

# Local Attachment and Competing Constraints

Colin Phillips  
Univ. of Delaware  
colin@udel.edu

Edward Gibson  
M.I.T.  
gibson@psyche.mit.edu

## 1. Introduction

It is widely assumed that some kind of locality or recency constraint influences decisions in the parsing of structurally ambiguous sentences. However, locality is usually assumed to be a relatively *weak* constraint, which may be overridden by a variety of other factors (e.g. argument-attachment biases, *Minimal Attachment*, frequency biases, semantic constraints etc.).

For example, although there is a consistent bias to associate the ambiguous adverbial *yesterday* in (1) with the embedded clause, i.e. locally,

(1) Bill thought it was fun at the beach *yesterday*.

numerous studies have shown that the presumed local attachment of the PP as a modifier in (2) is avoided by subjects. This suggests that locality is subordinate to other constraints in ambiguity resolution. Theories abound as to what these other constraints might be.

(2) The lifeguard saw the swimmer with inadequate experience.

This poster summarizes our ongoing work on a novel structural ambiguity which suggests that locality is a stronger constraint than normally assumed.

There are many different versions of locality constraints in the literature (e.g. Kimball 1973, Fodor & Frazier 1978, Gibson 1991, Stabler 1994, Phillips 1995, among many others), but the differences between them are irrelevant to the issues discussed here, so we refer to them all under the heading *local attachment*.

## 2. The matrix/embedded relative ambiguity

The underlined NP in a sentences like (3) is temporarily ambiguous between a **matrix clause subject** attachment and an **embedded relative clause subject** attachment.

(3) Although Erica hated the house she ...

In the relative clause parse, a gap is expected to occur at some later point in the sentence.

(4) Although Erica hated the house she had owned \_\_ for ten years ...

In the matrix clause parse, an overt NP occupies the position of the gap in the corresponding relative clause parse.

(5) Although Erica hated the house she had owned **it** for ten years ...

Figure 1 shows the alternative attachments for the ambiguous NP. Notice that the relative clause attachment is clearly more local to the immediately preceding lexical items, although the additional structure that must be built is somewhat more complex.

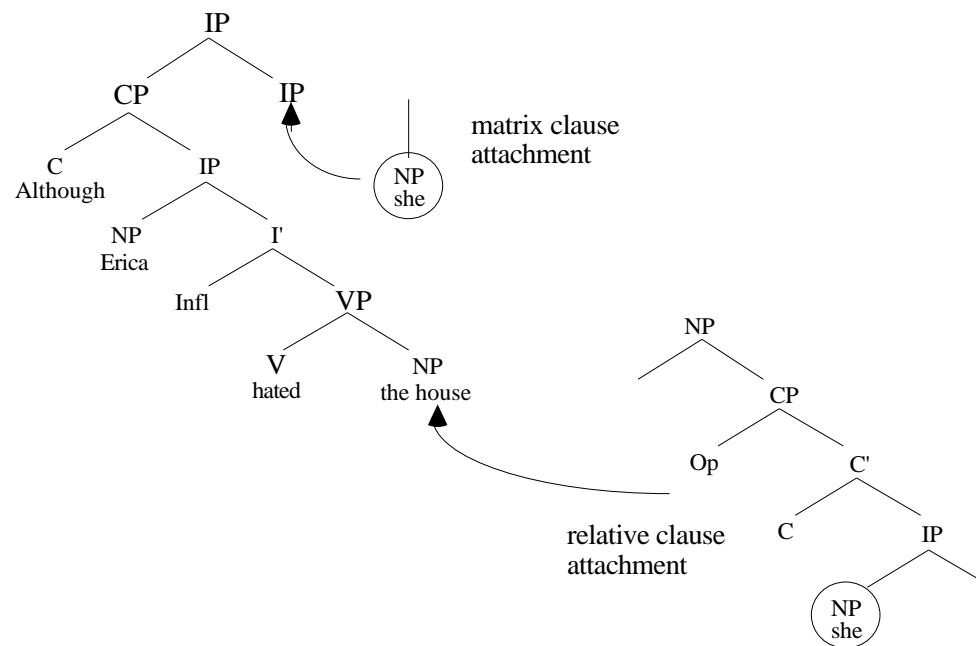


Figure 1: alternative structures in matrix/embedded-relative ambiguity

## Predictions

Although most constraints that we are aware of predict a bias for the matrix clause parse, locality clearly predicts a relative clause bias.

Factor	Preference	Reason
Minimal Attachment (Fodor & Frazier 1978)	matrix clause	fewer nodes required to build matrix Spec,IP than CP complex plu Spec,IP in relative clause
Attach arguments/ obligatory constituents	matrix clause	Matrix subject is obligatory, relative clause is not
Complete predicate– argument relations (Pritchett 1988, Gibson 1991)	matrix clause	matrix attachment has 1 argument missing a predicate; relative clause attachment has 2 arguments missing a predicate
Semantic selection	matrix clause	<i>because</i> or <i>although</i> clause selects antecedent matrix clause; relative clause not selected for
Discourse Accommodation (Crain & Steedman 1985)	matrix clause	In null context there's no contrast set for <i>the house</i> , so relative clause avoided
Local Attachment	relative clause	closer association to most recently attached lexical material

Table 1: predicted preferences in the matrix/embedded-relative ambiguity

Materials for both experiments followed the template in (6), with some minor additional changes in each experiment. In the ambiguous conditions disambiguation was provided by the presence or absence of a pronoun in object position of the ambiguous clause.

The unambiguous relative clause condition contained the overt complementizer *that* to signal the presence of the relative clause; the unambiguous matrix clause condition contained an NP with a possessor, which blocks restrictive modification.

- (6)
- a. Relative clause, ambiguous  
Although Erica hated the house she had owned for ten years since her grandmother died she had no desire to sell it.
  - b. Relative clause, unambiguous  
Although Erica hated the house that she had owned for ten years since her grandmother died she had no desire to sell it.
  - c. Matrix clause, ambiguous  
Although Erica hated the house she had owned it for ten years since her grandmother died and she had no desire to sell it.
  - d. Matrix clause, unambiguous  
Although Erica hated her house she had owned it for ten years since her grandmother died and she had no desire to sell it.

Figures 2–3 show the predicted results, according to which structure subjects prefer. If subjects opt for the matrix clause attachment, then a slowdown is expected at the point of disambiguation in the ambiguous relative clause condition; if subjects opt for the relative clause attachment, then a slowdown is expected at the point of disambiguation in the ambiguous matrix clause condition.




Figure 2: Matrix clause attachment

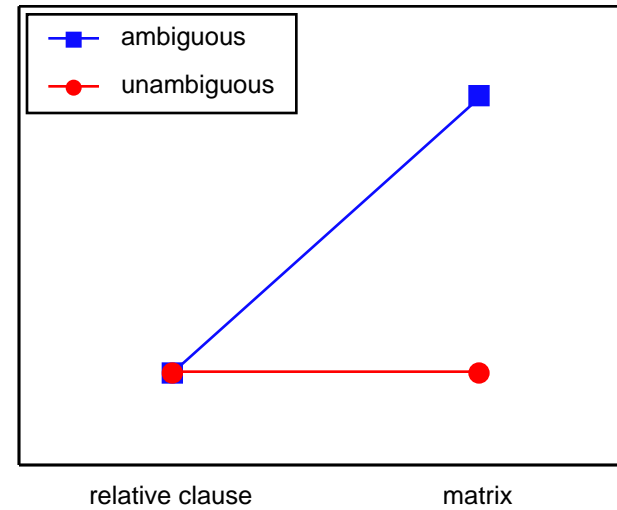


Figure 3: Local attachment

Both experiments described here used a word-by-word self paced reading paradigm with a moving window display. All results are based on residual reading times, which eliminate variability due to differing reading speeds across subjects and effects of word length. Each experiment used as subjects a different set of 48 MIT undergraduate students.



In the other half of our items the sentence-initial embedded clause described a temporal relation and the ambiguous NP was a full NP: these items showed the reverse effect: there was a significant slowdown following disambiguation in the ambiguous relative clause condition. [First word after disambiguation:  $F1(1,45) = 4.57, p < 0.05$ .  $F2(1,11) = 6.41, p < 0.05$ ]

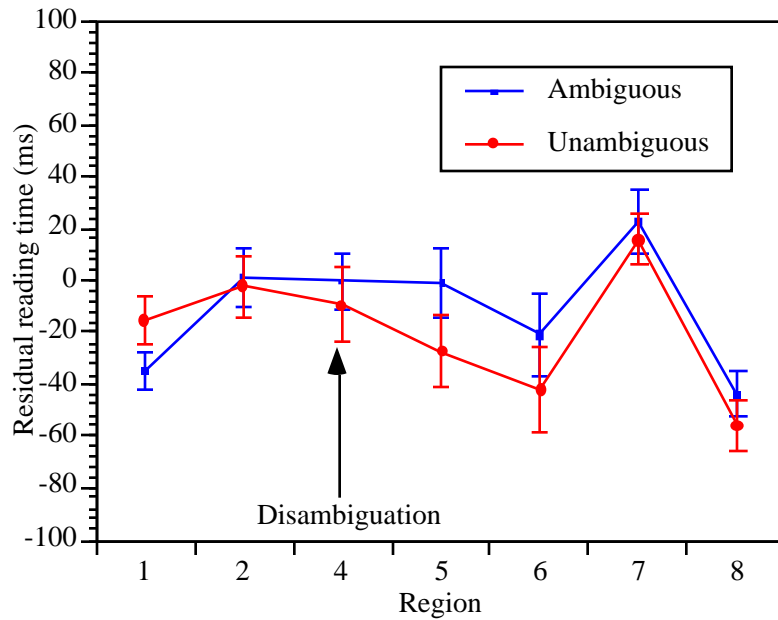


Figure 6: matrix clause, -pronominal/+temporal




Figure 7: relative clause, -pronominal/+temporal

Therefore the predictions of a strong locality bias were only partially confirmed. Moreover, since the factors  $\pm$ pronominal and  $\pm$ temporal were almost entirely correlated in the stimuli, it was impossible to determine which of these factors was responsible for the modulation of the local attachment preference.

Further work was clearly required in order to determine the scope of the local attachment preference, and to investigate the cause of the effects of NP-type and/or clause-type found in this experiment.

## 4. New experiment

This experiment aimed to clarify some of the conclusions from experiment 1. Two main changes were made to the stimuli: (a) The factors  $\pm$ pronominal and  $\pm$ temporal were independent factors; (b) the distance between the point of disambiguation and the next clause boundary was extended, to reduced potential contamination by later ambiguities in the materials. There were 32 items in the experiment.

- (7)
- a. **Relative clause, ambiguous: –pronominal/–temporal**  
Although Erica hated the house her father had owned for ten years since her grandmother died she had no desire to sell it.
  - b. **Relative clause, unambiguous: –pronominal/–temporal**  
Although Erica hated the house that her father had owned for ten years since her grandmother died she had no desire to sell it.
  - c. **Matrix clause, ambiguous: –pronominal/–temporal**  
Although Erica hated the house her father had owned it for ten years since her grandmother died and she had no desire to sell it.
  - d. **Matrix clause, unambiguous: –pronominal/–temporal**  
Although Erica hated her house her father had owned it for ten years since her grandmother died and she had no desire to sell it.

The results of this experiment conformed more closely to the predictions of a strong local attachment preference. There was a significant slowdown at disambiguation in the ambiguous matrix clause conditions, but not in the ambiguous relative clause conditions (figures 8–9).

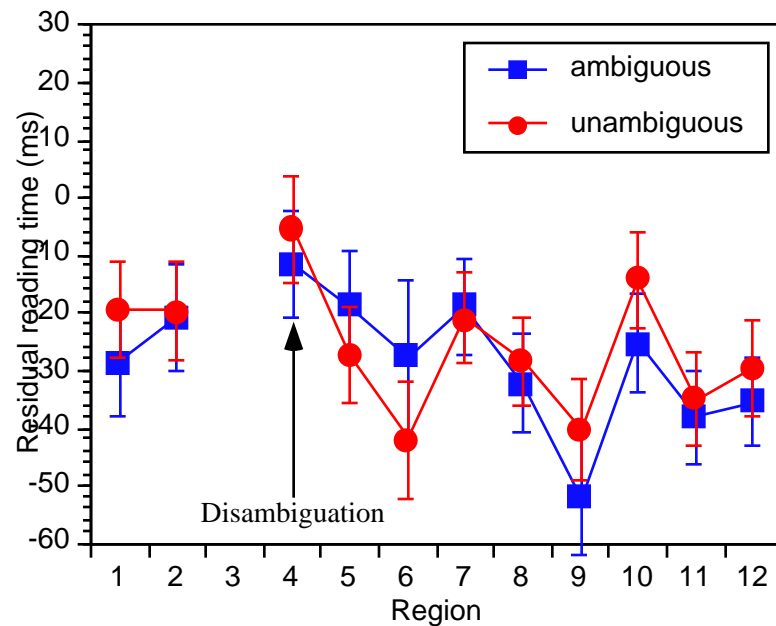


Figure 8: relative clause

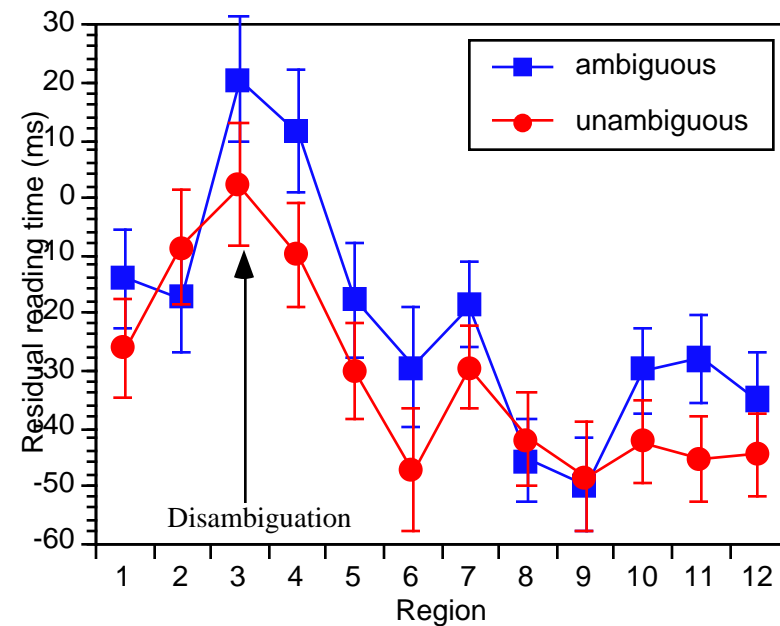


Figure 9: matrix clause



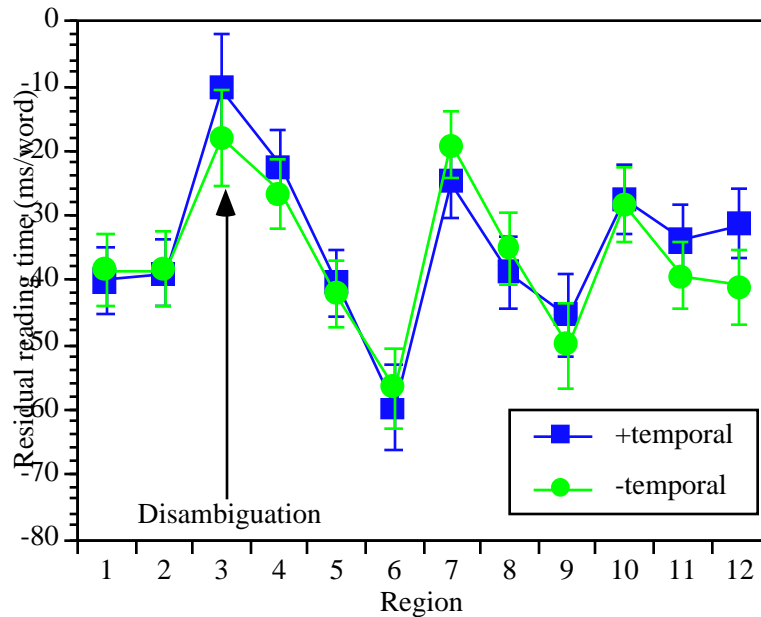


Figure 11: ±temporal conditions




Figure 12: ±pronominal conditions

Therefore, this study provides improved support for our original hypothesis, but fails to provide clear insight into the cause of the additional complications in the earlier study. One possibility is that the faster reading times in the +pronominal conditions could have contributed to a local attachment bias in the first experiment, but this remains to be confirmed.

## 5. A semantic explanation?

Existing accounts of the role of semantic/discourse factors in parsing predict a bias for the non-local matrix clause attachment: in the absence of a context providing a relevant contrast set for an NP, relative clause attachments are avoided (Crain & Steedman 1985, Altmann & Steedman 1988). The reasoning is that a restrictive relative clause is only felicitous if there is a set of alternative referents for the head noun which need to be chosen among.

Note, however, that it may be possible to give a semantic account of the local attachment effect if we adopt a slightly different view from Crain & Steedman.

### *Licensing condition on definite NPs*

Definite NPs must find a unique referent in the discourse model in which the NP is interpreted.

Under this view, relative clauses may serve either (a) to choose among multiple possible referents for the head noun in the discourse model, or (b) they may serve to provide an anchor in the discourse for an NP which otherwise would find no referent at all in the discourse model. This is why it is entirely felicitous to use an NP like (4) in the absence of any anchoring context:

(8) The surfer I met on the pier this morning ...

Once this function of relative clauses is taken into account, the possibility opens up that the local attachment bias observed in our experiments is due to the need for the definite NPs in the stimuli to provide their own anchors in the discourse model. Building a relative clause could achieve this.

We are currently investigating the interesting prediction of this approach that the local attachment preference will disappear if we alter our materials such that the first NP of the ambiguous NP-NP sequence is indefinite and the sentence-initial subordinate clause is generic.

(9) Although Erica has always hated dogs she owns...

## 6. A frequency bias?

Can corpus frequency data account for the observed local attachment preference? We conducted a number of different searches of the parsed Brown corpus (Penn Treebank I) to test this possibility.

A. Ambiguous NP-NP sequences in sentence-initial subordinate clauses headed by *because*, *since*, *when*, *after* etc., displaying the ambiguity tested in our experiments.

[Data based on hand-searched sample of 37% of parsed Brown corpus.]

B. Ambiguous NP-NP sequences in any VP.

C. NP-NP sequences in any sentential position.

D. Ambiguous NP-NP sequences in sentence-initial subordinate clauses, as in (A), but excluding pairs of NPs separated by a comma.

E. Ambiguous NP-NP sequences in any VP, as in (B), excluding pairs of NPs separated by a comma.

F. NP-NP sequences in any sentential position, as in (C), excluding pairs of NPs separated by a comma.

The results of our six different searches, conducted at different grains of analysis, yielded varying counts, as figure 13 shows.

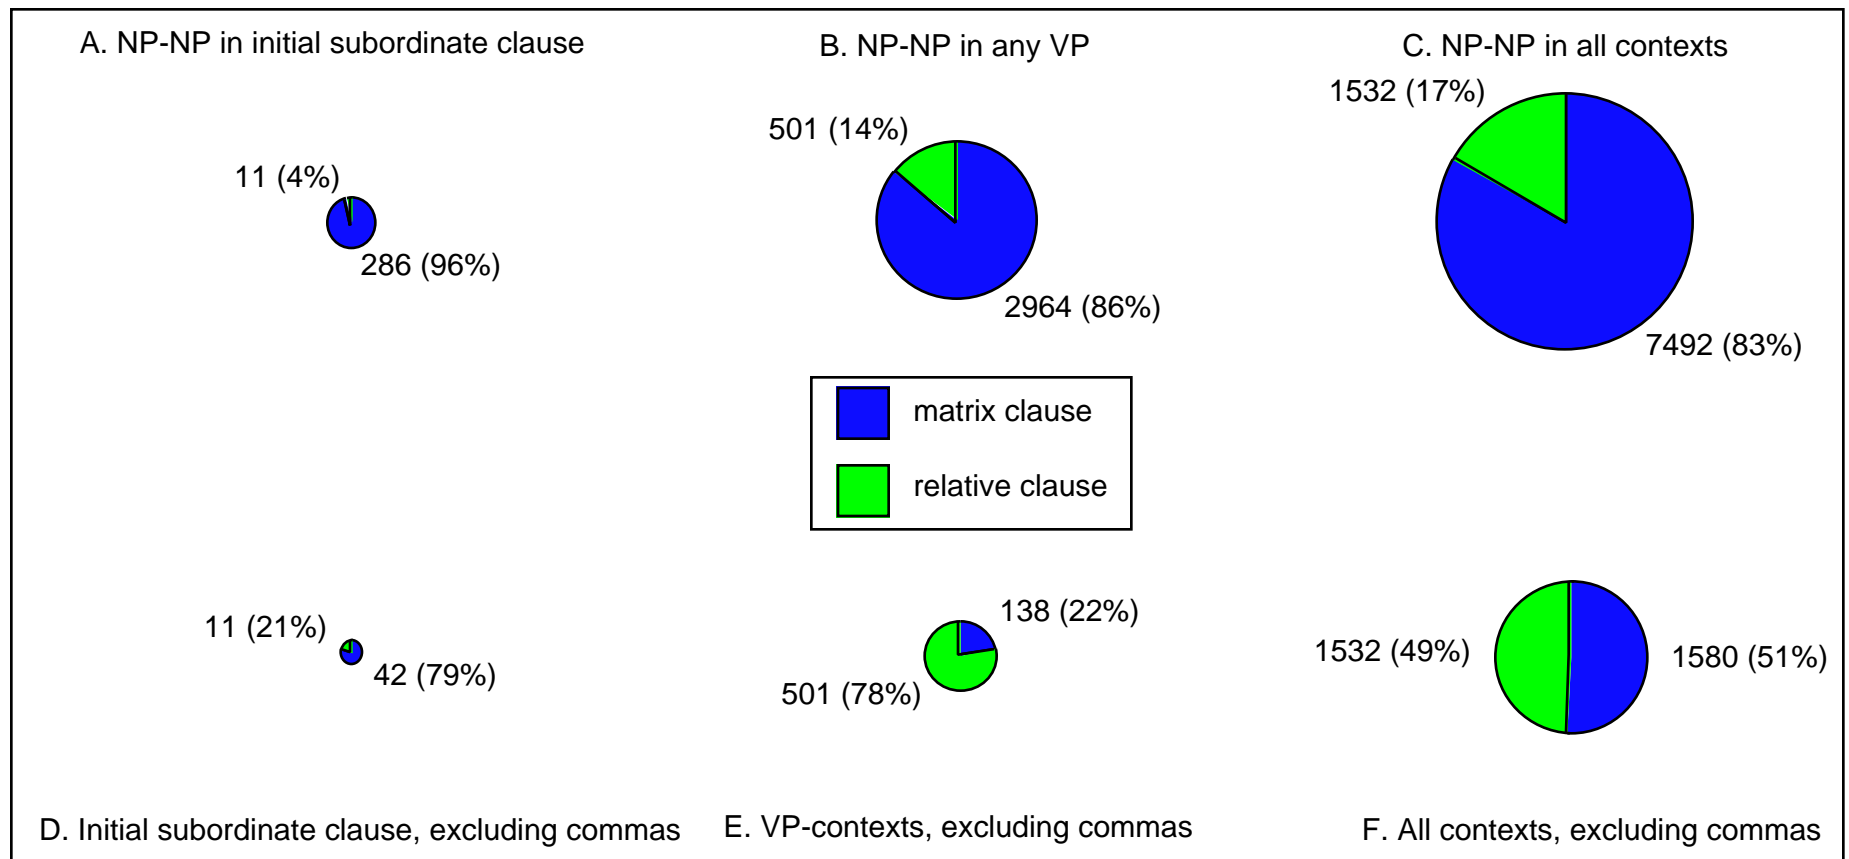


Figure 13: Multi-grained frequency analyses  
(Pie sizes reflect overall frequency in corpus of ~100,000 sentences)

In 4 of the 6 searches there was an overwhelming frequency bias towards the matrix clause resolution of the ambiguity; at one grain of analysis (F) there was an almost identical frequency of matrix and relative clause resolutions. Only at grain of analysis (E) was there a frequency bias towards the relative clause resolution of the ambiguity.

Given that little is known about the grains of analysis that speakers are most sensitive to in sentence-level frequency biases, we cannot exclude a frequency-based account of the local attachment preference. However, it should be clear from figure 13 that structural frequency information can *only* account for the observed local attachment bias if it is assumed that speakers selectively ignore most structural frequency information in the environment and attend to just one specific grain of analysis (E).

## 7. Conclusion

Our experiments support the existence of a preference to build a relative clause in the matrix/embedded-relative ambiguity. To our knowledge this is the first structure where a preference for a relative clause has been found, without additional manipulations of lexical or contextual factors.

The findings are consistent with the view that locality is a strong structural constraint in parsing. Alternative accounts of our findings, based on semantics or frequency information are still also possible, provided that they adopt certain very specific assumptions, as we have outlined.

### **Acknowledgment**

We gratefully acknowledge the help of San Tunstall and Ariel Salomon in the running of these experiments; CP was partially supported by an NSF RTG grant to MIT, and by the JST-sponsored “Mind Articulation Project” at MIT.