THE LIMITS OF UNIFICATION

Robert J. P. Ingria

BBN Systems and Technologies Corporation
10 Moulton Street, Mailstop 6/4C
Cambridge, MA 02138
Internet: ingria@BBN.COM

ABSTRACT

Current complex-feature based grammars use a single procedure—unification—for a multitude of purposes, among them, enforcing formal agreement between purely syntactic features. This paper presents evidence from several natural languages that unification—variable-matching combined with variable substitution—is the wrong mechanism for effecting agreement. The view of grammar developed here is one in which unification is used for semantic interpretation, while purely formal agreement involves only a check for non-distinctness—i.e. variable-matching without variable substitution.

1 Introduction

In recent years, a great deal of attention has been devoted to complex-feature based grammar formalisms—i.e. grammar formalisms in which syntactic elements are not atomic symbols, but rather complex elements, such as value-attribute or term structures; see Shieber (1986) for an overview. Typically such formalisms use a single mechanism—variable substitution—for all purposes, and the most widely used variable substitution mechanism is unification.[1] Such complex-feature based grammars, then, are viewed as systems in which partial feature structures are built up, by the process of unification, into successively more specified structures. While it is formally elegant to use a single mechanism for a number of purposes, this theoretical elegance is realized in practice only if the mechanism does not require the other modules of the system to be complicated to achieve this "elegance". Currently, unification is used for at least four purposes:

- to enforce formal agreement between purely syntactic features

- to "percolate" features between a pre-terminal category and the phrase which it heads

- to pass features between a dislocated element—such as a WH-phrase—and its trace

- to build up semantic representations

This paper will focus on the use of unification to enforce agreement and will present evidence from several natural languages which argues against its use in the case of purely formal syntactic features: when such features are lexically or morphologically underspecified, they remain so, even under agreement, contrary to the predictions of a system using unification for agreement. Moreover, it is worthwhile stressing at the outset that the main argument of this paper is not that there are certain constructions that present a problem for unification, and, hence, require some technical solution. The point is much stronger: even if some elaborate analysis can be devised that allows unification to be used to effect agreement, this would be the wrong tack to take. Rather, the argument will go, using unification to effect agreement is incorrect both for theoretical reasons—it presents a view of language which is contradicted by the facts—and for practical reasons—using unification to effect agreement can impede a system's robustness and transportability.

2 The Paradox

A typical paradigm that is presented to show the almost transparent application of unification to agreement phenomena is the following:

(1) a. The sheep is ready.
   b. The sheep are ready.
   c. The sheep is there.
   d. The sheep are there.
   e. *The sheep that is ready are there.

---

[1] In the rest of this paper, for convenience I will use the term "unification" instead of "variable substitution", since it is the most commonly used type of variable substitution, but it should be borne in mind that the point being made here holds for variable substitution, in general.

Sentences (1a) through (1d) are taken to indicate that "sheep" is underspecified with regard to number; it can be either singular or plural. (1e), on the other hand, shows that "sheep" cannot be both singular and plural at the same time. In the relative clause, "is" is marked as singular, and "sheep", interpreted as its subject via the relative connector "that", must also be singular. On the other hand, "are" in the matrix clause is marked as plural, and "sheep", its subject, must also be plural. Under a unification analysis, these facts are explained in the following way: "sheep" is syntactically unspecified for the feature *number*. The process of subject-verb agreement is effected by unification. Therefore, when "sheep" appears as the subject of a finite verb, unification will fix its number as singular or plural (unless the finite verb itself is ambiguous). (1e) is ungrammatical, then, since the values *singular* and *plural* cannot unify and the fact that "sheep" must agree with both "is" and "are" in number would require their unification.

This illegal feature configuration is shown in (2).

(2) [V [num:sg]] → [N [num:{sg,pl}]] ← [V [num:pl]]
    ("is")   ("sheep")   ("are")

Here, the arrows indicate the notional flow of information under agreement, but have no theoretical status. They indicate that agreement between "sheep" and "is" would set "sheep" 's *number* feature to *singular*, while agreement with "are" would set it to *plural*. More generally, the unification approach to agreement rules out the following configuration:

(3) *[X [F:α]] → [Y [F:{α,β}]] ← [Z [F:β]]

Here [F : x] denotes feature F, with value x, and [F : {x,y}] indicates feature F with value either x or y, x and y distinct. Thus, this schema indicates that a category which is specified for the values α and β for feature F cannot simultaneously agree in this feature with categories that specify distinct values for F. In the rest of this section, I will show cases of constructions which match this schema but are still grammatical.

2.1 Case 1: German Free Relatives

In German, as Groos and van Riemsdijk (1979) demonstrate, free relative clauses require that the relative pronoun agree in Case both with the position of the relative clause as a whole and also with the position with which the relative pronoun is construed (i.e. with the gap which the relative pronoun fills). This is shown

in (4) and (5), where the matrix verb and the verb in the free relative are annotated with the Case the relative pronoun must bear in that clause.

(4) a. Wer nicht stark ist,   muss klug sein.
        who not strong is   must clever be
        NOM   NOM NOM
        'Whoever isn't strong must be clever.'
    b. *Wen Gott schwach geschaffen hat, muss klug sein.
        ACC   ACC   NOM
       *Wer Gott schwach geschaffen hat, muss klug sein.
        NOM   ACC   NOM
        who God weak   created   has must clever be
        'Who(m)ever God has created weak must be clever.'

(5) a. Ich nehme, wen du mir empfiehlst.
        I   take   who you me recommend
                   ACC ACC   ACC
        'I take whomever you recommend to me.'
    b. *Ich nehme, wen du vertraust.
        ACC ACC   DAT
       *Ich nehme, wem du vertraust.
        ACC DAT   DAT
        I   take   who you trust
        'I take whomever you trust.'

Assuming that "Case assignment" is actually a form of agreement between a verb and a noun phrase that it governs, the data in (4)–(5) seems to fit nicely into a unification approach. However, the neuter free relative pronoun *was*, which is both nominative and accusative, can seemingly agree with both nominative and accusative Case assigning elements at the same time:

(6) Was du mir gegeben hast, ist prächtig.
    What you me given   have is wonderful
    NOM/ACC ACC   NOM
    'What you have given to me is wonderful.'
(7) Ich habe gegessen was noch übrig war.
    I   have eaten   what still left was
                      ACC NOM/ACC   NOM
    'I ate what was left.'

Note that sentences (6) and (7) are precisely instances of schema (3), just as (1e) is. Hence, if the explanation of the ungrammaticality of (1e) is correct, we should expect (6) and (7) to be ungrammatical.

(8) a. [V [case:ACC]] → [N [case:{N,A}]] ← [V [case:NOM]]
       ("gegeben")   ("was")   ("ist")
    b. [V [case:ACC]] → [N [case:{N,A}]] ← [V [case:NOM]]
       ("gegessen")   ("was")   ("war")

A possible solution to this seeming paradox, which still uses unification to effect agreement, is the following.[2] Assume that Case in German is not a single-valued feature, but rather an array of the different Cases of the language, each of which takes on one of the values T or NIL. We can then handle the data above with the following feature specifications. (The (a) representations use a "path" notation, consisting of attribute-value pairs, like that in Shieber (1986); while the (b) representations use a term notation, with positional features, like that in Definite Clause Grammars (Pereira and Warren (1980)).)

(9) wer:
    a. [case: [nom: T] [gen: NIL]
        [dat: NIL] [acc: NIL]]
    b. (CASE T NIL NIL NIL)
(10) wem:
    a. [case: [nom: NIL] [gen: NIL]
        [dat: T] [acc: NIL]]
    b. (CASE NIL NIL T NIL)
(11) wen:
    a. [case: [nom: NIL] [gen: NIL]
        [dat: NIL] [acc: T]]
    b. (CASE NIL NIL NIL T)
(12) was:
    a. [case: [nom: T] [gen: NIL]
        [dat: NIL] [acc: T]]
    b. (CASE T NIL NIL T)

Assuming that a verb is only specified for the Case it assigns and is unspecified for the others, the Case specifications for verbs that take nominal complements would be:

(13) geschaffen,nehme,empfiehlst,gegeben,gegessen:
    a. [case: [acc: T]]
    b. (CASE ?val ?val ?val T)
(14) vertraust:
    a. [case: [dat: T]]
    b. (CASE ?val ?val T ?val)

Similarly, the Case specification for nominative Case assignment, whether this is a property of syntactic structures or of particular lexical items, would be:

(15) a. [case: [nom: T]]
    b. (CASE T ?val ?val ?val)

This solution works, then, because *was*, and no other free relative pronouns, specifies the value T for

more than one element in its Case array and because verbs and other Case "assigning" elements only specify a value for the Case they "assign", and for no others. This solution of factoring out seemingly contradictory values for a single feature into values of different features allows us to get around the superficial violation of the schema in (3). However, there are other constructions which are harder to decompose in this fashion.

2.2 Case 2: Hungarian WH Movement and Topicalization

Let us now turn to a more complicated example, from Hungarian, described in Szamosi (1976). In Hungarian, WH words, like full NPs, are marked as either definite or indefinite. The verb in Hungarian is also marked as definite or indefinite, in agreement with its complement. When the complement is an accusative noun phrase, the definiteness marking on verb and noun phrase is the same.

(16) a.   Akart   egy könyvet.
        he-wanted   a   book
        −DEF   −DEF
    b. *Akarta   egy könyvet.
        he-wanted   a   book
        +DEF   −DEF
        'He wanted a book.'
    c. *Akart   a   könyvet.
        he-wanted   the book
        −DEF   +DEF
    d. Akarta   a   könyvet.
        he-wanted   the book
        +DEF   +DEF
        'He wanted the book.'
    e. Egy könyv amit   akart
        a   book which he-wanted
                  −DEF −DEF
    f. *Egy könyv amit   akarta
        a   book which he-wanted
                  −DEF +DEF
        'A book which he wanted
    g. *Ez az a könyv amelyiket akart
        this that the book   which   he-wanted
                      +DEF   −DEF
    h. Ez az a könyv amelyiket akarta
        this that the book   which   he-wanted
                      +DEF   +DEF
        'This book is the one which he wanted.'

When the complement is a finite clause, the verb bears definite agreement.

(17) a. János akarta, hogy elhozzak egy könyvet.
    John wanted that I-bring a book
                +DEF   +DEF
   b. *János akart, hogy elhozzak egy könyvet.
    John wanted that I-bring a book
                −DEF   +DEF
   'John wanted me to bring a book.'

Finally, WH phrases and topicalized constituents in Hungarian typically appear immediately preceding the verb; verb and WH word or topicalized noun phrase must agree in definiteness.[3] From these constraints, it follows that WH phrases and topicalized noun phrases extracted from complement clauses must be marked definite. Since the clausal complement forces the verb to bear definite agreement, and since the WH word or topicalized NP must agree with the verb in definiteness, the WH word or topicalized NP can only be definite. This is shown in the following examples:

(18) Ez az a könyv amelyiket akarta   hogy elhozzam.
      this that the book which   he-wanted that I-bring
                        +DEF   +DEF   +DEF +DEF
   'This is the book which he wanted me to bring.'
(19) *Egy könyv amit   akarta   hogy elhozzak.
     a   book which he-wanted that I-bring
               −DEF +DEF   +DEF −DEF
   'A book which he wanted me to bring.'

However, certain Hungarian verb forms[4] bear an ending which is ambiguous between definite and indefinite. In sentences involving such verbs, the WH word may be indefinite.

(20) A könyv amit   akarnánk,   hogy elhozzon.
     the book which we-would-want that he-brings
               −DEF ±DEF   +DEF −DEF
   'The book which we would want him to bring.'
(21) Egy könyv akartam, hogy elhozzon.
     a   book I-wanted that he-brings
               −DEF ±DEF +DEF −DEF
   'It was a book that I wanted him to bring.'

Once again, the grammatical (20) and (21) match the prohibited schema (3) ("c" = "complementizer"):

(22) a. [N [def:−] → [V [def:{+,−}] ← [c [def:+]
         ("amit")   ("akarnánk")   ("hogy")

b. [NP [def:−] → [V [def:{+,−}] ← [c [def:+]
     ("egy könyv") ("akartam")   ("hogy")

Let us consider the consequences of expanding out the definiteness feature into an array of separate values, analogous to the German example. First, this would require the underspecified verb forms to be represented as in (23).

(23) akarnánk,akartam:
    a. [definiteness: [definite: T] [indefinite: T]]
    b. (DEFINITENESS T T)

Next, it would require that the WH pronouns be specified as in (24) and (25):

(24) amelyiket:
    a. [definiteness: [definite: T]]
    b. (DEFINITENESS T ?val)
(25) amit
    a. [definiteness: [indefinite: T]]
    b. (DEFINITENESS ?val T)

Note that when either of these pronouns appeared with an underspecified definite and indefinite verb, such as those in (23), it would wind up with the definiteness specification in (23). This would totally neutralize the definiteness/indefiniteness contrast in such cases. But, in fact, no such ambiguity of interpretation is reported: a definite or indefinite WH phrase or topicalized noun phrase that appears in the suitable configuration with one of these ambiguously definite or indefinite verbs is interpreted as uniquely definite or indefinite, as is consistent with its overt marking, and not as ambiguous between definite and indefinite, as the proposed unification analysis would require. Thus, this unification based solution to a problem of a morphological ambiguity entails an ambiguity of interpretation that is not attested.

Moreover, aside from the empirically incorrect predictions about semantic interpretation, there is a more fundamental problem with the unification account of agreement. As was pointed out above, treating agreement as unification implies that structures meeting the schema in (3) should be superficially ungrammatical. In fact, this seems to be universally false: in every case in natural language in which an element does not morphologically distinguish between two or more values of a feature—a situation often referred to as *morphological neutralization*—it behaves as if this distinction is also neutralized for purposes of agreement. That is, instead of the configuration in (3) being universally ruled out,

---

[3]The situation is actually somewhat more complex; see Szamosi (1976) for full details.
[4]The first person singular past indicative and the first person plural present conditional.

it is universally attested. This creates a paradox, since the ungrammaticality of (1e) seems to depend on the ungrammaticality of structures matching the configuration in (3). To demonstrate that this seeming paradox is supported by the data, in the rest of this section, other examples will be presented to show that the configuration ruled out in (3) recurs again and again across the languages of the world.

2.3 Case 3: Objects of Conjoined VPs

In French, as Kayne (1975) points out, it is possible to conjoin past participles following the past auxiliary and a weak ("clitic") object pronoun which is the common object of the conjoined participles, under the requirement that the verbs of the conjuncts assign the pronoun the same Case. This is shown in (26) and (27):

(26) Paul l'a   insulté et mis à la porte.
     Paul him-has insulted and put to the door
          ACC   ACC   ACC
     'Paul insulted him and threw him out.'
(27) *Paul l'a   frappé et donné des coups de pied.
     Paul him-has struck and given blows   of foot
          ACC   ACC   DAT
     'Paul struck him and gave him some kicks.'

However, once again, if the object pronoun is marked for more than one Case, the conjunction of participles assigning those Cases is allowed.

(28) Paul nous   a frappé et donné des coups de pied.
     Paul us   has struck and given blows   of foot
          ACC/DAT ACC   DAT
     'Paul struck us and gave us some kicks.'
(29) On sait   que la police t'a   frappé
     one knows that the police you-has struck
                              ACC/DAT ACC
     et donné des coups de pied.
     and given blows   of foot
         DAT
     'Everybody knows that the police struck you and gave you some kicks.'

Similar facts hold for Icelandic, as well, as Zaenen and Karttunen (1984) point out.

(30) a. *Hann stal og borðaði kökuna/kökunni.
        he   stole and ate   the cookie
             ACC   DAT   ACC   DAT
        'He stole and ate the cookie.'

b. Hann stal og borðaði köku.
   he   stole and ate   a cookie
        ACC   DAT   ACC/DAT
   'He stole and ate a cookie.'

And German also has similar data, as Pullum and Zwicky (1986) show:

(31) a. *Sie findet und hilft Männer/Männem.
        she finds and helps men
                ACC   DAT ACC   DAT
        'She finds and helps men.'
     b. Er findet und hilft Frauen.
        he finds and helps women
                ACC   DAT ACC/DAT
        'He finds and helps women.'

The French, Icelandic, and German examples fall into the now familiar configuration.

(32) [V [case:A]] → [NP [case:{A,D}]] ← [V [case:D]]

2.4 Case 4: Elided Verbs in German

Eisenberg (1973) points out that in conjoined German subordinate clauses, the verb in all the non-final clauses can be elided, under identity of person and number agreement.

(33)   ...weil Hans Bier und Franz Milch trinkt.
        because Hans beer and Franz milk   drinks
               3rd   3rd   3rd
        '...because Hans drinks beer and Franz, Milk.'
(34) *...weil ich Bier und du   Milch trinkst/trinke.
      because I   beer and you milk drinks
             1st   2nd   2nd   1st
      '...because I drink beer and you, Milk.'

However, in forms which neutralize the person marking on the verb, elision is fine:

(35) ...weil wir das Haus und die Muellers den Garten kaufer
     because we the house and the Muellers the garden buy
            1st   3rd   1st/3rc
     '...because we buy the house and the Muellers, the garden
(36) ...weil Franz das Haus und ich den Garten kaufen könnt
     because Franz the house and I   the garden buy   could
            3rd   1st   1st/3ı
     '...because Franz could buy the house and I, the garden.'

This is yet another instance of our infamous schema:

(37) [NP [per:1]] → [V [per:{1,3}]] ← [NP [per:3]]

3 Resolving the Paradox

The previous section presented a paradox. There seems to be evidence, in the form of ungramatical utterances such as (1e), that the configuration in (3) is ungrammatical. However, the rest of the section presented evidence from different constructions and different languages which strongly indicates that (3) is the standard agreement configuration throughout the languages of the world. In this section, I will resolve this paradox by proposing that agreement is not effected by unification and but rather by a test for non-distinctness of feature values.

3.1 Neutralization versus Ambiguity

First, let us return to example (1e), repeated here for convenience

(38) *The sheep that is ready are there.

Recall that the analysis of this utterance which argued for the ungrammaticality of configuration (3) was based on the assumption that "sheep" is unspecified or underspecified for number. Note that this analysis is tenable if syntactic features alone are considered: syntactically, it seems plausible that "sheep" either has no *number* feature or that it has a variable, rather than a constant, as the value of this feature. However, when the ramifications of this analysis for semantics are considered, it becomes less tenable: while syntactic frameworks have been constructed in which features can take on underspecified values, most semantic frameworks require features such as singular/plural to be fully specified. That is, semantically, "sheep" can denote an individual or a set of individuals[5] but it cannot denote something indeterminate. This suggests that "sheep" is not underspecified, or vague, but rather ambiguous. That is, there is not a single representation for "sheep", which is underspecified for number, but rather two distinct entries, fully specified for number in both its syntactic and semantic aspects. If this is the case, the reason that (38) is ungrammatical is not that unification has filled in the underspecified value for *number*, but rather that subject-verb agreement disambiguates which of the two senses of "sheep" has been encountered and once one of

the fully specified entries is chosen, it naturally cannot agree with a constituent which bears a distinct number.

Once utterances like (38) are analyzed as not matching the agreement configuration in (3), it is possible to handle all the cases of morphological neutralization discussed in the previous section. Note that the feature involved in each example of neutralization discussed—Case in German and French, and definiteness on verbs in Hungarian—is either inherently formal, without semantic content (Case) or a feature that does not have any semantic ramification for the category in which it is neutralized: definiteness does affect the interpretation of noun phrases, but it serves purely as a formal agreement marker on verbs. If this observation is correct, then the solution to the apparent paradox runs along the following lines:

- Syntactic features which have semantic ramifications, such as number on nouns, tense on verbs, degree on adjectives, are never neutralized (underspecified). They are always fully specified and items which seem to be underspecified with regard to them are, in fact, ambiguous items with distinct, fully specified representations. (But see the discussion in the Section 4.)

- Purely formal syntactic features, on the other hand, can be neutralized, producing truly underspecified representations, either through the use of value disjunction or through the use of a variable, rather than a constant, as a feature value.[6]

- Agreement is effected not by unification but rather by a non-distinctness check.

That is, since we can view unification as logically composed of two parts—variable checking and variable substitution—agreement should be analyzed as involving only variable-matching, but not variable substitution. This would explain why constituents that neutralize a syntactic feature distinction are universally able to behave as if they are simultaneously marked for all the values of the feature that they neutralize: since agreement only involves variable matching, but not variable substitution, the original, underspecified representation is always available for agreement.

To make this proposal clearer, I will present an analysis of the German and Hungarian facts,[7] using a

---

[5]Nothing in the present argument hinges on this being the correct treatment of the singular/plural distinction. It does not matter which of the various proposals about the semantic interpretation of number is chosen. All that matters is that semantic theories require that singular and plural have different denotations, and do not allow indeterminate representations.

[6]Pullum and Zwicky (1986, p. 766) make a similar distinction between features "freely chosen" vs. those "syntactically imposed...by rules of agreement ... or government".

[7]For concreteness, I have analyzed nominative Case assignment as being a property of verbs in German. It is possible that this is a

term structure type notation, and adding the :OR operator, which introduces disjunctions of variable-free terms.[8]

German:	
Item:	Case:
wer	(nom)
wem	(dat)
wen	(acc)
was	(:or (nom) (acc))
empfiehlst...	(acc)
vertraust	(dat)
ist,war...	(nom)

Hungarian:	
Item:	Definite:
amit	(−)
amelyiket	(+)
hogy	(+)
akart	(−)
akarta	(+)
akarnánk,akartam	?val

These representations are matched by a non-distinctness check, which performs the same tests as unification. However, the non-distinctness check differs from unification in what it returns. Unification, when applied to two expressions, typically returns either a distinguished symbol, such as **Fail**, if they do not unify, or a single substitution expression, which is the *most general unifier* of its input; see e.g. Pereira and Shieber (1987, pp. 63–64). When two expressions are identical, this substitution expression is empty, since no substitutions need to be performed. In this case, then, unification effectively leaves its input unchanged. Thus, unification can be viewed as returning a single indicator of failure and an unbounded set of substitution expressions. Non-distinctness checking, on the other hand, returns a single indicator of failure but also a single indicator of success—an empty substitution expression. Alternatively, non-distinctness checking may be viewing as determining that two expressions are unifiable, without actually unifying them.

The following table contrasts the behavior of unification (⊔) and non-distinctness (≈):

Case:	⊔	≈
1. x, y are variable-free and non-disjunctive:		
a. $x = y$	NIL	NIL
b. otherwise:	Fail	Fail
2. x, y contain variables but are non-disjunctive:		
a. $\exists MGU(x, y)$	$MGU(x, y)$	NIL
b. otherwise:	Fail	Fail
3. x, y are both disjunctions:		
a. $x \cap y \neq \emptyset$	$(x \leftarrow x \cap y,$ $y \leftarrow x \cap y)$	NIL
b. otherwise:	Fail	Fail
4. x is a disjunction:		
a. y is a term in x	$(x \leftarrow y)$	NIL
b. otherwise:	Fail	Fail
5. y is a disjunction:		
a. x is a term in y	$(y \leftarrow x)$	NIL
b. otherwise:	Fail	Fail

where $MGU(x, y)$ is the *most general unifier* of x, y; NIL is the empty substitution expression; and $(\alpha \leftarrow \beta)$ indicates a substitution expression in which β substitutes for α.[9]

In examples (4) and (5) in German, and (16)—(19) in Hungarian, clause 1 applies. Since the terms involved in the agreements in all these examples are variable-free, the results are identical under the unification and non-distinctness analyses. In the German examples (6) and (7), which involve *was*, clauses 4 and 5 are the relevant ones, and it is here that the difference between the unification approach to agreement and the non-distinctness approach is apparent. Under the unification approach, once the disjunctive Case feature value associated with *was* unifies with a fully specified Case feature, a substitution list is produced that replaces the disjunction with one of its values:

(39) (:or (nom) (acc)) ⊔ (nom)
  ⇒ ((:or (nom) (acc))←(nom))

On the other hand, the non-distinctness check returns a null subsitution, so that the disjunction remains, allowing the Case feature of *was* to agree with distinct values of Case on different applications of non-distinctness.

(40) a. (:or (nom) (acc)) ≈ (nom)) ⇒ NIL
   b. (:or (nom) (acc)) ≈ (acc)) ⇒ NIL

A similar analysis holds for examples (20) and (21) in Hungarian; in these cases, it is clause 2 which is relevant.

The treatment of the conjoined verb phrase facts in Section 2.3 is analogous to that of the cases already discussed. However, one point is worth discussing here. It has not yet been made clear how it is that the object of the conjoined verb phrase is able to agree separately with each verb in the conjunct. While it might be possible to handle this mechanically by postulating some special percolation rule that combines the features of the conjuncts together into some underspecified or disjunctive form, there is a much more straightforward solution, namely, to postulate that the examples in Section 2.3 are generated by ellipsis. Certainly, given the strong lexical thrust of recent grammatical frameworks, in which syntactic structures, such as verbal complement structures, are projected from lexical representations, it is hard to see how such examples could *not* be analyzed as cases of ellipsis, at least in configurational languages. Thus, example (28) would be analyzed as in (42) rather than (41).

(41) nous a [$_{VP}$ [$_{VP}$ frappé] et
         [$_{VP}$ donné des coups de pied]]
(42) [$_{VP}$ [$_{VP}$ nous a frappé] et
         [$_{VP}$ [$_{NP}$ e] [$_{V}$ e] donné des coups de pied]]

In non-configurational languages, since complements may not be localized in any fixed position, some other mechanism for associating a head with its complements is needed, independent of these neutralization facts. In an active-objects approach to syntax, such as that outlined in Ingria and Pustejovsky (1990), message-passing would be the logical way of associating a head with its complements and would extend to the conjunction cases, as well. In any event, whatever mechanism is operative in the non-conjunctive case should also apply to the conjoined case.

3.2 Related Work

This paper is not the first to consider the problem that neutralization facts pose for theories of agreement. In particular, Zaenen and Karttunen (1984) and Pullum and Zwicky (1986) consider data of the type presented in Sections 2.3 and 2.4. However, the analysis of agreement proposed here seems more general in a number of ways.[10]

---

[10]In all fairness, Zaenen and Karttunen and Pullum and Zwicky also consider aspects of conjunction and agreement that fall outside the scope of the present paper.

While the earlier analyses only considered neutralization in the context of conjoined structures, as in Sections 2.3 and 2.4, this paper has examined the problem in general. In particular, the solutions proposed by Zaenen and Karttunen and Pullum and Zwicky crucially depend on the neutralized item standing in an agreement relation with a conjunction and, hence, cannot extend to cases of neutralization that do not involve conjunction.

While Zaenen and Karttunen and Pullum and Zwicky agree with the present analysis in associating the neutralized constistuent with each conjunct of the conjunction directly, rather than through the conjunction as a whole, both of their analyses require this association to be stated as a separate principle. If the brief sketch presented at the end of the preceding section is correct, no such stipulation is necessary. Rather, the behavior of neutralization with respect to conjunction follows from the interaction of the general agreement procedure with the way in which heads are associated with their complements.

Zaenen and Karttunen leave the bulk of the question of what features can be neutralized as a research topic. Pullum and Zwicky, on the other hand, limit neutralization to those features imposed by agreement. This is essentially the position argued for here, although there are subtle differences between the two proposals and some problematic data (which we will return to in Section 4). However, this proposal does seem to be fundamentally correct, and, combined with the view of agreement as non-distinctness, yields a more empirically valid theory of agreement than one which equates unification with agreement or which limits the effects of neutralization to conjoined structures.

Moreover, this view of agreement should contribute to the portability of natural language systems across languages. While it might be possible to reconcile the type of agreement behavior discussed here with a formalism in which unification is used for agreement by the use of arrays of feature values or some even more byzantine mechanism, such an approach would increase the fragility of any system embodying it. In a theory such as the one here, it should be possible to distinguish cases of ambiguity from cases of neutralization straightforwardly and to assign the appropriate representation accordingly. In a system that tried to maintain the use of unification for agreement by means of elaborated representations, the designer of a grammar for a new language would be faced with the problem of either using the elaborated representation for all cases of morphological underspecification, and, perhaps, blowing up the size and complexity of the grammar, or reserving

the elaborated representation for just those forms which enter into an agreement relation. This would require a thorough study of all the morphological forms of the language and the constructions they enter into before feature structures could be designed and might entail large scale changes later if previously unnoticed cases of neutralization were discovered.

3.3 The Place of Unification in Grammar

The proposal that agreement is not effected by unification does not, however, mean that unification plays no role in grammar. On the contrary: in most complex-feature based systems, semantic features are also full-fledged parts of syntactic representations and unification is used to build up more complex terms out of simpler or less specified terms and to build up formulas out of terms.[11] There is no argument at all in the data presented here that unification does not continue to play this role. In fact, there is a certain historical niceness in the picture of grammar that has been developed here: variable matching (non-distinctness) is used to effect agreement, and variable substitution (unification) is used to build up semantic representations. The reason why this view is historically satisfying is that it corresponds to views of agreement and semantic interpretation that were independently developed in theoretical and computational linguistics. In the earliest forms of generative grammar, it was recognized that certain constructions, such as the various types of ellipsis, depended on a notion of identity. Over the years, this notion of *identity* was refined into one of *non-distinctness*. Two linguistic representations agree if they are non-distinct from one another; they do not need to be identical (see Chomsky (1965, p. 181)). The view of agreement presented here accords with this well-established view. The use of unification for building up semantic representations, in turn, is based on Robinson's (1965) work on resolution theorem proving. Thus, using unification to build up semantic representations, but not for agreement, returns it to something close to its original use.

There are two other places where unification may play a role in grammar, although other mechanisms are also possibile in these cases. The first is feature percolation and the second is the use of empty categories, such as traces. Whereas agreement has been used here to mean matching of features between sister nodes, typi-

cally of distinct categories, feature-percolation involves the matching of features between one constituent and a constituent which it dominates, where the dominating constituent is a projection of the dominated, in the sense of the X-Bar theory of phrase structure (Chomsky (1970), Jackendoff (1977)). For example, a noun phrase has the same person and number features as its head noun, a verb phrase, the tense and mood of its head verb, etc. Unification has typically been used to effect feature-percolation and nothing in the data presented here suggests that it is wrong to use unification for this purpose. And while the proposal that agreement and feature-percolation are handled by different mechanisms is not usual in complex-feature based grammars, it is also not unprecedented. Ross's (1981) Local Grammar formalism is a complex-feature based grammar in which feature percolation and agreement are distinct.

Finally, unification has been used to "pass" features between a "dislocated" element and its trace. Here again, unification remains a viable mechanism. However, there are alternatives mechanisms for both these functions, such as inheritance and delegation, whose use should probably be investigated.

4   Future Research

There are a number of theoretical and practical issues that the analysis presented here raises. Their discussion will conclude this paper.

First of all, there is the question of how the non-distinctness test for agreement can be incorporated into a system in which unification is used for semantic interpretation and other purposes. Since non-distinctness returns a subset of the values returned by unification, interaction between non-distinctness and unification should be straightforward. However, a system using both these mechanisms would also need to contain some method for specifying which features of which constituents are subject to unification and which are subject to non-distinctness. This suggests the necessity of some sort of type declaration system, in which features are declared as semantically relevant or not for a particular category. The BBN ACFG formalism (Boisen 1989a,b), a form of Definite Clause Grammar, already includes a type declaration system, which has proven very useful for maintaining the consistency of large grammars. It should be possible to extend this kind of type system to the degree of delicacy required by a system incorporating both unification and non-distinctness.

---

[11] See Pereira and Warren (1980), Shieber (1986), and Pereira and Shieber (1987) for more detailed discussion of semantic interpretation in complex-feature based grammars.

202

A more problematic issue is the exact specification of the features which can be neutralized and those which can be ambiguous, and their contexts. In Section 3.1, it was suggested that semantically relevant features enter into ambiguity relations, while all others produce neutralization. However, the notion of semantic relevance may need to be refined. Zaenen and Karttunen (1984) produce examples such as the following:

(43) der Antrag   des oder der Dozenten
      the petition the or   the docent(s)
                    SG   PL GEN-SG/GEN-PL
      'the petition of the docent or docents'
(44) *Ich have den Dozenten gesehen und geholfen.
      I   have the docent(s) seen   and helped
              A-SG/D-PL   ACC   DAT
      'I have seen the docent and helped the docents.'

Example (44), which by the account presented here would involve the attempted neutralization of number, a semantically relevant feature, is ungrammatical, just as is predicted. However, (43), which also seems to involve the attempted neutralization of number, is unexpectedly grammatical. Zaenen and Karttunen also present an example from Finnish parallel to (43):

(45) He   lukivat hänen uusimman _ ja
     They read   his   newest   and
                        GEN-SG
     me hänen parhaat _ kirjansa.
     we his   best   book(s)
              NOM-PL   GEN-SG/NOM-PL
     'They read his newest book and we his best books.'

Here again, number, a semantically relevant feature, appears to be neutralized. Although Zaenen and Karttunen's treatment of neutralization is different from that suggested here is several respects, they suggest a crucial difference between (43) and (45) on the one hand and (44) on the other that may carry over. In (44), the constitutent level at which neutralization is attempted is that of the phrase (NP), whereas in (43) and (45) it is at the level of the pre-terminal (N). Zaenen and Karttunen (1984, p. 317) suggest that the neutralization is possible at the one level but not the other because "reference is assigned to noun phrases, not to common nouns." Or, in the terms we have been using here, number is semantically relevant for noun phrases, but not nouns.[12] Clearly, more research needs to be done to determine

if the proposed distinction is valid or not. Moreover, if it is valid, the theory of feature percolation needs to be modified to allow number to be neutralized at the level of N, but to produce ambiguity at the level of NP.

Finally, one issue that has not yet been mentioned is that of speaker preferences. While the discussion in Section 2 treated constructions involving the neutralized forms as perfectly grammatical, variation in speaker judgement has been reported. Thus, Zaenen and Karttunen (1984) comment that some Icelandic speakers reject (30b) as well as (30a). Pullum and Zwicky (1986) present similar sorts of judgements for other constructions. Moreover, there are also judgements in the opposite direction. For example, Modern Greek, unlike German, does not require that the relative pronoun in a free relative clause have a Case compatible with both its source and superficial positions; see, for example Mackridge (1985, pp. 259ff) for discussion. This means that the Modern Greek equivalents of (4b) and (5b) are grammatical. Nevertheless, some speakers[13], while accepting such sentences as grammatical, report that sentences containing a free relative pronoun which neutralizes the abstract Case conflict are somewhat more acceptable. These facts set us a broader research goal: that of proposing a theory of agreement which does not produce simple binary grammaticality statements but one which is capable of estimating degrees of relative grammaticality. Since the necessity of such a finer-grained theory of grammaticality is becoming more and more obvious in computational linguistics as a whole, it is no surprise to find it appearing in the study of agreement, as well.

5 Acknowledgments

---

[12]In our work on the BBN ACFG system (Boisen 1989a,b), we have also found that features such as number, degree, and tense seem to have their semantic effect at the phrasal level, rather than that of

the lexical head. Moreover, this distinction between the behavior of number on N and NP is reminiscent of Chomsky's (1965, pp. 171ff) claim that number is not an inherent feature of nouns.

[13]Sabine Iatridou, personal communication.

6 References

1. Boisen, S., Y. Chow, A. Haas, R. Ingria, S. Roucos, R. Scha, D. Stallard and M. Vilain (1989a) *Integration of Speech and Natural Language: Final Report*, Report No. 6991, BBN Systems and Technologies Corporation, Cambridge, Massachusetts.

2. Boisen, Sean, Yen-Lu Chow, Andrew Haas, Robert Ingria, Salim Roukos, and David Stallard (1989b) "The BBN Spoken Language System" In *Proceedings of the Speech and Natural Language Workshop February 1989*, Morgan Kaufmann Publishers, Inc., San Mateo, California, pp. 106–111.

3. Chomsky, Noam (1965) *Aspects of the Theory of Syntax*, The M.I.T. Press, Massachusetts Institute of Technology, Cambridge, Massachusetts.

4. Chomsky, Noam (1970) "Remarks on Nominalization", in R. A. Jacobs and P. S. Rosenbaum, eds., *Readings in English Transformational Grammar*, Ginn and Co., Waltham, Mass, pp. 184–221.

5. Eisenberg, Peter (1973) "A Note on 'Identity of Constituents' ", *Linguistic Inquiry* 4, pp. 417–420.

6. Groos, Anneke and Henk van Riemsdijk (1981) "Matching Effects in Free Relatives: A Parameter of Core Grammar", in A. Belletti, L. Brandi, and L. Rizzi, eds., *Theory of Markedness in Generative Grammar*, Proceedings of the 1979 GLOW Conference, Scuola Normale Superiore, Pisa, pp. 171–216.

7. Ingria, Robert J. P. and James Pustejovsky (1990) "Active Objects in Syntax, Semantics, and Parsing", in Carol Tenny, ed, *Papers from the Parsing Seminar*, MIT Center for Cognitive Science.

8. Jackendoff, Ray S. (1977) $\overline{X}$ *Syntax: A Study of Phrase Structure*, Linguistic Inquiry Monograph No. 2, The MIT Press, Cambridge, Massachusetts.

9. Johnson, Mark E. (1989) *Attribute-Value Logic and the Theory of Grammar*, Center for the Study of Language and Information.

10. Karttunen, Lauri (1984) "Features and Values", in *Proceedings of Coling84*, Association for Computational Linguistics, Morristown, NJ, pp. 28–33.

11. Kasper, Robert T. (1987) "A Unification Method for Disjunctive Feature Descriptions", in *25th Annual Meeting of the Association for Computational Linguistics: Proceedings of the Conference*, Association for Computational Linguistics, Morristown, NJ, pp. 235–242.

12. Kasper, Robert T. and William C. Rounds (1986) "A Logical Semantics for Feature Structures", in *24th Annual Meeting of the Association for Computational Linguistics: Proceedings of the Conference*, Association for Computational Linguistics, Morristown, NJ, pp. 257–266.

13. Kayne, Richard S. (1975) *French Syntax: The Transformational Cycle*, The MIT Press, Cambridge, Massachusetts, and London, England.

14. Mackridge, Peter (1985) *The Modern Greek Language*, Oxford University Press.

15. Pereira, Fernando C. N. and Stuart M. Shieber (1987) *Prolog and Natural-Language Analysis*, Center for the Study of Language and Information.

16. Pereira, Fernando C. N. and David H. D. Warren (1980) "Definite Clause Grammars for Language Analysis—A Survey of the Formalism and a Comparison with Augmented Transition Networks", *Artificial Intelligence* 13, pp. 231-278.

17. Pullum, Geoffrey K. and Arnold M. Zwicky (1986) "Phonological Resolution of Syntactic Feature Conflict", *Language* 62, pp. 751–773.

18. Robinson, J. A. (1965) "A Machine-oriented Logic Based on the Resolution Principle", *Communications of the ACM* 12, pp. 23–44.

19. Ross, Kenneth M. (1981) *Parsing English Phrase Structure*, Ph. D. Dissertation, University of Massachusetts at Amherst.

20. Shieber, Stuart M. (1986) *An Introduction to Unification-Based Approaches to Grammar*, Center for the Study of Language and Information.

21. Szamosi, Michael (1976) "On a Surface Structure Constraint in Hungarian", in James D. McCawley, ed, *Syntax and Semantics Volume 7: Notes from the Linguistic Underground*, Academic Press, New York, pp. 409–425.

22. Zaenen, Annie and Lauri Karttunen (1984) "Morphological Non-Distinctness and Coordination", in *ESCOL 84*, pp. 309–320.