

Right Attachment and Preference Semantics.

Yorick Wilks
 Computing Research Laboratory
 New Mexico State University
 Las Cruces, NM 88003, USA.

ABSTRACT

The paper claims that the right attachment rules for phrases originally suggested by Frazier and Fodor are wrong, and that none of the subsequent patchings of the rules by syntactic methods have improved the situation. For each rule there are perfectly straightforward and indefinitely large classes of simple counter-examples. We then examine suggestions by Ford et al., Schubert and Hirst which are quasi-semantic in nature and which we consider ingenious but unsatisfactory. We point towards a straightforward solution within the framework of preference semantics, set out in detail elsewhere, and argue that the principal issue is not the type and nature of information required to get appropriate phrase attachments, but the issue of where to store the information and with what processes to apply it.

SYNTACTIC APPROACHES

Recent discussion of the issue of how and where to attach right-hand phrases (and more generally, clauses) in sentence analysis was started by the claims of Frazier and Fodor (1979). They offered two rules :

(i) Right Association

which is that phrases on the right should be attached as low as possible on a syntax tree, thus

JOHN BOUGHT THE BOOK THAT I HAD BEEN TRYING
 TO OBTAIN (FOR SUSAN)

which attaches to OBTAIN not to BOUGHT.

But this rule fails for

JOHN BOUGHT THE BOOK (FOR SUSAN)


which requires attachment to BOUGHT not BOOK.

A second principle was then added :

(ii) Minimal Attachment

which is that a phrase must be attached higher in a tree if doing that minimizes the number of nodes in the tree (and this rule is to take precedence over (i)).

So, in :



as part of

JOHN CARRIED THE GROCERIES (FOR MARY)

attaching FOR MARY to the top of the tree, rather than to the NP, will create a tree with one less node. Shieber (1983) has an alternative analysis of this phenomenon, based on a clear parsing model, which produces the same effect as rule (ii) by preferring longer reductions in the parsing table; i.e., in the present case, preferring VP \leftarrow V NP PP to NP \leftarrow NP PP.

But there are still problems with (i) and (ii) taken together, as is seen in :

SHE WANTED THE DRESS (ON THAT RACK)

rather than attaching (ON THAT RACK) to WANTED, as (ii) would cause.

SEMANTIC APPROACHES

(i) Lexical Preference

At this point Ford et al. (1981) suggested the use of lexical preference, which is conventional case information associated with individual verbs, so as to select for attachment PPs which match that case information. This is semantic information in the broad sense in which that term has traditionally been used in AI. Lexical preference allows rules (i) and (ii) above to be overridden if a verb's coding expresses a strong preference for a certain structure. The effect of that rule differs from system to system: within Shieber's parsing model (1983) that rule means in effect that a verb like WANT will prefer to have only a single NP to its right. The parser then performs the longest reduction it can with the strongest leftmost stack element. So, if POSITION, say, prefers two entities to its right, Shieber will obtain :

THE WOMAN WANTED THE DRESS (ON THE RACK)

and

THE WOMAN POSITIONED THE DRESS (ON THE RACK).

But this iterative patching with more rules does not work, because to every example, under every rule (i, ii and lexical preference), there are clear and simple counter-examples. Thus, there is :

JOE TOOK THE BOOK THAT I BOUGHT (FOR SUSAN)

which comes under (i) and there is

JOE BROUGHT THE BOOK THAT I LOVED (FOR SUSAN)

which Shieber's parser must get wrong and not in a way that (ii) could rescue. Under (ii) itself, there is

JOE LOST THE TICKET (TO PARIS)

which Shieber's conflict reduction rule must get wrong. For Shieber's version of lexical preference there will be problems with :

THE WOMAN WANTED THE DRESS (FOR HER DAUGHTER)

which the rules he gives for WANT must get wrong.

(ii) Schubert

Schubert (1984) presents some of the above counter-examples in an attack on syntactically based methods. He proposes a syntactico-semantic network system of what he calls preference trade-offs. He is driven to this, he says, because he rejects any system based wholly on lexically-based semantic preferences (which is part of what we here will call preference semantics, see below, and which would subsume the simpler versions of lexical preference). He does this on the grounds that there are clear cases where "syntactic preferences prevail over much more coherent alternatives" (Schubert, 1984, p.248), where by "coherent" he means interpretations imposed by semantics/pragmatics. His examples are :

(MARY SAW THE MAN WHO HAD LIVED WITH HER (WHILE ON MATERNITY LEAVE))

(where full lines show the "natural" pragmatic interpretations, and dotted ones the interpretations that Schubert says are imposed willy-nilly by the syntax). Our informants disagree with Schubert : they attach as the syntax suggests to LIVE, but still insist that the leave is Mary's (i.e. so interpreting the last clause that it contains an elided (WHILE) SHE WAS (ON....)). If that is so the example does not split off semantics from syntax in the way Schubert wants, because the issue is who is on leave and not when something was done. In such circumstances the example presents no special problems.

JOHN MET A TALL SLIM AUBURN HAired GIRL FROM MONTREAL THAT HE MARRIED (AT A DANCE)

Here our informants attach the phrase resolutely to MET as commonsense dictates (i.e. they ignore or are able to discount the built-in distance effect of the very long NP). A more difficult and interesting case arises if the last phrase is (AT A WEDDING), since the example then seems to fall withing the exclusion of an "attachment unless it yields zero information" rule deployed within preference semantics (Wilks, 1973), which is probably, in its turn, a close relative of Grice's (1975) maxim concerned with information quantity. In the

(AT A WEDDING) case, informants continue to attach to MET, seemingly discounting both the syntactic indication and the information vacuity of MARRIED AT A WEDDING.

JOHN WAS NAMED (AFTER HIS TWIN SISTER)

Here our informants saw genuine ambiguity and did not seem to mind much whether attachment or lexicalization of NAMED AFTER was preferred. Again, information vacuity tells against the syntactic attachment (the example is on the model of :

HE WAS NAMED AFTER HIS FATHER

Wilks 1973, which was used to make a closely related point), but normal gendering of names tells against the lexicalization of the verb to NAME+AFTER.

Our conclusion from Schubert's examples is the reverse of his own : these are not simple examples but very complex ones, involving distance and (in two cases) information quantity phenomena. In none of the cases do they support the straightforward primacy of syntax that his case against a generalized "lexical preference hypothesis" (i.e. one without rules (i) and (ii) as default cases, as in Ford et al.'s lexical preference) would require. We shall therefore consider that hypothesis, under the name preference semantics, to be still under consideration.

(iii) Hirst

Hirst (1984) aims to produce a conflation of the approaches of Ford et al., described above, and a principle of Crain and Steedman (1984) called The Principle of Parsimony, which is to make an attachment that corresponds to leaving the minimum number of presuppositions unsatisfied. The example usually given is that of a "garden path" sentence like :

THE HORSE RACED PAST THE BARN FELL

where the natural (initial) preference for the garden path interpretation is to be explained by the fact that, on that interpretation, only the existence of an entity corresponding to THE HORSE is to be presupposed, and that means less presuppositions to which nothing is the memory structure corresponds than is needed to opt for the existence of some THE HORSE RACED PAST THE BARN. One difficulty here is what it is for something to exist in memory: Crain and Steedman themselves note that readers do not garden path with sentences like :

CARS RACED AT MONTE CARLO FETCH HIGH PRICES AS COLLECTOR'S ITEMS

but that is not because readers know of any particular cars raced at Monte Carlo. Hirst accepts from (Winograd 1972) a general Principle of Referential Success (i.e. to actual existent entities), but the general unsatisfactoriness of restricting a system to actual entities has long been known, for so much of our discourse is about possible and virtual ontologies (for a full discussion of this aspect of Winograd, see Ritchie 1973).

The strength of Hirst's approach is his attempt to reduce the presuppositional metric of Crain and Steedman to criteria manipulable by basic semantic/lexical codings, and particularly the contrast of definite and indefinite articles. But the general determination of categories like definite and indefinite is so shaky (and only indirectly related to "the" and "a" in English), and cannot possibly bear the weight that he puts on it as the solid basis of a theory of phrase attachment.

So, Hirst invites counter-examples to his Principle of Referential Success (1984, p.149) adapted from Winograd: "a non-generic NP presupposes that the thing it describes exists.....an indefinite NP presupposes only the plausibility of what it describes." But this is just not so in either case :

THE PERPETUAL MOTION MACHINE IS THE BANE OF LIFE IN A PATENT OFFICE

A MAN I JUST MET LENT ME FIVE POUNDS

The machine is perfectly definite but the perpetual motion machine does not exist and is not presupposed by the speaker. We conclude that these notions are not yet in a state to be the basis of a theory of PP attachment. Moreover, even though beliefs about the world must play a role in attachment in certain cases, there is, as yet, no reason to believe that beliefs and presuppositions can provide the material for a basic attachment mechanism.

(iv) Preference Semantics

Preference Semantics has claimed that appropriate structurings can be obtained using essentially semantic information, given also a rule of preferring the most densely connected representations that can be constructed from such semantic information (Wilks 1975, Fass & Wilks 1983).

Let us consider such a position initially expressed as semantic dictionary information attaching to the verb; this is essentially the position of the systems discussed above, as well as of case grammar and the semantics- based parsing systems (e.g. Riesbeck 1975) that have been based on it. When discussing implementation in the last section we shall argue (as in Wilks 1976) that semantic material that is to be the base of a parsing process cannot be thought of as simply attaching to a verb (rather than to nouns and all other word senses)

In what follows we shall assume case predicates in the dictionary entries of verbs, nouns etc. that express part of the meaning of the concept and determine its semantic relations. We shall write as [OBTAIN] the abbreviation of the semantic dictionary entry for OBTAIN, and assume that the following concepts contain at least the case entries shown (as case predicates and the types of argument fillers) :

[OBTAIN]	(recipient hum)	recipient case, human.
[BUY]	(recipient hum)	recipient case, human.
[POSITION]	(location *pla)	location case, place.
[BRING]	(recipient human)	recipient case, human.
[TICKET]	(direction *pla)	direction case, place.
[WANT]	(object *physob)	object case, physical object.
	(recipient hum)	recipient case, human.

The issue here is whether these are plausible preferential meaning constituents: e.g. that to obtain something is to obtain it for a recipient;

to position something is to do it in association with a place; a ticket (in this sense i.e. "billet" rather than "ticket" in French) is a ticket to somewhere, and so on. They do not entail restrictions, but only preferences. Hence, "John brought his dog a bone" in no way violates the coding [BRING]. We shall refer to these case constituents within semantic representations as semantic preferences of the corresponding head concept.

A FIRST TRIAL ATTACHMENT RULE

The examples discussed are correctly attached by the following rule :

Rule A : moving leftwards from the right hand end of a sentence, assign the attachment of an entity X (word or phrase) to the first entity to the left of X that has a preference that X satisfies; this entails that any entity X can only satisfy the preference of one entity. Assume also a push down stack for inserting such entities as X into until they satisfy some preference. Assume also some distance limit (to be empirically determined) and a DEFAULT rule such that, if any X satisfies no preferences, it is attached locally, i.e. immediately to its left.

Rule A gets right all the classes of examples discussed (with one exception, see below): e.g

JOHN BROUGHT THE BOOK THAT I LOVED (FOR MARY)

JOHN TOOK THE BOOK THAT I BOUGHT (FOR MARY)

JOHN WANTED THE DRESS (ON THE RACK) (FOR MARY)

where the last requires use of the push-down stack. The phenomenon treated here is assumed to be much more general than just phrases, as in:

PÂTÉ DE CANARD TRUFFÉ

(i.e. a truffled pate of duck, not a pate of truffled ducks!) where we envisage a preference (POSS STUFF) — i.e. prefers to be predicated of substances — as part of [TRUFFÉ]. French gender is of no use here, since all the concepts are masculine.

This rule would of course have to be modified for many special factors, e.g. pronouns, because of :

SHE WANTED { THE DRESS
IT } (ON THE SHELF)

A more substantial drawback to this substitution of a single semantics- based rule for all the earlier syntactic complexity is that placing the preferences essentially in the verbs (as did the systems discussed earlier that used lexical preference) and having little more than semantic type information on nouns (except in cases like [TICKET] that also prefers associated cases) but, most importantly, having no semantic preferences associated with prepositions that introduce phrases, we shall only succeed with rule A by means of a semantic subterfuge for a large and simple class of cases, namely:

JOHN LOVED HER (FOR HER BEAUTY)

or

JOHN SHOT THE GIRL (IN THE PARK)

Given the "low default" component of rule A, these can only be correctly attached if there is a very general case component in the verbs, e.g. some statement of location in all "active types" of verbs (to be described by the primitive type heads in their codings) like SHOOT i.e. (location *pla), which expresses the fact that acts of this type are necessarily located. (location *pla) is then the preference that (IN THE PARK) satisfies, thus preventing a low default.

Again, verbs like LOVE would need a (REASON ANY) component in their coding, expressing the notion that such states (as opposed to actions, both defined in terms of the main semantic primitives of verbs) are dependent on some reason, which could be anything.

But the clearest defect of Rule A (and, by implication, of all the verb-centered approaches discussed earlier in the paper) is that verbs in fact confront not cases, but PPs fronted by ambiguous prepositions, and it is only by taking account of their preferences that a general solution can be found.

PREPOSITION SEMANTICS: PREPLATES

In fact rule A was intentionally naive: it was designed to demonstrate (as against Shubert's claims in particular) the wide coverage of the data of a single semantics-based rule, even if that required additional, hard to motivate, semantic information to be given for action and states. It was stated in a verb-based lexical preference mode simply to achieve contrast with the other systems discussed.

For some years, it has been a principle of preference semantics (e.g. Wilks 1973, 1975) that attachment relations of phrases, clauses etc. are to be determined by comparing the preferences emanating from all the entities involved in an attachment: they are all, as it were, to be considered as objects seeking other preferred classes of neighbors, and the best fit, within and between each order of structures built up, is to be found by comparing the preferences and finding a best mutual fit. This point was made in (Wilks 1976) by contrasting preference semantics with the simple verb-based requests of Riesbeck's (1975) MARGIE parser. It was argued there that account had to be taken of both the preferences of verbs (and nouns), and of the preferences cued from the prepositions themselves.

Those preferences were variously called paraplates (Wilks 1975), preplates (Boguraev 1979) and they were, for each preposition sense, an ordered set of predication preferences restricted by action or noun type. (Wilks 1975) contains examples of ordered paraplate stacks and their functioning, but in what follows we shall stick to the preplate notation of (Huang 1984b).

We have implemented in CASSEX (see Wilks, Huang and Fass, 1985) a range of alternatives to Rule A: controlling both for "low" and "high" default; for examination of verb preferences first (or more generally those of any entity which is a candidate for the root of the attachment, as opposed to what is attached) and of what-is-attached first (i.e. prepositional phrases). We can also control for the application of a more redundant form of rule where we attach preferably on the conjunction of satisfactions of the preferences of the root and the attached (e.g. for such a rule, satisfaction would require both that the verb preferred a prepositional phrase of such a class, and that the prepositional phrase preferred a verb of such a class).

In (Wilks, Huang & Fass 1985) we describe the algorithm that best fits the data and alternates between the use of semantic information attached to verbs and nouns (i.e. the roots for attachments as in Rule A) and that of prepositions; it does this by seeking the best mutual fit between them, and without any fall back to default syntactic rules like (i) and (ii).

This strategy, implemented within Huang's (1984a, 1984b) CASSEX program, correctly parses all of the example sentences in this paper. CASSEX, which is written in Prolog on the Essex GEC-63, uses a definite clause grammar (DCG) to recognize syntactic constituents and Preference Semantics to provide their semantic interpretation. Its content is described in detail in (Wilks, Huang & Fass 1985) and it consists in allowing the preferences of both the clause verbs and the prepositions themselves to operate on each other and compete in a perspicuous and determinate manner, without recourse to syntactic preferences or weightings.

REFERENCES

- Boguraev, B.K. (1979) "Automatic Resolution of Linguistic Ambiguities." Technical Report No.11, University of Cambridge Computer Laboratory, Cambridge.
- Crain, S. & Steedman, M. (1984) "On Not Being Led Up The Garden Path: The Use of Context by the Psychological Parser." In D.R. Dowty, L.J. Karttunen & A.M. Zwicky (Eds.), *Syntactic Theory and How People Parse Sentences*, Cambridge University Press.
- Fass, D.C. & Wilks, Y.A. (1983) "Preference Semantics, Ill-Formedness and Metaphor," *American Journal of Computational Linguistics*, 9, pp. 178-187.
- Ford, M., Bresnan, J. & Kaplan, R. (1981) "A Competence-Based Theory of Syntactic Closure." In J. Bresnan (Ed.), *The Mental Representation of Grammatical Relations*, Cambridge, MA: MIT Press.
- Frazier, L. & Fodor, J. (1979) "The Sausage Machine: A New Two-Stage Parsing Model." *Cognition*, 8, pp.191-325.
- Grice, H. P. (1975) "Logic & Conversation." In P. Cole & J. Morgan (Eds.), *Syntax and Semantics 3: Speech Acts*, Academic Press, pp. 41-58.
- Hirst, G. (1983) "Semantic Interpretation against Ambiguity." Technical Report CS-83-25, Dept. of Computer Science, Brown University.
- Hirst, G. (1984) "A Semantic Process for Syntactic Disambiguation." *Proc. of AAAI-84*, Austin, Texas, pp. 148-152.
- Huang, X-M. (1984a) "The Generation of Chinese Sentences from the Semantic Representations of English Sentences." *Proc. of International Conference on Machine Translation*, Cranfield, England.
- Huang, X-M. (1984b) "A Computational Treatment of Gapping, Right Node Raising & Reduced Conjunction." *Proc. of COLING-84*, Stanford, CA., pp. 243-246.
- Riesbeck, C. (1975) "Conceptual Analysis." In R. C. Schank (Ed.), *Conceptual Information Processing*, Amsterdam: North Holland.
- Ritchie, G. (1978) *Computational Grammar*. Hassocks: Harvester.
- Shieber, S.M. (1983) "Sentence Disambiguation by a Shift-Reduced Parsing Technique." *Proc. of IJCAI-83*, Karlsruhe, W. Germany, pp. 699-703.
- Shubert, L.K. (1984) "On Parsing Preferences." *Proc. of COLING-84*, Stanford, CA., pp. 247-250.
- Wilks, Y.A. (1973) "Understanding without Proofs." *Proc. of IJCAI-73*, Stanford, CA.
- Wilks, Y.A. (1975) "A Preferential Pattern-Seeking Semantics for Natural Language Inference." *Artificial Intelligence*, 6, pp. 53-74.
- Wilks, Y.A. (1976) "Processing Case." *American Journal of Computational Linguistics*, 56.
- Winograd, T. (1972) *Understanding Natural Language*. New York: Academic Press.