Ellipsis

Intermediate article

Robert C May, University of California, Irvine, California, USA

CONTENTS

A linguistic ellipsis, most generally expressed, is a truncated or partial linguistic form. It is a linguistic form in which constituents normally occurring in a sentence are superficially absent, licenced by structurally present prior antecedents.

INTRODUCTION

A linguistic ellipsis, most generally expressed, is a truncated or partial linguistic form. This partiality is measured relative to a complete sentence; an elliptical sentence is one in which some of the constituent parts of a 'full' sentence are missing. For example, in answer to the question 'Who went to the store', one may answer 'Max went to the store'. In most contexts, however, speakers would avoid such prolixity and instead would employ an elliptical form: 'Max went', 'Max did', or even just 'Max' would suffice as answers to the question. The reason that they would suffice appears to be quite obvious; it is because they mean just what 'Max went to the store' means, except they express this more economically by leaving off at least part of what can otherwise be gleaned from the initial question. From such simple examples we can already observe a fundamental property of ellipsis that we would expect to be captured under any account of the relation between elliptical and non-elliptical forms: meaning is constant under ellipsis. Specifying this constancy is not as straightforward as it may initially appear, however, and trying to capture what it amounts to has been an on-going issue in linguistic theory. It has been an issue of particular importance because the accounts of elliptical phenomena have been used to bring empirical weight to fundamental claims about linguistic theory, including the extension of the notion of linguistic identity, the relation of syntax and semantics, and the abstractness of grammar.

To frame our discussion, we note the most well-known types of elliptical constructions that have been studied in the literature:

VP-ellipsis: Max went to the store, and Oscar did, too
gapping: Max went to the store, and Oscar to the arcade
sluicing: Max went to the store, but Oscar wondered why
stripping: Max saw Sally at the store, and Oscar, too
pseudo-gapping: Max loves Jane, and Harry does Sally
N′-ellipsis: Max's father went to the store, but Oscar's went to the arcade

Each of the these types of ellipsis have idiosyncratic properties. For example, with gapping there is a well-known correlation of the direction of gapping and word order, and with sluicing the restriction that the complementizer of a sluiced clause must be interrogative; these are among observations originating with Ross (1969, 1970). Our purpose here is not, however, to survey the differences between these various elliptical phenomena but to explore two fundamental properties that they all have in common: (1) ellipsis is of a syntactic constituent; (2) an antecedent occurrence of the elided constituent governs the ellipsis. For example, in VP-ellipsis the elided material is a verb phrase, and the ellipsis is licensed in the presence of a fully lexicalized antecedent: in the case above, 'went to the store'. The primary issue for ellipsis is how to properly characterize (1) and (2); what sorts of linguistic description are called for to capture these fundamentally grammatical aspects of ellipsis in their full generality? A number of subsidiary issues are implied by answers to the primary questions, including the nature of the grammatical mechanisms that account for the absence of the lexical material: is it a deletion of underlyingly present syntactic elements; or are these elements absent even at the underlying level? In what follows we will outline some of the main approaches that have been developed to these issues, focusing primarily on the case of VP-ellipsis, largely because the relevant

theoretical issues have been most clearly and widely discussed in this context.

VP-ELLIPSIS

Syntactic Reduction: The Transformational Theory

The initial approaches to ellipsis within contemporary linguistic theory attempted to account for the semantic constancy alluded to above by reducing it to syntactic constancy. The idea here is quite intuitive: if an elliptical form can be seen as a syntactic repetition of a corresponding non-elliptical form, i.e. as simply two occurrences of the same syntactic form, then it would follow immediately that they have the same meaning. The first systematic investigations along these lines is found in the work of J. R. Ross in the late 1960s (Ross, 1967, 1969). The approach plays itself out very naturally with respect to the deep structure/surface structure distinction: elliptical and non-elliptical forms have the same deep structure but different surface structures, the difference arising from the application of transformations that delete syntactic structure to give the elliptical forms. That (1) and (2) fall together was taken as a natural consequence of assuming that the transformational rules involved delete syntactic constituents and do so under identity with an antecedent constituent. So, for example, because there is an antecedent occurrence of the verb phrase 'went to the store', 'Oscar did, too' can be derived from 'Max went to the store' by VP-deletion.

During this period in the development of generative grammar, it was generally assumed that deletion transformations were constrained by a general theoretical condition that required the 'recoverability' of deletions; cf. Peters and Ritchie (1973) for discussion of the formal importance of this condition for the theory of transformations. Informally put, a deletion transformation satisfies this condition if it is possible to reconstruct the deleted material from within the structural context that the rule applies. Deletions that apply under identity obviously meet this criterion. For example, VP-deletion would appear to meet the condition because within the overall sentence 'Max went to the store, and Oscar did, too' we can reconstruct the deleted VP as a copy of the antecedent VP. However, two problems were noticed that indicated quite clearly that such deletion rules fail to meet the recoverability criterion.

The first problem is that the domains in which transformational rules apply are not the same as the domains in which deleted constituents can match up with their antecedents. While the domain of transformational rules is the sentence, the antecedent/deletion pairing is not restricted to this domain but transcends sentential boundaries. So, in our example of VP-deletion, we could replace the conjunction with a full stop, turning one sentence into two: (1) 'Max went to the store' and (2) 'Oscar did, too'. Moreover, we could imagine discourses in which other sentences would be interpolated between them or, even more telling, discourses in which the two sentences were uttered by different speakers. VP-deletion, it would appear, is not a rule of *sentence* grammar, but of *discourse* grammar.

Notice that broadening the applicability of syntactic deletion operations to discourse does not impact on what would seem a more fundamental aspect of recoverability, the reconstructive aspect of deletion given by the syntactic identity of the deleted constituent and its antecedent. The following observations, initially made by Ross (1967) do, however, and herein lies the second problem. Consider the following sentence:

Max saw his mother, and Oscar did, too (1)

Clearly, there is a reading of (1) that entails that Max and Oscar saw one and the same person; they each saw Max's mother. This interpretation could easily be obtained on the transformational view by taking the deep structure of (1) to be roughly (1′):

Max saw Max's mother, and Oscar saw Max's mother (1′)

(1) can derived from (1′) first by applying VP-deletion, applicable given the identity of the verb phrases, to be followed by pronominalization in the first clause. The construal of (1) characterized in this way is not, however, the only construal available of this sentence. It can also be understood in a manner comparable to (1″):

Max saw Max's mother, and Oscar saw Oscar's mother (1″)

The problem is that although (1″) ought to be a possible underlying source for (1), it is not a structure to which VP-deletion can apply, since the verb phrases are not identical. Insofar as (1) can be derived from an underlying form like (1″), it will be a nonrecoverable deletion.

Semantic Non-Reduction: The Property Theory

The ambiguity shown by (1) between a 'strict identity' reading (1′) and a 'sloppy identity' reading (1″)

thus scotches a transformational account of the sort envisaged as deletion under identity. But what went wrong? The answer that emerged is that the underlying problem lies with the initial presumption, that semantic constancy can be reduced to syntactic constancy. Deletions, in some sense, must still be recoverable, not least because knowing what has been deleted is necessary for understanding an elliptical sentence. But what is to be recovered is not syntactic information *per se*, but information about the logico/semantic roles played by syntactic expressions; it is in the identity conditions applicable to these roles that we are to look for the conditions that govern deletion. The semantic constancy of ellipsis, in this view, is not to be reduced to something else, but is to have a direct semantic analysis.

The nonreductionist approach found its fullest hearing in the work of Ivan Sag (1976) and Edwin Williams (1977) in the mid-1970s. The central observation animating this approach is that a sentence such as 'Max saw his mother' is ambiguous, depending upon whether the verb phrase 'saw his mother' expresses the property of seeing Max's mother or the property of seeing one's own mother. These two properties are not unrelated; the latter is an abstraction of the former, being general where the former is particular. This ambiguity, however, is masked in simple sentences because 'Max saw his mother' has the same truth conditions under either interpretation of the verb phrase. It becomes unmasked in elliptical contexts: depending on which property the ellipsis is taken as being identical with, different interpretations are obtained. Thus, if in 'Max saw his mother, and Oscar did, too', the ellipsis is understood as the property 'saw Max's mother' the strict reading ensues; if the ellipsis is understood as 'saw his own mother', the sloppy reading follows. In either case note that what is elided is identical with the antecedent; they each express the same property. In this view, the constancy of ellipsis is thus a matter of property identity, and this is an inherently semantic notion.

These ideas found a natural representation by assuming that the logical representation of natural language incorporate aspects of a λ-calculus. These sorts of logistic systems (first introduced by Church), incorporate an operation that abstracts properties from propositions. This operation, known as λ-abstraction, derives from the proposition expressed by 'Max read *Moby Dick*' the property: λx (x read Moby Dick). This is to be parsed as a λ-operator binding a variable in the following open sentence; it is interpreted as a characteristic function, taking an individual as argument and returning a truth value. If we supply an argument for this function, represented by placing it before the λ-expression: Max, λx (x read Moby Dick), we can return to our original proposition by the inverse operation, λ-conversion; we effect this by placing the argument in the place of the variable in the open sentence, and erasing the λ-operator.

In their analysis of ellipsis, Sag and Williams make two basic assumptions about the semantics of natural language. First, following a suggestion of Barbara Partee (1975), they assume that verb phrases are interpreted as λ-expressions, so that the logical form of 'Max read Moby Dick' would be represented as immediately above. Second, they assume that anaphoric pronouns can be interpreted as either constants or variables. Taking these together, it follows that 'Max saw his mother' is representationally ambiguous between the following logical forms: 'Max, λx (x saw Max's mother)' and 'Max, λx (x read x's mother)'. In the first form, the pronoun is represented as a constant specifying its anaphoric reference; in the latter, the pronoun occurs as a bound variable. Thus far these representations are only distinguished formally; via λ-conversion both convert to the same proposition. The distinction becomes more than this, however, when a third assumption comes into play: VP-ellipsis requires identity of λ-expressions. This gives two representations for 'Max saw his mother, and Oscar did, too':

Max, λx (x saw Max's mother) & Oscar, λy (y saw Max's mother) (2)

Max, λx (x saw x's mother) & Oscar, λy (y saw y's mother) (3)

The first representation is of the *strict* reading; the second clause means Oscar saw Max's mother. The second representation is of the *sloppy* reading; in it, the second clause means that Oscar saw his own mother; that is, Oscar saw Oscar's mother.

The part of the identity conditions in λ-expressions relevant to strict and sloppy identity is known as the alphabetic variance condition. This condition breaks down into two subcases. The first applies to λ-expressions if there are only bound occurrences of variables; in this case, the λ-expressions must be exactly the same up to alphabetic values of the variables. Thus, in (3) for example, the λ-expression on the left is nondistinct from that on the right because they are alphabetic variants, differing only in that where 'x' occurs on the left, 'y' occurs on the right. That is, it matters not that we have 'x' on the right and 'y' on the left so long as the

pattern of binding remains unaltered. When this is changed, the condition is not satisfied; a consequence of this is that a sloppy reading is unavailable in 'Max saw his mother, and Oscar believes Jane did, too'; i.e. the right-hand clause cannot mean that Oscar believes that Jane saw Oscar's mother. This is because in the following representation, which would represent this reading, there are no λ-expressions that are alphabetic variants: 'Max, λx (x saw x's mother)' and 'Oscar, λz (z believes Jane, λy (y saw z's mother))'. In particular, 'λx (x saw x's mother)' and 'λy (y saw z's mother)' are not alphabetic variants because 'z' is free within the λ-expression while the corresponding occurrence of 'x' is bound. Not all free occurrences of variables within λ-expressions are illicit however; this is the effect of the second case of the condition that permits λ-expressions to be alphabetic variants only if the free variables are all bound by the same operator. This is what we find in the representation of 'Max saw everyone before Oscar did':

$$\forall x \text{ (Max, } \lambda y \text{ } (y \text{ saw } x) \text{ before Oscar, } \lambda z \text{ } (z \text{ saw } x))$$

Although within each λ-expression there is a parallel occurrence of 'x' free, they are both bound by the universal quantifier, and hence are alphabetic variants.

The success of this account in overcoming the problems that plagued the prior transformational approach extends beyond the account of strict and sloppy identity. Because the notion of property identity on which the account depends is semantic, unlike in the syntactic account, which was limited by the structural extent of structural descriptions of transformational rules, there is nothing comparable that inherently restricts the context in which the identical properties may occur. Therefore, in the absence of some external constraint, the antecedent of an ellipsis may occur in positions quite detached in discourse from the ellipsis itself; the sentence in which the antecedent is expressed neither needs to be adjacent to the elliptical sentence, nor need it be uttered by the same speaker. All that is required is that the antecedent of the ellipsis be sufficiently salient in the surrounding context.

Property theory and the syntax of ellipsis

The property theory approach initiated by Sag and Williams has been highly influential in the study of ellipsis, and there have been any number of variations of this view. One important source of these variations arises from a changed perspective regarding what is the main syntactic issue raised by elliptical constructions. In the transformational deletion analysis, the concern was over what we may call the 'generation' problem: what are the syntactic operations that produce elliptical structures? But on the property theory, which assigns the explanatory role to semantics for the matters that were so troublesome *vis-à-vis* recoverability for the prior account, the focus is shifted to the mapping problem: how are syntactic structures translated into semantic structures? In particular, how are verb phrases translated into λ-expressions that satisfy the identity conditions (alphabetic variance)? Understanding the mapping problem in this way places a constraint on its solutions; however syntactic derivation is to be effected (i.e. whatever the solution to the generation problem is), it must be such that it allows for systematic translation. There are two broad approaches to the mapping problem falling within this constraint that differ in their views of the need to attribute syntactic structure to the ellipsis in order to generate the property it expresses.

The first, the rich syntax view, assumes that elliptical structures retain a syntactic relation to forms in which all constituents are structurally present. Thus, whatever procedures translate 'Oscar read *Moby Dick*' also translate 'Oscar did, too', because at the input to the translation, the latter has the same structure as the former. The rich syntax view thus calls for the 'reconstruction' of syntactic information in order to obtain a property that then serves as the basis of comparison for identifying salient antecedent properties in the context of an ellipsis. This is the view Sag and Williams take, although they differ on the derivational direction of this relation. Sag takes the elliptical structure to be derived from the non-elliptical by deletion, while Williams reverses the direction of the derivation, the non-elliptical arising from the elliptical by syntactic copying. The alternative, the poor syntax view, sees no need for such a syntactic relation; in this view, at all stages of derivation the elided phrase is missing, or if not actually missing is structurally noncomplex; i.e. an empty category with no internal constituents (cf. Hardt, 1993). Since there is no verb phrase, there is also no translation to a property, however, and hence, in contrast to the rich theory, there is no basis of comparison for determining the antecedent of an ellipsis. But, according to this view, there is in fact no need to derive this, for there is an independent model to draw upon, the anaphoric resolution of pronouns. In both the ellipsis/antecedent and the pronoun/antecedent relations, a possibly complex antecedent for a syntactically simple element is determined by the conditions on salience in

context. Thus, the relation of ellipsis to antecedent in 'Max read *Moby Dick*, and Oscar did, too' is comparable to the relation of pronoun to antecedent in discourse (e.g. 'Herman Melville wrote *Billy Budd*. He is more famous, however, for *Moby Dick*') or even more directly to VP-anaphors like 'so' and 'it' (e.g. 'Max read *Moby Dick*, and so did Oscar') for which elided verb phrases are, in this view, the covert analogues. In the poor syntax view, then, there is no reconstruction as in the rich view; information relevant to the resolution of the ellipsis is only that which is found in the antecedent.

Problems with the property theory

In subsequent research a number of problems have emerged with the property theory. One case, initially observed by Shalom Lappin (1984) calls into question the validity of the second clause of the identity condition in λ-expressions on the basis of examples such as (4):

I know which book Max read, and which book Oscar didn't (4)

Recall that the restriction that λ-expressions are alphabetic variants only if the free variables are all bound by the same operator is what accounted for 'Max saw everyone before Oscar did' as discussed above. But in allowing for this case, (4) should be disallowed, for here the λ-expression corresponding to the elided phrase – λz (z read x) – contains a free variable bound by a different operator (i.e. the *wh*-phrase in the second clause) than that which binds the free variable in the λ-expression corresponding to the antecedent phrase.

A second sort of case, derived from initial observations of Schiebe (1971) and Dahl (1974), is among a series of cases most extensively discussed in joint research by Robert Fiengo and Robert May (Fiengo and May, 1994), who label them the 'eliminative puzzles of ellipsis'. It does not pertain directly to the identity condition but to a systematic overgeneration problem in the translation of pronouns. This case, which Fiengo and May called the 'many-pronouns puzzle', arises when the number of pronouns is increased beyond the one found in the standard examples used to illustrate strict and sloppy identity:

Max said he saw his mother, and Oscar did, too (5)

Here, the expectation is that there should be a four-way ambiguity. This is because whether a pronoun is translated as a constant or a variable is independent of how any other pronoun is translated, predicting, for n-many pronouns, 2^n readings. Thus in (5) the pronouns could be (i) both variables, (ii) both constants, (iii) the first one a variable, the second a constant, or (iv), vice versa, the first a constant, the second a variable. Given the correspondence of variables with the sloppy reading, and constants with the strict reading, (i) and (ii) will result in 'across-the-board' sloppy and strict readings, respectively, while (iii) and (iv) will give readings mixed between sloppy and strict. The problem is that in (5) we observe only three, not four, readings: precluded is the reading corresponding to (iv). The ellipsis in (5) cannot be glossed as '… and Oscar said Max saw Oscar's mother'. What we actually observe in this case, as well as in those of increasing complexity, are only $n+1$ readings; readings do not grow exponentially.

Syntactic Reduction Redux: The Dependency Theory

What the 'many-pronouns puzzle' indicates is that the assumption of translational independence embedded in the property theory's account of anaphora is incorrect, suggesting instead that what is involved is some dependence relation between the pronouns. The most highly developed theory that seeks to capture these dependencies is the dependency theory developed by Fiengo and May (1994). Central to the dependency theory picture is that dependencies must have a syntactic characterization, and in establishing this result, Fiengo and May turn away from the nonreductive account of ellipsis embedded in the property theory, and return to a reductive syntactic approach. The dependency theory assumes that a sentence such as 'Max saw his mother' is structurally ambiguous, the ambiguity being attributed to a distinction in the representation of anaphoric pronouns that indicates whether the pronoun is formally dependent on its antecedent or not. By Fiengo and May's conventions, grammatically anaphoric pronouns are represented by co-indexing, with those that are formally dependent on their antecedent marked by 'β', those that are not by 'α'. 'Max saw his mother' thus has the following pair of representations: 'Max_1 saw his_1^{α} mother' and 'Max_1 saw his_1^{β} mother'. The dependency that a β-marked pronoun enters into is specified via a structural description of the sequence of categories that lies between the pronoun and its antecedent. So for the latter structure, this would be the dependency: $\langle (Max, his), 1, \langle NP, V, NP \rangle \rangle$, where the first member of the triple is the elements of the dependency, the (unique) antecedent and the dependent pronouns, the second

the index of the elements in the dependency, and the third the string of categories that links the co-indexed elements together.

The dependency theory is applied to ellipsis by holding that the identity conditions that allow for ellipsis are satisfied in the following two sorts of structure, where, given the co-indexings, the first represents the strict reading and the second represents the sloppy reading.

Max_1 saw his_1^{α} mother, and Oscar_2
saw $\textit{his}_1^{\alpha}$ *mother*

Max_1 saw his_1^{β} mother, and Oscar_2
saw $\textit{his}_2^{\beta}$ *mother*

While it is apparent that in the first structure the antecedent and elided verb phrases are simple syntactic copies, it is not in the second, for the pronouns are different in the two verb phrases. This discrepancy is reconciled by allowing for an identity condition such that dependencies are the same so long as there is the same sequence of categories, regardless of the index; dependencies that stand in this relation – same pattern, different index – Fiengo and May call '*i*-copies'. The dependencies in which the pronouns in the latter structure occur meet this criterion: since they are *i*-copies, they are sufficiently alike to allow for ellipsis even though the pronouns are syntactically distinct. On the other hand, where there are no dependencies to be calculated – where the pronouns are marked α, not β – there is no alternative but for the index of the pronoun to be unchanged.

Sloppy identity on the dependency theory view is thus the re-creation of an antecedent structural pattern of anaphora; strict identity, on the other hand, is the re-creation of the anaphora itself. Either may be extended to more complex structures, but when they are certain limitations arise. Thus, recall the 'many-pronouns puzzle' that swirled around examples such as (5) above – 'Max said he saw his mother, and Oscar did, too'. For this case, there are four possible combinations of indices for the antecedent; of these, only three give rise to well-formed elliptical structures:

Max_1 said he_1^{α} saw his_1^{α} mother, and Oscar_2
said $\textit{he}_1^{\alpha}$ *saw* $\textit{his}_1^{\alpha}$ *mother*

Max_1 said he_1^{β} saw his_1^{β} mother, and Oscar_2
said $\textit{he}_2^{\beta}$ *saw* $\textit{his}_2^{\beta}$ *mother*

Max_1 said he_1^{β} saw his_1^{α} mother, and Oscar_2
said $\textit{he}_2^{\beta}$ *saw* $\textit{his}_1^{\alpha}$ *mother*

*Max_1 said he_1^{α} saw his_1^{β} mother, and Oscar_2
said $\textit{he}_1^{\alpha}$ *saw* $\textit{his}_2^{\beta}$ *mother*

The first case is the across-the-board strict reading; since both pronouns are α, they have the same index in the antecedent and the ellipsis. The second case is across-the-board sloppy: the two pronouns are in a dependency, with 'Oscar' as the antecedent in the same way structurally as the pronouns in the prior clause are in a dependency with 'Max'. The third structure represents the mixed reading: the first pronoun is β, and thus may be in a dependency. The second pronoun, however, is α so it must be strict, i.e. co-indexed with 'Max', not 'Oscar'. The final case is the one that is excluded. This is because the dependency that reaches from the pronoun 'his' to 'Oscar' as antecedent is not structurally identical to any dependency in the prior clause. Insofar as there is a dependency in that clause, it must be to the closer possible antecedent, the pronoun 'he', but this is not structurally parallel to the dependency in the clause with the ellipsis that has a greater syntactic extent.

Antecedent-Contained Deletion

Thus far, the structural context of ellipsis we have considered has been that of a discourse; i.e. the ellipsis and its antecedent have each occurred in independent sentences. While the examples we have examined are ones in which these sentences are conjoined, this is not essential, for in all the examples 'and' could be replaced by a full stop. The one exception was 'Max saw everyone before Oscar did', but notice here that the ellipsis is contained in an adjunct, not a subordinate, clause, so that even this case falls under the generalization that the ellipsis and its antecedent are syntactically independent. The generalization, however, appears to clearly fail in the following case:

Dulles suspected everyone that Angleton did (6)

Sentence (6) is well formed, and is naturally understood to mean the same thing as its unelided counterpart, 'Dulles suspected everyone that Angleton suspected'. But in this case the elided VP is not independent of its antecedent – rather it is contained within it – and hence the name for this construction, 'antecedent-contained deletion'. Notice that not only does antecedent-contained deletion appear to run counter to the generalization, but it does so in a particularly curious way. We understand the antecedent of the elided VP to be that VP headed by the verb 'suspected'. But if we plug that VP into the ellipsis, it will again contain the ellipsis: 'Dulles suspected everyone that Angleton suspected everyone that Angleton did'.

Additional iterations of the process will continually give structures that still contain an ellipsis. Given this vicious regression, it is unclear how we are to establish the relation of ellipsis and antecedent.

The critical insight here was provided by Sag (1976). Sag observed that if we attend to the logical form of (6), that at the appropriate level of syntactic description it falls under the antecedence generalization. This is because (6) contains a quantified phrase, and this must be scoped out; if we assume that the logical structure of (6) is roughly as follows: '$\forall x$: Angleton *suspected* x (Dulles suspected x)', we then need only observe that here the ellipsis (filled in and indicated by italics) is no longer contained within the antecedent. (This logical form is comparable to the logical form that would be assigned to the fully lexicalized counterpart of (6).) Thus, the significance of antecedent-contained deletion, given the generalization regarding the relation of ellipsis and antecedent, is that the notion of structure relevant to this generalization must be sufficiently abstract so as to represent the logical form of sentences. In May (1985) it is argued that this structure is a form of syntactic structure, a result of the syntactic rule QR, which gives (7) as a representation of (6):

[everyone that Angleton did [Dulles suspected t]] (7)

In this structure the ellipsis is no longer contained within the antecedent, and so we can now plug in the antecedent VP without any regress:

[everyone that Angleton *suspected* t [Dulles suspected t]] (7′)

May's account of antecedent-contained deletion has been one of the main arguments that has been cited in support of the view that there is a syntactic level – *LF* – that represents the logical structure of natural language; see May (1985, chap. 1), as well as Hornstein (1995), for a contrary view.

VP-ANAPHORA

Previously, we briefly mentioned the phenomena of VP-anaphora, citing examples like 'Max hit Oscar, and Harry did it, too' or 'Max hit Oscar, and Harry did so, too'. These cases are closely akin to VP-ellipsis except that they have an anaphoric element – 'so' or 'it' – rather than an ellipsis. VP-anaphora is distributionally more restricted than VP-ellipsis. For example, with stative verbs, VP-ellipsis is possible, but not the 'it' form of VP-anaphora, and the 'so' form is marginal; compare 'Max knows French, and Oscar does, too' with 'Max knows French, and Oscar does it, too' and 'Max knows French, and Oscar does so, too'. However, what is more relevant to the present discussion is that with VP-anaphora we can find the same ambiguity of strict and sloppy identity that we observed with VP-ellipsis; compare 'Max hit his mother, and Oscar did, too' with (8):

Max hit his mother, and Oscar did it, too (8)

As before, the second clause can be taken to mean that Oscar hit Max's mother (strict) or that Oscar hit Oscar's mother (sloppy).

Examples like (8) pose issues as to how we are to understand ellipsis. If VP-anaphora and VP-ellipsis display uniform behavior, it would seem natural to posit a uniform analysis. One way to do this would be to take ellipses as anaphoric elements, silent counterparts, if you will, of 'it' or 'so'. (For one account along these lines, couched in a variant of the property theory, see Hardt (1993).) In this view, in which VP-ellipsis is reduced to VP-anaphora, that ellipsis involves some sort of syntactic reconstruction is effectively denied. The alternative would be to run the reduction in the opposite direction by maintaining that at an appropriately abstract syntactic level VP-anaphora, like VP-ellipsis, is syntactically complex, and that the overt pronominal elements are but superficial syntactic reflexes. In either account, the goal would be to isolate a level of representation in which VP-anaphora and VP-ellipsis are structurally nondistinct in order to account for their common behavior.

This search for a common analysis, however, needs to be weighed against ways in which VP-anaphora and VP-ellipsis do not cluster but diverge in properties. One sort of divergence has already been mentioned: namely, the distributional differences. Another can be gleaned from examples of antecedent-contained deletion, as in (9):

Dulles talked to everyone that Angleton did (9)

If VP-ellipsis were just a variant of VP-anaphora, then we would expect that the VP-anaphora version of (9) would also be grammatical. But, as was observed by Fiengo and May (1994), it is not:

*Dulles talked to everyone that Angleton did it (9′)

This case indicates (along with the distributional facts cited above) that there are substantial differences between VP-anaphora and VP-ellipsis. It remains an open research question how observations like these can be integrated with those about

strict and sloppy identity; until then it remains equivocal whether VP-ellipsis and VP-anaphora are a unified phenomenon.

SUMMARY

Returning to our initial observation, we have seen that capturing the basic intuition of semantic constancy under ellipsis devolves to the issue of the proper way to state the identity conditions that govern ellipsis. That ellipsis in general requires a notion of identity was the initial insight of the transformational account; its flaw, in a sense, was that the notion of identity it had available was not sufficiently abstract. What followed were attempts to find the right locus of abstractness for ellipsis. The property theory argues that the appropriate identity conditions are semantic and does not fundamentally question the relative lack of abstractness of the underlying syntax. The dependency theory, in contrast, argues for a more abstract notion of syntax by refining the criteria for identity of occurrences of syntactic categories. But regardless of how matters turn out, it is clear that in seeking to fix the identity conditions for ellipsis, issues fundamental to our conceptions of linguistic description have been raised. Phenomena relevant to these issues extend beyond what we have been able to consider here. Among them are further interactions with anaphora, such as the 'vehicle change' effect noticed by Fiengo and May (1994): the observation that 'Mary loves John, and he thinks that Sally does, too' is interpreted as comparable to 'Mary loves John, and he thinks that Sally loves him, too', not 'Mary loves John, and he thinks that Sally loves John, too'. Moreover, we have left aside discussion of the conditions on discourse that allow an ellipsis to be resolved, including conditions on discourse coherence with respect to antecedents that may occur in sentences at some degree of removal in the discourse (cf. Kehler, 2000), and on the abstractness of discourse: e.g. consider when Butch says to Sundance at the edge of the cliff before jumping: 'I will if you will', indicating that the antecedent need not even be uttered (cf. Chao, 1987). We have tried, however, to highlight some of the core issues that have animated the discussion and that have made understanding elliptical phenomena of continued interest within linguistic theory.

References

Chao W (1987) *On Ellipsis*. PhD dissertation. University of Massachusetts, Amherst, MA.

Dahl Ö (1974) How to open a sentence: abstraction in natural language, *Logical Grammar Reports*, no. 12. Göteberg, Sweden: University of Göteberg.

Fiengo R and May R (1994) *Indices and Identity*. Cambridge, MA: MIT Press.

Hardt D (1993) *VP Ellipsis: Form, Meaning, and Processing*. PhD dissertation, University of Pennsylvania.

Hornstein N (1995) *Logical Form*. Oxford, UK: Blackwell.

Kehler A (2000) Coherence and the resolution of ellipsis. *Linguistics and Philosophy* **23**: 533–575.

Lappin S (1984) VP-anaphora, quantifier scope, and logical form. *Linguistic Analysis* **13**: 273–315.

May R (1985) *Logical Form: Its Structure and Derivation*. Cambridge, MA: MIT Press.

Partee BH (1975) Deletion and variable binding. In: Keenan E (ed.) *Formal Semantics of Natural Language*, pp. 16–34. Cambridge, UK: Cambridge University Press.

Peters PS and Ritchie RW (1973) On the generative capacity of transformational grammars. *Information Sciences* **6**: 49–83

Ross JR (1967) *Constraints on Variables in Syntax*. PhD dissertation. Massachusetts Institute of Technology, Cambridge, MA.

Ross JR (1969) Guess who? In: Binnick RI, Davison A, Green GM and Morgan JL (eds) *Papers from the Fifth Regional Meeting of the Chicago Linguistic Society*, pp. 252–286. Chicago, IL: Department of Linguistics, University of Chicago.

Ross JR (1970) Gapping and the order of constituents. In: Bierwisch M and Heidolph K (eds) *Progress in Linguistics*. The Hague, Netherlands: Mouton.

Sag I (1976) *Deletion and Logical Form*. PhD dissertation. Massachusetts Institute of Technology, Cambridge, MA.

Schiebe T (1971) Zum Problem der grammatisch relevanten Identität. In: Kiefer F and Ruwet N (eds) *Generative Grammar in Europe*. Dordrecht, Netherlands: Reidel.

Williams E (1977) Discourse and logical form. *Linguistic Inquiry* **8**: 101–139.

Further Reading

Berman S and Hestvik A (eds) (1992) *Proceedings of the Stuttgart Ellipsis Workshop*. Heidelberg: Arbeitspapiere des Sonderforschungsbereichs 340, Bericht Nr. 29, IBM Germany.

Chung S, Ladusaw WA and McCloskey J (1995) Sluicing and logical form. *Natural Language Semantics* **3**: 239–282.

Dahl Ö (1973) On so-called 'Sloppy' identity. *Synthese* **26**: 81–112.

Dalrymple M, Shieber S and Perreira F (1991) Ellipsis and higher-order unification. *Linguistics and Philosophy* **14**: 399–452.

Hankamer J and Sag I (1984) Deep and surface anaphora. *Linguistic Inquiry* **7**: 391–426.

Kennedy C (1997) Antecedent contained deletion and the syntax of quantification. *Linguistic Inquiry* **28**: 662–688.

Lappin S (1996) The interpretation of ellipsis. In: Lappin S (ed.) *The Handbook of Contemporary Semantic Theory*, pp. 145–175. Oxford, UK: Blackwell.
Lappin S and Benmamoun EA (eds) (1999) *Fragments: Studies in Ellipsis and Gapping*. Oxford, UK: Oxford University Press.
Lobeck A (1995) *Ellipsis: Functional Heads, Licensing, and Identification*. New York, NY: Oxford University Press.
Merchant J (2001) *The Syntax of Silence – Sluicing, Islands, and the Theory of Ellipsis*. Oxford, UK: Oxford University Press.
Rooth M (1992) Ellipsis redundancy and reduction redundancy. In: Berman S and Hestvik A (eds) *Proceedings of the Stuttgart Ellipsis Workshop*. Heidelberg: Arbeitspapiere des Sonderforschungsbereichs 340, Bericht Nr. 29, IBM Germany.
Schwabe K and Zhang N (2000) *Ellipsis in Conjunction*. Tübingen, Germany: Max Niemeyer Verlag.
Steedman M (1990) Gapping and constituent coordination. *Linguistics and Philosophy* **13**: 207–264.
Zoerner CE (1995) *Coordination: The Syntax of &P*. PhD dissertation. University of California, Irvine, CA.

Embodiment

Intermediate article

Ronald Chrisley, University of Birmingham, Birmingham, UK
Tom Ziemke, University of Skövde, Skövde, Sweden

CONTENTS

An understanding of how cognition is realized or instantiated in a physical system, especially a body, may require or be required by an account of a system's embedding in its environment, its dynamical properties, its (especially phylogenetic) history and (especially biological) function, and its nonrepresentational or noncomputational properties.

INTRODUCTION

In recent years a number of researchers in cognitive science and artificial intelligence (AI) have criticized many traditional approaches to modeling, building and understanding cognitive systems as not placing sufficient emphasis on the body or physical realization of such systems. Non-embodied approaches to cognitive science typically involve some or all of the following features, to a greater or lesser extent:

- The belief that cognition is computation, and thus can be understood in an implementation-independent way, allowing cognitive science to proceed independently of biology and neuroscience.
- A search for general-purpose cognitive abilities, not relativized to any particular (biological, sensorimotor, physical) context or need.
- A method of analysis, modeling and design that for the most part ignores temporal aspects of cognition, in that it focuses on behaviors (e.g. chess playing) that are evaluated in terms of 'getting the right answer' rather than exhibiting a particular dynamic profile, and sees cognition as a module that mediates between the deliverances of a causally prior perceptual module and the inputs to an autonomous action system.

In contrast, embodied approaches to cognition typically involve some or all of the following features, again to varying extents:

- Acknowledgment of the role that the body and its sensorimotor processes can and do play in cognition. Some aspects of the system that would, on the traditional view, be considered mere matters of implementation, are instead taken to be crucial components.
- Understanding of cognition in the context of its (especially evolutionary) biological function: to support the activities of the body.
- A view of cognition as a real-time, situated activity, typically inseparable from and often fully interwoven with perception and action.

'Embodied' cognitive science or artificial intelligence, then, refers to a range of loosely affiliated philosophies, explanatory frameworks and design methodologies that strive to redress a perceived neglect of the body in cognitive science.

Since the mid-1980s there has been a rapid increase in interest in embodied cognition (and use of