(40)

Center-embedding and Self-embedding in Human Language Processing

by

James Davis Thomas

LIBRARIES

Submitted to the Department of Brain and Cognitive Sciences
in partial fulfillment of the requirements for the degree of

Master of Science in Brain and Cognitive Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1995

© Massachusetts Institute of Technology 1995. All rights reserved.

Author .....
Department of Brain and Cognitive Sciences
August 7, 1995

Certified by .....
Edward Albert Fletcher Gibson
Assistant Professor
Thesis Supervisor

Accepted by .....
Gerald Schneider
Chairman, Departmental Committee on Graduate Students

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

AUG 22 1995

LIBRARIES

SCHER-PLOUGH

Center-embedding and Self-embedding in Human Language Processing

by

James Davis Thomas

Submitted to the Department of Brain and Cognitive Sciences
on August 5, 1995, in partial fulfillment of the
requirements for the degree of
Master of Science in Brain and Cognitive Sciences

Abstract

It is one of the major enterprises of psycholinguistics to account for the notorious tendency of the human sentence parser to break down trying to process certain classes of embedded structures. Recently much work has been devoted to this topic; the theory that perhaps best accounts for the data is Gibson's [Gib91] account of thematic complexity. This thesis uses Gibson's theory as a starting point. Two major series of experiments were performed, and their results motivate us to propose a revision of Gibson's theory that depends crucially on the notion of self-embedding. Additionally, the possible relevance of work on processing certain kinds of ungrammatical sentences is discussed and some unresolved questions are presented.

Thesis Supervisor: Edward Albert Fletcher Gibson
Title: Assistant Professor

Acknowledgments

This thesis would not have happened if not for my advisor Ted Gibson, who, in addition to possessing the oft-underrated virtue of being a great guy, has proved to be an exceptional mentor. He has shown me where the interesting issues in psycholinguistics are, how to do experiments and build theories: in short, how to be a scientist. I cannot thank him enough, and I sincerely hope our collaboration continues after I leave. Thanks also go out to Ken Wexler, without whose encourgagement and often mysterious belief in my abilities I would not be here in the first place. I would also like to thank my first-year cohort of Jenny Ganger, Stephen Lines, Cristina Sorrentino, Josh Tenenbaum, and Emo Todorov for generalized intellectual stimulation and moral support. It would have been an honor to graduate with them.

At Wisconsin, my undergraduate thesis advisor Malcom Forster was kind enough to see promise in my rambling, half-formed ideas and to teach me how to shape them into a systematic path of philosophical investigation. And Kyle Johnson undertook the herculean task of making modern linguistics not only comprehensible but enjoyable, succeding with admirable panache.

Finally, I'd like to thank those who helped my frantic search for subjects in the final days, notably Mark Rousculp, Emily G. Wallis, Adee Matan, Ken Nwokeji, and Hoyt Bleakley and the kids down at the Boston Fed.

This thesis is dedicated to my grandmother E. S. Thomas, who first taught me to love language.

Contents

1 Introduction **8**

1.1 A Quick Note on the Performance/Competence Distinction 9

1.2 Chomsky & Miller 9

1.3 The Magic Number Two 10

1.3.1 Kimball and his Two Sentences 10

1.3.2 Lewis and Structural Position 12

1.3.3 Stabler and Case 12

1.3.4 The State of the Art for the Magic Number Two 12

1.4 Thematic Complexity 13

1.4.1 Pritchett 14

1.4.2 Gibson 14

1.5 How to Tell the Theories Apart 17

2 First Empirical Results **19**

2.1 Motivation 19

2.2 Methodology 21

2.3 Results 22

2.4 A Note on Unacceptability and Acceptability 24

3 Enter Self-embedding **26**

4 Second Empirical Results **32**

4.1 Purpose 32

4.2 Methodology 34

4.3 Results 35

5 The Missing Verb Effect **38**

5.1 The Alleged Phenomenon 38

5.2 The experiment 39

5.2.1 Methodology 39

5.2.2 Results 41

5.3 A Tentative Theory of NP Pruning 41

5.4 Future Experiments 42

5.4.1 First, Establish the Effect 42

5.4.2 Tie the Effect to Processing Overload 43

5.4.3 Where in the Structure Does Overload Occur? 44

5.5 Conclusion 45

List of Figures

List of Tables

2.1 Experiment 1, claim 1 23
2.2 Experiment 1, claim 2 23
2.3 Experiment 1, claim 3 24
2.4 Experiment 1, claim 4 24
2.5 Experiment 1, all conditions. 24

4.1 Experiment 2, claim 1 35
4.2 Experiment 2, claim 2 36
4.3 Experiment 2, claim 3 36
4.4 Experiment 2, claim 4 36

Chapter 1

Introduction

Why is it that a sentence with a single embedded relative clause, like:

(1) The child the dog bit developed rabies.

is easy to understand, while embedding just one more clause, as in:

(2) # The child the dog the man shot bit developed rabies

produces near incomprehensibility[1]? Similar results occur over a wide variety of constructions, and in all human languages. The search for a theory to explain these and other related facts has been one of the core questions of psycholinguistics for thirty years. This paper centers around two experiments designed to address this question. The first was designed primarily to test Gibson's [Gib91] original complexity theoretic theory; the second, to test a revision of that theory described in chapter three, crucially dependent on the notion of "self-embedding". But, as prelude, a brief review of previous work and competing theories is in order.

[1]The typographical symbol # is commonly used to mark sentences considered unacceptable, or diffucult to understand because of processing reasons; the symbol * is used to denote ungrammatical sentences.

1.1 A Quick Note on the Performance/Competence Distinction

This work lies solidly within the framework of modern generative linguistics innagurated by Chomsky [Cho65]. One of the key theoretical foundations of the approach is the performance/competence distinction. This posits that human language is split into two parts. Our knowledge of language is contained in a grammar composed of a set of generative rules. These rules are recursive, and theoretically allow us to produce and comprehend sentences of infinite length. However, there are a variety of factors, such as limited memory span, attention, outside noise, etc., that limit our actual comprehension and production. These are labelled performance factors. It is theorized that the problems in understanding sentences like (5) are caused by performance factors. Such doubly-embedded sentences are clearly grammatical in a deep way, and given sufficient time and resources can be understood fully. Therefore, the study of processing overload does not directly concern itself with the rules of grammar, rather, how that grammar is implemented and how that implementation leads to a finite-resources inspired breakdown.

There are many psychologists who would dispute these conclusions; much ink has already been spilt over it. This thesis is not the place to carry on the argument, and so the core assumptions of the Chomskyan framework are assumed without further discussion.

1.2 Chomsky & Miller

The beginning of systematic investigation of these phenomena begins with Chomsky. Chomsky & Miller's seminal work [CM63, MC63] established that problems with center-embedding were a function of performance limitations, not the grammar. They offered some suggestions as to the reasons behind processing breakdown.

First, they proposed that the ratio of nonterminal nodes to terminal nodes could serve as a metric to identify sentences too complex to process. The higher the ra-

tio, the more difficult the sentence. Unfortunately, this account predicts that left-branching and right-branching structures should be as difficult as center-embeddings.

(3) # The child the dog the man shot bit developed rabies

Chomsky & Miller also proposed a metric based on incomplete grammatical relations. In (3), "child", "dog", and "man" are all subjects, and stand in a certain grammatical relation to the corresponding verbs. Until we actually see those verbs, we cannot be sure what that relation is, so it is considered incomplete. Chomsky & Miller posited that too many incomplete relations produces overload.

Finally, Chomsky & Miller suggested that self-embedded structures were inherently hard. They noted that it was exactly the presence of self-embedding that made a grammar context-free, and not regular, thus requiring more computing power to parse[2]. In sentence (3) there is a object relative clause embedded inside another object relative clause, a clear case of self-embedded structures. Chomsky & Miller demonstrated that such self-similar structures led to an inevitable use of memory resources, potentially leading to processing difficulty.

While these suggestions were not worked out in enough detail to function as working theories of sentence processing, their basic insights underpin most contemporary accounts.

1.3 The Magic Number Two

1.3.1 Kimball and his Two Sentences

In [Kim73] Kimball established the Principle of Two Sentences:

The constituents of no more than two sentences can be parsed at one time [Kim73].

[2]Chomksy & Miller's notions of self-embedding were couched in the framework of the Chomsky hierarchy of pushdown automata. The notion of self-embedding presented in [CM63, MC63] is well defined in terms of the theory of computation; but, was not thoroughly developed in terms of the then nascent Transformational Grammar. Therefore, the use of Chomsky & Miller's original notion of self-embedding in this paper is more inspirational than technical.

This simple maxim covered a surprising amount of ground. It sucessfully predicts the distinction between singly and doubly center–embedded sentences:

(4) The child that the dog bit developed rabies.

(5) a. # The child that the dog that the man shot bit developed rabies.

b. # [$_{IP}$[$_{NP}$ The child [$_{IP}$ that [$_{NP}$ the dog [$_{IP}$ that [$_{NP}$ the man...]]]]]]

Note that at the point of "the man" in (5b), there are three incomplete IPs, and thus Kimball's account correctly identifies the sentence unacceptable.

Similarly, it accounts for the distinction between between single and double sentential subjects:

(6) a. That the child died bothered Suzie.

b. For John to date a teenager would bother Suzie.

(7) # That for John to date a teenager bothers Suzie surprised Cristina.

(8) # [$_{IP}$ That [$_{IP}$ for [$_{IP}$[$_{NP}$ John...]]]]

Here again, at "John" there are are three incomplete IPs, and Kimball again makes the right prediction.

However, there are some sentences with three incomplete IPs that are generally acceptable. One such acceptable example is a relative clause embedded inside of a sentential subject, as noted by Gibson [Gib91] and Cowper [Cow76].

(9) a. That the teenager that John dated was pretty annoyed Suzie.

b. [$_{IP}$ That [$_{IP}$[$_{NP}$ the teenager [$_{IP}$ that [$_{NP}$ John...]]]]]

At the point of "John" in (10b), there are three IPs, and so Kimball's account predicts this sentence to be unacceptable. But, empirically, this class of sentences seems to be easy to process.

Despite these problems, the parsimony of Kimball's theory has proved attractive to many. There are two current extensions of Kimball's basic account that attempt answer many of its problems while retaining its simplicity. Since they both make

similar predictions, they will be briefly introduced individually, and then discussed together in light of empirical results.

1.3.2 Lewis and Structural Position

Lewis' account [Lew95, Lew93] depends on structural position. It posits that the processer overloads if it has to store more than two heads occupying the same structural position before they can be attached, where structural position is the position of the head in X-bar theory, such as Spec IP, Topic, etc. The magic number two is here, but Lewis' account crucially differs in that there can be more than two incomplete sentences, so long as the subjects of those sentences do not all inhabit the same structural position. This allows Lewis to account for many of the facts that Kimball missed.

1.3.3 Stabler and Case

Similarly, Stabler [Sta94] has reworked Kimball's theory into an account dependent on case. He posits that processing overload occurs if the connectivity of a structure exceeds 2, where the connectivity of a structure is defined as the number of identical Cases assigned to or from it or the number of identical Chains going from it. Again, the magic number two is here, but crucially overload can be avoided in cases with more than two sentences, if the relevant subjects are not all of the same case.

1.3.4 The State of the Art for the Magic Number Two

These two theories account for much of the data that Kimball's original formulation missed. For example, they successufully predict the acceptability of relative clauses embedded inside sentential subjects:

(10)a. That the teenager that John dated was pretty annoyed Suzie.

b. [$_{IP}$ That [$_{IP}$[$_{NP}$ the teenager [$_{IP}$ that [$_{NP}$ John...]]]]]

In Lewis' theory, at the point of "John" in (10b) there are still three unattached

heads, but they occupy different structural positions: "teenager" and "John" both occupy spec-IP, but the head "that" is in topic position[3].

Similarly, in Stabler's account there are two unattached heads with nominative Case, "teenager" and "John", and the sentential subject is not assigned Case, so the sentence is predicted to be good.

In other cases, however, the two theories overpredict difficulty. In sentences consisting of embedded relative clauses embedded in a sentential complement of anNP[4], both theories incorrectly predict processing overload.

(11)a. The fact that the teenager that John dated was pretty annoyed Suzie.

b. [$_{IP}$[$_{NP}$ The fact that [$_{IP}$[$_{NP}$ the teenager [$_{IP}$ that [$_{NP}$ John...]]]]]]

For Lewis, all three of the subjects, "fact", "teenager", and "John" occupy Spec-IP position. And for Stabler, all three need nominative case. Both theories predict difficulty, yet strong intuitions, as well as empirical evidence (to be presented next chapter) show this class of sentences to be acceptable.

Before moving onto the "Thematic Complexity" approach, one note is in order. The "Magic Number Two" approaches never care about the *order* of the embeddings. They all depend crucially on incomplete relations between subject and verb, and the order in which those relations appear is not important.

1.4 Thematic Complexity

In addition to the "Magic Number Two" approach, there is a fundamentally different way to approach the question of processing overload, an approached dependent on the notion of thematic complexity. In order to introduce this topic, a brief detour into the topic of ambiguity resolution is required.

[3]Lewis' account depends on the analysis of sentential subjects that places them in topic position offered by Koster [Kos78].

[4]Sentential complement NPs are NPs that take whole sentences as complements, for example "the fact that the man dropped the ball". They differ from relative clauses, clauses like "the man who dropped the ball", in that they have no wh-operator and no trace in the clause.

1.4.1 Pritchett

Pritchett [Pri88] took a novel approach to the question of ambiguity resolution by appealing to the Θ-Criterion:

> Each argument bears one and only one Θ-role (thematic role) and each Θ-role is assigned to one and only one argument. [Cho81]

Pritchett combined this with the intuition that every part of syntax tries to be satisfied at every point during parsing, leading to a theory of ambiguity resolution that posited that given two possible structures for an ambiguous partial sentence, the structure with the least unfulfilled Θ-relations would be the preferred reading.

1.4.2 Gibson

Gibson [Gib91] took Pritchett's basic insight and refined it into a theory that handled both ambiguity resolution and processing overload[5]. Gibson's theory hypothesizes that incomplete thematic relations count as processing load. As long as these relations are incomplete, the sentence has a certain "cost" of resources, measured in PLUs[6] (processing load units). Gibson hypothesized, following Chomsky & Miller [CM63], that there is a limit to the cost associated with a structure; the limited was empirically determined to be four PLUs. Structures carrying five or more PLUs overload the limited resources of the parser and produce unacceptability.

The Theoretical Formalities

Gibson formalizes three crucial thematic relations whose lack of fulfillment "costs" PLUs. The formal definitions are presentend here, along with intuitive explanations and an example at the end.

[5] Gibson's theory uses the same constraints described here, in a framework of ranked parallelism. The details are irrelevant to the processing data, and will be ignored. See [Gib91] for the full story.

[6] In Gibson's original theory [Gib91] different contraints had different costs to violate. For the purpose of this paper, such extra complexity is unneeded and will be ignored.

The Property of Thematic Reception(TR):
Associate a PLU for the head of each chain that is in a position that can be associated with a thematic role, but which does not yet receive a thematic role.

This constraint means that stray NPs need Θ-roles. Each noun not assigned a Θ-role counts as processing load until its Θ-assigner is present.

The Property of Lexical Requirement(LR):
Associate a PLU to each lexical requirement position that is obligatory in the current structure, but is unsatisfied.

This constraint means that subcategorizers (usually verbs, but sometimes prepositions and complementizers) need complements. Processing load accumulates until this need is fulfilled.

The Property of Thematic Transmission(TT):
Associate a PLU to each semantically null C-node category in a position that can receive a thematic role, but whose lexical requirement is currently unsatisfied.

Primarily, this means that semantically null elements in a position to receive Θ-roles must pass them on to another category. Until they do, it counts as a violation.

A Helpful Example

Here is Gibson's theory applied to a relative clause embedded in a sentential complement, the same example as (12a) above. The "e" in (12b) represents the null form of the complementizer "that"[7].

(12)a. The fact that the teenager who John dated was pretty annoyed Suzie.

b. [$_{IP}$[$_{NP}$ The fact$_{TR}$ that [$_{IP}$[$_{NP}$ the teenager$_{TR}$[$_{IP}$ who$_{TR}$ e$_{LR}$[$_{IP}$[$_{NP}$ John dated...]]]]]]]

[7]Null complementizers are uncontroversial theoretical constructs of modern linguistics. Compare sentences like "I know that the man left" and "I know the man left". Both sentences mean the same thing, and are structurally identical. In the second sentence, the complementizer is still there, but simply not phonetically expressed.

Each relevant violation is marked here with the corresponding constraint. Note that some violations are not marked; they occur and are fully resolved before the point of maximum complexity, and so are not relevant to the question of whether or not Gibson's theory predicts the sentence to be acceptable.

In this sentence, there are at maximum four violations. The NPs "fact" and "teenager" need Θ-roles. Similarly, the relative pronoun "who" also needs a Θ-role. These are all violations of the Property of Thematic Reception. And finally, the null complementizer subcategorizes for an IP, violating Lexical Requirement. At the point of the null complementizer, there are four violations, still an accepbtable sentence. At the point of "John", the Lexical Requirement of the null complementizer is satisfied, and as "John" needs a Θ-role an additional violation of Thematic Reception occurs, keeping the total violations at four. The next word, "dated" provides "John" with its Θ-role, and the running total of violations drops to three. So, like the "Two Sentences" approaches, Gibson predicts this sentence to be acceptable.

Here is another example, an occurrence of the reverse embedding: sentential complement inside of a relative clause.

(13)a. The teenager who the fact that John dated Suzie annoyed was pretty.

b. # $[_{IP}[_{NP}$ The teenager$_{TR}$ who$_{TR}$ $[_{IP}[_{NP}$ the fact$_{TR}$ that$_{\#LR+TT\#}$ John...]]]]

Here, similar to (12a), the NPs "teenager" and "fact", and the relativizer "who" all need Θ-roles, violating Thematic Requirement. The "that" also subcategorizes for an IP, a temporary violation of Lexical Requirement. The fifth violation comes from Thematic Transmission. Here, "the fact" needs to assign a Θ-role to a category containing a thematic element, the embedded IP. However, the "that" immediately following cannot receive a Θ-role, and so must transmit it to the following IP, briefly incurring a violation of Thematic Transmission. At this point of the sentence there are five violations, and so the sentence is ruled unacceptable. Note that the a similar violation of Thematic Transmission occurs in sentence (12a) above, but it is resolved with the next word "the teenager", and the total violations at that point of the

sentence is less than four, so the sentence is still predicted to be acceptable.

1.5 How to Tell the Theories Apart

In much of science, theory building is a tradeoff: in order to account for the data better, you often pay a little additional theoretical complexity. So to with the phenomena noted here: as Lewis [Lew95] notes, Gibson's theory gets the most data, but is also the most complex. Although it should be noted that Gibson's theory has the virtue of not only getting the most of the processing overload data, but also serving as a very viable account of ambiguity resolution[8]. No other theory incorporates both domains.

While there are frameworks for resolving such trade-offs in between data and parsimony in some narrow domains (see [Ris89] for an explanation of Minimum Description Length techniques), in general science such things are a matter of personal taste and convention.

One traditional way to evaluate theories is to look at them in relation to other theoretical constructs. If a theory of processing overload uses ideas easily derivable from well-established theories of parsing or competence grammar, the theory becomes more plausible.

The "Magic Number Two" approaches have the virtue of being easy to motivate in terms of the parser – the parser is keeping track of certain kinds of unfulfilled grammatical relations, and doesn't have room to store more than two of any one kind of relation. And the relations in question are straightforward, well-motivated concepts established in previous theories – Case for Stabler, and structural position for Lewis.

Gibson's theory differs in that all the relevant incomplete relations are lumped together, and the processor can only store four of them in total. Two of these kinds

[8]Actually, Lewis' thesis [Lew93] is an attempt to build general theory of sentence processing that comprehensively covers many different phenomena including overload, ambiguity, and more; however, while Gibson's theory uses the same theoretical constructs in his accounts of both ambiguity resolution and processing overload, Lewis' general theory is composed of a number of subtheories which are somewhat independent of each other.

of incomplete relations, Lexical Requirement and Thematic Reception, are straighforward extensions of Θ-theory. The third, Thematic Transmission, is not, and this has been cause for concern in the past. The modifications proposed later in this paper will replace Thematic Transmission with a new property, dependent on the notion of self-embedding, that not only gets more data, but is more plausibly motivated in terms of the parser.

Chapter 2

First Empirical Results

One of the greatest problems with current work in processing overload is a lack of good empirical data. Despite the work of a few researchers, [BBMW86, KJ91] many researchers are drawn from non-experimental fields such as linguistics [Sta94] or computer science [Lew93], and as such have neither the interest nor the expertise to run carefully controlled experiments. The result has been a variety of differing intuitions used to argue for and against theories.

To alleviate this problem, a simple experiment was carried out, designed to test predictions of Gibson's [Gib91] thematic-complexity theory, primarily in relation to theories of the "magic number two" type.

2.1 Motivation

Specifically, the following four claims were tested:

Claim 1: That embedding a relative clause inside a sentential complement is easier than the opposite embedding:

(14)a. The hunch that the serial killer who the waitress had trusted might hide the body frightened the FBI agent into action.

b. The FBI agent who the hunch that the serial killer might hide the body had frightened into action had trusted the waitress.

The "Magic Number Two" approaches predict both of these sentences to be equally unacceptable. In both cases, there are three nouns that all have nominative case, and fill the Spec-IP position. In contrast, as explained in examples (12a) and (13a), while Gibson's account predicts (14b) to be bad, (14a) should be acceptable.

Claim 2: That embedding a relative clause inside a sentential subject is easier than the opposite embedding:

(15)a. Whether the serial killer who the waitress had trusted might hide the body frightened the FBI agent into action.

b. The FBI agent who whether the serial killer might hide the body had frightened into action had trusted the waitress.

To Gibson's theory, this is virtually identical to the sentential complement case above. (15a) is predicted to be good, and (15b) bad. The "Magic Number Two" theories treat it as somewhat of a special case. It was explained earlier how they predict (15a) as acceptable by positing that the initial sentential subject is in topic position, but to explain the unacceptability of (15b) they claim that embedded sentential subjects are just plain ungrammatical. (See [Lew95] for more).

Claim 3: That processing a doubly center-embedded relative clause structure whose most embedded relative clause is a subject relative clause is easier than processing one whose most embedded relative clause is an object relative clause:

(16)a. The serial killer who the FBI agent who the waitress trusted had frightened into action hid the body.

b. The serial killer who the FBI agent who trusted the waitress had frightened into action hid the body.

Gibson's theory predicts no difference between these two examples. In both "serial killer" and "FBI agent" violate Thematic Reception, as do the two relative pronouns "who". Finally, the null complementizer after the second "who" violates Lexical

Requirement. This totals to five violations in both sentences, ruling them both unacceptable. The "Magic Number Two" accounts rule clearly rule out (16a), with three unattached nominative case nouns in Spec-IP. In (16b) most embedded clause is a subject relative, the third noun in the sentence doesn't occur until after the verb "trusted". So, we never have more than two unattached NPs at one time, and the "Two Sentences" approaches predict the subject relative cases to be good.

Claim 4: That processing a doubly center-embedded relative clause structure in object position is easier than processing one in subject position:

(17)a. The serial killer who the FBI agent who the waitress trusted had frightened into action hid the body.

b. The body was hidden by the serial killer who the FBI agent who the waitress trusted had frightened into action.

Gibson predicts a contrast here, with (17a) bad and (17b) good. Note that (17a) is the same as (16a) in the previous claim, and the same analysis applies. However, for (17b), "serial killer" gets a Θ-role from the verb "hidden" through the preposition "by". When we reach the "who" aftter "FBI agent", the sentence is only carrying four PLUs: three from the property of Thematic Reception ("FBI agent" and the two "who"s) and one from Lexical Requirement (from the null complementizer immediately following the second "who"), and is predicted to be acceptable. Similarly, in the "Magic Number Two" accounts, (17a) is bad, but since "serial killler" in (17b) is the in object position, it neither resides in Spec IP nor receives nominative case, leading the "Magic Number Two" accounts to predict (17b) as acceptable.

2.2 Methodology

Subjects:

Forty-three native English speakers from MIT (primarily undergraduate students) participated. Subjects were paid $6.00 each.

Materials:

Thirty-five items with seven conditions (identical to those above–note that (16a) and (17a) are the same condition) were constructed. As illustrated by the examples above, the same NPs and verbs were used within each item across all conditions, and attempts were made to preserve, as much as possible, thematic relations across conditions. Given the diverse structure of the sentences, and the fact that some conditions had sentential complement NP-complement nouns that others lacked, this was not completely possible.

The thirty-five items were then combined with eighty-five filler items, roughly similar in length and complexity, to form seven lists. The experimental items were counterbalanced across each list such that each list contained, for each condition, five items, and items were never repeated within each list.

Procedure:

The lists were given as questionaires, in which subjects rated each sentence on a scale from 1 (the best) to 5 (the worst) according to how hard the sentences were on a quick first reading. The sentences were presented ten to a page in pseudo-random order for each list, and the pages of each individual list were randomized.

2.3 Results

The results were, as always, a mix of good and bad news.

Claim 1: That embedding a relative clause inside a sentential complement is easier than the opposite embedding.

The results in table 2.1 are straightforward. The relative clause inside sentential complement embedding is clearly easier than the reverse case. This counts as confirmation for Gibson's account and against the "Magic Number Two" models.

Claim 2: That embedding a relative clause inside a sentential subject is easier than the opposite embedding.

condition	mean unacceptabiliby	
rel-clause inside sent-complement	1.92	$p < .001$
sent-complement inside rel-clause	3.38	

Table 2.1: Experiment 1, claim 1

condition	mean unacceptabiliby	
rel-clause inside sent-subject	2.24	$p < .001$
sent-subject inside rel-clause	3.98	

Table 2.2: Experiment 1, claim 2

The results table 2.2 are also straightforward; the relative clause inside sentential subject is clearly better than the reverse. This follows from the predictions of both classes of theories.

Claim 3: That processing a doubly center-embedded relative clause structure whose most embedded relative clause is a subject relative clause is easier than processing one whose most embedded relative clause is an object relative clause:

The results in table (2.3) are a little surprising. Extensive previous empirical work on singly embedded relative clauses (see [KJ91] for just one example) has found a significant difference in reading times between the subject relatives and object relatives, with subject relatives clearly preferable. And the difference present in this experiment is close to significant; this is highly suggestive that there is some difference. However, it seems that the large gap in acceptability between the sentences in the first two claims and the small, potential difference between the examples here differ qualitatively. If there is a clean break between acceptability and unacceptability, the evidence here indicates that both doubly embedded subject relatives and doubly embedded object relatives are on the unacceptable side. So, the evidence again supports Gibson's account over the "Magic Number Two" theories.

Claim 4: That processing a doubly center-embedded relative clause structure in object position is easier than processing one in subject position.

The results in 2.4 are surprising. Both classes of theories had predicted that there would be a significant difference here, and none showed up. There was a small nu-

condition	mean unacceptabiliby	
most embedded clause subj-rel	3.17	$p < .066$
most embedded clause obj-rel	3.40	

Table 2.3: Experiment 1, claim 3

condition	mean unacceptabiliby	
double-embedded rel as obj	3.24	$p < .167$
double-embedded rel as subj	3.40	

Table 2.4: Experiment 1, claim 4

merical difference, but it was not statistically significant, and in the wrong direction. This sticks out as the only piece of data not predicted by the original Gibson account.

2.4 A Note on Unacceptability and Acceptability

What constitutes the boundary between "unacceptable" and "acceptable" sentences is unclear. It is easy to show that some sentences are more acceptable than others, but it is unclear we can draw a "line in the sand" between good and bad sentences.

In this experiment table 2.5 shows that all the mean acceptability of all seven classes of sentences conveniently patterned into two classes. The only possible exception was the sentential subject inside of relative clause embedding; this was worse than any of the others. As discussed before, many syntacticians think that this structure is fundamentally ungrammatical – this might explain the difference.

In this experiment, there does seem to be a pretty clear distinction–all the accept-

condition	mean unacceptability
rel-clause inside sent-complement	1.92
rel-clause inside sent-subject	2.24
sent-complement inside rel-clause	3.38
sent-subject inside rel-clause	3.98
most embedded clause subj-rel	3.17
most embedded clause obj-rel	3.40
double-embedded rel as obj	3.24

Table 2.5: Experiment 1, all conditions.

able sentences fall around 2, and the unacceptable ones around 3-3.4. It seems that this large gap does mark the boundary between acceptability and unacceptability, confirming strong introspective intuitions.

Chapter 3

Enter Self-embedding

The inability of Gibson's theory to account for the lack of difference in acceptability between doubly-center embedded object relatives in subject position versus those on object position is a problem. There was no easy way to adjust the theory to get the right empirical predictions.

Consider also the following sentences:

(18)a. The invasion plan which had exposed the aliens was analyzed by the conspiracy buff.

b. [$_{IP}$[$_{NP}$ The invasion plan$_{TR}$[$_{IP}$ which$_{TR}$ e$_{LR}$ had exposed the aliens...]]]

(19)a. The invasion plan which the aliens had exposed was analyzed by the conspiracy buff.

b. [$_{IP}$[$_{NP}$ The invasion plan$_{TR}$[$_{IP}$ which$_{TR}$ e$_{LR}$[$_{NP}$ the aliens had exposed...]]

Sentence (18a) contains a singly-embedded subject relative clause, while (19a) contains a singly-embedded object relative. As mentioned before in section 2.3, previous empirical work (e.g. [KJ91]), has shown reading times for object relatives to be significantly longer than for subject relatives. Here, since both sentences are only singly-embedded, they are both clearly acceptable, but they do differ.

Gibson's original theory, however, predicts no difference in processing difficulty between the two sentences. In both sentences, the point of maximal complexity is reached right after the null complementizer in the embedded clause. There are

three violations; one Thematic Reception violation for the matrix subject, one for the relative pronoun "which", and a violation of Lexical Requirement for the null complementizer. While Gibson's original theory predicts both to be acceptable, and both clearly are, it would be nice if we could give a more fine-grained account.

Almost identical concerns apply to the distinction between doubly center-embedded subject relative clauses and doubly center-embedded object relative clauses. The previous experiment seemed to show that there was some difference between the two; while it was not nearly as large as the distinctions between sentences on either side of the sharp distinction between unacceptable and acceptable, there might be something there. We would ideally like a theory that can account for these lesser differences. Gibson's original theory, while able in principle to make such fine distinctions (via the measure of maxmimum load of PLUs), predicts no difference.

Theory-internally, the Property of Theta Transmission has always seemed awkward. The other two properties, Thematic Reception and Lexical Requirement, are both straightforward extensions of the Θ-criterion[1]. But the Property of Theta Transmission had never been so solidly grounded, and always seemed stipulative.

This combination of empirical and theoretical concerns led to a careful re-examination of Gibson's original theory. The re-examination drew inspiration on an idea found in [CM63, MC63], spefically the concept of "self-embedding". They proposed that embedding an incomplete structure inside another incomplete structure of the same type causes processing difficulty, but this suggestion was analyzed in terms of the theory of computation and not developed into a functioning theory of sentence processing. The initial reaction is to think of self-embedding in terms of identical linguistic structures. Evidence from introspection suggests that self-embedding alone could not be the whole story. For example consider (25), the same example as (14b):

(20) # The FBI agent who the hunch that the serial killer might hide the body had frightened into action had trusted the waitress

[1]While the theories presented here work solidly within the Government and Binding framework establsihed by Chomsky [Cho81], it should be noted that all current linguistic theories have some correlate to the Θ-criterion, and the theory presented here could be adapted to work with any of them.

Embedding a sentential complement inside a relative clause produces processing overload, as shown in table (2.1). The two clause types are not *identical.* Thus, this type of self-embedding cannot be the sole cause of processing overload.

(21) The information that the evidence that the aliens were kidnapping the children would expose the invasion plan was analyzed by the conspiracy buff.

In contrast, in (21), there are identical embedded clauses: two sentential complements. Introspective judgements (and empirical evidence to be presented next chapter) suggest that this sentence is acceptable.

Clearly this strict notion of self-embedding tied to linguistically identical structures is inadequate to be the sole explanation of center-embedded processing overload. What if this particular definition of self-embedding is too strict? What if self-embedding were expressed not in terms of "identical" constructions, but constructions whose lists of features bore some sort of subset relation to each other? This intuition led to the following revision to the original Gibson theory.

The Property of Self-Embedding(SE) interference:
Associate a PLU for each predicted feature structure X_1 whose head has not yet appeared which is embedded inside another predicted feature structure X_1, whose head has also not yet appeared when the extened projection features of X_1 (the inner predicted category) are a subset of those of X_2 (the outer predicted category).

Exactly what the features involved are is crucial. Some of the features await further linguistic investigation, but two are relevant here. Since certain types of adjunction are allowed in matrix clauses that are not allowed in subordinate clauses, it is assumed that the feature sets of matrix clauses and relative clauses are distinct, ruling out a subset relationship[2]. So, a PLU from Self-embedding is never accrued

[2]Further evidence for this comes from the fact that some lanaguages (e.g. Korean) have differing morphology for matrix and relative clauses.

from embedding inside of a matrix clause. Also, relative clauses have operators, and sentential complements do not. Thus, the feature sets of sentential complements are assumed to be proper subsets of the feature sets of relative clauses. Thus, a Self-embedding PLU is accrued by embedding a sentential complement inside of a relative clause, but not vice versa.

This gives us a theory that depends on self-embedding, but differs from the original Chomsky & Miller account in two ways. It is not solely dependent on self-embedding; it is only one factor in a broader, thematic-complexity based approach. Also, exactly what "self-embedding" means is different – instead of the definition of "the same type of clause" used by Miller & Chomsky, this definition defines self-embedding in terms of subsets of features.

One additional change necessary to the theory at this point is to change the Property of Lexical Requirement so that it only applies to thematic words (nouns, verbs, adjectives, adverbs), as opposed to function words (determiners, complementizers, inflection).

(22)a. The invasion plan which had exposed the aliens was analyzed by the conspiracy buff.

b. $[_{IP}[_{NP}$ The invasion plan$_{TR}[_{IP}$ which$_{TR}$ e had exposed the aliens...]]]

(23)a. The invasion plan which the aliens had exposed was analyzed by the conspiracy buff.

b. $[_{IP}[_{NP}$ The invasion plan$_{TR}[_{IP}$ which$_{TR}$ e $[_{NP}$ the aliens$_{TR}$ had exposed...]]

We can now explain the perceived difference between sentence (23a) and (22a) in terms of maximum load. In (22a), the maximum load is two violations. It would be three, if the null complementizer violated Lexical Requirement, as in the old version of the theory. In (23a), again the null complentizer, as a function word, does not violate Lexical Requirement. But, the NP "the aliens" violates Thematic Reception until the following VP is reached. Thus, the maximal processing load differs between the two sentences, the new theory predicts mild processing differences. This matches the empirical data.

Similar analyses apply to the doubly embedded cases; the object relative case has a maximum load of five, and the subject relative case six. Since these are both over the sharp acceptable/unacceptable distinction at four violations, the sentences are both predicted to be unacceptable, but the object relative case to be slightly worse.

(24)a. # The body was hidden by the serial killer who the FBI agent who the waitress trusted had frightened into action.

b. # $[_{IP}[_{NP}$ The body was hidden $[_{PP}$ by $[_{NP}$ the serial killer $[_{IP}$ who$_{TR}$ e $[_{NP}$ the FBI agent$_{TR}[_{IP}$ who$_{TR+SE}$ e $[_{NP}$ the waitress$_{\#TR\#}$...]]]]]]

Here, when the embedded relative clause "who the waitress trusted" is processed, there is a violation of self-embedding; the relative clause is embedded inside of another clause with an identical feature set. This PLU from Self-embedding, together with the Θ-roles required by the nouns "FBI agent" and "waitress" and the two relative prounous, adds up to five PLUs. The sentence is now predicted to be unacceptable, as indicated in the previous experiment.

Additionally, with the definition of self-embedding used above, we no longer need the Property of Thematic Transmission.

(25)a. # The FBI agent who the hunch that the serial killer might hide the body had frightened into action had trusted the waitress.

b. # $[_{IP}[_{NP}$ The FBI agent$_{TR}$ who$_{TR}[_{IP}[_{NP}$ the hunch$_{TR}$ that$_{\#LR+SE\#}$ the...]]]]

Since the feature set of the embedded sentential complement is a subset of the relative clause in which it is embedded, a Self-embedding violation occurs, pushing the sentence over the four PLU limit. Every occurrence of a violation of Thematic Transmission (in the old theory) accrues a Self-embedding violation in the new theory. However, the Principle of Self-embedding Interference is considerably easier to motivate for outside reasons. Such a constraint could arise naturally out of the implementation of a parser; depending on how the parser represents incomplete clauses, it is plausible to posit that it is harder to keep track of two incomplete clauses that have similar features (are represented in similar ways) than two dissimilar clauses.

This would lead to an increase use of resources for self-similar embedded clauses, as described in the theory here.

In summary, this relatively simple revision preserves the assets of the old theory, while offering slightly better coverage of the empirical data as well as providing better outside motivation for its theory-internal machinery.

Chapter 4

Second Empirical Results

The new, self-embedding dependent theory makes predictions in relatively virgin territory for psycholinguistics. Thus, a new experiment was designed to test it.

4.1 Purpose

The following claims were tested[1]:

Claim 1: That embedding a sentential complement inside a sentential complement is acceptable:

(26)a. The information that the evidence that the aliens were kidnapping the children would expose the invasion plan was analyzed by the conspiracy buff.

b. The invasion plan which the evidence that the aliens were kidnapping the children had exposed was analyzed by the conspiracy buff.

So far, while strong evidence has been presented why the traditional notion of self-embedding (depending on identical constructions) is unable to account for many cases of center-embedded overload, nothing has shown that the new notion of self-embedding (depending on subsets of features) cannot fully account for the data.

[1]There were ten conditions tested altogether in the experiment, but some of them do not bear directly on the issues here and will not be discussed.

Sentence (26a) is a clear cut case of self-embedding; however, the new theory presented here still predicts it to be acceptable. Thus the acceptability of this sentence is crucial to justify the added complexity that a thematic complexity based theory requires. The sentence (28a) is present as a comparison; the last experiment showed it to be clearly unacceptable, and so a strong contrast between the two is predicted.

Claim 2: That a subject relative embedded inside of an object relative should be easier in object position than in subject position:

(27)a. The invasion plan which the aliens who were kidnapping the children had exposed was analyzed by the conspiracy buff.

b. The conspiracy buff analyzed the invasion plan which the aliens who were kidnapping the children had exposed.

Sentences like (27a) were discussed in detail in chapter two; the distinction between doubly center-embedded subject relatives and doubly center-embedded object relatives provided some of the motivation for revising Gibson's original theory. As explained before, (27a) has five PLUs at the point of maximum complexity; moving the embedded clauses to object position, as in (27b) reduces that load to four, since the Θ-role for the initial noun "invasion plan" is provided by the matrix verb. Thus the sentence is predicted to not overload the parser.

Claim 3: That a sentential complement embedded in a relative clause in object position should be easier than one in subject position.

(28)a. The invasion plan which the evidence that the aliens were kidnapping the children had exposed was analyzed by the conspiracy buff.

b. The conspiracy buff analyzed the invasion plan which the evidence that the aliens were kidnapping the children had exposed.

The predicted contrast here is similar to the previous claim. Sentence (28a), as discussed in chapter two, has five Θ-violations at the point of maximum complexity. But, in (28b), the embedded clauses are moved to object position, and the initial

nouns "invasion plan" gets a Θ role from the matrix verb, lessening the maximum load to four, leading to predicted acceptability.

Claim 4: That a VP gerund embedded in a subject relative is easier than a doubly center-embedded subject relative:

(29)a. The invasion plan which the aliens kidnapping the children had exposed was analyzed by the conspiracy buff.

b. The invasion plan which the aliens who were kidnapping the children had exposed was analyzed by the conspiracy buff.

This is a clear cut case of center-embedding that is not self-embedding. It is included here as a control for pragmatic and semantic effects; if the difficulty of center-embedding is caused by some sort of semantic or discourse phenomenon, then these two sentences should be equally difficult. But, theories based on syntactic factors all predict (29a) to be better than (29b).

4.2 Methodology

Subjects:

Forty-six native English speakers (primarily undergraduate students at MIT) participated. Subjects were paid $5.00 each.

Materials:

Forty items with ten conditions were constructed. As in the previous experiment, the same NPs and verbs were used within each item across all conditions, and attempts were made to preserve, as much as possible, thematic relations across conditions. As in the previous experiments, given the diverse structure of the sentences, and the fact that some conditions had NP-complement nouns that other conditions totally lacked, this was not completely possible.

condition	mean unacceptabiliby	
sent-comp inside sent-comp, subj-position	2.19	$p < .001$
sent-comp inside rel-clause, subj-postiion	2.95	

Table 4.1: Experiment 2, claim 1

The forty items were then combined with eighty filler items, roughly similar in length and complexity. These were combined to form ten lists, with the experimental items counterbalanced across lists such that each list contained, for each condition, four items, and items were never repeated within lists.

Procedure:

The lists were given as questionaires, in which subjects rated each sentence on a scale from 1 (the best) to 5 (the worst) according to how hard the sentences were on a quick first reading. The sentences were presented ten to a page in pseudo-random order for each list, and the pages of each individual list were randomized.

4.3 Results

Claim 1: That embedding a sentential complement inside a sentential complement is acceptable:

This data in 4.1 is straightforward; in fact, the doubly embedded sentential complement condition was the most acceptable of all the conditions.

Claim 2: That a subject relative embedded inside of an object relative should be easier in object position than in subject position:

The data in table 4.2 unfortunately do not support the predictions of the revised theory. There is no statistically significant difference between subject and object position, and what difference there is in the opposite direction from the theoretical prediction.

Claim 3: That a sentential complement embedded in a relative clause in object position should be easier than one in subject position.

condition	mean unacceptabiliby	
subj-rel inside obj-rel, subj position	2.55	$p > .2$
subj-rel inside obj-rel, obj position	2.69	

Table 4.2: Experiment 2, claim 2

condition	mean unacceptabiliby
sent-comp inside obj-rel, subj position	2.95
sent-comp inside obj-rel, obj position	2.94

Table 4.3: Experiment 2, claim 3

As in the previous experiment, table 4.3 shows no difference between subject and object position.

Claim 4: That a VP gerund embedded in a subject relative is easier than a doubly center-embedded subject relative:

The data in table 4.4 is straightforward; the VP gerund cases are clearly and significantly better.

Overall, the results are mixed. It has been clearly shown that a theory dependent on self-embedding alone cannot account for the data. However, in two cases, the theory predicts an asymmetry between embedded clauses in object versus subject position.

In the both of the experiments discussed here, the thematic complexity based theory has repeatedly predicted that various constructions unacceptable in subject position would improve when moved to object position. And, in every experimental condition, there was no statistically significant difference in acceptability between subject and object conditions. There might be a difference, but it is pretty clear that none of the object position examples are on the "acceptable" side of the border.

These facts are a problem for the theory in its current form. The problem is

condition	mean unacceptabiliby	
VP gerund inside obj-rel, subj position	2.19	$p < .014$
subj-rel inside obj-rel, subj position	2.55	

Table 4.4: Experiment 2, claim 4

systematic: difficult constructions do not appear to improve when moved from subject to object position. It is clear that the theory needs to be modified, and in what direction. However, it is beyond the scope of this thesis to propose such modications, and the discussion of the thematic-complexity based theory presented here and the question of its empirical verification will have to rest in this theoretical no-man's-land for now.

Chapter 5

The Missing Verb Effect

Most work done so far in processing overload has concentrated on building a theory of which classes of sentences produce overload, and have not made a systematic attempt to account for the location of processing breakdown in particular structures. Given the young state of the field, this is understandable, but ideally one would like an account of what is actually happening during overload. This chapter offers some speculative suggestions about how to use a phenomenon first noticed by Frazier [Fra85] and discussed by Gibson [Gib91] to explore the mechanism of processing breakdown. So far, the effect has only been explored and in the pilot experiment described here, unconclusive but promising enough to merit further experiments.

5.1 The Alleged Phenomenon

Consider the following sentences:

(30)a. The novel that the horror author had written in a burst of energy was banned by the local library.

b. * The novel that the horror author was banned by the local library.

c. # The novel that the horror author who the publishing company recently had fired had written in a burst of energy was banned by the local library.

d. * The novel that the horror author who the publishing company recently had fired was banned by the local library.

The sentence in (30a) contains a simple, singly-embedded object relative. In (30b), the inner verb is dropped, and unsurprisingly the sentence becomes ungrammatical. Similarly, (30c) is a doubly center-embedded object relative, shown to cause processing difficulty in both experiments. In (30d), the center verb is dropped. Crucially, in some interesting sense (30d) does not turn to garbage. This intuition was noted by Lyn Frazier, in [Fra85]. Further judgements by a variety of native speakers seemed to confirm the intuition, as well as show that the effect was most prominent when the second verb is removed. Experimentally, Dickey [Dic95] has performed reading time experiments that show a speedup in reading times when an ungrammatical resumptive pronoun is inserted in the second of three noun gaps. While not identical to the phenomenon here, it is certainly suggestive of a similar effect. An experiment was performed in an attempt to empirically verify these intuitions.

5.2 The experiment

5.2.1 Methodology

Subjects

Forty-six native English speakers (primarily undergraduate students at MIT) participated, and were paid $5.00 each.

Materials

Twelve items with five conditions were constructed. These twelve items were contained with one-hundred and eight filler items[1], roughly similar in length and complexity. These were combined to form five lists, with the experimental items counterbalanced across lists such that each list contained, for each condition, three items, and items were never repeated within lists.

Procedure

The lists were given as questionaires, in which subjects rated each sentence on a scale from 1 (the best) to 5 (the worst) according to how hard the sentences were on a quick first reading. The sentences were presented ten to a page in pseudo-random order for each list, and the pages of each individual list were randomized.

We tested five types of sentences:

(31)a. # The novel that the horror author who the publishing company recently had fired had written in a burst of energy was banned by the local library.

b. * The novel that the crazed horror author who the publishing company wrote in a burst of energy was banned by the local library.

c. * The novel that the crazed horror author who the publishing company recently fired was banned by the local library.

d. * The novel that the crazed horror author who the publishing company recently fired wrote in a burst of energy.

e. The publishing company fired the crazed horror author who wrote the novel that was banned by the local library.

Sentence (31a) is a doubly center-embedded object relative, discussed at length in the previous two chapters. It has repeatedly been shown to be unacceptable. Sentences (31b), (31c), and (31d) are the same sentence, after the removal of the the

[1]Actually, the lists from the previous experiment and the list from this experiement were combined. Since the sentences for the two experiments were roughly the same length and complexity, and only one conditioned overlapped, it was felt that they could serve as "filler" for each other.

condition	mean unacceptability
doubly center-embedded base	2.63
first verb missing	3.54
second verb missing	2.89
third verb missing	3.30
right-branching control	1.66

first, second, and third verbs respectively. Sentence (31e) is a simple right-branching sentence constructed of the same phrases and theta-relations as the other examples. This sentence was inserted as a clearly grammatical, easy to parse control.

5.2.2 Results

The results in table 5.2.1 are inconclusive, but suggestive. First of all, it is unsurprising that the right-branching conditon (sentence (31e)) is considerably easier than any of the other conditions. The first verb missing and third verb missing conditions are the worst; they are significantly worse than the doubly center-embedded condition ($p < .004$ for both cases), and not significantly different from each other ($p < .391$).

Unexpectedly, the second verb missing condition was worse than the doubly center-embedded base condition, but the difference was not close to statistical significance ($p < .25$). And, the second verb missing condition was statistically better than the first verb missing condition ($p < .005$), and statistically better than the third verb missing condition, on a subject analysis ($p < .002$), but not on an items analysis ($p < .101$). While hardly conclusive, the fact that a clearly ungrammatical sentence was rated almost as acceptable as a grammatical doubly center-embedded object relative clause condition justifies further examination.

5.3 A Tentative Theory of NP Pruning

Given the working assumption that some version of this effect is real, that doubly center-embedded sentences process better when the middle verb is dropped, the fol-

lowing hypothesis[2] suggests itself:

The Overloaded NP Pruning Hypothesis:

At the point of overload while processing a sentence, prune the second unattached NP from the current representation and continue.

This would explain why the sentences somehow get better; in some part of the parser, only two NPs are present, and the parser is only expecting two verbs.

This hypothesis posits some sort of disassociation between the parsing of the sentence and its semantic interpretation. Subjects find the missing verb sentences acceptable, but when asked to paraphrase the sentence, they realize that something was wrong with it. Clearly, the second NP does not disappear completely from the mind of the subject, but only from the representation used by the parser.

As it stands, the NP Pruning Hypothesis is necessarily vague. Why the drop the second NP? On first blush, it seems reasonable to guess that there is some sort of recency/primacy effect, similar to the well documented effect in short-term memory, that allows the parser to preserve only the first and third verbs. This tentative working hypothesis can be better investigated once a better understanding of the phenomenon has been reached.

5.4 Future Experiments

5.4.1 First, Establish the Effect

The first thing to be done is to try to get the effect to show up more clearly. There are a number of things that could have spoiled the effect in the pilot experiement. First, there were two examples of each missing verb condition in the survey, for a total of six, or fully five percent of the survey. It is very possible that the subjects noticed one of the conditions with a verb missing and either consciously or unconsciously started

[2]Gibson [Gib91] proposes a similar account of these phenomena; however, his account suggests that the least recent (i.e. the first) NPs get pruned. The data does not support this hypothesis, but it should be regarded as the starting point for the ideas presented in this chapter

to watch for similar sentences. Since the pilot has clearly established that the effect is associated with the middle verb, the next experiment will use only two or three missing verb conditions in a survey, helping to alleviate this possible problem. Also, the judgements on the missing verb cases varied tremendously between items; careful examination of specific items might explain factors that lessened the effect.

Also, as explained before, the effect depends on subjects not "thinking too much" about the meaning of the sentence. Such reflection is unavoidable in the seconds between the reading of the sentence and the decision on the rating. Hopefully, on-line reading time experiments will alleviate this difficulty; they have proved to work on what must be a similar phenomenon in recent work by Dickey [Dic95].

5.4.2 Tie the Effect to Processing Overload

Further experiments would attempt to tie the NP pruning hypothesis more strongly to processing overload. For example:

(32)a. The hunch that the serial killer who the waitress had trusted might hide the body frightened the FBI agent into action.

b. * The hunch that the serial killer who the waitress had trusted frightened the FBI agent into action.

c. # The FBI agent who the hunch that the serial killer might hide the body had frightened into action had trusted the waitress.

d. * The FBI agent who the hunch that the serial killer might hide the body had trusted the waitress.

Sentence (32a) is a object relative clause embedded in a sentential complement; (32c) is a sentential complement embedded in a relative clause. These constructions, and the strong empirical evidence that (32a) is acceptable while (32c) is not, were discussed in chapter two. Sentences (32b) and (32d) are the same sentences, with the middle verbs dropped.

If the Overloaded NP Pruning hypothesis is true, than (32d) ought to be more

acceptable than (32c), and (32b) ought to be noticably problematic, since no overloard occurs in the its grammatical correlate (32a), and thus no NP pruning is predicted to occur. However, if both (32b) and (32d) are either both clearly rated as bad , or both are rated as acceptable, then some explanation not dependent on overload must be developed. Intuitions are tricky here; empirical data is needed. A an off-line survey experiment, similar in methodology to those presented here, will be developed.

5.4.3 Where in the Structure Does Overload Occur?

Many psychologists would dispute the validity of the whole experimental paradigm used here; the gold standard seems to be on–line reading times. It is imperative that many of the empirical findings reported here be replicated with reading time studies. However, no experimental paradigm currently exists comparable to the one used to test ambiguity resolution. The missing verb phenomenon might offer a guide to how such experiments could be constructed.

To start with, an on–line experiment will be designed, consisting of doubly center-embedded relative clauses constructed so that each verb is obligatorily transitive and takes inanimate subjects, so that they are reversible and can take any of the NPs as subject or object. For example:

(33)a. # The specialist who the surgeon who the nurse was criticizing had defended talked to the administrator on Saturday afternoon.

b. # The specialist who the surgeon who the nurse was talking to had criticized had defended the administrator on Saturday afternoon.

c. # The specialist who the surgeon who the nurse was defending had talked to had criticized the administrator on Saturday afternoon.

Reading times will be compared for each of the three verb regions, allowing us to examine two hypotheses. If, when a representation of a sentence gets too expensive, it is completely thrown out, then reading times for all three verbs should show no difference. On the other hand, if the NP Pruning hypothesis is correct, we should

expect reading times to slow down at the third verb. If this turns out to be the case, then since we know how to determine where in a sentence the effects of overload manifest themselves, we can do on–line experiments on similar structures, and look for similar third-verb reading time slowdowns to establish the existence of overload.

5.5 Conclusion

Although currently speculative, if this "missing verb" phenomenon can be explored and analyzed, it could lead us from building theories than merely detail the conditions under which processing overload occurs to theories of what actually happens during processing overload. This could dramatically increase our understanding of how the human language parser works, and expecially how it interacts with the limitations of human memory. For now, however, we can only dimly glimpse a possible path to such an theory; much work remains to be done.

Bibliography

[BBMW86] E. Bach, C. Brown, and W. Marslen-Wilson. Crossed and nested dependencies in german and dutch: A psycholinguistic study. *Language and Cognitive Processes*, 1:249–262, 1986.

[Cho65] Noam Chomsky. *Aspects of the Theory of Syntax.* MIT Press, Cambridge, MA, 1965.

[Cho81] Noam Chomsky. *Lectures on Government and Binding.* Foris, Dordrecht, The Netherlands, 1981.

[CM63] Noam Chomsky and George A. Miller. Introduction to the formal analysis of natural languages. In R.D. Luce, R.R. Bush, and E. Galanter, editors, *Handbook of Mathematical Psychology*, volume 2, pages 269–321. John Wiley, New York, NY, 1963.

[Cow76] E. A. Cowper. *Constraints on Sentence Complexity: A model for Syntactic Processing.* PhD thesis, Brown University, 1976.

[Dic95] Mike Dickey. Constraints on the sentence processor and the distribution of resumptive pronouns. Technical report, University of Massachesetts-Amherst, 1995.

[Fra85] Lynn Frazier. Syntactic complexity. In D. Dowty and A. Karttunen, L. Zwicky, editors, *Natural Language Processing: Psychological, Computational and Theoretical Perspectives.* Cambridge University Press, Cambridge, U.K., 1985.

[Gib91] Edward Albert Fletcher Gibson. *A Computational Theory of Human Linguistic Processing: Memory Limitations and Processing Breakdown.* PhD thesis, Carnegie Mellon University, 1991.

[Kim73] J. Kimball. Seven principles of surface structure parsing in natural language. *Cognition*, 2:15–47, 1973.

[KJ91] J. King and M.A. Just. Individual differences in syntactic processing: the role of working memory. *Journal of Memory and Language*, 30:580–602, 1991.

[Kos78] J. Koster. Why subject sentences don't exist. In S. J. Keyser, editor, *Recent Transformational Studies in European Languages.* MIT Press, Cambridge, MA, 1978.

[Lew93] Richard Lawrence Lewis. *An Architecturally-based Theory of Human Sentence Comprehension.* PhD thesis, Carnegie Mellon University, 1993.

[Lew95] Richard L. Lewis. A theory of grammatical but unacceptable embeddings. Technical report, Princeton University, 1995.

[MC63] George A. Miller and Noam Chomsky. Finitary models of language users. In R.D. Luce, R.R. Bush, and E. Galanter, editors, *Handbook of Mathematical Psychology*, volume 2, pages 419–491. John Wiley, New York, NY, 1963.

[Pri88] Bradley L. Pritchett. Garden path phenomena and the grammatical basis of language processing. *Language*, 64:539–576, 1988.

[Ris89] Jorma Rissanen. *Stochastic Complexity in Statistical Inquiry.* World Scientific, Singapore, 1989.

[Sta94] Edward P. Stabler. The finite connectivity of linguistic structure. Technical report, UCLA, 1994.