

Recent developments in the NL-Soar garden path theory

Richard L. Lewis

May 21, 1992

CMU-CS-92-141

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

This report describes a theory of garden path phenomena that is emerging from work on NL-Soar, a computational model of language comprehension embedded within the Soar architecture. The theory is constrained by a corpus of two kinds of sentences: garden paths (GP), which reveal the limitations of human comprehension in dealing with local ambiguities, and non-garden-paths (NGP), which reveal its power in handling local ambiguities. NL-Soar is a single-path comprehender with a limited capacity to *repair* misanalysed input. A space of repair mechanisms is explored by hand simulation on the corpus. The importance of phonology and plausible search control is established, leading to a theory up to 95% accurate (76% worst case) in predicting performance on 37 distinct GP and NGP types.

This research was sponsored by the Avionics Laboratory, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U.S. Air Force, Wright-Patterson AFB, Ohio 45433-6543 under Contract F33615-90-C-1465, ARPA Order No. 7597.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. government.

Keywords: artificial intelligence, cognitive simulation, natural language processing, parsing and understanding, Soar

This report describes a theory of garden path phenomena that is emerging from work on building a computational model of language comprehension. The model, NL-Soar (Lehman et al., 1991b; Lehman et al., 1991a), is embedded within Soar, a theory of the human cognitive architecture (Newell, 1990; Laird et al., 1987; Lewis et al., 1990). NL-Soar is thus intended to develop into a psycholinguistic theory of comprehension. The garden path theory reported on here is a step in that direction.

1 The phenomena

A striking feature of adult comprehension is its real-time ability to comprehend language with local ambiguities. A sentence may have a number of *choice points*, where each choice point represents a set of possible grammatical interpretations of the sentence up to that point. These choices occur even after the application of all knowledge sources (apart from the knowledge contained in the rest of the utterance). A *path* through a sentence corresponds to a series of choices made at each choice point. Successful comprehension of a sentence requires the computation of one (or more) complete grammatical paths.

Garden path phenomena can arise when a reader or listener attempts to comprehend a grammatical sentence and takes a wrong path at one of the choice points. The partial path is grammatical, but is not part of a complete path corresponding to a correct interpretation of the sentence. If the comprehender is unable to recover the correct interpretation, the resulting impression is that the sentence is ungrammatical. Thus garden paths provide important clues about human sentence processing, since they directly reveal some of its limitations. Garden path sentences are typically determined by rapid-grammaticality judgment experiments (e.g., Warner and Glass (1987)), or by intuitions of the theorist (in the tradition of linguistic evidence—e.g., Pritchett (1988)). The classic garden path example is (1), due to Bever (1970):

- (1) The horse raced past the barn fell. (cf. *The horse that was raced past the barn fell.*)

Complementing the garden path data is evidence from sentences with local ambiguities that do not cause any difficulty. This evidence takes the form of sentence pairs¹, in which the two sentences are identical up to some choice point, but end up requiring different interpretations. Subjects will often take the wrong path on such sentences, since the disambiguating material is not available at the choice point. For example, the sentences

- (2) (a) John believed Mary.
(b) John believed Mary was lying.

¹They need not be pairs; k -way ambiguities with corresponding non-garden-path k -tuples are possible.

are both identical up to the ambiguous word *Mary*, which may be taken as the object of *believed* or the subject of an incoming clause. (2a) requires the object reading, and (2b) requires the subject reading. However, neither sentence causes a garden path effect.

Together, evidence from *garden paths* (GP) and *non-garden-path pairs* (NGP) constrains any theory of human sentence comprehension: a theory must not be so powerful that it has no difficulty with garden path sentences, yet it must predict the ease of comprehending the non-garden-path pairs. I have collected a set of distinct types of garden paths and non-garden-path pairs (primarily from the literature and my own invention). The types are distinguishable because they involve different syntactic constructions. Appendix A lists the current set of 19 garden path types and 19 non-garden-path types².

2 The structure of the NL-Soar theory

NL-Soar is a set of problem spaces in Soar that perform word-by-word comprehension. A *problem space* (Newell, 1990, Chap. 2) is a formulation of a task as an *initial state*, a *goal state*, and a set of *operators* that apply to states to produce new states. Any series of operator applications that leads to the goal state is taken as a solution to the problem. As NL-Soar proceeds through a set of sentences, applying operators in a *comprehension space*, it maintains a partial comprehension state in the form of two data structures: the *utterance model*, which is a dependency graph representing the syntactic relations in the current sentence (Hays, 1964; Mel'čuk, 1988), and the *situation model*, which is a semantic structure representing what the discourse is about. For a detailed description of the NL-Soar system, see (Lehman et al., 1991b).

Before examining the structure of the NL-Soar garden path theory in detail, it will be useful to consider some simple baseline theories and how they fare against the data.

2.1 Baseline theories

Given the impressive power of the human sentence processor, one possibility is to assume that at each choice point all interpretations are computed and carried forward in the analysis of a sentence. Such an *all-paths* theory perfectly predicts the ease of comprehending the NGP pairs. However, it incorrectly predicts that the GP sentences will be comprehended with equal ease. An alternative is to assume that only one interpretation at a time may be maintained. This *single-path* theory, coupled with assumptions about how the choices are made, can potentially predict the garden path effects. But regardless of how the choices are made, a single path theory incorrectly predicts that one sentence in each NGP pair will produce a garden path effect. Therefore both single path and all-paths theories must be modified to account for the data. An all-paths theory must limit in some way the interpretations carried forward to account for the GP sentences (Gibson, 1990b; Gibson,

²That there are currently 19 GP types and 19 NGP types is a coincidence.

1991; Just and Carpenter, 1991), and a single-path theory must posit some limited recovery mechanism to account for the NGP pairs (Frazier and Rayner, 1982; Warner and Glass, 1987; Pritchett, 1988; Abney, 1989)³.

The single path assumption has dominated theories in psycholinguistics. Most work has focused on identifying the knowledge sources and processes that determine how the local ambiguities are resolved (Altmann, 1989; Frazier, 1987). Relatively little attention has been given to specific recovery mechanisms. The theory presented here focuses on the *recovery* process rather than the choice point.

2.2 Simple destructive repair

Many varieties of recovery processes are possible. We consider only a particular kind of process called *simple destructive repair*. It can be defined generally in the following way. Assume a problem space for comprehension contains a set of *constructor* operators that produce new partial comprehension states by combining existing objects in the state (introducing new relationships) and/or introducing new objects. Assume the set of constructors is sufficient to reach the goal state (a successful comprehension). If the available knowledge to control the search in this space is incomplete (as it inevitably is in sentences with local ambiguities), then it may be possible to reach a deadend—a state in which no constructors apply. A *destructor* operator is one that applies to a state and removes existing relations or objects so that the state may be reconstructed in a different way. There is no additional communication between constructors and destructors other than the partial comprehension state assumed by the set of constructors alone. A *repair* consists of the application of a destructor operator followed by one or more applications of constructor operators. The actual reconstruction is effected by the same constructors that build all states. Repair is limited because it must work with the given state, unlike backtracking, which maintains a memory of previous states that may be returned to.

NL-Soar is a single-path comprehender with a simple destructive repair mechanism. The constructors are *link* operators that join together words with syntactic relations, building the utterance model. The destructor is the *snip* operator that breaks a link in the utterance model, permitting the model to be reconstructed in a different way⁴.

To illustrate repair, consider the comprehension of (3a):

³Lookahead parsers (Marcus, 1980) are a third class of theories; they get their power by expanding the available knowledge at each choice point to include incoming words in the utterance. They get their limitations from a fixed window of lookahead (defined by Marcus (1980) in terms of constituents rather than words). Conceivably, lookahead parsers could provide the basis for a garden path theory. However, no existing systems or fully specified theories in this class have been applied to any appreciable set of GP/NGP sentences. Problems with Marcus's original theory have been noted by Pritchett (1988) and Gibson (1991).

⁴ It's actually a little more complicated than this: the constructors also include the *merge* operator, which matches an incoming word against an expectation, and the *refer* operator, which builds the situation model. Snip breaks down the situation model as well the utterance model. But this complication need not concern us here.

- (3) (a) John likes green dragons.
(b) John likes green.

Assume that upon comprehending *green*, the system commits to an interpretation consistent with (3b); that is, it takes *green* as the direct object of *likes*. At *dragons*, a repair is necessary: snip breaks the object link between *likes* and *green*, permitting link to attach *green* to *dragons* as a describer, and then to attach *dragons* as the object of *likes*.

To make the theory complete, we assume a knowledge level (Newell, 1990, Chap. 2). component that guides the processing at the initial choice points. This component applies whatever knowledge is necessary to guide the comprehender down the path consistent with the data. Despite this degree of freedom, the theory is still constrained by the data in the following way. For each GP sentence, there must exist a grammatical partial path in the sentence such that the system cannot repair from that path to the correct interpretation. For each NGP pair, there must exist a single grammatical partial path such that the system can obtain a correct interpretation in both cases, either because it chooses the correct path or because it chooses the wrong path but can recover.

3 Requirements for a repair mechanism

In addition to accounting for the GP and NGP data, a repair mechanism must satisfy additional requirements, summarized in the list below (the NGP and GP data constraints appear as items 1 and 2):

1. **Must permit comprehension of NGP pairs.** The primary functional requirement for repair is that it be effective on the NGP sentences. Two functions that must be supported to perform repair are *lexical reinterpretation* and *syntactic reinterpretation*. Lexical reinterpretation changes the sense of a word⁵. For example, in (3), *green* must be reinterpreted as an adjective after first being interpreted as a noun. Syntactic reinterpretation changes the syntactic relations of the utterance model. (3) is also an example of syntactic reinterpretation, since the utterance model is restructured to include the describer relation for *green*.
2. **Must fail on garden path sentences.** The repair mechanism must not be so powerful that it cannot explain the garden path effects.
3. **Must be correct.** Repair must not produce interpretations inconsistent with the utterance. Repair must ultimately be bound by the same grammatical constraints that guide comprehension generally.

⁵Exactly what constitutes a unique sense is determined by the structuring of the lexicon, a matter to be returned to shortly.

4. **Must work without reprocessing input.** Repair must work without rereading or rehearing the misanalyzed utterance. This requirement distinguishes repair from a recovery mechanism based on regressive eye movements (Frazier and Rayner, 1982)⁶.
5. **Must be real-time.** Repair must work as part of the online comprehension process—within a few hundred milliseconds. Given the constraints of the Soar architecture (Newell, 1990), this requirement makes explicit that the repair must be *recognitional* (within a single Soar operator) or happen via a very few operator applications.

4 Searching for the right repair mechanism

Many varieties of simple destructive repair are possible. This section discusses exactly how various repair mechanisms might fail (and thus produce garden path effects), describes a large space of possible repair mechanisms, and gives the results of the search in this space for a theory consistent with the data.

4.1 How repair can fail

The partial comprehension state includes two sets of pairs of the form (*syntactic-relation, word*) that specify which words are available for linking via which syntactic relations. The two sets correspond to relations that can be assigned, and relations that can be received. The dynamic contents of the sets, collectively referred to as the A/R set, are determined by the grammar⁷. For example, after comprehending *The horse*, the A/R set is as follows:

Assigns: (restrictive-qualifier, *horse*)
 Receives: (subject, *horse*)
 (object, *horse*)

Thus if the sentence began *The horse that ate the oats...*, *horse* would assign the restrictive-qualifier relation. The crucial ambiguity in (1) is precisely whether *horse* assigns the restrictive-qualifier relation to *raced*, or receives the subject relation from *raced*.

Assume that the knowledge level component selects the main verb reading for *raced*. Upon encountering *fell*, why does repair fail to yield the correct interpretation?

One plausible assumption is given below:

Assumption 1 *The snip operator can only snip a link between two words in the A/R set.*

⁶Of course, such recovery mechanisms also exist.

⁷Currently, the grammatical components that generate the A/R set include the lexicon (which specifies the syntactic relations that each word can assign or receive), and structural knowledge similar to that captured by traditional phrase structure rules, but without the ordering information. The remainder of the grammar is cast as a set of constraints (various agreement and order tests) on the possible syntactic relations permitted by the A/R set. The precise nature of the grammar and its representation in the system is still under development.

If we further assume for independent syntactic reasons that *horse* is no longer in the A/R set after comprehending *The horse raced past the barn*, then the crucial subject link between *horse* and *raced* cannot be snipped, and the repair fails.

One problem with this explanation is that this syntactic assumption is suspect. Consider sentence (4):

(4) The horse raced past | wearing a saddle.

The most straightforward analysis is to have *horse* assign the non-restrictive-qualifier relation to *wearing*. Then we would expect at the point marked by | that *horse* indeed appears in the assigns set, even after comprehending the main verb. If *horse* is not in this set, then we must either posit an additional repair mechanism to account for comprehension of these sentences, or an alternative syntactic structure.

NGP type 1, repeated here as (5), presents a more serious problem.

(5; NGP1) (a) The defendant examined | the evidence.

(b) The defendant examined | by the lawyer shocked the jury.

At the point marked by | in (5), the A/R set is as follows, under the assumptions above:

Assigns: (object, *examined*)
Receives: empty

Thus, *examined* can assign a direct object as in (5a). However, the word *by* in (5b) forces reinterpretation of *examined* as a reduced relative, which requires snipping the subject link between *defendant* and *examined*. This would be impossible under Assumption 1, since *defendant* is no longer in the A/R set.

Consider relaxing the constraint on snip as follows:

Assumption 2 *Snip may work on any relation of a word that is part of the A/R set.*

Under Assumption 2 the subject link from *examined* in (5) is available regardless of the status of *defendant*, because *examined* is in the assigns set. However, a problem now shows up in (1), repeated here as (6).

(6; GP1) The horse raced past the barn | fell.

At the point marked by | in (6), the A/R set is:

Assigns: (clausal-modifier, *raced*)
(restrictive-qualifier, *barn*)
Receives: empty

(If the sentence was *The horse raced past the barn quickly*, *raced* assigns a clausal-modifier; if it was *The horse raced past the barn with the broken door*, *barn* assigns a restrictive-qualifier.) Since *raced* is in the A/R set, Assumption 2 permits the subject link between *horse* and *raced* to be snipped, allowing *raced* to be relinked as the restrictive-qualifier. No garden path effect arises.

Consider one more possible theory. Notice that three pieces of the utterance model become available for linking after a snip: the two pieces formed by splitting the old model, and the incoming word. Furthermore, these pieces are ordered according to the word order of the utterance; after snipping the subject link between *horse* and *raced*, *the horse* is the left piece, *raced past the barn* is the middle piece, and *fell* is the right piece. To make these concepts precise, we need the following definitions:

- An *utterance* consists of a pair $(W, <_{wo})$, where $W = \{w_1, w_2, \dots, w_n\}$ is a set of words, and $<_{wo}$ is a total ordering on the words, such that $w_1 <_{wo} w_2$ iff w_1 comes before w_2 in the utterance.
- A *dependency graph* D is a connected graph (W, R) , where W is a set of words, and R is a set of directed syntactic relations between the words.
- The *utterance model* U for an utterance $(W, <_{wo})$ consists of a set of one or more dependency graphs, called *submodels*, $\{D_1, D_2, \dots, D_n\}$, where $D_i = (W_i, R_i)$ and W_1, \dots, W_n form a partition of W .
- An ordering $<_{mo}$ on submodels is defined as follows. If $D_1 = (W_1, R_1)$ and $D_2 = (W_2, R_2)$, then $D_1 <_{mo} D_2$ iff for all $w \in W_1$ and $v \in W_2$, $w <_{wo} v$.
- A submodel D is *rightmost* if there is no submodel D' in the utterance model such that $D <_{mo} D'$. Two submodels D_1 and D_2 are *adjacent* if there is no submodel D' in the utterance model such that $D_1 <_{mo} D' <_{mo} D_2$.
- The *partial comprehension state* includes the utterance model U and an A/R set for each submodel in U .

Thus, link joins two submodels together to form a single submodel, and snip splits a submodel into two new submodels. We can now make explicit an implicit assumption in all the theories considered to this point:

Assumption 3 *Link may join only two adjacent submodels.*

Returning to the theory we were about to examine, consider constraining repair in the following way:

Assumption 4 *Link may join only the two rightmost submodels (the rightmost model and the model adjacent to it).*

If this is the case, (6) cannot be repaired, since there is no way of linking *fell* to *raced*. However, (5b) can be repaired, since *examined* can be linked to *by*, and then linked again to *defendant*.

Assumption 4 appears to be a better solution, but it too has problems. Consider (7) below:

- (7; NGP12) (a) Is the block in the box?
- (b) Is the block in the box red?

Upon comprehending *red* in (7b), the prepositional phrase *in the box* must be snipped from *is* and relinked as a modifier to *box*. However, Assumption 4 prevents this repair, because it forces *in the box* to first be linked to *red*.

4.2 RM: A space of repair mechanisms

The preceding discussion illustrated how different variations of repair mechanisms give rise to different effects, conveyed the nature of these variations, and demonstrated how the GP/NGP data set constrains the theorizing. There are many such variations, determined by decisions such as the two described above: what relations are available to snip (Assumptions 1 and 2), and what words are available to link (Assumption 4). The list below describes a dozen such relevant dimensions along with their possible values, defining a large space of repair mechanisms called *RM*.

1. **Available destructor operators.** In addition to snip, *peel* is another possible destructor operator. Peel removes a word but leaves behind an *expectation*, which is a node in a dependency graph just like a word, except it does not specify a particular lexical item⁸. Peel therefore preserves the structure of the existing utterance model.
 - (a) Only snip is available.
 - (b) Only peel is available.
 - (c) Both snip and peel are available.
2. **Availability of relations for snipping.**
 - (a) Just the relation connecting the last word to the utterance model⁹.
 - (b) Any relation between words in an A/R set (Assumption 1).
 - (c) Any relation to or from a word in an A/R set (Assumption 2).
 - (d) Any relation in the utterance model.

⁸Expectations are independently motivated parts of NL-Soar; they exist with or without the peel operator (see also footnote 4). They are similar to the empty categories in some linguistic theories (Chomsky, 1986) but exist in response to processing considerations, not in response to grammatical issues.

⁹More precisely, the last word that has not been snipped.

3. **Availability of words for peeling.**
 - (a) Only words in an A/R set.
 - (b) Any word.
4. **Availability of words for linking.**
 - (a) Words in two rightmost submodels only (Assumption 4).
 - (b) Words in two leftmost submodels only.
 - (c) Words in any two adjacent submodels (Assumption 3).
5. **Local relaxation of order constraint for link.**
 - (a) The relation order constraint is always locally enforced for link. The *relation-order* constraint is a constraint on the ordering of sister relations to a head. For example, the constraint that determiners come before describers in a noun phrase (**Red the box*) is a relation-order constraint¹⁰.
 - (b) Relation-order constraint relaxed for link after snip.
6. **Relaxation of word order constraint for snip.**
 - (a) Snip preserves original ordering. I.e., snip preserves the definition of $<_{mo}$ based on $<_{wo}$.
 - (b) Snip not constrained to preserve original ordering. I.e., snip may introduce a new ordering $<_{mo}$.
7. **Number of possible snips per word.**
 - (a) One snip only.
 - (b) No more than one consecutive snip. Snips are *consecutive* if a link operator does not occur between them. This constraint is equivalent to saying that an utterance model consists of no more than three submodels.
 - (c) Unlimited number of consecutive snips. I.e., an unlimited number of submodels.
8. **Preservation of snipped relations.** By requiring snipped relations to be reassigned, the original structure of the utterance model can be preserved as much as possible while accommodating the new input.
 - (a) A snipped relation must be reassigned.
 - (b) A snipped relation need not be reassigned.

¹⁰Relaxing the constraint locally means that *red* can be linked to *box* before *the*, but the final structure must still be grammatical.

9. **Computation of the new A/R sets after a snip.** This parameter specifies constraints on what words appear in the A/R sets for the new submodels created by a snip.
- (a) Only the words being unlinked by snip (the local arguments of the operator) appear in the new A/R sets.
 - (b) Only the words being unlinked by snip may be added to or deleted from the new A/R sets.
 - (c) No constraints (other than the grammar) for computation of the new A/R sets.
10. **Structure of the lexicon.** The nature of the lexical entries determines the level of commitment made by choosing a lexical entry, and thus affects when a word must be reaccessed. Two dimensions are relevant: words may be ambiguous with respect to syntactic properties (*syntactic ambiguity*), and ambiguous with respect to semantic properties (*semantic ambiguity*). Examples of syntactic ambiguity include the regular past/past participle interpretations of *raced*, and the noun/verb categorial ambiguity of *block*. Examples of semantic ambiguity include the wealthy versus food interpretations of *rich*, and the toy versus city area ambiguity of *block*. Each unique syntactic interpretation of a word is called a *syntactic sense*, and each unique semantic interpretation is called a *semantic sense*¹¹.

The possibilities for structuring the lexical entries are:

- (a) Lexical entries are a combination of a unique syntactic sense and a unique semantic sense.
 - (b) Lexical entries have a unique semantic sense, but may combine multiple syntactic senses.
 - (c) Lexical entries have a unique syntactic sense, but may combine multiple semantic senses.
 - (d) Lexical entries combine both semantic and syntactic senses.
11. **Reaccess of the lexicon.** This parameter specifies when lexical reinterpretation may happen. Permitting free reinterpretation of any word in the utterance model at any time would require a complete reconstruction of the model to maintain its consistency. The parameter values below restrict the reaccess to contexts without such negative computational consequences.
- (a) Only *free words* (words not linked to any other words) may be reaccessed.
 - (b) Only free words and words linked only to expectations may be reaccessed.

¹¹Of course, such a definition still leaves open the question of exactly what constitutes a unique semantic interpretation. For the moment, the analysis rests on relatively clear cases such as the examples given here. A precise and complete response to this question requires uncovering the right theory of lexical semantics for NL-Soar.

12. Cardinality of assigns and receives sets.

- (a) Only one word per relation.
- (b) Unlimited number of words per relation.

These parameters and values are motivated by a need to provide the basic functionality required to handle the NGP cases, and a desire to keep the processing and mechanism as cheap and simple as possible. For example, several of the parameters (4, 7, 9, and 12) embody constraints which can reduce working memory size and the match cost in the recognition memory.

Twelve theories in a promising part of the space have been tested against the data set by hand-simulation. Table 1 summarizes the parameter settings for these theories, and Appendix B gives the detailed descriptions. The starting point for the hill-climbing search

	Parameters											
	1	2	3	4	5	6	7	8	9	10	11	12
Theory A	c	d	b	a	a	b	b	a	b	c	a	b
Theory B	c	d	b	c	a	b	b	a	b	a	a	b
Theory B'	c	d	b	c	a	a	b	b	b	a	a	b
Theory C	c	d	b	b	a	a	b	b	c	a	a	b
Theory D	c	d	b	a	b	a	b	b	b	a	a	b
Theory E	c	d	b	a	b	a	b	b	a	a	a	b
Theory F	c	b	a	a	b	a	b	b	b	a	a	a
Theory G	c	c	a	a	b	a	b	b	b	a	a	a
Theory H	b	<i>n.a.</i>	b	a	b	a	b	b	c	a	b	b
Theory J	c	d	b	c	b	a	b	b	c	b	a	b
Theory K	c	d	b	c	b	a	b	b	c	b	b	b
Theory L	c	d	b	c	a	a	b	b	c	b	b	b

Table 1: **Parameter settings for Theories A-L.**

that produced these theories was fixing the value of parameter 7 to restrict the number of possible submodels. This constraint, combined with restrictions on the availability of the submodels for linking, appeared to account for many of the GP sentences. Even more importantly, loosening these restrictions resulted in mechanisms with excessive power. Most of the remaining parameter values were chosen to ensure enough power to account for the NGP sentences.

Table 2 shows how the 12 theories did against the data set. Most of the theories account for between 70% and 80% of the data, with the best at 78%. (Many of the theories were not tested against all of the data because the data set continued to expand as the theories

were being developed.) Although this search explored only a tiny fraction of the RM space, it appears that no theory in RM will deal with the data satisfactorily. The problem is the fundamental tradeoff in predictive power over the two types of sentences (recall the discussion of the baseline theories in Section 2.1). Any change in a parameter that increases the power of the mechanism to account for additional NGP sentences usually results in additional incorrect predictions for GP sentences, and vice versa. Figure 1 graphically shows this tradeoff. It thus seems unlikely that any theory in the RM space will do much better than 80%¹². To provide a better account of the data, we must look outside RM.




Figure 1: **The tradeoff between predicting GP and NGP sentences.**

¹²Of course, there is the remote possibility that we are stuck in a local maximum in RM, and jumping to a completely different part of the space will result in a large improvement. However, just such a radical jump is considered in Section 6, and it demonstrates the problem tradeoff even more dramatically.

	A	B	B'	C	D	E	F	G	H	J	K	L
GP 1	●	●	●		●	●	●	●	●	●	●	●
GP 2	●	●	●		●	●	●	●	●	●	●	●
GP 3	●	●	●		○	●	●	○	○	○	○	○
GP 4	●	●	●		●	●	●	●	●	●	●	●
GP 5	●	●	○		●	●	●	●	●	●	●	●
GP 6	●	●	●		○	○	●	○	○	○	○	●
GP 7	●	●	●		●	●	●	●	●	●	●	●
GP 8	●	○	○		○	○	○	○	●	●	●	●
GP 9	●	●	●		○	○	○	○	○	●	●	●
GP 10	●	●	●		●	●	●	●	●	●	●	●
GP 11	●	○	○		●	●	●	●	●	●	●	●
GP 12												
GP 13	●	●	○		○	○	●	○	○	●	●	●
GP 14	○	○	○		○	○	●	●	○	○	○	○
GP 15	●	○	○		○	○	●	●		●	●	●
GP 16	●	●	●		●	●	●	●	●	●	●	●
GP 17	●	○	○		●	●	●	●	●	○	○	○
GP 18												
GP 19												
NGP 1	○	●	●	●	●	●	○	●	●	●	●	●
NGP 2	○	○	○	○	●	○	○	●	●	●	●	●
NGP 3	●	●	●	○	●	●	○	●	●	●	●	●
NGP 4	●	●	●	○	●	●	●	●	●	●	●	●
NGP 5	●	●	●	○	●	●	●	●	●	●	●	●
NGP 6	●	●	●	●	●	●	●	●	●	●	●	●
NGP 7	●	●	●	●	●	●	●	●	●	●	●	●
NGP 8	○	●	●	○	●	●	●	●	○	○	○	○
NGP 9	●	●	●	●	●	●	●	●	●	○	●	●
NGP 10	○	●	○	○	○	○	○	○	○	○	○	○
NGP 11	●	●	●	●	●	●	○	○	●	●	●	●
NGP 12	○	○	○	●	○	○	○	○	○	●	●	●
NGP 13	○	●	●	○	●	●	○	●	○	●	●	●
NGP 14	●	●	●	●	●	●	○	○	●	○	○	○
NGP 15											●	●
NGP 16											○	○
NGP 17											●	●
NGP 18												
NGP 19												
Percent correct	77%	77%	67%	50%	70%	70%	67%	70%	69%	73%	76%	78%

●= correct prediction, ○= incorrect prediction

Table 2: Performance of Theories A-L against the 38 sentence types.

5 The role of phonology

Phonology is an important knowledge source in comprehension. Phonology is not captured in the existing utterance and situation models, and it is a reasonable conjecture that some of what is missing in the prior space of theories resides in the phonology. A separate *phonological model* is required to represent this information. In this section we consider the role such a model might play in the comprehension process, and why it is important for a garden path theory.

5.1 Phonological constraints

Phonology is part of the grammar of a language (Chomsky and Halle, 1968; Fromkin and Rodman, 1983). Like syntax, it provides constraints on how language is interpreted. For example, stress may be used to distinguish given and new information. Only (8a) below is an appropriate response to the question *What happened to the ball?*

- (8) (a) The ball was **hit**.
- (b) The **ball** was hit.

(8b) is an appropriate response to *What was hit?*

Pitch, or intonation, is another important aspect of phonology. Consider question (9) from Fromkin and Rodman (1983):

- (9) What did you put in my drink, Jane?

If the pitch rises on *drink* and then falls off, the questioner is asking what Jane put in the drink. If the pitch rises sharply on *Jane* instead, without decreasing, the questioner is asking if someone put Jane in the drink.

Intonation *breaks* also provide crucial constraints on interpretation. In the examples below, / is used to mark a break. (10), adapted from Marcus and Hindle (1990), shows how different breaks in an identical word sequence lead to different meanings:

- (10) (a) We only suspected / they all knew that a burglary had been committed.
- (b) We only suspected / they all knew / that a burglary had been committed.

Note that (10b) would be written as:

- (11) We only suspected—they all knew—that a burglary had been committed.

5.2 How phonology interacts with repair in speech comprehension

The available constructor operators can be categorized according to what they take as input and what they construct as output. The phonological model will be abbreviated as *P*, the utterance model as *U*, the situation model as *S*, and the auditory signal as *A*. Then we can assume the following constructor type for the phonological model:

$$(12) A \dots \rightarrow P$$

where the \dots represent possible input from higher level models (taken here to be U and S)¹³. Constructors that perform lexical access are of type:

$$(13) P \dots \rightarrow U$$

where again \dots represents possible input from higher level models. Constructors that build the utterance model based on the existing partial utterance model (including the results of lexical access) and the phonological model, are of type:

$$(14) P, U \dots \rightarrow U$$

To take into account some of the constraints discussed in Section 5.1, the constructors that build the situation model must be of the form:

$$(15) P, U, \dots \rightarrow S$$

The utterance and situation models may be repaired by applying destructor operators (snip or peel) and rebuilding the models with constructor types (13), (14) and (15). Since we require that repair work without rehearing the utterance (Section 3), the auditory signal is not available. Thus, the phonological model itself may not be repaired, since the required input is not available to its constructors at repair time.

These simple assumptions outline the beginnings of a mechanism that supports the basic functional requirements of repair (Section 3). It ensures correctness because no knowledge sources are ignored (e.g. the utterance model may not be reconstructed in a way inconsistent with the phonology.) It permits reaccess of the lexicon and syntactic restructuring via constructor (14). It works without reprocessing the input (i.e., requiring A to be again present for reanalysis).

5.3 How phonology interacts with repair in reading

The phonological model plays an important role in reading as well as in listening¹⁴. Written text may contain explicit cues about the intended phonological model: for example, the boldface in (8) or the dashes in (11). Thus, we assume the reader builds a phonological model as part of the online comprehension process, even though some aspects of phonology are not explicitly represented orthographically. Using V to represent the visual signal, we then have constructors of the following type¹⁵:

¹³Whether these inputs exist is a version of the modularity question, but this issue seems not to affect the discussion here.

¹⁴The precise role of phonology in reading is still a matter of some controversy in psychology and neuropsychology (e.g., see footnote 15 concerning lexical access). The theory emerging here is most consistent with those theories that assign a critical role to phonological processes, even in fluent reading; e.g., (Perfetti, 1985).

¹⁵The presence or absence of constructors of type $V \dots \rightarrow U$ is at the heart of the debate over whether lexical access proceeds directly from the visual form, or from a phonological code, or both (Jared and Seidenberg, 1991).

(16) $V \dots \rightarrow P$

Repair of the utterance and situation models may then proceed in the same fashion as described in Section 5.2: by applying a destructor operator and rebuilding the models with constructor types (13), (14) and (15). As in the speech case, these operators support lexical and syntactic reinterpretation. Since we require repair to work without rereading (Section 3), the visual input is not available to repair the phonological model.

5.4 Consequences of adding a phonological model

We can examine the consequences of a *phonological addition* (Theory PA) independent of the RM space by making the theory operational with a knowledge level repair mechanism. That is, without specifying the precise details of how repair works (i.e., choosing values for the parameters in Section 4.2), assume that, if a repair can be made *within the constraints of the phonological model*, it will be made. In Section 6 the specifics of the repair mechanism will be re-examined in light of the phonological addition.

Theory PA requires that the repaired interpretation be phonologically consistent with the initial incorrect interpretation. Said differently, whenever the correct interpretation is phonologically *inconsistent* with the initial interpretation, a garden path effect arises. This is because the phonological model cannot itself be repaired, and the constructors are constrained by the content of the phonological model.

In fact, this is the case in the data: in *all* NGP pairs, the repaired interpretation is phonologically consistent with the initial incorrect interpretation, and where the initial interpretation is phonologically inconsistent, a garden path effect arises. Consider (17a) below.

- (17; GP8) (a) While Mary mended a sock fell on the floor.
(b) While Mary mended, a sock fell on the floor.

The garden path in (17a) is avoided by the introduction of the comma in (17b). The hypothesis is that the correct intonation of this sentence has an obligatory break between *mended* and *sock*, or else *sock* must be taken as the direct object of *mended*. Without the comma, (17b) is intonated without the break, and is therefore phonologically inconsistent with the correct interpretation.

If the phonological model is represented in the short-term phonological buffer (Baddeley, 1990)¹⁶, a distance-to-disambiguation effect (Warner and Glass, 1987) may arise. As the phonological model becomes unavailable, the constructors will be unable to complete the repair. Consider (18):

¹⁶There have been at least three audition-based memories hypothesized in work on short-term memory: a short-term auditory buffer, a short-term phonetic buffer, and a short-term phonological buffer (Clark and Clark, 1977). The initial hypothesis is that the phonological model is represented in the short-term phonological buffer identified in rehearsal theories of verbal memory, estimated to hold 1.5–2 seconds worth of speech (Baddeley, 1990).

- (18; GP14) The girls believe the man who believes the very strong ugly boys struck the dog killed the cats.
(cf. *The man who believes the boys struck the dog is believed by the girls to have killed the cats.*)

If the distance between *the man* and *killed* (the disambiguating word) is great enough, the phonological encoding of *the man* will no longer be in the buffer to permit the repair. Unfortunately, without further specification of the nature of the phonological code, the relationship of the buffer size to articulation rate is unclear, making predictions of garden paths due to distance-to-disambiguation tenuous¹⁷.

Finally, Theory PA makes interesting predictions regarding homographs and homophones (which are not currently part of the GP/NGP data set). For example, it predicts difficulty if a homograph is initially assigned a phonological encoding inconsistent with its correct interpretation, as in (19):

- (19) After the symphony, we had the bass player over for dinner, but the bass we picked up at the market yesterday didn't seem fresh and we had to order out.

Since lexical reaccess occurs via the phonological model, a word may be substituted by a homophone with a different spelling, without seriously affecting comprehension, as in (20)¹⁸:

- (20) (a) I wish the Penguins had one the game against the Capitals.
(b) I found a pair of keys today . . . describe them and there yours.

Table 3 shows how theory PA does against the entire data set. This evaluation represents a subjective judgment on my part on the phonological similarity of the interpretations or the distance-to-disambiguation. Some cases are less clear than others (as noted above, the distance effect is quite difficult to gauge); these are indicated in the table. The great performance on the NGP pairs is due, of course, to the knowledge level repair component. Finally, there are still several critical garden paths left unexplained, most interestingly, GP1, *The horse raced past the barn fell*.

6 Re-examining the repair mechanism

Given that Theory PA predicts many garden path effects even with a powerful knowledge level repair component, the underlying repair mechanism may actually be more powerful than those in Theories A-L. In particular, since none of the theories permitted more than

¹⁷This is also an obvious place where individual differences will show up; once the theory is developed further it might be interesting to see if there is a correlation between buffer size as measured by standard STM experiments and buffer size as measured by long GP sentences, for individual subjects.

¹⁸(20b) was recently seen on an electronic bboard.

one consecutive snip (i.e., no theory had value (c) for parameter 7), it is worth considering relaxing this restriction. Theory M does just that:

Theory M. Only snip is available (1a). Only a relation to or from the last word may be snipped (2a). Link works between any adjacent submodels (4c). Link need not respect relation-order after a snip (5b). Snip respects word order (6a). Unlimited number of consecutive snips (7c). A snipped syntactic relation need not be refilled (8b). No constraints except grammar for computation of A/R sets after a snip (9c). Lexical entries may combine syntactic senses (10c). Only free words may be reaccessed (11a). Unlimited number of words per relation in A/R sets (12b).

Besides eliminating the constraint on consecutive snips, Theory M differs from Theories A-L in one other important way. The only available destructor operation is snipping the last word. (Theory M is therefore similar to Theory H in positing only one destructor). Eliminating the structural restriction on submodels permits simplifying the destructor in this way, without a loss of functionality.

As expected, however, Theory M exhibits the problematic tradeoff plaguing the RM space. Theory M accounts for all the NGP pairs, but misses almost all of the GPs (Table 3).

We can combine Theories M and PA, so that Theory M specifies the actual repair mechanism described at the knowledge level in Theory PA. As the table shows, the resulting Theory M+PA does somewhat better than any of the previous theories considered, but still leaves several garden paths unexplained. There must be more to the story.

Theories in the RM space specify aspects of a problem space: defining the structure of the state, conditions on operator application, and so on. For a system to behave in a problem space, the *search control* knowledge must be specified as well. An implicit assumption in the RM theories is that perfect search control knowledge is available: if there is a path of constructor and destructor operators that leads to the goal, then that path will be taken. Such an assumption can be clearly formulated as a knowledge level component. At the symbol level, the assumption can be realized by exhaustive, recursive, lookahead search; this is the mechanism used in the NL-Soar system. However, such exhaustive search is psychologically implausible, and in fact violates NL-Soar's single-path assumption, since multiple paths in the space must be maintained to permit backtracking¹⁹.

One alternative is to eliminate exhaustive lookahead and guide the search with a fixed body of search control rules. If the rules are insufficient to guide the comprehension down the right path, comprehension will fail and a garden path occurs. To explore this alternative, Theory N was constructed. Theory N is identical to Theory M in all respects, except it posits a fixed set of control rules to replace the lookahead search. The set consists of eight relatively simple rules, such as *prefer link to snip, ceteris paribus*²⁰.

¹⁹Note that these are multiple paths being explored during the comprehension of a *single* word, not paths maintained across words. Nevertheless, it violates the single-path assumption.

²⁰Note that such control rules are not alternatives to syntactic heuristics such as minimal attachment (Frazier and Rayner, 1982), which are concerned with the initial choice point, not the repair.

As shown in Table 3, Theory N is still powerful enough to handle most of the NGP sentences. When combined with the phonological addition, the resulting Theory PA+N appears to be the most promising yet, making predictions with up to 95% accuracy (further work is required to tighten the predictions on the unclear cases).

7 Conclusion

The garden path theory described here has been motivated by an attempt to realize a functional comprehension capability in Soar that matches human performance. This required working out the details of the problem space representation, resulting in the formulation and exploration of the RM space. This exploration has led us to consider the integration of a new knowledge source (phonology), and the effects that integration might have on the garden path predictions. In addition, we have reevaluated an implausible component of the existing system (the lookahead search), and proposed an alternative that is preferred on both computational and empirical grounds.

There is much work yet to be done. An implementation of the theory is required to verify the hand simulation and fully demonstrate the effectiveness of the fixed search control on a large corpus. Other schemes for managing the search control are worth exploring as well. For example, we are currently considering strategies for learning search control rules that distribute the exhaustive search across multiple sentences, permitting local processing that maintains only one or two states.

Theory PA sketched the broad outlines of the integration of a phonological model, but the precise nature of the phonological representation is still an open issue. Furthermore, introducing a phonological buffer in Soar has far-reaching implications, impacting general issues of short term memory and learning. That is a consequence of working within a theory of the cognitive architecture: architectural changes cannot be proposed in isolation just to handle the demands of a local theory.

Acknowledgements: Allen Newell and Jill Lehman have provided invaluable input during the development of this theory (and tech report). Brad Pritchett's comments also helped improved the quality of this report. Any errors, of course, are mine.

	PA	M	M+PA	N	N+PA
GP 1	○	○	○	●	●
GP 2	○	○	○	⊙	⊙
GP 3	⊙	○	⊙	○	⊙
GP 4	○	○	○	●	●
GP 5	●	○	●	●	●
GP 6	○	○	○	⊙	⊙
GP 7	○	○	○	⊙	⊙
GP 8	●	○	●	●	●
GP 9	○	○	○	○	○
GP 10	●	○	●	●	●
GP 11	●	●	●	●	●
GP 12					
GP 13	⊙	○	⊙	●	●
GP 14	⊙	○	⊙	●	●
GP 15	⊙	○	⊙	●	●
GP 16	●	○	●	●	●
GP 17	○	○	○	●	●
GP 18	●	○	●	●	●
GP 19	⊙	○	⊙	⊙	⊙
NGP 1	●	●	●	●	●
NGP 2	●	●	●	●	●
NGP 3	●	●	●	●	●
NGP 4	●	●	●	●	●
NGP 5	●	●	●	●	●
NGP 6	●	●	●	⊙	⊙
NGP 7	●	●	●	●	●
NGP 8	●	●	●	●	●
NGP 9	●	●	●	●	●
NGP 10	●	●	●	●	●
NGP 11	●	⊙	⊙	⊙	⊙
NGP 12	●	●	●	●	●
NGP 13	●	●	●	●	●
NGP 14	●	●	●	○	○
NGP 15	●	●	●	●	●
NGP 16	●	●	●	●	●
NGP 17	●	●	●	●	●
NGP 18	●	●	●	●	●
NGP 19	●	●	●	●	●
Percent correct (best case [†])	81%	54%	81%	91%	95%
Percent correct (worst case [‡])	68%	51%	65%	72%	76%

●= correct prediction, ⊙= unclear correct prediction, ○= incorrect prediction
[†]all unclear predictions correct, [‡]all unclear predictions incorrect

Table 3: Performance of Theories PA, M, and N against the 38 sentence types.

Appendix A: GP and NGP corpus

The following four tables give the current corpus of garden paths and non-garden-path pairs. For each type, I have attempted to acknowledge the original source, as well as the source for the specific example listed. (The particular type labels given to the sentences are not those of the original authors'.)

	Type and source	Example
GP1	Matrix-verb/reduced-relative (Bever, 1970)	The horse raced past the barn fell. (cf. <i>The horse that was raced past the barn fell.</i>)
GP2	NP-modifier/NP (Marcus, 1980; Pritchett, 1988; Gibson, 1990a)	The Russian women loved died. (cf. <i>The Russian that women loved died.</i>)
GP3	Object/reduced-relative (Gibson, 1991; Pritchett, 1988)	John gave the boy the dog bit a dollar. (cf. <i>John gave a dollar to the boy that the dog bit.</i>)
GP4	Oblique-comp/NP-modifier (Gibson, 1990a)	I put the candy on the table in my mouth. (cf. <i>I put the candy that was on the table into my mouth.</i>)
GP5	Embedded-object/object (Pritchett, 1988)	Sue gave the man who was reading the book. (cf. <i>Sue gave the book to the man who was reading.</i>)
GP6	Verb/noun (Milne, 1982)	The building blocks the sun faded are red. (cf. <i>The building blocks that the sun faded are red.</i>)
GP7	Clausal-comp/relative-clause (Crain and Steedman, 1985; Gibson, 1990a)	John told the man that Mary kissed that Bill saw Phil. (cf. <i>The man that Mary kissed was told by John that Bill saw Phil.</i>)
GP8	Object/subject (Frazier and Rayner, 1982; Pritchett, 1988)	While Mary mended a sock fell on the floor. (cf. <i>While Mary mended, a sock fell on the floor.</i>)
GP9	Predicate-comp/noun (Pritchett, 1988)	The boy got fat melted. (cf. <i>The boy got some fat melted for the cook.</i>)
GP10	Object/subject w/relative (Warner and Glass, 1987)	Before the boy kills the man the dog bites strikes. (cf. <i>Before the boy kills, the man that the dog bites strikes.</i>)

Table 4: **Current set of garden path types (part 1 of 2).**

Type and source	Example
GP11 †Reduced-relative/clausal-comp (Warner and Glass, 1987)	When the horse kicks the boy the dog bites the man. (cf. <i>When the horse kicks the boy, the dog bites the man.</i>)
GP12 Tense ambiguity (Warner and Glass, 1987)	The boys put out the dogs that are strong when the man who is very ugly strikes the clock. (cf. <i>The boys put out the dogs when the man struck the clock.</i>)
GP13 †Reduced-relative/matrix verb (Warner and Glass, 1987)	The men believed to strike the dog is ugly. (cf. <i>The men believed that striking the dog is ugly.</i>)
GP14 Clausal-object ambiguity (Warner and Glass, 1987)	The girls believe the man who believes the very strong ugly boys struck the dog killed the cats. (cf. <i>The man who believes the boys struck the dog is believed by the girls to have killed the cats.</i>)
GP15 †Relative-clause/clausal-object (Crain and Steedman, 1985)	The psychologist told the wife that he was having trouble with her husband. (cf. <i>The psychologist let the wife know that he was having trouble with her husband.</i>)
GP16 Complementizer/pronoun	Before she knew that she went to the store. (cf. <i>Before she knew that, she went to the store.</i>)
GP17 Matrix-verb/relative (short) (Pritchett, 1988; Abney, 1989)	The boat floated sank. (cf. <i>The boat that was floated sank.</i>)
GP18 <i>Throughout</i> ambiguity Discovered in (Allen, 1987)	Throughout the plan structure that serves as the expectation will be called the e-plan. (cf. <i>Throughout, the plan structure that serves as the expectation will be called the e-plan.</i>)
GP19 Distant particle (Chomsky, 1965)	I called the man who wrote the book that you told me about up. (cf. <i>I called up the man who wrote the book that you told me about.</i>)

†These garden paths depend on an initial interpretation counter to what is normally assumed.

Table 5: **Current set of garden path types (part 2 of 2).**

	Type and source	Example
NGP1	Matrix-verb/reduced-relative (Ferreira and Clifton, 1986)	The defendant examined the evidence. The defendant examined by the lawyer shocked the jury.
NGP2	Clausal-modifier/noun-modifier (Taraban and McClelland, 1988)	The spy saw the cop with the binoculars. The spy saw the cop with the revolver.
NGP3	Direct-object/clausal-object (Pritchett, 1988)	I knew the man. I knew the man hated me passionately.
NGP4	Plural/possessive (Pritchett, 1988)	The woman kicked her sons. The woman kicked her sons' dogs' houses' doors.
NGP5	Noun/Noun-modifier (Pritchett, 1988)	Without her we failed. Without her contributions we failed.
NGP6	Theta-role switch (Pritchett, 1988)	I gave the dogs to Mary. I gave the dogs some bones.
NGP7	Noun/verb (Milne, 1982)	The building blocks are red. The building blocks the sun.
NGP8	Object/subject (Warner and Glass, 1987)	When the boys strike the dog kills. When the boys strike the dog the cat runs away.
NGP9	Have-question/imperative (Marcus, 1980)	Have the boys taken the exam today? Have the boys take the exam today.

Table 6: **Current set of non-garden-path pair types (part 1 of 2).**

	Type and source	Example
NGP10	Noun-modifier/noun (Gibson, 1990b)	I gave her earrings to Sally. I gave her earrings on her birthday.
NGP11	Adjective sense ambiguity	The deep pit was scary. The deep philosopher was kind.
NGP12	Question-predicate/NP-modifier (Marcus, 1980)	Is the block in the box? Is the block in the box red?
NGP13	Coordinate ambiguity	I went to the mall and the drugstore. I went to the mall and the drugstore was closed.
NGP14	Direct-obj/clausal-obj (long) (Warner and Glass, 1987)	The girls believe the man who struck the dog. The girls believe the man who struck the dog killed the cats.
NGP15	Matrix-verb/reduced-relative	The defendant carefully examined the evidence. The defendant carefully examined by the prosecutor looked nervous.
NGP16	Pred-comp/describer	The boy got fat. The boy got fat mice for his pet snake.
NGP17	Object/prep-object gap	John saw the ball the boy hit. John saw the ball the boy hit the window with.
NGP18	Singular/plural noun	The sheep seem very happy. The sheep seems very happy.
NGP19	Verb/verb+particle	John picked the boy for this team. John picked the boy up yesterday.

Table 7: **Current set of non-garden-path pair types (part 2 of 2).**

Appendix B: Descriptions of Theories A-L

The numbers and letters in the descriptions refer to the dimensions and values of space RM, given in Section 4.2.

Theory A. Both snip and peel are available (1c). Any relation available for snipping (2d). Any word available for peeling (3b). Link works only between two rightmost submodels (4a). Link respects local relation order constraint (5a). Snip need not respect word order (6b). Not more than one snip in succession (7b). A snipped syntactic relation must be refilled (8a). Only the words being snipped may be added to or deleted from the new A/R sets. (9b). Lexical entries may combine syntactic senses (10c). Only free words may be reaccessed (11a). Unlimited number of words per relation in A/R sets (12b).

Theory B. Both snip and peel are available (1c). Any relation available for snipping (2d). Any word available for peeling (3b). Link works between any two adjacent submodels (4c). Link respects local relation order constraint (5a). Snip need not respect word order (6b). Not more than one snip in succession (7b). A snipped syntactic relation must be refilled (8a). Only the words being snipped may be added to or deleted from the new A/R sets. (9b). Lexical entries may not combine multiple syntactic senses or semantic senses (10a). Only free words may be reaccessed (11a). Unlimited number of words per relation in A/R sets (12b).

Theory B'. Both snip and peel are available (1c). Any relation available for snipping (2d). Any word available for peeling (3b). Link works between any two adjacent submodels (4c). Link respects local relation order constraint (5a). Snip respects word order (6a). Not more than one snip in succession (7b). A snipped syntactic relation need not be refilled (8b). Only the words being snipped may be added to or deleted from the new A/R sets. (9b). Lexical entries may not combine multiple syntactic senses or semantic senses (10a). Only free words may be reaccessed (11a). Unlimited number of words per relation in A/R sets (12b).

Theory C. Both snip and peel are available (1c). Any relation available for snipping (2d). Any word available for peeling (3b). Link works between two leftmost submodels only (4b). Link respects local relation order constraint (5a). Snip respects word order (6a). Not more than one snip in succession (7b). A snipped syntactic relation need not be refilled (8b). No constraints except grammar for computation of A/R sets after a snip (9c). Lexical entries may not combine multiple syntactic senses or semantic senses (10a). Only free words may be reaccessed (11a). Unlimited number of words per relation in A/R sets (12b).

Theory D. Both snip and peel are available (1c). Any relation available for snipping (2d). Any word available for peeling (3b). Link works between two rightmost submodels

only (4a). Link need not respect relation-order after a snip (5b). Snip respects word order (6a). Not more than one snip in succession (7b). A snipped syntactic relation need not be refilled (8b). Only the words being snipped may be added to or deleted from the new A/R sets. (9b). Lexical entries may not combine multiple syntactic senses or semantic senses (10a). Only free words may be reaccessed (11a). Unlimited number of words per relation in A/R sets (12b).

Theory E. Same as Theory D, except: Snip adds only the words being unlinked to the new A/R sets (9a).

Theory F. Same as Theory D, except: Snip only works on relations between words in the A/R set (2b). Peel only works on a word in the A/R set (3a). Only one word per relation in A/R sets (12a).

Theory G. Same as Theory F, except: Snip only works on any relation to or from a word in an A/R set (2c).

Theory H. The only destructor is peel (1b). Any word available for peeling (3b). Link works between two rightmost submodels only (4a). Link need not respect relation-order after a snip (5b). Snip respects word order (6a). Not more than one snip in succession (7b). A snipped syntactic relation need not be refilled (8b). No constraints except grammar for computation of A/R sets after a snip (9c). Lexical entries may not combine multiple syntactic senses or semantic senses (10a). Only free words and free words attached to expectations may be reaccessed (11b). Unlimited number of words per relation in A/R sets (12b).

Theory J. Both snip and peel are available (1c). Any relation available for snipping (2d). Any word available for peeling (3b). Link works between any two adjacent submodels (4c). Link need not respect relation-order after a snip (5b). Snip respects word order (6a). Not more than one snip in succession (7b). A snipped syntactic relation need not be refilled (8b). No constraints except grammar for computation of A/R sets after a snip (9c). Lexical entries may combine syntactic senses (10b). Only free words may be reaccessed (11a). Unlimited number of words per relation in A/R sets (12b).

Theory K. Same as Theory J, except: Only free words and free words attached to expectations may be reaccessed (11b).

Theory L. Both snip and peel are available (1c). Any relation available for snipping (2d). Any word available for peeling (3b). Link works between any two adjacent submodels (4c). Link respects local relation-order constraint (5a). Snip respects word order (6a). Not more than one snip in succession (7b). A snipped syntactic relation need not be refilled (8b). No constraints except grammar for computation

of A/R sets after a snip (9c). Lexical entries may combine syntactic senses (10b). Only free words and free words attached to expectations may be reaccessed (11b). Unlimited number of words per relation in A/R sets (12b).

References

- Abney, S. P. (1989). A computational model of human parsing. *Journal of Psycholinguistic Research*, 18(1):129–144.
- Allen, J. (1987). *Natural Language Understanding*. Benjamin/Cummings, Menlo Park, CA.
- Altmann, G. (1989). Parsing and interpretation: An introduction. *Language and Cognitive Processes*, 4:SI 1–19.
- Baddeley, A. D. (1990). *Human Memory: Theory and Practice*. Allyn and Bacon, Boston.
- Bever, T. G. (1970). The cognitive basis for linguistic structures. In Hayes, J. R., editor, *Cognition and the Development of Language*. Wiley, New York.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. MIT Press, Cambridge, MA.
- Chomsky, N. (1986). *Barriers*. MIT Press, Cambridge, MA.
- Chomsky, N. and Halle, M. (1968). *The Sound Pattern of English*. Harper and Row, New York.
- Clark, H. and Clark, E. (1977). *The Psychology of Language: An introduction to psycholinguistics*. Harcourt Brace Jovanovich, New York.
- Crain, S. and Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological syntax processor. In Dowty, D. R., Karttunen, L., and Zwicky, A. M., editors, *Natural Language Parsing*. Cambridge University Press, Cambridge, U.K.
- Ferreira, F. and Clifton, Jr., C. (1986). The independence of syntactic processing. *Journal of Memory and Language*, 25:348–368.
- Frazier, L. (1987). Sentence processing: A tutorial review. In Coltheart, M., editor, *Attention and Performance XII: The Psychology of Reading*. Lawrence Erlbaum Associates, Ltd., East Sussex, U.K.
- Frazier, L. and Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, 14:178–210.

- Fromkin, V. and Rodman, R. (1983). *An Introduction to Language, Third Edition*. Holt, Rinehart and Winston, New York.
- Gibson, E. (1990a). Memory capacity and sentence processing. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*, pages 39–46.
- Gibson, E. (1990b). Recency preference and garden-path effects. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, pages 372–379.
- Gibson, E. A. F. (1991). *A Computational Theory of Human Linguistic Processing: Memory Limitations and Processing Breakdown*. PhD thesis, Carnegie Mellon. Available as Center for Machine Translation technical report CMU-CMT-91-125.
- Hays, D. G. (1964). Dependency theory: A formalism and some observations. *Language*, 40(4):511–525.
- Jared, D. and Seidenberg, M. S. (1991). Does word identification proceed from spelling to sound to meaning? *Journal of Experimental Psychology: General*, 120(4):358–394.
- Just, M. A. and Carpenter, P. A. (1991). A capacity theory of comprehension: Individual differences in working memory. Department of Psychology, Carnegie Mellon University. In press, *Psychological Review*.
- Laird, J. E., Newell, A., and Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, 33:1–64.
- Lehman, J. F., Lewis, R. L., and Newell, A. (1991a). Integrating knowledge sources in language comprehension. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pages 461–466.
- Lehman, J. F., Lewis, R. L., and Newell, A. (1991b). Natural language comprehension in Soar: Spring 1991. Technical Report CMU-CS-91-117, School of Computer Science, Carnegie Mellon University.
- Lewis, R. L., Huffman, S. B., John, B. E., Laird, J. E., Lehman, J. F., Newell, A., Rosenbloom, P. S., Simon, T., and Tessler, S. G. (1990). Soar as a unified theory of cognition: Spring 1990. In *Twelfth Annual Conference of the Cognitive Science Society*, pages 1035–1042.
- Marcus, M. and Hindle, D. (1990). Description theory and intonation boundaries. In Altmann, G. T. M., editor, *Cognitive Models of Speech Processing: Psycholinguistic and Computational Perspectives*. The MIT Press, Cambridge, MA.
- Marcus, M. P. (1980). *A Theory of Syntactic Recognition for Natural Language*. MIT Press, Cambridge, MA.

- Mel'čuk, I. A. (1988). *Dependency Syntax: Theory and Practice*. State University of New York Press, Albany, NY.
- Milne, R. W. (1982). Predicting garden path sentences. *Cognitive Science*, 6:349–373.
- Newell, A. (1990). *Unified Theories of Cognition*. Harvard University Press, Cambridge, Massachusetts.
- Perfetti, C. A. (1985). *Reading ability*. Oxford, New York.
- Pritchett, B. L. (1988). Garden path phenomena and the grammatical basis of language processing. *Language*, 64:539–576.
- Taraban, R. and McClelland, J. L. (1988). Constituent attachment and thematic role assignment in sentence processing: influences of content based expectations. *Journal of Memory and Language*, 27:597–632.
- Warner, J. and Glass, A. L. (1987). Context and distance-to-disambiguation effects in ambiguity resolution: Evidence from grammaticality judgments of garden path sentences. *Journal of Memory and Language*, 26:714–738.