one \*CASE constraint has been replaced by a lower-ranked one).

Compare the dispreferred continuation of example (9), in which the initial fragment:

*John put the candy on the table*

violates LOC<sup>1</sup> (since *on* attaches above the NP *the candy* as the second argument of *put*), as well as \*NOM, \*ACC and \*DAT for the Cases assigned to the NPs. The dispreferred continuation, *John put the candy on the table into*, causes noticeable processing difficulty for the human sentence processor (Gibson 1991). In our model, the new parse violates the same constraints as the previous parse, with the addition of a LOC<sup>2</sup> violation for the new PP (which c-commands and follows *the candy on the table*). Thus here the change in markedness is more severe than above, as there is the additional violation of a higher ranked constraint, LOC<sup>2</sup>.

Finally, the preferred interpretation of example (18):

*John gave the child the dog*

similarly violates LOC<sup>1</sup> and the three \*CASE constraints \*NOM, \*ACC and \*DAT. The dispreferred continuation, *John gave the child the dog bit*, causes a severe garden

path effect in people (Pritchett 1992). In this parse, there is a violation of OB-HD,<sup>8</sup> ASSIGN- $\theta$ , \*ACC, and two violations of \*NOM. Note that although the \*CASE constraints have improved, the violation of the very highly ranked OB-HD decreases the Harmony greatly, accounting for the difficulty.

We conclude then that a promising notion for building a theory of reanalysis difficulty is the change in Harmony between the optimal structures determined successively during incremental parsing. The approach has the advantage of applying a uniform mechanism for initial preferences and reanalysis difficulty, and strengthens the evidence for the grammar-parser relation we hypothesized in (1).

## 5. CONCLUSION

The approach to a theory of sentence processing begun here has, in our view, three primary advantages over alternative types of theories.

### (44) Advantages of the OT theory

- a. *Basis in grammar*: We look directly to the grammar for the cause of preferences. The assumption is that surface-level phenomena which seem to influence processing (such as word-order flexibility; cf. Gibson et al. 1996; Miyamoto 1998) are themselves a reflex of grammar (likely the Case system, which we exploit here).
- b. *Uniformity*: All our constraints have grammatical motivation and play a role in both grammar and processing. This is possible in part because we adopt a grammaticalized notion of locality that captures effects that have typically been cast as extra-syntactic in other work (e.g., Gibson 1991; Gibson et al. 1996; Hemforth et al. 2000; Stevenson 1994b).
- c. *Restrictiveness*: The theory has the flexibility to yield general preference patterns observed cross-linguistically (due to constraint re-ranking), while at the same time formally generating a restricted typology that excludes unattested patterns (due to the general grammatical structure of OT – such as universal markedness hierarchies and strict domination, rather than numerical weighting). We have demonstrated this behavior for one important class of ambiguities (relative clause attachment), and future work will investigate the generality of this result.

Of course, fundamental conceptual questions remain open. How do the rankings employed here for a “performance” theory relate to those of the ordinary “competence” grammar? It appears that, very generally speaking, the ranking needed for the performance theory of a language will be more

---

<sup>8</sup> The OB-HD violation is due to the highest phrase XP of the relative clause, which holds an (understood) empty operator in specifier position, but whose head has no lexical material; see the parse tree for (18b).

---

constrained than that for the competence grammar. For instance, in the competence theory, it makes no difference how the grammar ranks two constraints  $\div$  and  $\div'$  which are both unviolated in the surface forms of the language. Both these constraints can be simultaneously satisfied given a *complete* sentence. But prior to completion of a sentence, it may not be possible to satisfy both  $\div$  and  $\div'$ , so their relative priority may guide the preference among competing incremental analyses. Thus these same constraints may need to be ranked in a particular way to make the right predictions concerning on-line processing.

Now if the hierarchy needed for the processing theory has additional rankings relative to the hierarchy needed for the competence theory, the question arises of how the learning process determines not only the rankings of the competence grammar, but also the additional rankings relevant to processing. We will not speculate on this intriguing question here. We note, however, that there are suggestive parallels between this question, pertaining to syntactic comprehension, and the problem discussed in [Chapter 17](#) of learning the 'hidden rankings' apparently involved in phonological production.

The set of preliminary empirical results reported here suggest that the general approach may be a productive one. OT provides a conception of grammar that incorporates one of the central computational insights of sentence processing theories, competitive optimization. It thus raises the possibility of a unified theory in which a single grammar can explain both the "competence" data of theoretical linguistics and the "performance" data of psycholinguistics.

---

**References**

- Aissen, Judith. 2001. Markedness and subject choice in Optimality Theory. In *Optimality-Theoretic Syntax*, eds. Géraldine Legendre, Sten Vikner and Jane Grimshaw, 61-96. Cambridge, MA: MIT Press.
- Artstein, Ron. 2000. Case Constraints and Empty Categories in Optimality Theory Parsing. Ms. Rutgers University.
- Artstein, Ron, and Stevenson, Suzanne. 1999. Case Constraints and Empty Categories in Optimality Theory Parsing. Talk presented at the Conference on Architectures and Mechanisms in Language Processing (AMLaP 1999). Edinburgh, Scotland.
- Bader, Markus. 1996. Reanalysis in a Modular Architecture of the Human Sentence Processing Mechanism. Poster presented at the 10th Annual CUNY Conference on Human Sentence Processing.
- Chomsky, Noam. 1981. *Lectures on Government and Binding*. Dordrecht, Holland: Foris.
- Cuetos, Fernando, and Mitchell, Don. 1988. Cross-linguistic differences in parsing: Restrictions on the use of the Late Closure strategy in Spanish. *Cognition* 30:73-105.
- Fanselow, Gisbert, Cavar, Damir, Kliegl, Reinhold, and Schlesewsky, Matthias. 1999. Optimal Parsing: Syntactic Parsing Preferences and Optimality Theory. Ms. University of Potsdam.
- Fodor, Janet. 1998. Learning to parse? *Journal of Psycholinguistic Research* 27:285-319.
- Frazier, Lyn. 1978. On comprehending sentences: syntactic parsing strategies. IULC. Bloomington, IN.
- Frazier, Lyn, and Clifton, Charles. 1996. *Construal*. Cambridge, MA: MIT Press.
- Frazier, Lyn, and Rayner, Keith. 1982. Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology* 14:178-210.
- Gibson, Edward. 1991. *A Computational Theory of Human Linguistic Processing: Memory Limitations and Processing Breakdown*. Doctoral dissertation, Carnegie-Mellon University.
- Gibson, Edward. 1995. Optimality Theory and human sentence processing. Talk presented at the Conference on Optimality and Competition in Syntax. MIT, Cambridge, MA.
- Gibson, Edward, Pearlmutter, Neal, Canseco-Gonzalez, Enriqueta, and Hickok, Greg. 1996. Recency Preference in the Human Sentence Processing Mechanism. *Cognition* 59.
- Gibson, Edward, and Broihier, Kevin. 1998. Optimality Theory and human sentence processing. In *Is the Best Good Enough? Optimality and Competition in Syntax*, eds. Pilar Barbosa, Danny Fox, Paul Hagstrom, Martha McGinnis and David Pesetsky, 157-191. Cambridge, MA: MIT Press and MITWPL. Rutgers Optimality Archive 117.
- Gilboy, Elizabeth, Sopena, Josep-Maria, Clifton, Charles, and Frazier, Lyn. 1995. Argument Structure and Association Preferences in Spanish and English Complex NPs. *Cognition* 54:131-167.
- Grimshaw, Jane. 1997. Projection, heads, and optimality. *Linguistic Inquiry* 28:373-422.
- Grimshaw, Jane. 2001. Optimal clitic positions and the lexicon in Romance clitic systems. In *Optimality-Theoretic Syntax*, eds. Géraldine Legendre, Sten Vikner and Jane Grimshaw, 205-240. Cambridge, MA: MIT Press.
-



- Hemforth, Barbara, Konieczny, Lars, and Scheepers, Christoph. 2000. Syntactic attachment and anaphor resolution: Two sides of relative clause attachment. In *Architectures and Mechanisms of Language Processing*, eds. Matthew Crocker, Martin Pickering and Charles Clifton. Cambridge University Press.
- Kamide, Yuki. 1998. *The Role of Argument Structure Requirements and Recency Constraints in Human Sentence Processing*. Doctoral dissertation, University of Exeter, UK.
- Kayne, Richard S. 1994. *The Antisymmetry of Syntax*. Cambridge, MA: MIT Press.
- Legendre, Géraldine, Smolensky, Paul, and Wilson, Colin. 1998. When is less more? Faithfulness and minimal links in *wh*-chains. In *Is the Best Good Enough? Optimality and Competition in Syntax*, eds. Pilar Barbosa, Danny Fox, Paul Hagstrom, Martha McGinnis and David Pesetsky, 249–289. Cambridge, MA: MIT Press and MITWPL. Rutgers Optimality Archive 117.
- Legendre, Géraldine, Vikner, Sten, and Grimshaw, Jane eds. 2001. *Optimality-Theoretic Syntax*. Cambridge, Mass.: MIT Press.
- Legendre, Géraldine, Wilson, Colin, Smolensky, Paul, Homer, Kristin, and Raymond, William. 1995. Optimality in *wh*-chains. In *University of Massachusetts Occasional Papers in Linguistics 18: Papers in Optimality Theory*, eds. Jill Beckman, Laura Walsh Dickey and Suzanne Urbanczyk, 607–636. Amherst, MA: GLSA, University of Massachusetts.
- Lewis, Richard L. 1998. Reanalysis and limited repair parsing: Leaping off the garden path. In *Reanalysis in Sentence Processing*, eds. Janet Dean Fodor and Ferreira Fernanda, 247–285. Kluwer Academic.
- MacDonald, Maryellen, Pearlmutter, Neal, and Seidenberg, Mark. 1994. Lexical Nature of Syntactic Ambiguity Resolution. *Psychological Review* 101:676–703.
- Meng, Michael, and Bader, Markus. 1997. Syntax and Morphology in Sentence Parsing: A new look at German subject-object ambiguities. Ms. University of Jena.
- Miyamoto, Edson. 1998. A low attachment preference in Brazilian Portuguese relative clauses. Talk presented at the Conference on Architectures and Mechanisms in Language Processing (AMLaP-98). Freiburg, Germany.
- Phillips, Colin. 1996. *Order and Structure*. Doctoral dissertation, Linguistics Department, MIT, Cambridge, MA.
- Prince, Alan, and Smolensky, Paul. 1993/2004. *Optimality Theory: Constraint Interaction in Generative Grammar*. Malden, MA: Blackwell. Technical Report, Linguistics Department, Rutgers University, New Brunswick, NJ and Department of Computer Science, University of Colorado at Boulder (1993). Rutgers Optimality Archive 537 (2002).
- Pritchett, Bradley. 1992. *Grammatical Competence and Parsing Performance*. Chicago, IL: University of Chicago Press.
- Rizzi, Luigi. 1990. *Relativized Minimality*. Linguistic Inquiry Monograph 16. Cambridge, MA: MIT Press.
- Sauerland, Uli, and Gibson, Edward. 1998. How to Predict the Relative Clause Attachment Preference. Talk presented at the 11th Annual CUNY Conference on Human Sentence Processing.
- Stevenson, Suzanne. 1994a. Competition and recency in a hybrid network model of syntactic disambiguation. *Journal of Psycholinguistic Research* 23:295–322.
- Stevenson, Suzanne. 1994b. *A Competitive Attachment Model for Resolving Syntactic Ambiguities in Natural Language Parsing*. Doctoral dissertation, Department of Computer Science, University of Maryland, College Park, MD.
- Stevenson, Suzanne. 1995. Arguments and Adjuncts: A Computational Explanation of Asymmetries in Attachment Preferences. *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, 748–753.
- Stevenson, Suzanne. 1998. Parsing as Incremental Restructuring. In *Reanalysis in Sentence Processing*, eds. Janet Fodor and Fernanda Ferreira, 327–363. Kluwer Academic.

- 
- Stevenson, Suzanne. in preparation. Agreeing to Attach: Structure Building and Feature Processing. Ms. University of Toronto.
- Sturt, Patrick, and Crocker, Matthew. 1996. Monotonic syntactic processing: A cross-linguistic study of attachment and reanalysis. *Language and Cognitive Processes* 11:449-494.
- Traxler, Matthew , Pickering, Martin, and Clifton, Charles. 1996. Architectures and mechanisms that process prepositional phrases and relative clauses. Talk presented at the Conference on Architectures and Mechanisms of Language Processing (AMLaP 1996). Torino, Italy.
- Trueswell, John, and Tanenhaus, Michael J. 1994. Toward a Lexicalist Framework for Constraint-Based Syntactic Ambiguity Resolution. In *Perspectives on Sentence Processing*, eds. Charles Clifton, Lyn Frazier and Keith Rayner, 155-179. Hillsdale, NJ: Lawrence Erlbaum.
- Weinberg, Amy. 1995. Licensing constraints and the theory of language processing. In *Japanese Sentence Processing*, eds. Reiko Mazuka and Noriko Nagai. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Woolford, Ellen. 2001. Case patterns. In *Optimality-Theoretic Syntax*, eds. Geraldine Legendre, Sten Vikner and Jane Grimshaw, 509-543. Cambridge, MA: MIT Press.
-