Applying Lexical Rules Under Subsumption

Erhard W. Hinrichs
Dept. of Linguistics
University of Tübingen
Wilhelmstr. 113
D-72074 Tübingen
Germany
eh@sfs.nphil.uni-tuebingen.de

Tsuneko Nakazawa
NTT Laboratories
1-2356 Take, Yokosuka 238-03
Japan
tsuneko@nttnly.ntt.jp

Abstract

Lexical rules are used in constraint-based grammar formalisms such as Head-Driven Phrase Structure Grammar (HPSG) (Pollard and Sag 1994) to express generalizations among lexical entries. This paper discusses a number of lexical rules from recent HPSG analyses of German (Hinrichs and Nakazawa 1994) and shows that the grammar in some cases vastly overgenerates and in other cases introduces massive spurious structural ambiguity, if lexical rules apply under unification. Such problems of overgeneration or spurious ambiguity do not arise, if a lexical rule applies to a given lexical entry iff the lexical entry *is subsumed by* the left-hand side of the lexical rule. Finally, the paper discusses computational consequences of applying lexical rules under subsumption.

1 Introduction

Current linguistic theories place an increasing amount of grammatical information in the lexicon and employ a variety of mechanisms to express generalizations across lexical entries: templates (Flickinger 1987, Shieber 1986), inheritance hierarchies (Flickinger 1987, Pollard and Sag 1994), and lexical rules (Bresnan 1982, Dowty 1982, Gazdar et al. 1985, Pollard and Sag 1994). Lexical rules (henceforth: LRs) have been subjected to particularly close scrutiny. This research has focused on two important issues: 1. how the use of LRs affects the generative power of grammar formalisms and the computational complexity of parsing algorithms (Uszkoreit and Peters 1986, Carpenter 1991), and 2. how to provide a denotational semantics for LRs (Calcagno and Pollard 1995, Meurers 1995). In this paper we address neither of these two issues. Instead we will concentrate on a question that we consider to be of equal importance, but that has received surprisingly little attention: Under what conditions should an LR be applicable to a given lexical entry (henceforth: LE)? For grammar formalisms that employ the notion of unification of attribute-value structures, two criteria for applicability naturally suggest themselves:

1. *Hypothesis A*: A lexical rule applies to a lexical entry iff the lexical entry *unifies with* the left-hand side of the lexical rule.

2. *Hypothesis B*: A lexical rule applies to a lexical entry iff the lexical entry *is subsumed by* the left-hand side of the lexical rule.[1]

Without much argument, it is commonly assumed that Hypothesis A is correct (cf. Pollard and Sag 1994, Calcagno and Pollard 1995, and Meurers 1995). This paper argues that Hypothesis A should be rejected on empirical grounds. We discuss a number of LRs that have been used in HPSG analyses of German (Hinrichs and Nakazawa 1994) and show that the grammar will either vastly overgenerate and accept ungrammatical sentences or introduce spurious structural ambiguity for grammatical sentences, if Hypothesis A is adopted. However, no such problems of overgeneration or spurious ambiguity arise, if one adopts Hypothesis B, instead.

It would go beyond the scope of this paper to present a fully worked-out proposal on how to process LRs in a computational system for HPSG. However, as discussed in section 6, it is worth noting that the subsumption test for LR application can be integrated straightforwardly into two recent proposals by van Noord and Bouma (1994) and by Meurers and Minnen (1995) of how to implement LRs in a processing system for HPSG.

[1] This paper will not provide a formal definition of subsumption or unification for typed feature structures. Instead, we refer the reader to the standard definitions of Kasper and Rounds (1986) and Carpenter (1992), among many others. Informally speaking, two feature structures are unifiable iff they do not contain incompatible information. One feature structure subsumes another iff the information contained in the former is less specific than in the latter.

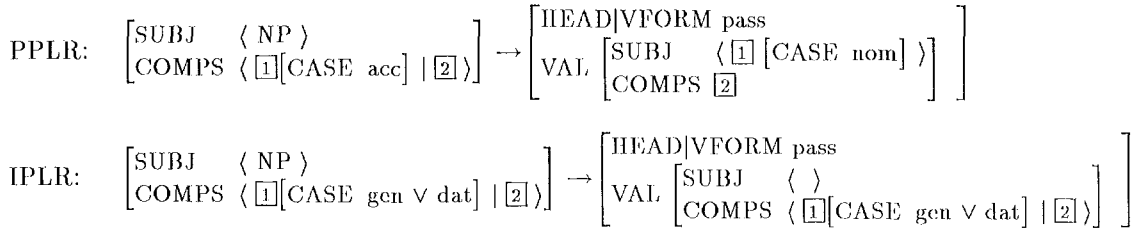

Figure 1: Passive Lexical Rules for German – Kiss 1992

2 Passive by Lexical Rule

It has been assumed in a variety of syntactic frameworks that the active/passive alternation should be treated as a lexical process: Bresnan (1982) in LFG, Dowty (1982) in Categorial Grammar, and Pollard and Sag (1987) for HPSG. German exhibits two types of passives: personal passives, as in (1b), and impersonal passives, as in (2b,3b).

(1) a. Peter sah den Mann.
'Peter saw the man.'
b. Der Mann wurde gesehen.
'The man was seen.'

(2) a. Peter half dem Mann.
'Peter helped the man.'
b. Dem Mann wurde geholfen.
'The man was helped.'

(3) a. Die Veteranen gedachten der Toten.
'The veterans commemorated the dead.'
b. Der Toten wurde gedacht.
'The dead were commemorated.'

For personal passives the accusative object NP of a transitive verb, e.g. *den Mann* in (1a), in its active form corresponds to an NP with nominative case, e.g. *der Mann* in (1b). In impersonal passives a dative or genitive NP complement of a transitive verb, e.g. the dative NP *dem Mann* in (2), exhibits the same case assignment in the active and passive forms. To formally capture the relationship between the LEs for active and passive forms of German verbs in HPSG, Kiss (1992) formulates the LRs in Fig. 1 (henceforth referred to as the PPLR and the IPLR, respectively).[2]

The LRs in Fig. 1 employ the HPSG feature geometry of Pollard and Sag (1994).[3] The specification of a syntactic category (CAT) in HPSG includes the feature VALENCE (abbreviated as VAL in Fig. 1), which in turn specifies for verbs under the features SUBJ and COMPS the subject and non-subject complements. COMPS takes a list of categories, called *synsem* objects, as its value. If the list is non-empty, the leftmost category in the list represents the direct object. The intended effect of the PPLR is to promote the direct object of a transitive verb in the LE for the active form to the subject of the passive form. For impersonal passives the COMPS list of any transitive verb whose leftmost complement is marked by genitive or dative case remains unchanged, while the singleton list of the subject value of the active form becomes the empty list in the LE for the passive form. Note that the case specifications on the left-hand side of the rules are crucial since they condition which classes of transitive verbs appear in personal and impersonal passives.

Viewed procedurally, the PPLR is meant to apply to LEs for transitive verbs such as *kaufen* as shown in Fig. 2.[4] For transitive verbs it makes no difference whether we use unification or subsumption as the test of applicability for the PPLR. The LE for any transitive verb is more specific than the feature structure of the input description of the PPLR, since the PHON value and the COMPS and SUBJ values will be further instantiated than in the input description of the PPLE. In particular, the COMPS list of the PPLR is entirely schematic for any non-empty list of categories whose first element is an accusative NP, while the COMPS value for *kaufen* is a list with exactly one element that has the same category and case specification as the NP in the PPLR. Therefore, the LE will both unify with and be subsumed by the input feature structure of the PPLR. However, as we will see in the next section, the choice of subsumption or unification makes a crucial difference when LEs are themselves highly schematic and underspecified.

[2] However, see Kathol (1994) and Pollard (1994) for an account of German passive without LRs.

[3] In order to reduce the size of the feature structures, prefixes of paths that begin with the SYNSEM attribute have been omitted as much as possible in Fig. 1 and all other feature structures that are shown in this paper.

According to Kiss, specifying two different case values under one reentrancy (cf. tags 1 in the PPLR) is a shorthand notation for identity of the two categories in all respects, except for the case value. The original formulation of the LRs by Kiss (1992) differs from the version presented here in minor definitional details. However, these differences are entirely orthogonal to the theoretical issues discussed in this paper.

[4] Following abbreviatory conventions in HPSG, the subscripted tags in Fig. 2 stand for the index values contained in the CONTENT specifications of the NPs.

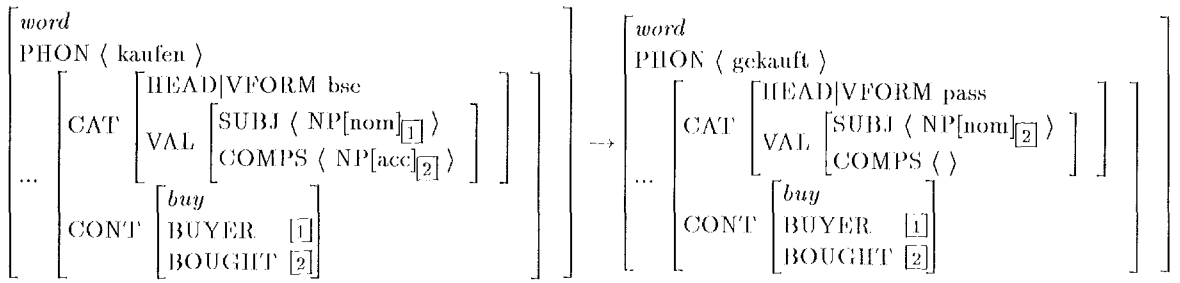

Figure 2: Applying the PPLR to the abbreviated LE for the verb *kaufen* ('buy')

3 Argument Composition in the Lexicon

The notion of argument composition was first introduced into the HPSG literature by Hinrichs and Nakazawa (1989) to account for the topicalization of verbal constituents and for the auxiliary flip construction in German.[5]

In German VERBAL COMPLEXES, i.e. main verbs together with non-finite auxiliaries, such as *lesen können* can be topicalized, as shown in (4).

(4) Lesen können wird er es.
read can will he it
'He will be able read it.'

In addition, the verbal complex serves as the domain over which auxiliaries can be fronted. This so-called AUXILIARY FLIP construction positions finite auxiliaries such as *wird* in (5) to the left in the verbal complex, instead of the customary sentence-final position for subordinate clauses.

(5) Ich glaube nicht, daß er es wird lesen können.
I believe not that he it will read can
'I don't believe that he will be able to read it.'

Topicalization and auxiliary-flip, thus, provide crucial evidence for treating main verbs and auxiliaries as constituents. The proposed constituent structure requires that subcategorization information about non-verbal complements is propagated from the main verb to the top of the verbal complex. In HPSG this can be achieved by structure sharing the complements of the main verb with the subcategorization information of each auxiliary in the sentence. This leads to LEs for auxiliaries such as *können* 'can' shown in Fig. 3.

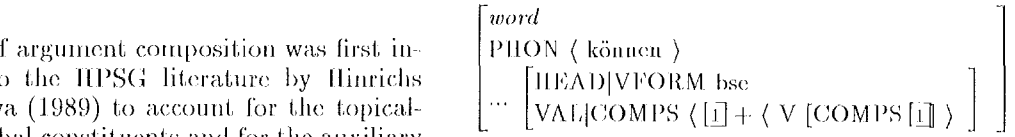

where 1 = *list(→ [SYNSEM|...|HEAD verb])*

Figure 3: Abbreviated LE for auxiliary *können*

können requires a base infinitive complement, as indicated in the COMPS value in Fig. 3. The COMPS value of *können* also contains a (possibly empty) list of non-verbal categories (identified by the tag 1) that the governed verb subcategorizes for.[6] In other words, the COMPS value of the governed verb is merged with the COMPS list of *können* itself. Formally this merging is achieved by appending the COMPS list of the governed verb with the one-element list that consists of the governed verb itself. Since auxiliaries have to be able to combine with different types of verbs (e.g. intransitives, transitives, ditransitives, etc.), the COMPS list of the governed verb has to remain underspecified. It is this underspecification in the LE for auxiliaries that makes a crucial difference in the applicability of LRs.

4 Subsumption and Unification Revisited

For interactions between the LEs for auxiliaries such as *können* and the passive LR for German, it makes a crucial difference whether the LR applies under unification or under subsumption. First consider unification as the test for applicability of LRs. Since the LE for *können* does not contain any information that is inconsistent with the input specification of the PPLR in Fig. 1, the rule is applicable and will produce the derived LE for

[5] HPSG practicioners have adopted the notion of *argument composition* to account for a variety of syntactic constructions in different languages, including clitic-climbing in Italian (Monachesi 1993), the syntactic properties of auxiliaries in French (Abeillé and Godard 1994), and cross-serial dependencies in Dutch (Rentier 1994).

[6] For further details as to why the list of raised elements has to be restricted to contain only non-verbal categories see Hinrichs and Nakazawa (1994).

können in Fig. 4.

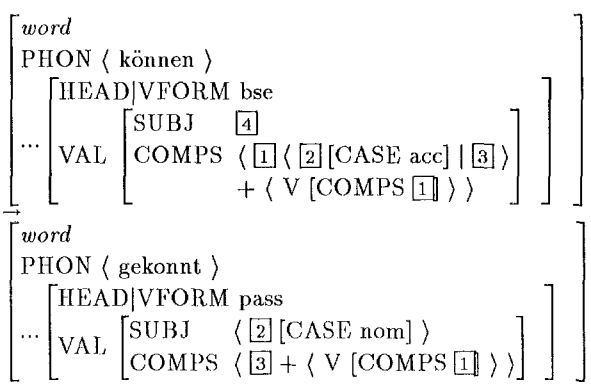

Figure 4: Applying the PPLR to *können* under unification

When the LE for *können* is unified with the left-hand side of the PPLR, the COMPS list of *können* – and, via structure sharing, the COMPS value of the governed verb – becomes further instantiated. This COMPS list now contains as its leftmost element a category with accusative case. In accordance with the PPLR, this leftmost element is promoted to the SUBJ value of *können*, while the remainder of the COMPS list of the verb governed by *können*, identified by tag 3, is retained. However, this derived LE would have the undesirable consequence that it admits ill-formed sentences such as (6).

(6) * Das Auto wurde kaufen gekonnt.
the car was bought can

In (6) the auxiliary *können* has been passivized and the direct object of the transitive verb has been promoted as the subject of *können*. However, since in German only main verbs can be passivized, the sentence is ungrammatical.[7]

Ungrammatical sentences such as (6) can be successfully ruled out if the PPLR is applied to an LE only if the input specification of the LR subsumes the LE (Hypothesis B). The subsumption requirement for LR application is based on the insight that LRs should apply only to LEs that are instantiated at least to the extent that the input description of the LR minimally requires. In the case at hand, the list of raised arguments in the LE for *können* in Fig. 3 is totally unspecified – it can be any list of non-verbal *synsem* objects, including the empty list. The COMPS list of the left-hand side of the PPLR, on the other hand, requires the leftmost element to carry accusative case. Therefore the COMPS list of *können* does not subsume the COMPS list on the left-hand side of the PPLR, or vice versa. Accordingly, no subsumption relationship exists between the input specification of the PPLR as a whole and the LE for *können*. Hence if Hypothesis B is assumed, the LR can be successfully blocked.[8] However, even under subsumption nothing blocks the PPLR from applying to the transitive verb *kaufen*, as discussed in section 2. Therefore the grammatical sentence in (7) can be derived successfully.

(7) Das Auto wurde gekauft.
'The car was bought.'

5 Avoiding Spurious Ambiguity

This section will consider the interaction of highly schematic LEs with another LR commonly used in HPSG: the Complement Extraction LR, formulated for English by Pollard and Sag (1994). Fig. 5 shows the version of the Complement Extraction Lexical Rule for German (henceforth: CELRG) that has been proposed by Hinrichs and Nakazawa (1994).

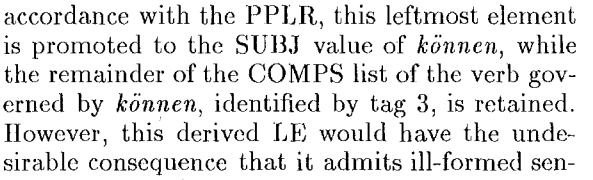

Figure 5: Complement Extraction Lexical Rule for German – Hinrichs and Nakazawa 1994

The CELRG moves an element from the COMPS list of a verb to its SLASH set. The value of the feature SLASH contains those items that are realized in left dislocated position, e.g. as a topicalized constituent in sentence-initial position. Analogous to the applicability of the PPLR, the CELRG is applicable to LEs of transitive verbs such as *kaufen* shown in Fig. 2, under both unification and subsumption. The result of such an application is the same as shown in Fig. 2, except that the direct object is placed on the SLASH set instead of the SUBJ list. For the analysis of sen-

[7] There are some cases of "long distance" passives, i.e. passives which involve the complement of an embedded verb, that at least some German speakers accept, e.g. *Der Hund wurde vergessen zu füttern* ('It was forgotten to feed the dog'): However, we are not aware of any German speakers that would allow passives with raising verbs such as *können*.

[8] The reader may wonder whether one couldn't get around the requirement that the LR for passive apply under subsumption by restricting the rule to apply only to main verbs. At first glance this seems like a possible way-out, since the cases that we have considered problematic for applying the LR under unification involve LEs of auxiliaries. However, the class of argument raising verbs is not restricted to auxiliaries. Verbs such as *verstehen* ('know how'), as in *Er versteht Parser zu implementieren* ('He knows how to implement parsers.'), fall into the same class. But *verstehen* is a main verb, not an auxiliary. Thus, even if the LR would be reformulated to apply 'only to auxiliaries, the following ungrammatical sentence could not be excluded: * *Parser werden verstanden zu implementieren.*

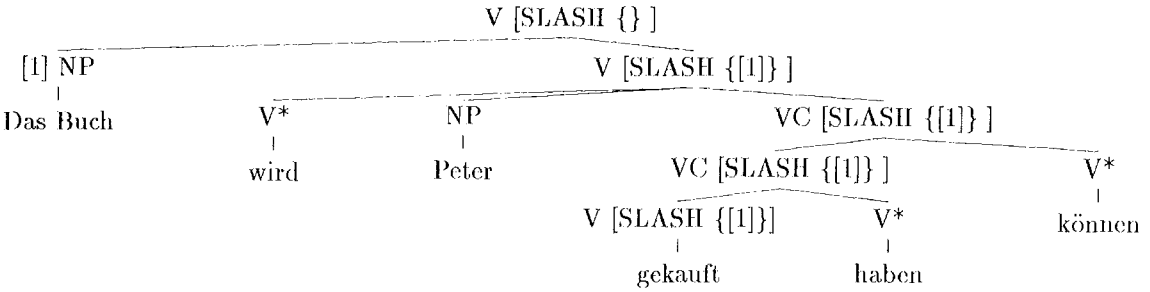

Figure 6: Analysis Tree for Sentence (8)

tence (8), the tree in Fig. 6 illustrates the percolation of the relevant SLASH value that is introduced via the CELRG in the LE for *kaufen*.[9]

(8) Das Buch wird Peter gekauft haben können.
the book will Peter bought have can
'Peter will have been able to buy the book.'

The SLASH value in Fig. 6 is percolated from the non-terminal node for the verb *kaufen* by the Nonlocal Feature Principle to the sister node of the topicalized constituent *das Buch*. The top local tree is licensed by the Head-Filler ID Schema which binds off the SLASH value so that the sentence node has an empty SLASH set.[10]

The CELRG restricts topicalized constituents to phrasal categories. This restriction is necessary to rule out sentences such as (9) in which a single lexical item, i.e. a *word* in terms of the type hierarchy of HPSG, is topicalized.

(9) * Können wird Peter das Auto gekauft haben.
can will Peter the car bought have

As in the case of the PPLR, the difference between Hypotheses A and B comes into play when we consider the interaction of the CELRG with highly schematic entries such as the ones Hinrichs and Nakazawa (1994) assume for auxiliaries in German. If Hypothesis A is assumed, then the CELRG will be applicable to the type of LE shown for *können* in Fig. 3, since such an LE will unify with the input description of the LR. One of the possible outputs of the LR to such an auxiliary entry would look identical to the putative output of the PPLR shown in Fig. 4, expect that one of the elements from the COMPS list of the auxiliary is assigned to the SLASH set instead of the SUBJ list. However, this would have the undesirable consequence that the SLASH value in the analysis of topicalized sentences that contain auxiliaries could originate in the LE for the main verb, but also in the LE for each auxiliary present in the sentence. Hence, in addition to the tree in Fig. 6, three additional trees are admitted for sentence (8) in which the SLASH value originates in one of the pre-terminal nodes for the auxiliaries. These nodes are marked for emphasis by an asterisk in Fig. 6. This ambiguity is, of course, totally spurious since it does not correlate with a difference in semantics or any other relevant linguistic property of the sentence. From a computational perspective, such spurious ambiguities are highly undesirable since they force the parser into considering multiple analyses where a single analysis suffices. The spurious ambiguity that we have just identified is particularly pernicious, since it would affect a wide range of sentences in any grammar of German that employs argument composition and the CELRG: all assertion main clauses that contain auxiliaries would be affected since in assertion clauses the initial constituent is the result of topicalization.

Once again undesirable consequences of overapplying an LR under unification can be avoided if applicability of LRs instead requires subsumption (Hypothesis B). Since the CELRG limits extracted constituents to phrases and since the COMPS list of an auxiliary does not restrict its elements in the same way, the LE for auxiliaries and the input description of the CELRG do not stand in a subsumption relation. Hence under Hypothesis B the CELRG applies only to main verbs, and the tree in Fig. 6 is the only tree admitted.

6 Computational Consequences

Finally, we will consider the computational implications that the adoption of Hypothesis B has for the processing of LRs in a computational system. Since consensus on how to provide an adequate denotational semantics for LRs has not yet been reached, it would go well beyond the scope of this paper to develop a fully worked-out proposal on how to process LRs.[11] A very promising approach

[9]The tree in Fig. 6 assumes the flat constituent structure for German clause structure proposed by Hinrichs and Nakazawa (1994). However, the issue raised with respect to the CELRG in this paper is orthogonal to overall assumptions of German clause structure.

[10]See Pollard and Sag (1994) and Hinrichs and Nakazawa (1994) for further details on the Nonlocal Feature Principle and the Head-Filler ID Schema.

[11]See Calcagno and Pollard (1995) and Meurers (1995) for further discussion.

on how to integrate LRs into processing systems for HPSG has been developed independently by van Noord and Bouma (1994) and by Meurers and Minnen (1995). It turns out that the conclusions reached in this paper can be easily integrated in the general framework that these authors provide.

Common to these two proposals is the idea of treating LRs as (Horn clause) constraints on LEs of the kind shown schematically in Fig. 7.

derive-lexical-entry(X_0,X_n) :-
base-lexical-entry(X_0),
lex-rule$_1$(X_0,X_1), ... , lex-rule$_n$(X_{n-1},X_n).

Figure 7: Lexical rules as Horn Clause constraints on relations between lexical entries

Definite clause schemata as in Fig. 7 (for $n \geq 0$) define relations between base LEs listed in the lexicon and derived LEs that are obtained via the application of a sequence of LRs. Each LR is then viewed as a two-place definite relation as illustrated in Fig. 8 for the PPLR of Fig. 1.

$$\text{lex-rule}\left(\begin{bmatrix}\text{SUBJ} & \langle\,\text{NP}\,\rangle \\ \text{COMPS} & \langle\,\boxed{1}\begin{bmatrix}\text{CASE acc}\end{bmatrix} \mid \boxed{2}\,\rangle\end{bmatrix},\right.$$
$$\left.\begin{bmatrix}\text{HEAD|VFORM pass} \\ \text{VAL}\begin{bmatrix}\text{SUBJ} & \langle\,\boxed{1}\begin{bmatrix}\text{CASE nom}\end{bmatrix}\,\rangle \\ \text{COMPS} & \boxed{2}\end{bmatrix}\end{bmatrix}\right)$$

Figure 8: The PPLR encoded as a definite relation

It is important to distinguish two tasks that need to be performed in computing with lexical rules:[12]

1. the algorithm that decides for a given lexical entry whether a lexical rule is applicable to it, and
2. the algorithm that computes for a given lexical entry the output specification of the lexical rule, i.e. the derived lexical entry.

The subsumption test for lexical rule application that we have argued for in this paper pertains to the first task. The execution of the definite clauses of the kind shown schematically in Fig. 7, which encode the possible relations between base and derived lexical entries, pertains to the second task.

Regarding the first task, in van Noord and Bouma's approach the sequences of lexical rules that are applicable to a given base lexical entry have to be specified by the grammar developer along with delay statements, which allow goal freezing at run time of not sufficiently instantiated relations. In Meurers and Minnen's approach such sequences are automatically generated from the set of base lexical entries and the set of lexical rules specified by the grammar. The allowable sequences of lexical rules are compiled into finite state automata which are in turn encoded as definite clause attachments to base lexical entries.

Note that both approaches are general enough to accommodate different assumptions about the applicability of lexical rules to lexical entries, i.e. they are compatible with both Hypotheses A and B. Whether or not a given lexical rule applies to a lexical entry in van Noord and Bouma's approach needs to be stipulated by the grammar writer who is in theory free to use either a unification or subsumption test. In Meurers and Minnen's approach the deduction rules for automatically generating a finite-state encoding of lexical rules can likewise be based on a subsumption check or a unification check. In this paper we have argued on empirical grounds that subsumption should be the relevant operative criterion. The theoretical results of the present paper can therefore be straightforwardly integrated into a lexical rule compiler of the sort described by Meurers and Minnen in which applicability of lexical rules is checked automatically under subsumption.

7 Conclusion

This paper has discussed a number of lexical rules from recent HPSG analyses of German (Hinrichs and Nakazawa 1994) and has shown that the grammar in some cases vastly overgenerates and in other cases introduces massive spurious structural ambiguity, if lexical rules apply under unification.[13] However, no such problems of overgeneration or spurious ambiguity arise, if a lexical rule applies to a given lexical entry iff the lexical entry *is subsumed by* the left-hand side of the lexical rule. Finally we have shown that the subsumption test for the applicability of lexical rules can be integrated straightforwardly into the proposals by van Noord and Bouma (1994) and by Meurers and Minnen (1995) of how to implement lexical rules in a processing system for HPSG.

[12] This distinction is rightly emphasized by Calcagno and Pollard (1995).

[13] It is worth pointing out that the importance of subsumption has been noted for other linguistic phenomena as well. Building proposals originating with Gazdar et al. (1985), Bayer and Johnson (1995) have pointed out that the grammar of (cross-categorial) coordination for English will make the right predictions if the mother category of the conjunction schema is required to subsume the category of each conjunct. By contrast, if the relation between the mother category and each daughter category is that of unifiability, then the resulting grammar vastly overgenerates.

References

Anne Abeillé and Danièle Godard. 1994. The Complementation of French Auxiliaries. *Proceedings of the Thirteenth West Coast Conference on Formal Linguistics.* CSLI Publications/SLA, Stanford University.

Sam Bayer and Mark Johnson. 1995. Features and Agreement. *Proceedings of the 33rd Annual Meeting of the Association of Computational Linguistics.* MIT, Cambridge, Mass.

Joan Bresnan. 1982. Passive. Bresnan, J. ed. *The Mental Representation of Grammatical Relations.* Cambridge, MA: The MIT Press.

Michael Calcagno and Carl Pollard. 1995. Lexical Rules in HPSG: What are they. *Proceedings of the ACQUILEX II Workshop on Lexical Rules.* Cambridge, UK.

Bob Carpenter. 1991. The Generative Power of Categorial Grammars and Head-Driven Phrase Structure Grammars with Lexical Rules. *Computational Linguistics.* Vol. 17.3, pp. 301-314.

Bob Carpenter. 1992. *The Logic of Typed Feature Structures.* Cambridge Tracts in Theoretical Computer Science. Cambridge University Press.

David R. Dowty. 1982. Grammatical relations and Montague grammar. Jacobson, P. and G. K. Pullum eds. *The Nature of Syntactic Representation* Dordrecht: D. Reidel.

Dan Flickinger. 1987. *Lexical Rules in the Hierarchical Lexicon.* Ph.D. thesis, Stanford University, Stanford, CA.

Gerald Gazdar, Ewan Klein, Geoffrey Pullum and Ivan Sag. 1985. *Generalized Phrase Structure Grammar.* Oxford: Basil Blackwell.

Erhard W. Hinrichs and Tsuneko Nakazawa. 1989. Subcategorization and VP Structure in German. Paper presented at the *Third Symposium on Germanic Linguistics*, Purdue University: Lafayette, Indiana. Published in E. Hinrichs and T. Nakazawa. *Aspects of German VP Structure.* SfS Technical Report 01-93, Dept. of Linguistics: University of Tübingen.

Erhard W. Hinrichs and Tsuneko Nakazawa. 1994. Partial-VP and Split-NP Topicalization in German - An HPSG Analysis. E. Hinrichs, D. Meurers and T. Nakazawa eds. *Arbeitsberichte des Sonderforschungsbereichs 340.* Stuttgart and Tübingen, pp. 1-46.

Robert Kasper and William Rounds. 1986. A Logical Semantics for Feature Structures. *Proceedings of the 24th Annual Meeting of the Association of Computational Linguistics.* New York, N.Y.

Andreas Kathol. 1994. Passives Without Lexical Rules. In Nerbonne et al. eds. *German in Head-Driven Phrase Structure Grammar.* CSLI Lecture Notes No. 46, University of Chicago Press and CSLI, Chicago.

Paul John King. 1989. *A logical formalism for head-driven phrase structure grammar.* Ph.D. thesis, University of Manchester.

Tibor Kiss. 1992. Variable Subkategorisierung - Eine Theorie der unpersönlichen Konstruktionen im Deutschen. *Linguistische Berichte* Vol. 140, pp. 256-293.

W. Detmar Meurers. 1995. Towards a Semantics for Lexical Rules as Used in HPSG. *Proceedings of the ACQUILEX II Workshop on Lexical Rules.* Cambridge, UK.

W. Detmar Meurers and Guido Minnen. 1995. A Computational Treatment of HPSG Lexical Rules as Covariation in Lexical Entries. *Proceedings of the Fifth International Workshop on Natural Language Understanding and Logic Programming.* Lisbon, Portugal.

Paola Monachesi. 1993. Object Clitics and Clitic Climbing in Italian HPSG Grammars. *Proceedings of the Sixth Conference of the European Chapter of the Assocation of Computational Linguistics.* Utrecht, The Netherlands.

Carl Pollard. 1994. Toward a Unified Account of Passive in German. In Nerbonne et al. eds. *German in Head-Driven Phrase Structure Grammar.* CSLI Lecture Notes No. 46, University of Chicago Press and CSLI, Chicago.

Carl Pollard and Ivan Sag. 1987. *Information-Based Syntax and Semantics: Volume I - Fundamentals* CSLI Lecture Notes, Volume 13. Stanford: Center for the Study of Language and Information.

Carl Pollard and Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar.* Chicago: University of Chicago Press and Stanford: CSLI Publications.

Gerriet M. Rentier. 1994. A Lexicalist Approach to Dutch Cross-Serial Dependencies. Papers from the 30th Regional Meeting of the Chicago Linguistics Society. Chicago, Illinois.

Stuart Shieber. 1986. *An Introduction to Unification-based Approaches to Grammar.* CSLI Lecture Notes No. 4. Chicago: University of Chicago Press and Stanford: CSLI Publications.

Hans Uszkoreit and Stanley Peters. 1986. On some Formal Properties of Metarules. *Linguistics and Philosophy* Vol. 9.4, pp. 477-494.

Gertjan van Noord and Gosse Bouma. 1994. The Scope of Adjuncts and the Processing of Lexical Rules. *Proceedings of COLING 94* Kyoto, Japan.