RIGHT ASSOCIATION REVISITED *

Michael Niv

Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA, USA
niv@linc.cis.upenn.edu

Abstract

Consideration of when Right Association works and when it fails lead to a restatement of this parsing principle in terms of the notion of heaviness. A computational investigation of a syntactically annotated corpus provides evidence for this proposal and suggest circumstances when RA is likely to make correct attachment predictions.

1 Introduction

Kimball (1973) proposes the parsing strategy of Right Association (RA). RA resolves modifiers attachment ambiguities by attaching at the lowest syntactically permissible position along the right frontier. Many authors (among them Wilks 1985, Schubert 1986, Whittemore et al. 1990, and Weischedel et al. 1991) incorporate RA into their parsing systems, yet none rely on it solely, integrating it instead with disambiguation preferences derived from word/constituent/concept co-occurrence based. On its own, RA performs rather well, given its simplicity, but it is far from adequate: Whittemore et al. evaluate RA's performance on PP attachment using a corpus derived from computer-mediated dialog. They find that RA makes correct predictions 55% of the time. Weischedel et al., using a corpus of news stories, report a 75% success rate on the general case of attachment using a strategy Closest Attachment which is essentially RA. In the work cited above, RA plays a relatively minor role, as compared with co-occurrence based preferences.

The status of RA is very puzzling. consider:

(1) a. John said that Bill left yesterday.
     b. John said that Bill will leave yesterday.

(2) In China, however, there isn't likely to be any silver lining *because the economy remains guided primarily by the state.*
     (from the Penn Treebank corpus of Wall Street Journal articles)

On the one hand, many naive informants do not see the ambiguity of 1a and are often confused by the putatively semantically unambiguous 1b – a strong RA effect. On the other hand (2) violates RA with impunity. What is it that makes RA operate so strongly in 1 but disappear in 2? In this paper I argue that it is an aspect of the declarative linguistic competence that is operating here, not a principle of parsing.

2 Heaviness

Quirk et al. (1985) define *end weight* as the tendency to place material with more information content after material with less information content. This notion is closely related with *end focus* which is stated in terms of importance of the contribution of the constituent, (not merely the quantity of lexical material.) These two principles operate in an additive fashion. Quirk et al. use heaviness to account for a variety of phenomena, among them:

- genitive NPs: the shock of his resignation,
  * his resignation's shock.

- it-extraposition: It bothered me that she left quickly. ? That she left quickly bothered me.

Heaviness clearly plays a role in modifier attachment, as shown in table 1. My claim is that what is wrong with sentences such as (1) is the violation, in the high attachment, of the principle of end weight. While violations of the principle of end weight in unambiguous sentences (e.g. those in table 1) cause little grief, as they are easily accommodated by the hearer, the on-line decision process of disambiguation could well be much more sensitive to small differences in the degree of violation. In particular, it would seem that in (1)b,

	John sold it today.	*	John sold today it.
	John sold the newspapers today.	?	John sold today the newspapers.
	John sold his rusty socket-wrench set today.		John sold today his rusty socket-wrench set.
?	John sold his collection of 45RPM Elvis records today.		John sold today his collection of 45RPM Elvis records.
?	John sold his collection of old newspapers from before the Civil War today.		John sold today his collection of old newspapers from before the Civil War.

Table 1: Illustration of heaviness and word order

the heaviness-based preference for low attachment has a chance to influence the parser before the inference-based preference for high attachment.

The precise definition of heaviness is an open problem. It is not clear whether *end weight* and *end focus* adequately capture all of its subtlety. For the present study I approximate heaviness by easily computable means, namely the presence of a clause within a given constituent.

3 A study

The consequence of my claim is that light adverbials cannot be placed after heavy VP arguments, while heavy adverbials are not subject to such a constraint. When the speaker wishes to convey the information in (1)a, there are other word-orders available, namely,

(3) a. Yesterday John said that Bill left.
   b. John said yesterday that Bill left.

If the claim is correct then when a short adverbial modifies a VP which contains a heavy argument, the adverbial will appear either before the VP or between the verb and the argument. Heavy adverbials should be immune from this constraint.

To verify this prediction, I conducted an investigation of the Penn Treebank corpus of about 1 million words syntactically annotated text from the Wall Street Journal. Unfortunately, the corpus does not distinguish between arguments and adjuncts – they're both annotated as daughters of VP. Since at this time, I do not have a dictionary-based method for distinguishing (VP asked ($\overline{S}$ when...)) from (VP left ($\overline{S}$ when...)), my search cannot include all adverbials, only those which could never (or rarely) serve as arguments. I therefore restricted my search to subgroups of the adverbials.

1. $\overline{S}$s whose complementizers participate overwhelmingly in adjuncts: *after although as because before besides but by despite even lest meanwhile once provided should since so though unless until upon whereas while.*

2. single word adverbials: *now however then already here too recently instead often later once yet previ-*

ously especially again earlier soon ever first indeed sharply largely usually together quickly closely directly alone sometimes yesterday

The particular words were chosen solely on the basis of frequency in the corpus, without 'peeking' at their word-order behavior[1].

For arguments, I only considered NPs and $\overline{S}$s with complementizer *that*, and the zero complementizer.

The results of this investigation appear the following table:

adverbial:	single word		clausal	
arg type	pre-arg	post-arg	pre-arg	post-arg
light	760	399	13	597
heavy	267	5	7	45
total	1027	404	20	642

Of 1431 occurrences of single word adverbials, 404 (28.2%) appear after the argument. If we consider only cases where the verb takes a heavy argument (defined as one which contains an S), of the 273 occurrences, only 5 (1.8%) appear after the argument. This interaction with heaviness of the argument is statistically significant ($\chi^2 = 115.5, p < .001$).

Clausal adverbials tend to be placed after the verbal argument: only 20 out of the 662 occurrences of clausal adverbials appear at a position before the argument of the verb. Even when the argument is heavy, clausal adverbials appear on the right: 45 out of a total of 52 clausal adverbials (86.5%).

(2) and (4) are two examples of RA-violating sentences which I have found.

(4) Bankruptcy specialists say Mr. Kravis set a precedent for putting new money in sour LBOs *recently* when KKR restructured foundering Seaman Furniture, doubling KKR's equity stake.

To summarize: light adverbials tend to appear before a heavy argument and heavy adverbials tend to appear after it. The prediction is thus confirmed.

---

[1] Each adverbial can appear in at least one position before the argument to the verb (sentence initial, preverb, between verb and argument) and at least one post-verbal-argument position (end of VP, end of S).

286

RA is at a loss to explain this sensitivity to heaviness. But even a revision of RA, such as the one proposed by Schubert (1986) which is sensitive to the size of the modifier and of the modified constituent, would still require additional stipulation to explain the apparent conspiracy between a *parsing* strategy and tendencies in *generator* to produce sentences with the word-order properties observed above.

4 Parsing

How can we exploit the findings above in our design of practical parsers? Clearly RA seems to work extremely well for single word adverbials, but how about clausal adverbials? To investigate this, I conducted another search of the corpus, this time considering only ambiguous attachment sites. I found all structures matching the following two low-attached schemata[2]

low VP attached: [vp ... [s * [vp * adv *] *] ...]
low S attached: [vp ... [s * adv *] ...]

and the following two high-attached schemata

high VP attached: [vp v * [... [s]] adv *]
high S attached: [s * [... [vp ... [s]]] adv *]

The results are summarized in the following table:

adverb-type	low-attached	high-att.	% high.
single word	1116	10	0.8%
clausal	817	194	19.2%

As expected, with single-word adverbials, RA is almost always right, failing only 0.8% of the time. However, with clausal adverbials, RA is incorrect almost one out of five times.

5 Toward a Meaning-based account of Heaviness

At the end of section 3 I stated that a declarative account of the ill-formedness of a heavy argument followed by a light modifier is more parsimonious than separate accounts for parsing preferences and generation preferences. I would like to suggest that it is possible to formalize the intuition of 'heaviness' in terms of an aspect of the meaning of the constituents involved, namely their givenness in the discourse.

Given entities tend to require short expressions (typically pronouns) for reactivation, whereas new entities tend to be introduced with more elaborated expressions.

---

[2] By * I mean match 0 or more daughters. By [x ... [y]] I mean constituent x contains constituent y as a rightmost descendant. By [x ... [y] ...] I mean constituent x contains constituent y as a descendant.

In fact, it is possible to manipulate heaviness by changing the context. For example, (1)b is natural in the following dialog[3] (when appropriately intoned)

A: John said that Bill will leave next week, and that Mary will go on sabbatical in September.
B: Oh really? When announce all this?
A: He said that Bill will leave yesterday, and he told us about Mary's sabbatical this morning.

6 Conclusion

I have argued that the apparent variability in the applicability of Right Association can be explained if we consider the heaviness of the constituents involved. I have demonstrated that in at least one written genre, light adverbials are rarely produced after heavy arguments – precisely the configuration which causes the strongest RA-type effects. This demarcates a subset of attachment ambiguities where it is quite profitable to use RA as an approximation of the human sentence processor.

The work reported here considers only a subset of the attachment data in the corpus. The corpus itself represents a very narrow genre of written discourse. For the central claim to be valid, the findings must be replicated on a corpus of naturally occurring spontaneous speech. A rigorous account of heaviness is also required. These await further research.

References

[1] Kimball, John. Seven Principles of Surface Structure Parsing. *Cognition 2(1)*. 1973.

[2] Quirk, Randolph, Sidney Greenbaum, Geoffrey Leech and Jan Svartvik. *A Comprehensive Grammar of the English Language*. Longman. London. 1985.

[3] Schubert, Lenhart. Are there Preference Trade-offs in Attachment Decisions? AAAI-86.

[4] Wilks Yorick. Right Attachment and Preference Semantics. ACL-85

[5] Weischedel, Ralph, Damaris Ayuso, R. Bobrow, Sean Boisen, Robert Ingria, and Jeff Palmucci. Partial Parsing: A Report on Work in Progress. Proceedings of the DARPA Speech and Natural Language Workshop. 1991.

[6] Whittemore, Greg, Kathleen Ferrara, and Hans Brunner. Empirical study of predictive powers of simple attachment schemes for post-modifier prepositional phrases. ACL-90.

---

[3] I am grateful to Ellen Prince for a discussion of this issue.