THE EFFECT OF ESTABLISHING COHERENCE IN ELLIPSIS AND ANAPHORA RESOLUTION

Andrew Kehler
Harvard University
Aiken Computation Laboratory
33 Oxford Street
Cambridge, MA 02138
kehler@das.harvard.edu

Abstract

This paper presents a new model of anaphoric processing that utilizes the establishment of coherence relations between clauses in a discourse. We survey data that comprises a currently stalemated argument over whether VP-ellipsis is an inherently syntactic or inherently semantic phenomenon, and show that the data can be handled within a uniform discourse processing architecture. This architecture, which revises the dichotomy between *ellipsis* vs. *Model Interpretive Anaphora* given by Sag and Hankamer (1984), is also able to accommodate divergent theories and data for pronominal reference resolution. The resulting architecture serves as a baseline system for modeling the role of cohesive devices in natural language.

1 Introduction

There has been much debate concerning the appropriate level of language processing at which to treat VP-ellipsis resolution. Syntactic accounts (Fiengo and May, 1990; Haïk, 1987; Hellan, 1988; Hestvik, 1993; Lappin, 1993; Lappin and McCord, 1990) claim that syntactic material is copied from the antecedent clause and reconstructed in the elided clause, whereas semantic accounts (Dalrymple, 1991; Dalrymple *et al.*, 1991; Gawron and Peters, 1990; Hardt, 1992; Kehler, 1993; Klein, 1987) claim this material is retrieved from semantic representations. This debate is currently deadlocked; indeed a survey of the data seems to indicate that ellipsis must be both a syntactic and semantic phenomenon. In Section 2, we examine five types of ellipsis contexts, and show a pattern that has gone unnoticed in the literature. In Section 3 we break the deadlock by presenting a discourse processing architecture from which the correct predictions naturally result. We present further evidence for this architecture from pronominal NP reference resolution data in Section 4, and conclude in Section 5.

2 VP Ellipsis Data

A simple example of verb phrase (VP) ellipsis is given in sentence (1):

(1) Ross likes his mother, and Bill does too.

The stranded auxiliary in the second clause (the *target* clause) indicates the deletion of a verb phrase, a representation for which is to be located from another clause (the *source* clause), in this case, the first clause. Sentence (1) displays a *strict/sloppy* ambiguity: Bill may like *Ross's* mother (the strict reading) or *his own* mother (the sloppy reading).

In this section we examine five types of elliptical contexts, and show the following heretofore unnoticed pattern.[1] When the relationship between the source clause A and the target clause B is what we term *parallel* (as exemplified by the sentential structure *"A and B too"*, as in sentence (1)), the data indicates that a syntactic representation is reconstructed at the target site. That is, a syntactically-parallel source VP must be available, and reconstruction of this VP in the target clause is subject to syntactic constraints. However, in *non-parallel* constructions (for example, contrastive conjunctions (*"A but B"*), subordinates (*"A because B"*), and comparatives (*"A better than B"*)) neither a syntactic source VP nor compliance with syntactic constraints is necessary, instead only a suitable semantic source representation is required. These data strongly suggest a discourse processing architecture that is sensitive to the establishment of coherence relations between clauses, which is described in Section 3.

2.1 Passive/Active Alternation

Active elided target VPs can receive interpretations from passive source VPs:

(2) In March, four fireworks manufacturers asked that the decision be reversed, and on Monday the ICC did. [reverse the decision]
(from text of Rosenthal (1988), cited in Dalrymple (1991))[2]

[1]Although space precludes us from discussing past work in detail, no approach known to us addresses (let alone accounts for) all of the phenomena discussed herein.

[2]This use of *and* in this example does not signify a parallel relationship between the source and target, as

(3) This problem was to have been looked into, but obviously nobody did. [look into the problem]
(Vincent Della Pietra, in conversation)

Less frequent, but still extant, are cases of passive targets with active sources:

(4) In addition to inducing lethality during the first instar, it[3] retards embryonic development, but not to the extent that the growth cones were. [retarded]
(from text of (Jarecki, 1992))

(5) Actually I have implemented it[4] with a manager, but it doesn't have to be. [implemented with a manager]
(Steven Ketchpel, in conversation)

These examples are problematic for syntactic analyses of ellipsis, because the source clause is not a surface VP constituent. In contrast, semantic analyses cannot account for the unacceptability of similar examples with parallel constructions:

(6) * This problem was looked into by John, and Bob did too.

(7) * This agent retards embryonic development, and the growth cones were too.

Syntactic parallelism is apparently required in the parallel construction, but not in the non-parallel ones.

2.2 Condition C Violations

Fiengo and May (1990) and Lappin (1993) note the unacceptability of examples such as (8) and (9):

(8) * I hit Bill$_i$, and then he$_i$ did too.

(9) * Mary introduced John$_j$ to everyone, and he$_j$ did too.

Under a syntactic reconstruction analysis, this is predicted by a Condition C violation within Binding Theory (Chomsky, 1981):

(10) * I hit Bill$_i$, and then he$_i$ hit Bill$_i$ too.

(11) * Mary introduced John$_j$ to everyone, and he$_j$ introduced John$_j$ to everyone too.

Because syntactic constraints do not apply at the semantic level, semantic accounts predict these cases to be acceptable. However, Dalrymple (1991) gives examples where Condition C should apply but apparently does not:[5]

(12) I expected Bill$_i$ to win even when he$_i$ didn't.

(13) The lawyer defended Bill$_j$ against the accusations better than he$_j$ could have.

As these examples are non-parallel constructions, again it appears that syntactic constraints apply in parallel constructions but not in non-parallel ones.

2.3 Condition A Violations

As predicted by Condition A of Binding Theory, it is generally difficult to obtain a strict reading when the source clause contains a reflexive pronoun:

(14) ?? John$_i$ defended himself$_i$, and Bob$_j$ did too. [defended John$_i$]

(15) ?? Fred$_i$ voted for himself$_i$, and Gary$_j$ did too. [voted for Fred$_i$]

Given appropriate semantic context, judgements improve but still remain somewhat stilted:[6]

(16) ? The alleged murderer$_i$ defended himself$_i$, and his lawyer$_j$ did too. [defended the alleged murderer$_i$]

(17) ? Bill Clinton$_i$ voted for himself$_i$, and his campaign manager$_j$ did too. [voted for Bill Clinton$_i$]

The stiltedness of reflexives under a strict reading disappears, however, in non-parallel constructions (from Dalrymple (1991)):

(18) Bill$_i$ defended himself$_i$ against the accusations better than his lawyer$_j$ did. [defended Bill$_i$]

(19) John$_i$ voted for himself$_i$ even though no one else$_j$ did. [voted for John$_i$]

In these cases, the strict reading is readily available and perhaps preferred. Again, there appears to be a syntactic dependency in the parallel cases that is absent from the non-parallel ones.

2.4 Non-VP Antecedents

In the following examples, the source representation is not a syntactic VP but instead comes from a nominalization:[7]

(20) This letter deserves a response, but before you do, ... [respond]
(Gregory Ward, p.c.)

(21) Meanwhile, they sense a drop in visitors to the city. Those who do, they say, are not taking cabs. [visit the city]
(Chicago Tribune, courtesy Gregory Ward)

---

the use of *too* would be inappropriate under the desired reading. We might term this relationship to be *result*.

[3] Here, *it* refers to a previously mentioned chemical agent.

[4] Here, *it* refers to a previously mentioned computer system.

[5] These particular cases also violate the *Vehicle Change* proposal of Fiengo and May (1990), which adequately handles other acceptable cases that violate Condition C.

[6] There appears to be a dialect that readily allows strict readings with reflexives. However, even for those speakers, the non-parallel constructions given below are more acceptable under the strict interpretation than the parallel cases.

[7] Some speakers find these cases slightly awkward or stilted. Again, however, most find these better than the parallel construction cases given below.

Semantic analyses can account for these cases if nominalizations are assumed to evoke event representations into the discourse model. However, in parallel constructions, discourse-model events from nominalizations are apparently not available:

(22) * This letter provoked a response from Bush, and Clinton did too. [responded]

(23) * There is a rise in American visitors to the city, and Canadians do too. [visit the city]

A similar pattern is seen in cases where the antecedent of the ellipsis is evoked by an adjectival phrase:

(24) First person pronouns aren't very shiftable, although the plural ones can be. [shifted] (Barbara Partee, in conversation)

The acceptability of example (24) is to be compared with the relative unacceptability of an otherwise similar parallel construction case:

(25) * First person pronouns aren't very shiftable, and the plural ones also don't. [shift]

Again, non-syntactic source representations appear to only be available in the non-parallel construction cases.

2.5 Subjacency Violations

Haïk (1987) gives examples of apparent subjacency violations in antecedent contained deletion (ACD):

(26) John read everything which Bill believes he did.

(27) * John read everything which Bill believes the claim that he did.

(28) * John read everything which Bill wonders why he did.

This data is problematic for a purely semantic theory of ellipsis, as there should be no syntactic dependency at the ellipsis site. However, sentence (29), which has a subordinate conjunction, does not display the expected subjacency violation (from Rooth (1981)):[8]

(29) Which problem did you think John would solve because of the fact that Susan did?

Without ellipsis, the gap remains and a subjacency violation results:

(30) * Which problem did you think John would solve because of the fact that Susan solved?

In our account, the relative pronoun *which* does not specify an interclausal coherence link, and therefore sentences (26-28) are parallel constructions. Consequently, again the source representation for the parallel construction is apparently syntactic, whereas that for the non-parallel construction (e.g., example (29)) is semantic.

[8] I thank Stuart Shieber for bringing this example to my attention.

3 An Architecture that Exploits Coherence

The data given in Section 2 suggests that VP-ellipsis resolution copies a syntactic representation in parallel constructions and a semantic representation in non-parallel ones. In this section, we present a discourse processing architecture from which these predictions naturally result. We first describe Sag and Hankamer's (1984) (henceforth S&H) earlier work, the representations from which our analysis will utilize.

3.1 Sag and Hankamer's Architecture

S&H give a performance-based view of anaphoric processing that utilizes two forms of representation, a *propositional representation* and a *discourse model*. Propositional representations, which they conjecture are held by short-term registers in memory (henceforth *propositional registers*), maintain the surface syntactic constituent structure of an utterance as well as binding relations; however, discourse anaphors are resolved.[9] These representations are built in tandem with a discourse model. S&H claim that the ellipsis resolution process obtains referents from propositional representations, whereas what they term *Model Interpretive Anaphora* (MIA) (e.g., 'do it' anaphora) obtains referents from the discourse model. They give the following example to illustrate (judgements are theirs):

(31) The children asked to be squirted with the hose, so

a. they were. [ellipsis]

b. * we did. [ellipsis]

c. we did it. [MIA]

In their theory, example (31a) is acceptable because the source representation is a surface VP and therefore is retrievable from its propositional representation. Example (31b) is unacceptable because the source *squirt the children with the hose* is not a surface VP in the propositional representation.[10] Sentence (31c) is acceptable because 'do it' anaphora is an MIA process, and therefore obtains referents from the discourse model, in which a representation for *squirt the children with the hose* is assumed to exist. One problem with this account is that it does not explain the dichotomy of judgements for the data given in Section 2. For each of these phenomena, the S&H approach predicts that all cases are

[9] In fact, they suggest that propositional representations (as opposed to classical surface structures) are what the parser constructs, a view consistent with our account.

[10] Ellipsis, in their formulation, is subject to an *identity-of-logical-form* constraint on propositional representations. See Dalrymple et al. (1991) and Hardt (1992) for arguments that this condition is flawed.

either acceptable or unacceptable without regard to the type of construction. The data instead suggests that VP-ellipsis within the parallel construction behaves like S&H's definition of ellipsis, and in non-parallel constructions it behaves like their MIA. We believe that their dichotomy of phenomena is somewhat illusory, and a more compelling and elegant treatment is possible by appealing to discourse structure and coherence.

3.2 A Revised Architecture

We follow S&H and earlier arguments by Johnson-Laird (1983) in specifying a discourse processing architecture that includes both propositional representations and a discourse model. We propose that a main function of the propositional representations is to serve as a "bridge" in establishing clause-to-clause coherence. That is, as clauses are parsed, propositional representations are generated and held in the short-term registers. When it is established that a non-parallel coherence relationship is present between a clause (or set of clauses) and a subsequent clause, the propositional representation(s) for the first is integrated into the discourse model so that coherence can be computed, thereby freeing the short-term propositional register for subsequent representations. Non-parallel constructions specify a non-parallel coherence relationship overtly (e.g., 'contrast' (*because*), 'comparison' (*better than*), 'result' (the use of *and* in example 2))[11], thereby identifying the first clause as a unit for coherence establishment and invoking integration of its propositional representation into the discourse model.

Parallel constructions, on the other hand, do not invoke this integration. Hobbs (1979), in discussing the parallel coherence relation he defines, suggests one basis for this distinction:

> *One frequent function of the Parallel relation is to suggest or support the generalization which the two segments are specific instances of. The relation often links segments which together function as an Exemplification or Elaboration of a more general statement.*

In our terms, clauses conjoined by a parallel relation will form a group and cohere *as a unit* with prior and subsequent statements. Therefore, this relation in itself does not cause the representation for the first clause to be integrated into the discourse model. Instead, the integration of *both* representations into the discourse model as a unit is invoked by subsequent coherence establishment.

[11] This use of *and* raises the question of how the listener knows which meaning of *and* is present during processing. We assume that the listener can consider multiple possibilities in parallel, although it may also be that in these cases the interclausal relationship has been established by the time the ellipsis site is processed.

The preceding analysis makes the predictions necessary to account for the ellipsis data given in Section 2. Under our account, the representation of an utterance *either* exists in a propositional register *or* in the discourse model; these are not created in tandem as S&H suggest. An elided VP then receives an interpretation from whichever representation is available. The parallel construction cases in Section 2 are unacceptable because the source clause has not been integrated into the discourse model when the ellipsis site is processed. The source must therefore be retrieved from the propositional representation, where surface syntactic structure and binding relations are maintained. Reconstructing this representation requires syntactic parallelism (ruling out passive/active and nominalization cases) and can result in violation of syntactic constraints (such as Condition C violations, Condition A violations, or subjacency violations). The non-parallel construction cases in Section 2 are acceptable because the antecedent for the ellipsis has been integrated into the discourse model when the ellipsis site is processed. Because the vestiges of syntactic information are no longer present in the representation, syntactic constraint violations do not occur; furthermore source representations from nominalizations and clauses of differing voice now exist in the discourse model.[12]

3.3 Examples

In this section, we work through two examples to illustrate the proposals that have been made thus far. For resolution at the syntactic level, we assume an S&H-like propositional representation and a straightforward mechanism whereby a VP representation is copied from the source clause representation into the target clause representation. For resolution in the discourse model, we will use the event representation and resolution algorithm defined in (Kehler, 1993). The focus of this section, however, is to offer a general illustration of the architecture rather than to make specific claims concerning the

[12] Differences remain between the distribution of S&H's ellipsis and MIA phenomena that need to be accounted for in a comprehensive treatment of event anaphora, as examples (31a-c) show (although judgements as to the unacceptability of sentence (31b) vary). Interestingly, contra S&H, MIA phenomena also appear to be sensitive to syntactic constraints in certain contexts, as the following example from Dalrymple (1991) shows:

* I hit Sue$_i$, and then she$_i$ did it.

One hypothesis is that VP-ellipsis is actually event anaphora with an empty pronoun; it may then be that distributional differences between "do ϕ", "do it", and "do that" are due only to the anaphoric properties of the event pronoun involved, and not due to a fundamental difference between ellipsis and MIA phenomena. This hypothesis is the subject of ongoing work.

nature of the representations involved.

Examples (32) and (33) exhibit the contrast between parallel and non-parallel constructions with active target clauses that have passive source clauses, as discussed in Section 2.1:

(32) * John was shot by Bob, and Bill did too.

(33) John was shot by Bob because Bill wouldn't.

The propositional representation for the source clause in these sentences is shown in representation (34), where P denotes the passive voice:

(34) [P [was_shot′(by(Bob′))] (John′)]

Because the two clauses in sentence (32) stand in a parallel relationship, the source representation has not been integrated into the discourse model at the time the ellipsis site is processed; therefore the ellipsis is resolved at the propositional level of representation. A representation is constructed with the information present in the target clause:

(35) [A [...] (Bill′)]

Here A denotes the active voice, a feature which is indicated in the target clause through the use of *did*. When the elided VP is retrieved, a mismatch occurs: the passive VP cannot be brought into a representation marked as having active voice. The copying can therefore not be completed, resulting in the unacceptability of the sentence.

Sentence (33) also has representation (34) for its source clause. However, because the two clauses stand in a non-parallel relationship, representation (34) has already been integrated into the discourse model when the ellipsis site is processed, and thus resolution occurs in the discourse model. The representation for the source clause is:

(36) e_1: [predicate: shot
time: past
polarity: positive
modality: necessity
agent: Bob
theme: John]

Because this representation is based upon thematic (and not syntactic) roles, the representations for active and passive forms of a sentence are identical. For the target clause, a parallel event representation is created with empty roles, and the role fillers present in the target clause are filled in:

(37) e_2: [predicate:
time: past
polarity: negative
modality: volitional_possibility
agent: Bill
theme:]

Representations for the the remaining role fillers are retrieved from the source clause representation:

(38) e_2: [predicate: shot
time: past
polarity: negative
modality: volitional_possibility
agent: Bill
theme: John]

This resolution successfully yields the correct representation for the target clause in sentence (33).

Examples (39) and (40) illustrate the contrast between parallel and non-parallel constructions with respect to potential Condition C violations, as described in Section 2.2:

(39) * The lawyer defended Bill$_i$ and he$_i$ did too.

(40) The lawyer defended Bill$_i$ better than he$_i$ could have.

In each sentence, the propositional representation for the source clause takes the form given in (41):

(41) [[defend′(Bill′)] (lawyer1′)]

Because the two clauses in sentence (39) stand in a parallel relationship, the source representation has not been integrated into the discourse model at the time the ellipsis site is processed. The ellipsis is then resolved at the propositional level of representation. After filling in the information present in the target clause and copying the representation of the source VP, representation (42) results:[13]

(42) [[defend′(Bill$_i$′)] (he$_i$′)]

A manifestation of Condition C applying at this level rules out this representation as ill-formed, because the pronoun *he* c-commands the coreferential NP form *Bill*.

Sentence (40) also has representation (41) for its source clause. Because the two clauses stand in a non-parallel relation, representation (41) has already been integrated into the discourse model when the ellipsis site is processed. Resolution then occurs in the discourse model. The representation for the source clause is:

(43) e_3: [predicate: defend
time: past
modality: necessity
agent: lawyer1
theme: Bill]

After creating a parallel event representation, filling in role fillers present in the target clause, and retrieving remaining role fillers from the source representation, representation (44) results:

(44) e_4: [predicate: defend
time: past
modality: possibility
agent: Bill
theme: Bill]

[13] Recall that pronouns have been resolved at this level of representation; we indicate this by coindexing.

Because no syntactic constraints apply at this level of representation, representation (44) is well-formed and yields the correct interpretation for the target clause.

In summary, our architecture accounts for data supporting both the syntactic and semantic accounts of ellipsis resolution in an elegant and intuitive way. Section 4 examines pronominal NP resolution in this model.

4 Pronoun Resolution

There are also dichotomous views in the literature concerning the process by which pronoun resolution is performed. Theories incorporating a notion of *local focus* generally utilize syntactic information in their specifications. For example, the *Centering* framework of Grosz et al. (1986) utilizes grammatical role information in specifying the accessibility of potential referents on an ordered forward-looking center list. Kameyama's work (Kameyama, 1988) contains rules for property sharing utilizing grammatical roles. Passonneau (1991), in looking at the distribution of 'it' and 'that' for NP reference in naturally-occurring texts, concludes that both syntactic form of the antecedent and its grammatical role are needed to adequately account for the data. Furthermore, she suggests that the function of the propositional register discussed by S&H is appropriate for accommodating her rules.

Alternatively, some researchers (Hobbs, 1979; Wilensky, 1978) have suggested that coreference is determined as a by-product of coherence determination between sentences. In Hobbs' account, for example, pronouns are modeled as free variables and are assigned to objects during the process of establishing coherence relations.

However, Hobbs himself acknowledges the power of grammatical role-based heuristics,[14] noting that upon hearing example (45),

(45) John can open Bill's safe. He ...

one is likely to assume that *John* is the referent of *He*. The existence of a garden-path effect in example (46), where *He* refers to *Bill* instead of *John*, suggests that pronominal reference resolution is not guided by coherence considerations alone:

(46) John can open Bill's safe. He's going to have to get the combination changed soon.

As focus-based theories would predict, the reader assigns *John* as the referent of *He*, and double-takes when semantic information later suggests otherwise.[15] Our architecture provides an explanation for this phenomenon. Since a coherence relation has not been established at the time the pronoun is processed, the propositional representation for the first clause (which preserves information that focus-based theories utilize, such as surface-string ordering and depth-of-embedding of potential referents) is the representation available to the reference resolution algorithm.[16] However, when a non-parallel coherence link is overt, our architecture would predict that a semantically-based resolution process would be used because the propositional representation containing potential referents has already been integrated into the discourse model at the time the pronoun is processed. This prediction is borne-out empirically; consider the following two sentence prefixes (complete sentences taken from (Ehrlich, 1980)):

(47) Steve blamed Frank and he ... [spilt the coffee].

(48) Steve blamed Frank because he ... [spilt the coffee].

Focus-based theories predict the strong bias toward the referent of *he* in example (47) being the subject (i.e., *Steve*), even though *he* is consistent with both potential referents. Because this sentence is a parallel construction (i.e., the meaning of "and" is not *result*), our architecture also makes this prediction in accordance with those theories. The heuristic preferring subjects does not apply in example (48), where *Frank* is perhaps the preferred referent of *he*, seemingly as a result of reasoning using semantic features of the verb *blame*. Our architecture correctly predicts that the pronoun in sentence (48) does not cause processing problems while the one in example (46) does, because only in sentence (48) has the clause containing the referent of *he* been integrated into the discourse model at the time the pronoun is processed.

Ehrlich (1980) gives experimental evidence supporting this view. Ehrlich's goal was to test the biasing effect of the so-called "implicit causality" feature (Garvey and Caramazza, 1974) of verbs such as *blame* in pronoun reference assignment in two-clause sentences with conjunctions other than *because* (which was the only conjunction used in previous work (Garvey *et al.*, 1976)). In her experiments, subjects were tested for their ability to identify correct referents of pronouns in three versions of six two-clause sentences (such as those in sentences (47) and (48)), where each of the sentences contained one of the conjunctions *and*, *but*, and *because*. It was found that subjects were significantly more accurate in determining correct referents of

[14] Hobbs (1976) found that a heuristic favoring subjects over objects was 90% accurate for written texts.

[15] This effect causes Hobbs to admit that "this strongly suggests that some psychological reality underlies the heuristic [favoring subjects over objects]."

[16] After garden-pathing, "off-line" reasoning apparently allows the reader of example (46) to identify the correct referent of the pronoun. This reasoning may allow propositional representations to be integrated into the discourse model so that Hobbs-like coherence determination can be performed.

pronouns when the conjunction used was *because* or *but*, and therefore that the effect of implicit causality was not constant with respect to the conjunction used. While a detailed analysis of her work is beyond the scope of this paper, two generalizations that she draws as a result of her experiments are: (1) that subjects were more able to utilize 'general knowledge' in determining the referent when the conjunction used was *because* or *but* than when it was *and*; and (2) that hearers analyze language a clause at a time. The first of these results supports our view that semantic information required for reasoning is primarily available in the discourse model (since the representation for the first clause is integrated into the discourse model when the conjunction used is *but* or *because*); the second point supports our claim that the propositional registers hold clause-level representations.[17]

In summary, our architecture also accommodates evidence supporting competing theories of how pronominal NP resolution is performed.

5 Conclusions

This paper presents a model for anaphoric processing that incorporates the role of establishing coherence relationships between clauses in a discourse. By postulating the existence of propositional representations in addition to a discourse model, we account for ellipsis data that has gridlocked work on the topic. Furthermore, arguments for dichotomous approaches to pronoun resolution are resolvable within this framework.

It should be noted that coherence establishment is not likely to be the only discourse factor involved in integrating propositional representations into the discourse model. Therefore, the analysis described herein only indicates tendencies, as opposed to predicting cut-and-dry judgements on the basis of type of construction alone. For instance, example (49) has been judged by some speakers to be acceptable under a strict reading:[18]

(49) I voted for myself, and I hope you did too!

Our account predicts that this case would be at least somewhat stilted due to a Condition A violation. One factor distinguishing this example from others we have discussed is the use of first and second person pronouns, and a second is the fact that the pronominal referent necessary to yield the strict reading is also present in the target clause. Future work is needed to further analyze the effects of these differences.

The theory presented here evokes many other questions for future study. One such question is how the postulated representations should be further formalized, and how reasoning with these formalizations is to be performed. A second question is how this conception of discourse processing may be integrated with theories of discourse structure (Grosz and Sidner, 1986; Scha and Polanyi, 1988; Webber, 1991). While we have looked primarily at two-clause structures, the ramifications that the claims have on multi-clause discourse structure require further investigation. Such studies will form the basis for further characterization of the role of coherence establishment in anaphoric processing.

Acknowledgments

This work was supported in part by National Science Foundation Grant IRI-9009018, National Science Foundation Grant IRI-9157996, and a matching grant for the latter from the Xerox Corporation. I would like to thank Mary Dalrymple, Barbara Grosz, Shalom Lappin, Karen Lochbaum, Christine Nakatani, Stuart Shieber, and two anonymous reviewers for valuable discussions and comments on earlier drafts.

References

(Chomsky, 1981) Noam Chomsky. *Lectures in Government and Binding.* Foris, Dordrecht, 1981.

(Dalrymple *et al.*, 1991) Mary Dalrymple, Stuart M. Shieber, and Fernando Pereira. Ellipsis and higher-order unification. *Linguistics and Philosophy*, 14:399–452, 1991.

(Dalrymple, 1991) Mary Dalrymple. Against reconstruction in ellipsis. Technical Report SSL-91-114, Xerox, 1991.

(Ehrlich, 1980) Kate Ehrlich. Comprehension of pronouns. *Quarterly Journal of Experimental Psychology*, 32:247–255, 1980.

(Fiengo and May, 1990) Robert Fiengo and Robert May. Anaphora and ellipsis. ms., City University of New York and University of California at Irvine, 1990.

(Garvey and Caramazza, 1974) C. Garvey and A. Caramazza. Implicit causality in verbs. *Linguistic Inquiry*, 5:549–564, 1974.

(Garvey *et al.*, 1976) C. Garvey, A. Caramazza, and J. Yates. Factors underlying assignment of pronoun antecedents. *Cognition*, 3:227–243, 1976.

[17] Ehrlich's results with the conjunction *and* are mixed with respect to our theory, as in some cases her participants preferred a non-subject position referent over a subject position one. In particular, she notes that this happens when the main verb of the second clause is the stative verb *be*, as in *Sue criticized Penny and she was gloomy.* These sentences contain the *result* meaning of *and* as opposed to the parallel one. Unfortunately, Ehrlich's original data was not available at the time of this writing so an analysis distinguishing between uses of *and* could not be performed.

[18] I thank an anonymous reviewer for this example.

(Gawron and Peters, 1990) Mark Gawron and Stanley Peters. *Anaphora and Quantification in Situation Semantics.* CSLI/University of Chicago Press, Stanford University, 1990. CSLI Lecture Notes, Number 19.

(Grosz and Sidner, 1986) Barbara Grosz and Candace Sidner. Attention, intentions, and the strcture of discourse. *Computational Linguistics*, 12(3):175–204, 1986.

(Grosz *et al.*, 1986) Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. Towards a computational theory of discourse interpretation. Unpublished manuscript, 1986.

(Haïk, 1987) Isabelle Haïk. Bound variables that need to be. *Linguistics and Philosophy*, 11:503–530, 1987.

(Hardt, 1992) Daniel Hardt. VP ellipsis and contextual interpretation. In *Proceedings of the International Conference on Computational Linguistics (COLING-92)*, Nantes, July 1992.

(Hellan, 1988) Lars Hellan. *Anaphora in Norwegian and the Theory of Grammar.* Studies in Generative Grammar 32. Foris, Dordrecht, 1988.

(Hestvik, 1993) Arild Hestvik. Strict reflexives and the subordination effect. In S. Berman and A. Hestvik, editors, *Proceedings of the Stuttgart Workshop on Ellipsis: Arbeitspapiere des Sonderforschungsbereich 340, Bericht Nr. 29-1992, SFB 340.* University of Stuttgart, University of Tuebingen, and IBM Germany, 1993.

(Hobbs, 1976) Jerry Hobbs. Pronoun resolution. Technical Report 76-1, Department of Computer Science, City University of New York, 1976.

(Hobbs, 1979) Jerry Hobbs. Coherence and coreference. *Cognitive Science*, 3:67–90, 1979.

(Johnson-Laird, 1983) P. N. Johnson-Laird. *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness.* Cambridge University Press, 1983.

(Kameyama, 1988) Megumi Kameyama. Japanese zero pronominal binding: Where syntax and discourse meet. In William J. Poser, editor, *Papers from the Second International Workshop on Japanese Syntax*, pages 47–74. CLSI, 1988.

(Kehler, 1993) Andrew Kehler. A discourse copying algorithm for ellipsis and anaphora resolution. In *Proceedings of the Sixth Conference of the European Chapter of the Association for Computational Linguistics (EACL-93)*, Utrecht, the Netherlands, April 1993.

(Klein, 1987) Ewan Klein. VP-Ellipsis in DR theory. In Groenindijk and Stokhof, editors, *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers.* Foris, 1987.

(Lappin and McCord, 1990) Shalom Lappin and Michael McCord. Anaphora resolution in slot grammar. *Computational Linguistics*, 16:197–212, 1990.

(Lappin, 1993) Shalom Lappin. The syntactic basis of ellipsis resolution. In S. Berman and A. Hestvik, editors, *Proceedings of the Stuttgart Workshop on Ellipsis: Arbeitspapiere des Sonderforschungsbereich 340, Bericht Nr. 29-1992, SFB 340.* University of Stuttgart, University of Tuebingen, and IBM Germany, 1993.

(Passonneau, 1991) Rebecca Passonneau. Persistence of linguistic form in discourse processing. In *Proceedings of the Fall Symposium on Discourse Structure in Natural Language Understanding and Generation*, Monterey, CA, November 1991.

(Rooth, 1981) Mats Rooth. A comparison of three theories of verb phrase ellipsis. University of Massachusetts Occasional Papers in Linguistics, Volume 7, 1981.

(Sag and Hankamer, 1984) Ivan Sag and Jorge Hankamer. Toward a theory of anaphoric processing. *Linguistics and Philosophy*, 7:325–345, 1984.

(Scha and Polanyi, 1988) Remko Scha and Livia Polanyi. An augmented context free grammar for discourse. In *Proceedings of the International Conference on Computational Linguistics (COLING-88)*, pages 573–577, Budapest, August 1988.

(Webber, 1991) Bonnie Lynn Webber. Structure and ostension in the interpretation of discourse deixis. *Language and Cognitive Processes*, 6(2):107–135, 1991.

(Wilensky, 1978) Robert Wilensky. *Understanding Goal-Based Stories.* PhD thesis, Yale, 1978. Research Report #140.

Sources of Examples

(Jarecki, 1992) Jill Jarecki. The role of synaptic activity during embryonic neuromuscular junction development in *drosophila melangestr.* Unpublished Yale University Prospectus, August 1992.

(Rosenthal, 1988) Harry F. Rosenthal. Agency orders UPS to resume fireworks deliveries–after July 4th. Associated Press Newswire, June 1988.