Simplifying Analyses by Fitting the Semantic Representation Closely to Language

David D. McDonald
Brandeis University
14 Brantwood Road, Arlington MA 02174-8004
mcdonald@cs.brandeis.edu

Abstract

The best prospects for robust, extensible Natural Language Understanding systems will come through the use of a semantics with a sound fit to way language is organized into units. Many single words correspond to relations rather than terms; many full phrases (maximal projections) correspond to lambda abstractions rather than wffs. In this paper we explore the consequences of some of these common but syntactically unconventional phrase types on the representations that could be used to represent their meanings.

These phrases—conjunction reductions without semantically natural heads, and displaced constituents that must be given interim representations until the parse reaches a point where they can be completed—are common in real texts, and must be accounted for in any successful language understanding system. We use them to illustrate some of the unusual properties of the representational system we work with, including first class representations of the links between relations and their arguments and first class representations of partially saturated relations.

1. Introduction

The work described here has been done as part of the Sparser natural language comprehension system and its knowledge representation system Krisp (McDonald 1992, 1994a). Sparser is designed for the task of information extraction, which is to say that its goal is to locate and represent specific, literal information wherever it may appear in an unrestricted corpus of texts such as newswire articles, while systematically ignoring whatever other information the texts might contain that is outside the domain of interest.

This is a approachable task today, and one that allows us to begin to study real texts: texts written by professional writers to be read by people with a serious interest in knowing what they have to say. One thing that we quickly learned about these texts is that even when we try to restrict them to a sub-language they exhibit a strikingly large variety of phrasal patterns, far more than had been considered a 'large' set in the past when developing grammars for restricted, short-utterance tasks like question answering.

We expect that it will not be possible to develop grammars and interpretation functions for these texts using the style that has typified work in the past. We cannot expect to succeed if what we do is attempt to catalog all the individual syntactic constructions and give each of them their own compositional, rule-by-rule interpretation. The enormous number of constructions and the variations imposed on them by the peculiarities of the semantics of diferent words will swamp any such effort.

Instead, we think that what is necessary is to move ever more of the work into operations in semantics rather than syntax. By developing more powerful compositional mechanisms and more expressive and flexible representations, it should be possible to keep pace with the increase in variations that we see at the surface of the language now that we are starting to deal seriously with real texts.

In this paper we will look at a number of new problems that we have been studying that seem to us to speak to the issues of this workshop. We will start with problems and move on from them to the probably unusual representation devices that we have developed to solve them. This will not be an exposition of the representation system used with Sparser (for that see McDonald 1994a), but we will touch on most of its elements in the course of the exposition.

2. Non-standard constituents

Our experience has been that the real demands on expressiveness in representations do not come from traditional problems like quantification but from trying to make sense of the abbreviated and strange constituents that one finds in actual texts, particularly in conjunctions. Consider the following text, excerpted from the joint ventures corpus of the Tipster project and representative of the texts it contains.

"The new firm will be 50 percent owned by Takarabune, 45 percent by Merciries and the rest by a Japanese advertisement concern, Cybac Co., the company said."

What does it mean to understand a text like this or any other? Variations in notation and epistemology aside, we can probably all agree that understanding requires determining all of the relationships among the individuals—express or implied—that the text references. When the understanding is complete, all of the identified relationships will be completely saturated, which is to say that all of the arguments (variables, parameters) that the relations require in order to be individuated will have been bound to individual values that were either expressly specified in the wording of the text or arrived at through the application of default knowledge.[1]

The usual means of identifying and populating a domain relationship during a parse is to start with the head of the phrase as the basis of the type of object or relation to be instantiated, e.g. "*firm*" or "*owned*", and then to interpret the syntactic arguments to the head as values for the relation's arguments or the object's properties, e.g. "*new*" or "*by Takarabune*".

The problem of non-standard constituents comes (a) from constituents without semantically sensible heads, such as "*45 percent by Merciries*" (there is no sensible way that a company can directly provide an argument to a percentage); and (b) from displaced constituents.

3. Displaced constituents and partially saturated objects

A displaced constituent is one like this instance of "*50 percent*" within the full verb group "*will be* [ ] *owned*". It is 'displaced' in the sense that it is not in construction with (roughly speaking, adjacent to) the syntactic argument or head that will be needed to denote a full (saturated) relation. In the era of transformational grammar we would have said it had 'moved'; today we would be more circumspect.

Functionally, the usual reason for displacing a phrase is to realize one of the speaker's presentational goals: the use of displacement does not change the information that the utterance conveys but affects the process of how that information is assimilated by the hearer in the course of his or her incremental understanding of it.[2]

A version of this example text that did not displace the percentage phrase would be

"50 percent of the new firm will be owned by Takarabune, 45 percent by Merciries and the rest ..."

This version of the text conveys the same information in terms of its content, but its presentation is shifted in emphasis and in what is implied about the continuity or shift in topic with respect to earlier and later parts of the news article. (An even larger shift would occur in a version of the text that started with the company Takarabune; less drastic than that would be to start with something like "*The shares in the new firm will be divided ...*". All of these different patterns and more have been observed in this corpus.)

To see the implications of displaced constituents for representation, we should start by looking at the semantic elements involved. To begin with, a percentage by itself is semantically just a measurement, a kind of quantifier whose mention implies the existence of a thing being measured; compare "*a half dozen*", "*third*", or "*all*".[3]

A measurement can only be taken as a term once it has been composed with the thing being measured to yield a quantified instance of stuff. In this case this will happen once the verb group and the subject have composed (giving us 'a 50 percent share in the new firm').

Sparser is an incremental language comprehension system. It tries to determine the denotation of each word/minimal phrase at the very moment it is scanned/composed during Sparser's single beginning-to-end pass through the text.[4] This means that Sparser recognizes the percentage phrase and identifies a referent for it before it has even finished processing the verb group. Ipso facto, by the end of the verb group it will have appreciated that the

---

1 Note that this kind of representation show no particular regard for the manner in which that information had been presented: such information as what elements were thematic in the utterance; what had been presumed as already 'given' for the audience; what perspectives were imposed; what was most salient. Such *presentational information* can in principle play a role during the course of the parsing and interpretation. It would unquestionably have to be captured and represented if the goal was to be able to repeat (re-generate) the information to someone else, but it is not part of what we tend count as the 'meaning' of the text.

2 This is an area of linguistic research that some people call 'information structure', see e.g. Vallduví (1992) Prince (1988).

3 A closely related class of measurement phrases are headed by words that characterize the stuff being measured: "*10 ton*" implies a measurement of something's mass; "*10 foot tall*" implies a linear measurement of something's height; etc. As relationships, these too are arguably unsaturated, since while they do provide sometimes crucial information about the aspects of the thing being measured (cp. "*high*" "*wide*" "*deep*"), they would not stand by themselves as individuals in a denotational model of a text's meaning but only would be components of, say, 'quantified instances of stuff".

4 A similar kind of incremental parsing system is extensively described in Mellish 1985.

percentage phrase is stranded—displaced from its canonical phrasal context where the thing being measured would be part of the same phrase.

Following a general principle that any observation that cannot be acted on immediately must be given some representation so that it can be remembered and acted on later, this means that we must have some object type in our epistemological inventory that can serve this purpose, i.e. an object type that can represent the fact that we have identified an incomplete object that, from the point of view of the parser, is looking for an object to bind to its open variable.

As described in McDonald 1994, Sparser uses an object type called a '*partially-saturated individual*' for this purpose. This is a reification of the information at hand, the anticipated larger relation, and an open variable. Identifying what the larger relation and its open variable are is presently done by hand while designing the semantic model; we anticipate that it ought to be calculable from an automatic examination of the rules of the grammar.

The larger relation in this case is taken to be 'amount of stuff', a relationship between some individual (which in this case will be 'the new firm') and a measurement taken on it (the '50%' that has already been identified[5]). The open variable is the variable of the amount of stuff relation to which the individual will eventually be bound. (We elaborate on the notion of 'binding a variable' later.)

Sparser forms the needed partially saturated individual by (a) instantiating an individual with the indicated relation-type as its domain type; (b) binding its variable for the measurement to the individual that represents the percentage; and (c) posting the indicated open-variable in the appropriate place in the resulting structure.

There are any number of ways one could imagine to implement this sort of object. We chose typed structured objects. Expanding out all of the embedded objects, the printed representation of the partially saturated individual needed to capture what is known about the percentage phrase at the point when the parsing process has finished the verb group it is embedded in looks like this:

```
#<psi amount-of-stuff 1
 :type ( #<category amount-of-stuff> )
 :open-in ( #<variable amount-of-stuff.stuff> )
 :binds
   ( #<measurement =
       #<individual  percentage 1
          :type ( #<category percentage> )
          :binds
            ( #<value =
                 #<individual number 23
                    :type ( #<category
                                 number> )
                    :binds
                      ( #<value = 50>>
                        )>)>)>
```

This print form probably raises more questions than it answers. We will begin to take them up in the next section. First however we should finish the discussion of the function of this kind of object in the parsing/interpretation process.

When the parsing of the verb group is finished it will point to a compound structure as its denotation. The primary element of that denotation is the open predicate denoted by the verb "*to own*", specialized as picking out an event in the future. The secondary element is the partially saturated individual denoted by the percentage phrase given that it occurred in this context. If there were other such objects corresponding to other stranded elements or relational adverbs, they would be part of the compound as well.

The finished verb group, syntactically marked as passive, is then taken by the parser and combined with is neighbor phrases as dictated by the grammar. Being passive, its next composition is with the subject, "*the new firm*". When the phrases are combined their denotations are as well, in this case following the dictates of the rule for 'company + be owned'.

What happens to the predicate that is the primary element of the verb group's compound denotation during this composition is what one would expect: it binds the company as the thing being owned and becomes the denotation of the clause as a whole. By contrast. the activity of a secondary element is opportunistic: It will be carried up the head line of the growing phrasal analysis with each composition until the composition adds a phrase whose denotation is compatible with the constraints on its open variable. As it happens in this

5 There is a domain-level question of course about what is actually being measured here: 50% of the company's workforce, 50% of its physical plant, or what? We can form the relationship between the new firm and this measurement, but what does it actually mean? In general one cannot know this independently of the context. In this example the context is 'ownership', from which we can know (can incorporate into the semantic model) that what is being measured is the shares of stock in the company, ultimately pieces of elegantly printed paper that will be in the physical possession of the legal officers of the three company that have formed the joint venture in this new firm. Just how much of this real world detail belongs in the denotation of a phrase like "*50 percent of the new firm*" is ultimately a matter of judgment and convention. We expand it out as far as representing the notion of shares of stock (as a mass term), since in other comparable texts in this corpus the words "*stock*" and "*shares*" are often used, and we want all of the instances to be represented in comparable detail.

example, this occurs immediately with this first composition of subject and verb group.

With its target object located, the partially saturated relation binds its open variable to the object and converts to a regular, individuated relation that is added to the model containing the information so-far gleaned from the text on an equal status with the (still-growing) ownership relation, the new firm, or all the information from the earlier part of the article.

This idea of allowing partially saturated relations to be carried up the headline in parallel with the primary denotations and then taking the very first available compatible complement to be the object they are 'looking for' is experimental. We have been adding it case by case in selected examples as we develop our grammars and cannot at this point say whether it will ultimately pan out as an effective comprehension device in its present form: It might turn out, for instance, that it binds too quickly to intermediate unintended objects along the composition path; or that the detailing of the taxonomy and value restrictions needed to adequately constrain the bindings unduly baroque and unintuitive. Only further experiment will tell.[6]

4. Explicitly representing variable bindings

All representational systems form instances of object types (henceforth 'individuals') through some kind of association between what we can informally agree to call the 'parameters' of a predicate and other objects. In the printed notation of the propositional calculus, the parameters are reified as ordered locations in the parentheses following the predicate, delimited by commas, e.g.

```
partial-ownership-of-a-company
        (new-firm, Takarabune, 50%)
```

Here the identity of the different parameters is a matter of convention since no particular ordering of the parameters has any sensible priority over any other. This is carried over to the implementation of these logics in Prolog, with specific parameters implicitly and conventionally identified with positions in a sequence. And in set-theoretical models, individuals are again taken to be ordered tuples of objects. If it is useful to manipulate instances of a particular parameter, then access functions over the tuples (or other ordered representations) are defined that are equivalent to Lisp's ordinal access functions ('first', 'second', 'nth', etc.).

In such systems, the access functions are the closest one gets to actually treating parameters as objects in their own right.

In representation systems like the KL-ONE family or typed feature structures, parameters are realized as 'slots' or 'features'. They have names, are objects in the epistemology, have domains of applicability (in the sense that in KL-ONE, lower predicates ('concepts') in the taxonomy inherit slots from the predicates they inherit from), and often are grouped with other slots in their own taxonomic hierarchies.

The semantics of variable binding in such systems is essentially equivalent to the positional treatment of regular logics however. Each individual satisfying a given predicate or predicates is allocated a number of locations in the implementational equivalent of a record structure, and a binding is just the occupation of one of those locations by an object.

In our work, we have wanted a more explicit treatment of binding. It is often important to know what linguistic structure or observation was responsible for a given binding, which prompts the desire for a simple place to record such information. Also we prefer to treat a denotational model not as a bag of sets but as a semantic net: a system of nodes and links that processes can navigate as part of their reasoning. This means that one needs an object type to represent the links, and we identify this with the binding of variables.[7]

Perhaps the most familiar analog to our notion of a binding can be found in the lambda calculus. In the lambda calculus, a function is represented by a 'lambda expression' that indicates what variables—formal parameters—are available to be bound when the function is applied to arguments. A lambda expression can have some of its variables bound and some open. Bound variables are usually cached out through so called 'beta reductions' by the substitution of the value of the variable for the variable symbol throughout the formula portion of the expression. The partially saturated object for "*50 percent*" discussed above would look like this in the lambda calculus:

```
Lambda (stuff)  .
     amount-of-stuff (50%, stuff)
```

A 'function application' in the lambda calculus is an explicit representation of the binding of an argument to a variable, e.g.

6 Complementary examples and discussion of this compositional device can be found in McDonald 1994b

7 Other semantic network formalisms presumably do something analogous (for example SNePs, Shapiro & Rapaport 1992), but we are not familiar enough with what these systems look like when one goes beyond the diagrams given in papers for it to be worth going into further comparisons.

```
( Lambda (stuff)
      amount-of-stuff (50%, stuff) )
the-new-firm
```

Our notion of a 'binding'—as a type of object at the epistemological level—captures just this three-part relationship: namely the identification of a variable ('stuff'), a body or formula to which the binding will apply ('amount-of-stuff (50%, stuff)'), and the value to which the variable is to be bound ('the-new-firm').

An individual in our representational system is effectively a bag of bindings and a representation of a type (the predicates that license the bindings).

Because they are first class objects (the equivalent of how slots are treated at KL-ONE's meta-level), the implementation of our bindings allows them to be treated as links that can be traversed in either direction: from body to value or value to body. They are also required to be unique. The interpretation function that takes linguistic structures into their denotations consults indexes that require it to retrieve the very same binding object every time the same body, value, and variable are referenced.

Given this view of individuals as unordered bags of the variable bindings they are comprised of, we can perform some interesting manipulations that have turned out to be useful in the treatment of some non-standard constituents. This is the topic we will now turn to.

5. Constructing functions dynamically to solve problems in conjunction

If we look again at the text we have been using as our example, we see two examples of the other non-standard constituent types that we are concerned about in this paper: constituents without semantically sensible heads. These constituent types appear most frequently in conjunctions, but can also appear in answers to questions, or in any other context where what would once have been called a 'reduction under identity' transformation could apply.

> *"The new firm will be 50 percent owned by Takarabune, 45 percent by Merciries and the rest by a Japanese advertisement concern, Cybac Co., the company said."*

In texts written by journalists, it is quite common for full sentence conjunctions like this to have these irregular phrase combinations, e.g. *"45 percent by Merciries"*, *"the rest by a Japanese advertisement concern, Cybac Co."*. Syntactically they must be treated as constituents—two of the three constituents that are conjoined by *"and"* in this text.[8] However they cannot be given an interpretation in the usual way since there are no heads within the constituents on whose denotations to anchor the other elements. (The percentage phrases are nominally the syntactic heads of the phrases, with the agentive 'by' clauses as their adjuncts; however there is no rule in any sensible, extensible grammar that would allow us to make semantic sense out of the combination of those two phrase types—hence our assessment that they are non-standard constituents.)

Conjunctions have long been a steady source of theoretical inspiration for linguists. The stipulation that, by definition, a conjunction combines phrases of the same type has prompted some computational linguistics, notably Steedman (1989) to adopt syntactic formalisms such as categorial grammar. The flexibility in how phrasal categories in a categorial grammar can be assembled permits these linguists to employ operations such as type-raising to take the two later constituents in this example, which are ostensibly noun phrases, and convert them to clauses, thereby permitting the text to be seen as a conjunction of three clauses and parsed in the customary way.

However we are concerned not just with constructing a syntactic description of this text but with assembling an informative model of its meaning. We are required to go further and ask what is it about these independent percentages that relates them to agentive roles of the two companies they are in construction with? Even a cursory examination of the corpus will reveal that we cannot answer this question with any sort of fixed schema: There are too many variations in what can appear in the reduced, non-standard constituents (what can be elided from the first 'full' clause of the conjunction) for any such treatment to be possible. Instead we need a dynamic procedure that will formulate the needed operators case by case as the elements of particular reduced phrases and full clauses require.

To do this, we begin by assuming that all three conjoined phrases denote instances of the same complex type, and that the first conjunct, the full clause, is fully instantiated—every variable defined by the complex type has a value in the object

---

8 An issue we will not discuss here since our treatment has yet to be validated by extensive tests is exactly how the parsing process correctly identifies the scope of this conjunction. In short, the preference of the conjunction process is to link two (or more) adjacent minimal constituents that denote individuals of the same type. That is not possible here, and in such cases the process's heuristic is to go back to the most recent phrase that contains an object of the same type, i.e. one containing a percentage phrase. This takes it back to the entire first clause and thereby forces the two later constituents to be construed as containing the same amount and type of information as the first clause.

denoted by the full clause. We proceed by constructing a function from this object, and will then apply the function to the elements of the two reduced conjuncts to assemble comparable fully instantiated objects.

In constructing this function, the question is what variables it should be open in. This is the core of what makes this procedure dynamic, since we determine these variables by examining what values were explicitly given in the two reduced conjuncts.

In this example, the reduced conjuncts each provide a percentage and a company in an agentive role. To identify the variables, we back off from the identity of these individuals to their types, and then look to the object denoted by the full clause to see which of its variables are bound to objects of those types. In this case this will give us a measurement of the amount of stock in the new firm and the owner of the new firm respectively.

Constructing the function is a matter of assembling all of the bindings of the other variables in the full clause and forming abstractions from them comprised of the bindings' values and variables. These 'var+value' objects, plus the two open variables, constitute the function.

Applying the function is a matter of (1) constructing (for each of the two reduced conjuncts) the shell of a new individual, (2) giving this individual the same complex type as the full clause's, (3) reconstituting bindings from the var+value objects with this new individual as the bindings' body, and (4) creating bindings for the two open variables combining the individual and the two individuals (values) provided in the reduced conjuncts.

This is a completely general procedure, and we think it holds great promise. We have yet, however, to have done enough testing and application to assert this definitively. That will come during the course of the next few quarters.

6. Specializing individuals transparently to the relations they are part of

Type inference is a standard element of modern representational systems. All of the predicates in a system are typically organized into a subsumption lattice according to their airity and the type restrictions on their variables, and descriptions of new predicates are fit into these lattices by algorithms like 'least specific subsumer' which move a predicate up or down as more of its properties are examined. A comparable kind of inference is a crucial element in the language comprehension process.

To achieve the kind of flexibility in writing grammar/semantic rules that we need to cope with the range of variations found in real texts, it is important to be able to start with fairly general definitions of relational words, especially verbs, and to have their denotations specialized dynamically in the course of understanding a text.

The verb "*own*", for example, can start out a just a 'generalized possession' predicate relating two objects, one for each of its syntactic arguments in a transitive clause. In the example text, composing the verb in the passive with a company subject gives us 'possess (have) a company'. Further composition with the agentive by-phrase gives us 'company–possesses–company'. This last predicate is a core notion in the joint ventures domain, one that we would invariably predefined. The point, however, is that we do not need to also predefine the corresponding rule(s) of semantic interpretation provided that have a means of allowing the individuals that one constructs during a parse to have their types specialized—moved down in the lattice—as the parse proceeds and the types of their complements become apparent.

In Sparser, the type of an individual is established separately from its identity. Identity rests in the object one constructs. With the judicious use of indexes to insure that textual descriptions of the same individual are always taken to the same object, we can allow what is known about an individual to vary without having to update all of the other objects that reference this one every time a new property (binding) is added.

The type of an individual is simply one of its fields that contains a list of predicate terms (what we calls 'categories'—the referents of most common nouns, verbs, or other relational words). As the parse proceeds, we can change the contents of this field to reflect a more specific type as the evidence for it surfaces.

A similar thing happens when further parsing elaborates or specializes the objects involved in already established relations. A clear example of this is what happens to the individual representing "*the new firm*" in our ongoing example. Of course the object denoted by this phrase will not be literally an instance of a 'new firm': Even before the parsing of the next phrase has begun that wording will have been used to 'de-reference' the phrase and identify the company (the joint venture firm) to which it refers. This company will be the actual denotation, and be the individual that is bound as the 'thing owned' when the verb group is combined with the subject.

Recall that a part of this same moment of composition is the combination of the partially saturated object based on the percentage phrase in the verb group with that same subject ("*the new firm will be 50 percent owned by ...*"). An important consequence of this composition is that it establishes that what we are dealing with in this example is not a company per se, but one of its necessary parts, its shares of stock.

At the moment when this specialization becomes known, the binding linking the company to the individual representing this instance of 'something (will) own a company' is already in place. We can still make the substitution however, since what that amounts to is changing the value of this binding from the company to this specialization of it; the identity of the binding never changes, and there is no impact on the rest of the ownership relation.

7. Final Remarks

We firmly believe that now that the study of the representation of the meaning (semantics) of natural language texts has reached a substantial level of sophistication, it is crucial to ground further work in the problems that arise from the study of the real texts that real people encounter every day. We are all too adroit at the technical details of representation to continue much longer following the paths that the logics themselves suggest—we should be more linguist than mathematician at this point.

To this end, we have in this paper pointed to some of the actual problems that we have been trying to deal with that have forced us to expand our representations and consider whether we are employing appropriate designs. This has led us to focus on the interpretation function that links a text with its denotation, rather than with issues of inference once the denotation has been established and, for example, we (or rather our systems) then need to know the consequences of what they have just read for the actions they might take next.

Our guiding principle has been to try to simplify the interpretation function at the possible expense of making the epistemology of the underlying representation more elaborate. This was, for example, the motivation behind using an unusual kind of referent for displaced constituents, the partially saturated object. The alternative would have been to make the syntax more complex, a programme that we expect would ultimately fail.

It should also be said that throughout this work we have always been dealing with only two levels of representation: the surface linguistic analysis of the text, and the model-level representation to which its interpretation is projected. We have never found a use for an intermediate level of analysis such as 'logical form', in large part this is because we have always worked incrementally on very large continuous texts (hundreds to thousands of words), and the questions that need answering moment to moment during such analyses have invariably been best couched at one of our two levels. Logical forms seem to us to make more sense when working in successive passes sentence by sentence rather than incrementally with whole texts.

A further influence is the fact that Sparser's design has always assumed that many of the individuals that it will be called on to recognize will already have a representation in a long-standing model (companies, people, recent events, calendars of time, etc.). If this is the case, it makes the most sense in an analysis to move to that level of representation as quickly as possible, if for no other reason than for the help it will provide in disambiguation.

The work we have described is very much an ongoing effort where small changes are being made month by month. Many of the techniques are experimental as we have said. Nearly all of them are likely to evolve as this work continues.

8. References

McDonald, David (1992) "An Efficient Chart-based Algorithm for Partial-Parsing of Unrestricted Texts". In: Proceedings of the 3d Conference on Applied Natural Language Processing (ACL). Trento, Italy. April 1992. 193 -- 200.

_____ (1994a) "'Krisp': A Representation for the Semantic Interpretation of Texts", *Mind and Machines* **4**: 59-73.

_____ (1994b) "Trade-offs Between Syntactic and Semantic Processing in the Comprehension of Real Texts", in the Proceedings of RIAO-94, Rockefeller Univ. October 11-13, 1994

Mellish, C. S. (1985) *Computer Interpretation of Natural Language Descriptions*, John Wiley & Sons, New York.

Prince, Ellen F. (1988) "Information-Status", in Bright (ed.) *International Encyclopedia of Linguistics*, Oxford University Press.

Shapiro, Stuart C. & William J. Rapaport (1992) "The SNePs Family", *Computers Math. Applic.*, (23) 2-5, pgs 243-275.

Steedman, Mark (1989) "Constituency and Coordination in a Combinatory Grammar", in Baltin & Kroch (eds) *Alternative Conceptions of Phrase Structure*, Univ. of Chicago Press, pgs. 201-231.

Vallduví, Enrich (1993) *The Informational Component*, Garland Publishing, New York.