THEORETICAL LINGUISTICS PROGRAMME, BUDAPEST UNIVERSITY (ELTE)

LEXICON MATTERS

PÉTER DIENES, KÁLMÁN DUDÁS, LÁSZLÓ KÁLMÁN, ÁGNES LUKÁCS,
ATTILA NOVÁK, GÁBOR RÁDAI, VIKTOR TRÓN, KÁROLY VARASDI

THEORETICAL LINGUISTICS PROGRAMME, BUDAPEST UNIVERSITY
RESEARCH INSTITUTE FOR LINGUISTICS, HUNGARIAN ACADEMY OF SCIENCES

RESEARCH INSTITUTE FOR LINGUISTICS, HUNGARIAN ACADEMY OF SCIENCES

WORKING PAPERS IN THE THEORY OF GRAMMAR, VOL. 5, NO. 2

RECEIVED: DECEMBER 1998

LEXICON MATTERS

PÉTER DIENES, KÁLMÁN DUDÁS, LÁSZLÓ KÁLMÁN, ÁGNES LUKÁCS,
ATTILA NOVÁK, GÁBOR RÁDAI, VIKTOR TRÓN, KÁROLY VARASDI

THEORETICAL LINGUISTICS PROGRAMME, BUDAPEST UNIVERSITY (ELTE)
RESEARCH INSTITUTE FOR LINGUISTICS, HUNGARIAN ACADEMY OF SCIENCES
P.O. BOX 19, 1250 BUDAPEST 1, HUNGARY
E-MAIL: kalman@nytud.hu

SUPPORTED BY THE OTKA, PROJECT NUMBER T 018001, AND THE HUNGARIAN MINISTRY FOR CULTURE AND EDUCATION, GRANT NO. MKM FKFP 0589 /1997

THEORETICAL LINGUISTICS PROGRAMME, BUDAPEST UNIVERSITY (ELTE)
RESEARCH INSTITUTE FOR LINGUISTICS, HUNGARIAN ACADEMY OF SCIENCES

BUDAPEST I., P.O. BOX 19. H-1250 HUNGARY
TELEPHONE: (36-1) 175 8285; FAX: 212 2050

Part I

Lexical Phonology

Chapter 1

Finite State Devices in Phonology

1.1 Introduction

Phonology – or at least part of phonology – is claimed to be the component of the language most closely attached to the lexicon. Within a widespread framework, Lexical Phonology (cf. Kiparsky (1985), Kenstowicz (1994)), a number of phonological processes apply in the lexicon itself. Hence, the theory of phonological rule systems and phonological representation is crucial in the investigation of the lexicon-grammar interactions. This paper intends to give a brief introduction to the formal (mathematical and computational) handling of phonological rules, based on three influential articles: Kaplan and Kay (1994), Antworth (1990) and Bird and Ellison (1994).

In traditional – SPE-based (Chomsky and Halle (1968)) – framework, the phonological component of the grammar consists of a set of ordered rewrite rules, which connect underlying and surface representations. These representations are viewed as linearly ordered feature matrices; the two representations differ in that respect that the former allows not fully specified (underspecified) matrices whereas the latter does not. This means that underlying representations permit abstract segments while surface representations exclude them. (On the controversial issue of abstractness, cf. Kenstowicz (1994, pp. 107–114.)) However, the situation is less straightforward when rules are taken into account. In principle, nothing prevents a rule from applying on another rule's output (or even on its own output, cf. below), and indeed, phonologists often claim this to be the case. This leads to the – in principle unbounded – proliferation of intermediate representations.

At a first glance, phonological rules have the context-sensitive format:

(1) $\phi \rightarrow \psi / \lambda\_\rho$

where each symbol refers to a feature matrix, probably with finite feature variables, such as a vowel harmony rule:

(2)

$$\begin{bmatrix} \text{+ syllabic} \\ \text{– consonantal} \end{bmatrix} \rightarrow [\alpha\text{back}] \; / \begin{bmatrix} \alpha\text{back} \\ \text{+ syllabic} \\ \text{– consonantal} \end{bmatrix} [\text{+ consonantal}]^{*} \_\_\_$$

This kind of formalism, however, seems to be too powerful, especially taking the fact into account that syntax relies solely on context-free rewrite rules. This overgeneralisation made many theorists and computational linguists suspicious, till Johnson (1972) proved that phonological rules of this formalism generally are equivalent in power to finite-state devices. The only exceptions to this claim are cyclic rules (though cf. next section) and the rules which can apply to their own output, such as:[1]

(3) $e \rightarrow ab/a\_b$

This rule generates the language $\{a^n b^n | n \in \mathbf{N}\}$, which is not regular. However, phonologists tend to disfavour such rules. Hence, this finding has the straightforward consequence that phono-

[1] e stands for the empty string.

logical rule systems can be modelled by finite-state devices, such as Finite State Automata or Finite State Transducers. Furthermore, as regular relations are closed under serial composition (cf. below), a single Finite State Transducer (FST) can be constructed algorithmically to represent a whole phonological rule system. This reduces the number of representations to two again, i.e. to lexical and surface forms.

A further problem for traditional SPE-type rule systems is their application in recognition. If we want to revert these phonologically well-motivated rules, we end up in an amount of indeterminacy which is impossible to handle effectively. Consider, for example, the following two ordered rules accounting for nasal place-assimilation:

(4) **Rule 1** N → m / __ [+ labial]
Rule 2 N → n

These two rules produce one surface form for an underlying form. On the other hand, application of these two rules in finding the underlying representation of the surface form *intractable* results in two forms: *intractable* or *iNtractable*. If our system contains more rules, the number of possible underlying forms of a given input multiplies. This requires so much memory (every path needs to be pursued, maybe for a considerable distance) and computational time that this cannot be implemented. However, finite state devices, especially FSTs, help overcome this problem as well.

The overall presentation of the paper is as follows: in Section 2, we sketch the mathematical tools and concepts important for the further parts of the paper. In Section 3, we illustrate Kaplan and Kay (1994)'s method to translate SPE rules into regular relations. The following section gives a brief introduction to the Kimmo-formalism (cf. Antworth (1990), Karttunen (1993)), and the way it views the organization of phonology. Finally, Section 5 deals with a possible incorporation of autosegmental phonology into the finite-state paradigm, following Bird and Klein (1990) and Bird and Ellison (1994).

1.2 Mathematical Background

This section gives the definitions of mathematical and computational tools and devices crucial for the understanding of the remainder of the paper. One may, however, use this section as a reference section, turning back some pages when necessary. Throughout this section, we adopt the notation of Kaplan and Kay (1994).

One of the most basic concepts of the articles this paper is dealing with is (binary) *string relation*. This is a set of ordered pairs of strings, namely the subsets of $\Sigma^* \times \Sigma^*$,[2] where Σ^* denotes the alphabet. If $X = \langle x_1, x_2 \rangle$ and $Y = \langle y_1, y_2 \rangle$ are string relations, then we define their concatenation as:

$$X \cdot Y = XY =_{df} \langle x_1 y_1, x_2 y_2 \rangle$$

With these definitions at hand, we can construct a family of string relations similar to that of formal languages, namely the family of regular relations. The definition is parallel to the recursive definition of regular languages. (e is the empty word in the definitions, $\Sigma^e = \Sigma \cup \{e\}$, whereas superscript i indicates concatenation repeated i times.)

(5) **Definition**. Regular Relations

i. The empty set and $\{a\}$ for all $a \in \Sigma^e \times \Sigma^e$ are regular relations.

ii. If L_1 and L_2 are regular relations, then so are

$L_1 L_2 = \{xy | x \in L_1, y \in L_2\}$ (concatenation)
$L_1 \cup L_2$ (union)
$L^* = \cup_{i=0}^{\infty} L^i$ (Kleene closure)

iii. There are no other regular relations.

Another important device in computational linguistics is the (nondeterministic) Finite State Automaton (FSA), which is a quintuple $\{\Sigma, Q, q, F, \delta\}$, where Σ is a finite alphabet, Q is a finite set of states, $q \in Q$ is the initial state and $F \subseteq Q$ is the set of final states. The transition function δ is a total

[2]We only defined binary relations here. *n-ary* relations, however, may be defined in the same vein; they might be applied in autosegmental representations (cf. Kay (1987)).

$Q \times \Sigma^e \to \mathcal{P}(Q)$ function, and for every state $s \in Q, s \in \delta(s, e)$ vacuously holds. The definition of δ is usually extended to δ^* on Σ^* as follows: for all $r \in Q$: $\delta^*(r, e) = \delta(r, e)$ and for all $u \in \Sigma^*$ and $a \in \Sigma^e$: $\delta^*(r, ua) = \delta(\delta^*(r, u), a)$, where in the case of $P \subseteq Q$ and $a \in \Sigma^e$, $\delta(P, a) = \cup_{p \in P}\delta(p, a)$. Now, the machine accepts a string $x \in \Sigma^e$ just in case $\delta^*(q, x) \cap F$ is nonempty. The machine blocks in a state if there is no possible transition defined by δ. The definition of Finite State Transducers is the same, except for the difference in the transition function, which is in the latter case is a $\delta : Q \times \Sigma^e \times \Sigma^e \to \mathcal{P}(Q)$ function. Thus an FST can be viewed as an FSA defined over the product alphabet $\Sigma^e \times \Sigma^e$.

The basic theorem connecting regular languages and FSAs, on the one hand, and regular relations and FSTs, on the other, is the following:

(6) Every regular language is accepted by an FSA and every FSA accepts a regular language. Every regular relation is accepted by an FST and every FST accepts a regular relation.

This theorem emphasises that regular languages and regular relations (or FSAs and FSTs) are basically the same. Both families are closed under union ($X \cup Y$), concatenation ($X \cdot Y$), Kleene-star (X^*), inversion (X^{-1}) and serial composition ($X \circ Y$). There is, however, a crucial difference between the two families: regular languages are closed under intersection, while regular relations are not. An example is the following: suppose that $R_1 = \{\langle a^n, b^n c^* \rangle | n \geq 0\}$ and $R_2 = \{\langle a^n, b^* c^n \rangle | n \geq 0\}$. These relations are clearly regular, whereas their intersection $\{\langle a^n, b^n c^n \rangle | n \geq 0\}$ is not. Nevertheless, a subclass of regular relations, namely the *same-length* regular relations, is closed under intersection. These relations contain only string-pairs $\langle x, y \rangle$, where the length of x is the same as the length of y. It can be proved that

(7) R is a same-length regular relation iff it is accepted by an e-free FST.

Finally, as an illustrative example, consider how a generative rule, such as **Rule 1** can be implemented by a FST. The details of such an implementation will be discussed in the next section summing up the methods of Kaplan & Kay.

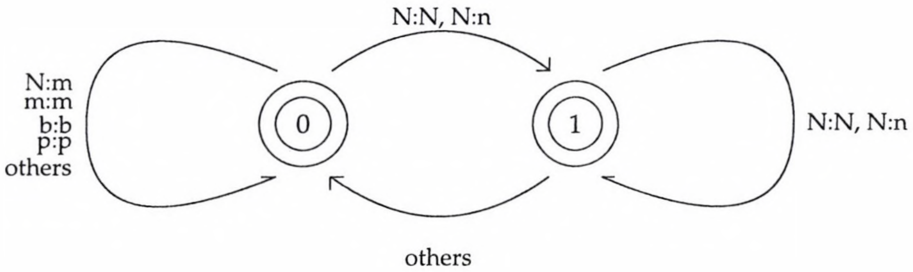

The circled numbers are the states of the FST, double circles denote final states (in the present case, both states are final), state 0 is, by convention, the initial state. The transitions are indicated by labelled arrows, the two symbols on the two tapes are separated by the colon ':'. The term *others* represents all other feasible pairs not explicitly mentioned in the diagram (such as a:a, n:n, t:t, d:d, t:D, d:D, etc.).[3] Let us see how this machine accepts the relation *iNpractical:impractical*:

$$\xrightarrow{0} \substack{i\\i} \xrightarrow{0} \substack{m\\N} \xrightarrow{0} \substack{p\\p} \xrightarrow{0} \ldots$$

On the other hand, the machine blocks in the case of *iNpractical:inpractical* relation:

$$\xrightarrow{0} \substack{i\\i} \xrightarrow{0} \substack{n\\N} \xrightarrow{1} \text{BLOCK, no } p{:}p \text{ transition in state 1}$$

[3]Note that we assume that N can only participate in three feasible pairs, namely: *N:N*, *N:n*, *N:m*.

However, FSTs can be interpreted in a different way, which interpretation is more fruitful in linguistic application. Namely, the two tapes of an FST can be viewed as the input and output tapes, thus the machine can be regarded not only as an accepting device, but as a generating device as well. If the input string of the FST is the underlying representation of a string (*iNpossible*), then the machine works in a generative fashion. If, on the other hand, the input is a surface string (*input*), then the machine is applied for recognition:

(8) Generation of *impossible*

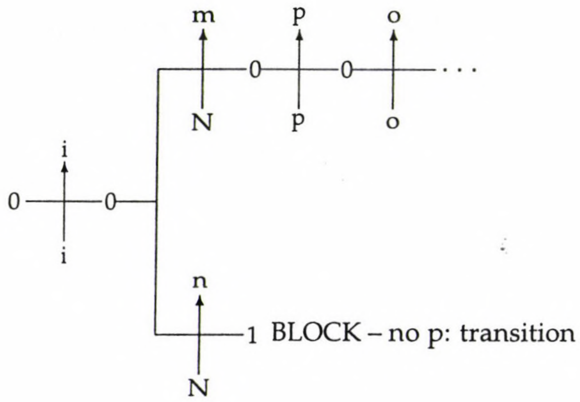

(9) Recognition of *input*

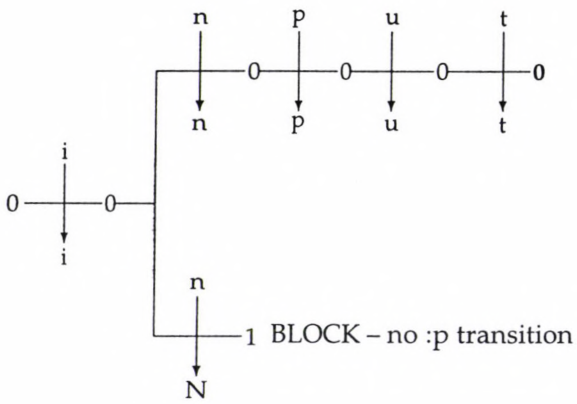

As we can see, the major goal of this treatment is that we no longer differentiate between generation and recognition, since the same algorithm (and the same FST) can account for both procedures.

1.3 Kaplan and Kay (1994)

In their article, Kaplan and Kay prove that every non-cyclic phonological rule system is regular, i.e. it can be modelled by one (though very complex) FST or, to put it in an other way, every non-cyclic phonological rule system defines a regular relation between underlying and surface forms. This treatment has two important benefits:

(i) intermediate representations can be done away with;

(ii) it provides for an effective machine for recognition as well as for generation.

The paper shows an algorithm how context-sensitive generative rules can be interpreted as regular relations. After this implementation, the authors rely on the closure properties of regular relations (such as closure under serial composition), which gives the desired result, namely that a whole rule-system also defines a regular relation. In the present overview, we would like to give only a taste how this implementation can be achieved.

Before we go into the necessary details, some definitions concerning regular relations and languages are in order. Let R be a regular relation. The image of a string x under the relation (x/R) is the set of all strings y for which $\langle x, y\rangle \in R$. If X is a regular language or relation, let $Opt(X)$ be $X \cup \{e\}$, where e is the empty string. If L is a regular language, then the relation $Id(L) = \{\langle l, l\rangle | l \in L\}$ carrying every member of L onto itself is regular. Moreover, if L_1 and L_2 are regular languages, then the relation $L_1 \times L_2 = \{\langle l_1, l_2\rangle | l_i \in L_i\}$ is also regular. Note the difference between $Id(L)$ and $L \times L$: if $L = \{a, b\}$ then $Id(L) = \{\langle a, a\rangle, \langle b, b\rangle\}$, while $L \times L = \{\langle a, a\rangle, \langle b, b\rangle, \langle a, b\rangle, \langle b, a\rangle\}$.

We need five other operators as well, which preserve regularity.

(10) i.Let S be a designated finite set. The $Intro(S)$ relation, defined by the expression $[Id(\Sigma) \cup [\{e\} \times S]]^*$ freely introduces symbols from S.

ii.Let S be a finite set and L a regular language. The language L_S (read 'L ignoring S') is also a regular language, where $L_S = Range(Id(L) \circ Intro(S))$. Language L_S differs from L in that occurrences of symbols in S may be freely interspersed in a string.

iii.If L_1 and L_2 are regular languages, the language *If-P-then-S*(L_1, L_2) ("if prefix then suffix") contains an string a if each of its prefixes in L_1 are followed by a suffix in L_2. Formally:

$$\textit{If-P-then-S}(L_1, L_2) = \{a | \text{ for every partittion } a = x_1 x_2,$$
$$\text{if } x_1 \in L_1 \text{ then } x_2 \in L_2\} = \overline{L_1 \overline{L_2}}$$

iv.The definition of *If-S-then-P*$(L_1, L_2) = \overline{\overline{L_1} L_2}$ is analogous.

v.Finally we can combine these two operators to impose that the prefix be in L_1 if and only if its suffix is in L_2:

$$\textit{P-iff-S}(L_1, L_2) = \textit{If-P-then-S}(L_1, L_2) \cap \textit{If-S-then-P}(L_1, L_2).$$

Now, we can set out to define the regular relation properly accounting for the following SPE-rule:

(11) $\phi \rightarrow \psi / \lambda \_\rho$

We assume, for the time being, that all letters represent strings in Σ^*, and that this rule is optional. As an initial step, we can define the relation *Replace* modelling the rule as follows:

(12) $Replace = [Id(\Sigma^*) Opt(\phi \times \psi)]^*$

The asterisk allows for repetitions of the $\phi \times \psi$ relation (multiple application of the rule), whereas $Id(\Sigma^*)$ accepts identical string pairs between the replacements. The image of the input string under the *Replace* relation is identical with the input except for possible replacements of ϕ substrings in x with ψ. However, this relation is not sensitive to the contextual part of the rule. As a second step we may simply add context requirements to the operator:

(13) $Replace = [Id(\Sigma^*) Opt(Id(\lambda)\phi \times \psi Id(\rho))]^*$

At a first glance, this relation includes pairs of strings which differ in that respect that some occurrences of ϕ in the input may be replaced with ψ in the output in the context required by the rule; this, indeed, is the thing we need. A closer inspection, however, shows that this relation, in fact, undergenerates. Consider, as an example, the following simple rule:

(14) B → b / V __ V

Our *Replace* relation accepts the pair on the left but not the one on the right:

V	B	V	B	V	V	B	V	B	V
V	b	V	B	V	V	b	V	b	V

The problem is that, in general, the right context of a rewrite rule in one application may serve as the left context of the same rule in another application on the same string. The *Replace* relation of (14) does not allow for such an application. We should incorporate this requirement into the definition.

Kaplan and Kay (1994)'s solution to the problem is ingenious. In order to keep track of the contexts, let us introduce two markers: $<$ and $>$. If we put a marker $<$ *after* each left-context (λ) and the other one, $>$ *before* each right context, then the possible sites for replacement are bracketed ($< >$), as it is shown below:

>V<B>V<B>V<	>V<B>V<B>V<	>V<B>V<B>V<
>V<b>V<b>V<	>V<b>V<B>V<	>V<B>V<b>V<

Now, the *Replace* relation would be:

(15) $Replace = [Id(\Sigma_m^*)Opt(Id(<)\phi_m \times \psi_m Id(>))]^*$

where $m = \{<, >\}$ is the set of markers and the subscript m stands for the $Ignore$ operation.[4] But how do the markers get into the string? They are introduced freely by the $Prologue$ operator:

(16) $Prologue = Intro(m)$

Our final task now is to construct an appropriate filter which only allows for properly placed left-context markers $<$. This filter would require that the left-context bracket $<$ appears if and only if it is preceded by the left context λ. Such an operator is $P\text{-}iff\text{-}S(\Sigma^*\lambda, <\Sigma^*)$. However, the situation is slightly more complex. Suppose that $\lambda = ab^*$. In this case, this operator accepts the bracketing $a < b <$, but refuses the bracketings $ab <$ and $a < b$. The brackets of the shorter prefixes must not prevent the proper bracketing of longer prefixes either. A possible reaction to this requirement is to ignore the occurrences of the left-context bracket $<$ in λ (and Σ^*). Furthermore, right-context brackets $>$ should also be disregarded. Hence, the operator:

(17) $P\text{-}iff\text{-}S(\Sigma^*_<\lambda_<, <\Sigma^*_<)_>$

Unfortunately, this relation disregards slightly too many brackets, since the left-context $\lambda_<$ followed by a bracket $<$ is also an instance of $\lambda_<$, so it should be followed by another $<$, which should also be followed by another $<$ and so forth. To identify left-contexts properly, we should not disregard left-context brackets $<$ following an instance of $\lambda_<$. Hence, the correct left-context identifier operation (filter) is:

(18) $Leftcontext(\lambda, <, >) = P\text{-}iff\text{-}S(\Sigma^*_<\lambda_< - \Sigma^*_< <, <\Sigma^*_<)_>$

A parallel definition of the filter $Rightcontext(\rho, <, >)$ can be obtained.

So, the following regular relation implements the optional, left-to-right SPE-type rule (12), where $\circ$ stands for (serial) composition:

(19)

$$\begin{array}{l} Prologue \circ \\ Id(Rightcontext(\rho, <, >)) \circ \\ Replace \circ \\ Id(Leftcontext(\lambda, <, >)) \circ \\ Prologue^{-1} \end{array}$$

In fact, there are some minor issues not addressed here, such as the problems of empty contexts, of obligatory rule applications, of various directions of application (left-to-right, right-to-left, simultaneous), of feature matrices and variables and of unordered rules. These are discussed thoroughly in Kaplan and Kay (1994), and the authors prove that they do not constitute a problem for the regularity of the model. Since this section is planned to give just a taste of the methods and tools of Kaplan and Kay (1994), we refer the interested reader to the original source.

The main theorem in the center of investigation in the section was:

---

[4]The markers must be ignored in the replacement proper, since they may occur within ϕ or ψ.

(20) Every non-cyclic phonological rule system defines a regular relation between the underlying and surface representations.

This is a very important result, however, the term 'non-cyclic' imposes certain restrictions on the grammar. First of all, it cannot contain a rule which can freely apply to its own output, such as (3) repeated here for the sake of convenience:

(21) $e \rightarrow ab/a\_b$

Fortunately, this kind of rule is rare (if not non-existent) in phonological rule-systems, so it does not constitute a real problem for phonologists. We also saw that multiple applications of a rule on the same string (but not on its own output!) is not problematic either for the regular approach (cf. 14). However, in principle, nothing prohibits that the output of a rule *R1* is the input of another rule *R2* whose output is the input of the first rule, etc. This kind of application is indeed possible, and, in fact, it is one of the basic tenets of a widely-accepted phonological framework, Lexical Phonology.

In Lexical Phonology, ordered lexical (cyclic) rules are claimed to apply in consecutive cycles. The morphological information is reflected by brackets when a word enters the phonological module: *[un[[en[force]]able]]*.[5] The rules first apply on the string between the innermost brackets. Then these brackets are deleted (*Bracket Erasure*), and the rules reapply on the material between the appropriate brackets, and so on until all brackets are deleted. This process results in – in principle unbounded – reapplication of the rules; the relation thus obtained is no longer regular.[6] Let us see the authors' opinion on the problem:

(22) 'The cycle has been a major source of controversy ever since it was first proposed by Chomsky and Halle (1968), and many of the phenomena that motivated it can also be given noncyclic descriptions. Even for cases where nonrecursive, iterative account has not yet emerged, there may be restrictions on the mode of reapplication that limit the formal power of the grammar without reducing its empirical or explanatory coverage.' (Kaplan and Kay (1994, p. 365.))

In their article, Kaplan and Kay prove that – in general – phonological rule systems are equivalent with two-tape FSTs, i.e. phonology of a given language can be viewed as a two-level regular relation, or equivalently, as an FST. This treatment has the advantage in generation and especially in recognition. It is criticised, however, for two reasons:

(i) the algorithms for construction of the single FST are very slow and ineffective, hence no real implementation of a whole phonological system has ever been given;

(ii) the FST constructed from the rules is so complex that it loses its explanatory power for phonologists.

In the next section, we shall see another two-level model, the KIMMO formalism, which allows for both straightforward, efficient implementation and phonological plausibility.

1.4 The KIMMO formalism

This section is a gentle introduction to the formalism of Kimmo Koskenniemi (1983) based on Antworth (1990) and Karttunen (1993). The formalism relies on the knowledge that phonological rule sytems are in fact regular relations between underlying and surface forms. This relation can be decomposed into individual rules but in a different fashion that has been seen so far. In this model, rules (individual FSTs) do not apply serially but they work as parallel constraints on the input-output pairs. Such a treatment has two advantages over the single complex FST approach: (i) it is easy to implement (in fact, this implementation has already been carried out for languages such as Finnish, English, Russian and French); (ii) it is phonologically plausible.

---

[5] Under other views, phonological processes apply within the lexicon itself, after each morphological process.

[6] In fact, it cannot be modelled by any relation, since we cannot know how many items the composition of rule will consist of.

The KIMMO system consists of two major parts: the set of *lexical character:surface character* correspondence pairs (*feasible pairs*)[7] and the set of the rules, like:

(23) t:c ⇒ __ i

Feasible pairs have two subtypes: *default* correspondences, such as *i:i, t:t,* (i.e. the elements of $Id(\Sigma)$) and *special* correspondences like *t:c* in the example above.

The rules, at a first glance, seem to be similar to the traditional transformational rules. However, there are important differences. Transformational rules are rewrite rules in the sense that they *change* lexical (underlying) representations into surface forms (via unbounded number of intermediate representations). So the rule:

(24) t → c / __ i

actually turns *t* into *c*, thus the lexical *t* no longer exists after the rule applied. In contrast, two-level rules of the KIMMO formalism are declarative rules expressing correspondences between lexical and surface forms, and not changing the former into the letter. A further difference between the two formalisms is that two-level rules do not apply sequentially (like transformational rules do) but in a parallel fashion.

A two-level rule is made up of three parts: the *correspondence*, the *rule operator* and the *environment* or *context*. The notation allows a number of possible shorthands in the rules:

(25) **i** This stands for the correspondence $i : i$.

i:@ The notation means '*any* possible feasible pair with lexical i regardless of how it is realised on the surface.' This is usually simplified to i :.

@:i The same as above but now the surface realisation must be i. The shorthand for this is $: i$.

The rules can also refer to subsets of sounds such as:

	SUBSET D	t, d, s, z	(dental stops)
(26)	SUBSET P	c, j, š, ž	(palatal stops)
	SUBSET Vhf	i, e	(high, front vowels)

A palatalization rule can now be expressed as:

(27) D:P ⇒ __ Vhf

An important feature of two-level rules is that they do not allow deletion in the literal sense, rather they introduce the symbol 0 to express material not phonologically or lexically realised. These can be stress marks (′) in lexical representation and morpheme boundaries (+) in the surface form:

(28)	LR:	0	t	a	t	+	i
	SR:	′	t	a	c	0	i

The rule operator expresses special logical relations between the correspondence and the environment it may occur. There are four types of rule operators, expressing conditional or implicational relationship.

(29) t:c ⇒ __ i:i

The operator ⇒ means 'only but not always.' That is, the rule can be translated as 'lexical t corresponds with surface c only if it is followed by $i : i$' but not necessarily always does so in that environment. This is roughly the optional rule application in the SPE-formalism. When implementing this rule, we first have to construct an FST modelling it. It is quite easy:

[7] In fact, this set need not be constructed, since the correspondence pairs can be learned from the rules. Nevertheless, this approach does not make any difference.

(30) 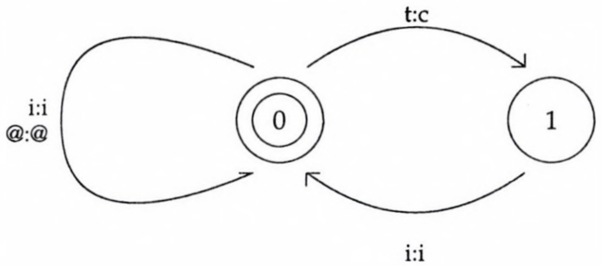

	t	i	@
	c	i	@
0:	1	0	0
1.	B	0	B

The arc @ : @ indicates any feasible pairs except for the ones explicitly mentioned in the transducer ($t : c$ and $i : i$ in this case). The table on the right is the corresponding *state table*: capital B means that there is no transition from the given state, i.e. the machine blocks. The colon : next to the states indicates final states, the dot . non-final states.

The second type of operator ⇐ expresses obligatory rule application ('always but not only'):

(31) t:c ⇐ __ i:i

Thus, the rule means: lexical t always corresponds with surface c if it is followed by $i : i$. However, such a correspondence may exist in other environments as well. The corresponding state table is:

(32)

	t	t	i	@
	c	@	i	@
0:	0	1	0	0
1:	0	1	B	0

Here, the column t : @ expresses any feasible pair with lexical t except for the one explicitly mentioned, i.e. $t : c$. The machine blocks if and only if a pair t : @ is followed by $i : i$, or equivalently, if $i : i$ is preceded by a lexical t which does not correspond with surface c.

The two operators seen so far can be combined into one ⟺ operators, meaning 'always and only':

(33) t:c ⟺ __ i:i

	t	t	i	@
	c	@	i	@
0:	2	1	0	0
1:	2	1	B	0
2.	B	B	0	B

The rule can be formulated as: lexical t corresponds with surface c if and only if it is followed by a lexical i which is realised on the surface as i. The transducer is simply the combination of the two FSTs seen before.

The last operator is /⇐ can be translated as 'never':

(34) t:c /⇐ __ i:y

	t	i	@
	c	y	@
0:	1	0	0
1:	1	B	0

The rule prohibits the occurence of the pair $t : c$ if followed by a lexical i realised as y on the surface.

Finally, let us see how a two-level rule such as:

(35) **R** t:c ⇒ __ i

works in the case of, for example, generating the surface form. The correspondence of the rule is a special correspondence; the two-level description of the language must also contain the default correspondences such $a : a$, $i : i$ and $t : t$.[8] Beginning the first character of the input, the generator finds a correspondence with lexical *t* (our rule):

(36)

LR:	t	a	t	i
	|			
Rule:	R			
	|			
SR:	c			

At this point, the generator entered the rule R, which states, however, that a $t : c$ correspondence must be followed by the correspondence $i : i$, which is not the case at present. Hence, the generator must back up, and try the default correspondence $t : t$. It in fact works so the generator can continue with the next lexical character a. The only feasible pair defined for lexical a is $a : a$, thus the generator proceeds onto the next segment t. At this point, it can enter the rule R again, positing a surface c:

(37)

LR:	t	a	t	i
	|	|	|	
Rule:	|	|	R	
	|	|	|	
SR:	t	a	c	

Now, the generator encounters a lexical i, which must correspond with a surface i by the rule R. Hence, the generator produces the surface form *taci*. However, the generator is not done yet. It will continue backtrackig, trying to find alternative realisations. First, it will undo the correspondences $i : i$ and $t : c$ and try the default $t : t$:

(38)

LR:	t	a	t	i
	|	|	|	
Rule:	|	|	|	
	|	|	|	
SR:	t	a	t	

Now, it will continue with the default $i : i$ correspondence, generating the other correct surface form *tati*. Since there are no other backtracking paths, the generator exits. If the system contains more than one rule, the procedure is similar. Parallel application means that in each step the possible correspondences are the intersection of the correspondences posited by the rules involved.[9]

A final note is in order concerning the formalism developed in this section. Two-level rules have the advantage of simplicity and plausibility over the single-FST approach of the previos section. The regularity of the system, however, seems to be questionable: two-level rules, in fact, define regular relations between lexical and surface representations. Parallel application means intersection of these regular relations, which – as we saw above – may no longer be a regular relation. That is, the KIMMO formalism seems to have more formal power than the traditional SPE rule-sytems, which have been critcised for overgeneration. A closer inspection, however, reveals that the relations of this formalism are *same-length* regular relations, which family of regular relations has been proved to be closed under intersection by Kaplan and Kay (1994). Thus, the two approaches have equal power.

The SPE-formalism has been criticised for overgenerating. Rewrite rules can express phonologically implausible changes with the same simplicity as plausible events can be stated:

(39) (i) [– sonorant] → [+ nasal] / __# *vs.*
(ii) [– sonorant] → [- voiced] / __#

[8] Note, if $Id(\Sigma)$ is necessarily a subset of the correspondence set, then this formalism does not permit absolute neutralisation, i.e. underlying segments that do not appear on the surface.

[9] An alternative approach might be generating the sets of possible outputs for each rule and the output of the system will be the intersection of these sets.

There is no language which would work in the way the first rule requires, whereas the change expressed by the second rule is very frequent. This problem along with similar cases led to the emergence of a new framework, i.e. Autosegmental Phonology in the late 1970s and 1980s (eg. Goldsmith (1976)). The formalisms considered so far are, however, unable to model autosegmental representations. In the following section we shall see an approach coping with the problems raised by modelling Autosegmental Phonology, namely the one-level approach of Bird and Ellison (1994).

1.5 One-Level Phonology

Phenomena such as vowel harmony, the regularity of tonal patterns and the behaviour of tones in Bantu languages and problems like in (39) gave a rise to Autosegmental Phonology. Within this framework, representations are no longer viewed as linear sequence of feature matrices but rather as autonomous segments or *autosegments* – like [+ nasal], [+ high], etc. – associated with skeletal positions. Thus, under this view, the representation of i bearing high tone would be something like:[10]

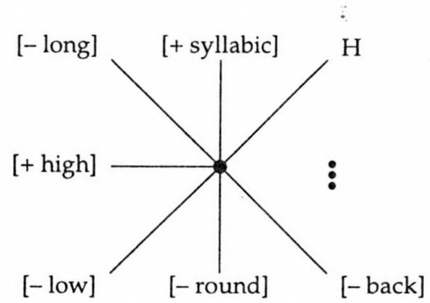

The possible associations are constrained by the *No Crossing Constraint*, which ensures that association lines cannot cross. In this framework, the most important (if not the only) phonological processes are *spreading*, *delinking* and *deletion* of autosegments.

Autosegmental representations are more or less accepted in all current phonological frameworks; there is, however, a frequently debated issue, the *Obligatory Contour Principle*, stated as:

(40) **OCP**: At the melodic level of the grammar, any two adjacent [autosegments] must be distinct. Thus HHL is not a possible melodic pattern; it automatically simplifies to HL. Bird and Ellison (1994, p. 59.)

Autosegmental representations have always been a challenge for computational linguist, trying to interpret and formalise autosegmental charts. Several authors, such as Kay (1987) or Kornai (1991), have also tried to incorporate them into the finite-state phenomena. One of the most successful approaches was that of Bird and Ellison (1994), especially when measured against Kornai's four desiderata (Bird and Ellison (1994, p. 72.)):

i. **Computability** The number of terms in the encoding is equal to the number of autosegments, and each term has a fixed size. Therefore, the encoding can be computed in linear time.

ii. **Compositionality** If D_1 and D_2 are two autosegmental diagrams then $\mathcal{E}(D_1 D_2) = \mathcal{E}(D_1)\mathcal{E}(D_2)$, where concatenation of encodings – $\mathcal{E}$ – is done in a tier-wise manner. Thus the encoding is compositional.

iii. **Invertibility** A representation can be reconstructed from its encoding.

[10] In fact, autosegments are hierarchically ordered; they have a certain geometry, called *feature geometry* (cf. Clements (1985)). Note also that it is not customary to use the feature [+ long]; rather, length is represented by association to two timing units.

iv. **Iconicity** If an autosegment in the diagram is changed, the effect on the encoding is local, since only one term is altered. However, if an association line is added or removed, two terms must be altered.

In this section, we will introduce the notation and methods of this article, which model incorporates the OCP as well.

First we have to formalise what the association in (41) exactly means:

(41)
$$\begin{array}{ccc} \ldots & A & \ldots \\ & | & \\ \ldots & B & \ldots \end{array}$$

It can be interpreted as a partial overlap (cf. Bird and Klein (1990)) of the *intervals* representing A and B. Thus, any of the following strings is described by the diagram above:

(42)

A	A	A	•	•
•	•	B	B	B

A	A	A	•	•
•	•	B	B	•

•	•	A
B	B	B

This treatment has two important advantages over any other approach to autosegmental representation. First, the problem of line-crossing is evaded, since intervals are sequenced linearly on their tier. Secondly, but more importantly, the intervals can extend freely, thus the OCP is rendered trivial in this model: two adjacent intervals of the same feature is in fact one – though longer – interval of the feature.

For the representation of such an association, however, we must introduce a new device, the *State-Labelled Finite Automaton* (SFA),[11] which is a septuple $(V, \Sigma, \lambda, \delta, S, F, \alpha)$, where:

V is a finite set of *states;*

Σ is a finite *alphabet;*

$\lambda \subseteq V \times \Sigma$ is the *labelling relation* (i.e. the states are labelled with subsets of the alphabet);

$\delta \subseteq V \times V$ is the *transition relation;*

$S \subseteq V$ is the set of *start states;*

$F \subseteq V$ is the set of *final states;*

α is a Boolean *flag* that is true if and only if the null string e is accepted by the automaton.

A *situation* of an SFA A is a triple $\langle x, T, y \rangle$ where $T \subseteq V$ is the set of currently active states, and x and y are the portions of the input string to the left and right of the reading head, respectively. If $\langle x, T, y \rangle$ and $\langle x', T', y' \rangle$ are two situations, then $\langle x, T, y \rangle \vdash_A \langle x', T', y' \rangle$ iff there is a $\sigma \in \Sigma$ such that:

(i) $y = \sigma y'$ and $x' = x\sigma$ (σ is the first symbol in the string x and the last in y'),

(ii) for each $t' \in T'$ there is a $t \in T$ such that $\langle t, t' \rangle \in \delta$ (the new situation must be reachable from the previous one), and

(iii) $\langle t', \sigma \rangle \in \lambda$ for each $t' \in T'$ (σ is the label of each currently active state).

The transitive closure of $\vdash_A$ is $\vdash_A^*$. Finally, the automaton A accepts a string w iff either $w = e$ and α is true, or $\langle \sigma, \{s\}, \beta \rangle \vdash_A^* \langle \sigma\beta, F', e \rangle$, for some $\langle s, \sigma \rangle \in \lambda$, $s \in S$, $\beta \in \Sigma^*$ and $F \cap F' \neq \emptyset$, and where $w = \sigma\beta$.

The SFAs are, in fact, equivalent to FSAs in formal power, but – Bird and Ellison (1994) claim – they are empirically more adequate. Consider, as an illustrative example, the following automata that prohibit two adjacent occurences of any symbol (out of the four: $\{a, b, c, d\}$), the constraint of OCP. The machine on the left is the FSA, the one on the right is the corresponding SFA. Note the differences of notation from what we have seen so far: bullets stand for states, circles final states, whereas initial states are marked by $\succ$. The state labelled with $\emptyset$ indicates that the automaton accepts the empty string e.

[11] Bird and Ellison (1994, p. 59ff.)

(43) The OCP automata

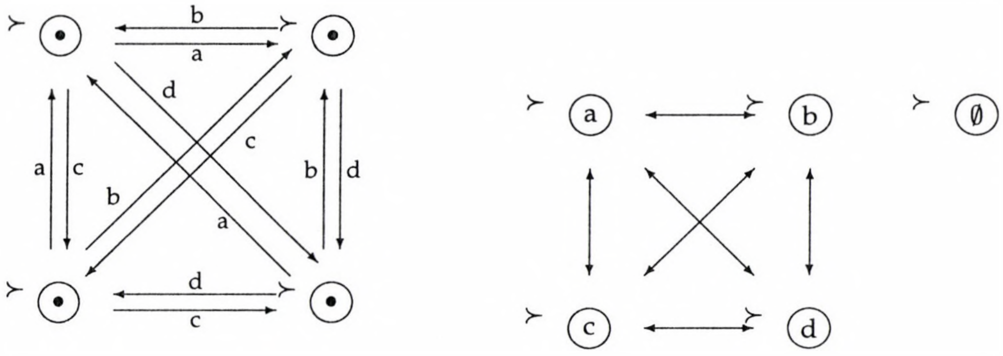

How can an autosegment P be modelled by this new machine? It is fairly easy:[12]

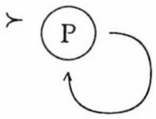

(44)

Now, consider the association:

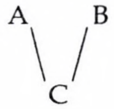

(45)

How can we model this diagram with an automaton? First, we construct a *synchronised* SFA: this is in fact two automata running in parallel, but the states linked by vertical lines must be active simultaneously.

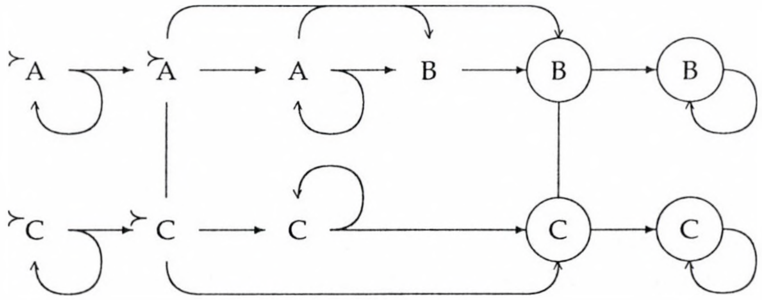

(46)

As a second step, we simulate this synchronised SFA with adding indices to each state: 0 is associated with unsynchronized states, whereas 1 with synchronised states. Then we erase the lines.

[12]Here, the interpretation of P is in fact the set of all segments bearing the autosegment P. Thus, for example, [+ high] stands for the set $\{i, u, y\}$.

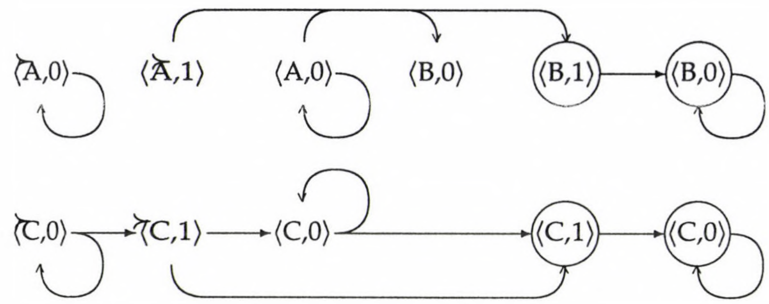

(47)

Now, let us see the intersection (⊓) of these two automata (defined over the alphabet $\Sigma \times \{0,1\}$):

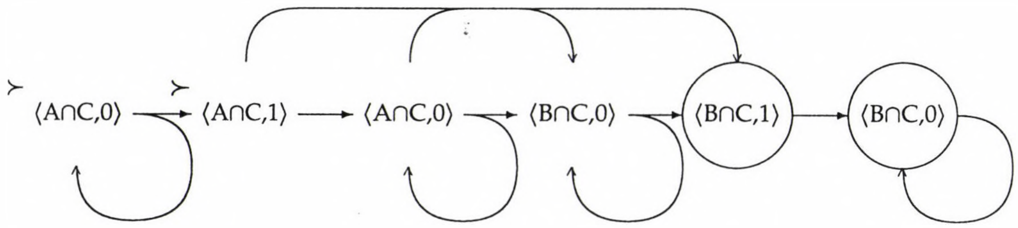

(48)

The function of indices was to rule out certain states in the intersection. Now, they have served their purpose, so they can be omitted. Thus we arrive at our final SFA representing the autosegmental association in (49):

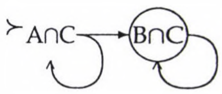

(49)

Thus, we managed to construct an automaton, representing the autosegmental diagram in (45). In fact, we can construct automata for more elaborate charts with the same ease. Furthermore, it is obvious that this treatment can be easily extended to multiple tiers. Thus, the autosegmental representation of a string can be viewed as an SFA. The only thing left we have to cope with is the account for autosegmental rules in this model.

Any generative rule has the format: $SD \rightarrow SC$, i.e. any string that meets the *structural description* of the rule must undergo the *structural change*. In an equivalent formula: $\neg\exists s \subseteq S$, SD($s$) ∧ ¬SC($s$). This can be expressed as: $\neg(\bullet^*(\mathrm{SD} \sqcap \overline{\mathrm{SC}})\bullet^*)$. Autosegmental rules have the same format, though this is not necessarily transparent for the first glance. Consider, for example, our familiar nasal place-assimilation rule (**Rule 1**):

(50) 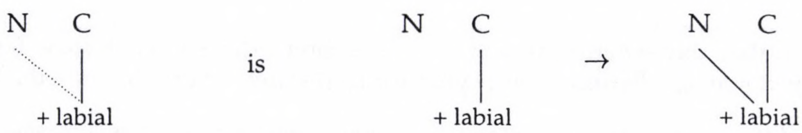

Now, the corresponding rule format is the following:

(51)

$$\neg\left(\bullet^{*}\left[\left[\begin{matrix}\bullet:0^{*} & \text{N:0} \quad \text{C:1} & \bullet:0^{*}\\ \bullet:0^{*} & \text{+labial:1} & \bullet:0^{*}\end{matrix}\right] \sqcap \neg\left[\begin{matrix}\bullet:0:0^{*} & \text{N:0:1} \quad \text{C:1:0} & \bullet:0:0^{*}\\ \bullet:0:0^{*} & \text{+labial:1:1} & \bullet:0:0^{*}\end{matrix}\right]_{SC}\right]_{SD}\bullet^{*}\right)$$

In evaluating (51), we intersect the two tiers of the SC part, and then delete the second index of each tuple. Next, the complement of this automaton is intersected with the SD part. Then, we delete the first index of the tuples. Finally, we add the $\bullet^*$ wildcards and form the complement. Thus, we obtain the SFA rejecting all nonhomorganic nasal-labial clusters:

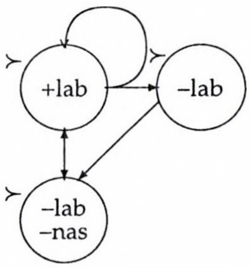

(52)

The algorithm for more complex rules is the same; thus, all autosegmental rules can be represented by an SFA. Note that the rules are modelled in the same way as lexical representations in this framework. Indeed, both the lexical form of a string and the rules are inviolable *constraints* on the *surface* forms. Thus, these constraints work in a parallel fashion, i.e. we have to intersect of the SFAs to account for the strings on the surface.

1.6 Conclusion

In this paper, we have discussed three important finite-state approaches to phonological representations and rule systems. The first model, that of Kaplan and Kay (1994), the single-FST approach differed from the others in that respect that it implemented serial generative rule-systems. On the other hand, the KIMMO formalism and Bird and Ellison (1994) view the grammar as parallel constraints. Though these two approaches (the serial and the parallel one) are equivalent in formal power (both define regular languages and relations), the question – which is more appropriate for the description of phonological processes? – may well arise. This is briefly discussed in Karttunen (1993, p. 186ff.).

Another difference between the three approaches groups the KIMMO formalism and the single-FST model together: they both differentiate between lexical and surface representations on the one hand, and representations and rules, on the other. In contrast, in the one-level phonology of Bird and Ellison (1994), rules and representations are viewed as the same: they both are SFAs, and they both represent constraints on the surface realisations of a string. This roughly means that there is no such a thing as underlying form (hence the label 'one-level' phonology).

Finally, the KIMMO formalism differs from the other two in one important respect. While Kaplan and Kay (1994) and Bird and Ellison (1994) both encode the traditional generative rule format: $SD \to SC$, the KIMMO formalism relies on rules of a different kind, the two-level rules. Here, apart from the implication $SD \to SC$, $SD \leftarrow SC$, $SD \leftrightarrow SC$ and $SC/ \leftarrow SD$ can also hold in the rules. This results in easier computational implementation and greater phonological transparency, though this approach bears the same formal power as the others, namely that all non-cyclic (or having a bounded number of cycles) phonological rule-systems define a regular relation between lexical and surface forms, i.e. they can be modelled by finite-state devices.

References

Antworth, E. 1990. *PC-KIMMO: A Two-level Processor for Morphological Analysis*. SIL, Dallas.

Bird, Steven and T. Mark Ellison. 1994. One-level phonology: Autosegmental representations and rules as finite automata. *Computational Linguistics*, 20(1):55–90.

Bird, Steven and Ewan Klein. 1990. Phonological events. *Journal of Linguistics*, 26:33–56.

Chomsky, Noam and Morris Halle. 1968. *The Sound Pattern of English*. Harper and Row.

Clements, George N. 1985. The geometry of phonological features. *Phonology Yearbook*, 2:225–252.

Goldsmith, John. 1976. *Autosegmental Phonology*. Ph.D. thesis, MIT.

Johnson, C. Douglas. 1972. *Formal Aspects of Phonological Description*. Mouton.

Kaplan, Ronald M. and Martin Kay. 1994. Regular models of phonological rule systems. *Computational Linguistics*, 20(1):331–378.

Karttunen, Lauri. 1993. Finite state constraints. In John Goldsmith, editor, *The Last Phonological Rule*. The University of Chicago Press, Chicago and London, pages 173–194.

Kay, Martin. 1987. Nonconcatenative finite-state morphology. In *Proceedings of the Third European Conference of the Association for Computational Linguistics*, pages 2–10.

Kenstowicz, Michael. 1994. *Phonology in Generative Grammar*. Blackwell.

Kiparsky, Paul. 1985. Some consequences of lexical phonology. *Phonology Yearbook*, 2.

Kornai, András. 1991. *Formal Phonology*. Ph.D. thesis, Stanford University.

Koskenniemi, Kimmo. 1983. Two-level morphology: A general computational model for word-form recognition and production. Publication No. 11, Department of General Linguistics, University of Helsinki.

Chapter 2

Three-level Phonology

2.1 Introduction

Classical generative phonology, as formulated in the *Sound Pattern of English* (Chomsky and Halle, 1968) has received much criticism and has by now been developed into several new theories within generative phonology itself: lexical phonology, for example, was an attempt to integrate morphology and phonology, while autosegmental phonology announced the break with linear forms of representation. Yet all generative approaches maintained the idea that phonological processes are a set of symbol-manipulating rules which mediate between two levels of representation – the underlying or phonological and the surface or phonetic levels – and which apply sequentially, in a language-specific, often extrinsic order. These core concepts of mainstream derivational phonology are challenged by several new non-generative frameworks all belonging to the broader perspective of Declarative Phonology (DP). Common to every new theory under this name is the pursuit of a restricted phonology which, describing processes in a radically different way making its theoretical constructs better approximations of real-time mechanisms, is of greater psychological plausibility than its generative predecessors. In compliance with such an endeavour, DP accounts share the conception that constraints should replace rules. Some of them use formalisms that can be naturally implemented on connectionist networks, which are now thought to be a model of the physical structure and the parallel operation of the brain. Here we would like to introduce a three-level declarative approach, which is termed Cognitive Phonology (Lakoff, 1995) or Harmonic Phonology (Goldsmith, 1995a) in the literature, but the name Construction Phonology, in the light of its similarities to Construction Grammar, would also be appropriate. Throughout the paper we will use all these terms as synonyms. After outlining the general principles of declarative approaches and Cognitive Phonology, we will focus on the question concerning the number of levels. In the second part of the paper, armed with the formalism of the theory, we test the framework, and hope to prove that it stands the trial: we present three-level declarative reanalyses of two textbook examples of the need for rule ordering, and also show through a third example that iterative rules receive a simpler treatment in DP.

2.2 Declarative phonology

One of the main criticisms concerning generative phonology is levelled at the concepts of rules (rewrite rules of the form $A \rightarrow B / C \_ D$) and rule ordering. Declarative Phonology's central objective is to rid itself of both mechanisms. Declarative grammars replace transformational rules with constraints on well-formedness. Ordered rules are symptoms of a theory of many levels, and, more importantly, of levels which cannot claim the status of being mentally real: in an ordered set of rules the application of each rule would result in a representation on a new stratum, but there is no evidence of the existence of multiple levels of intermediate representations in the mind; besides, speakers of languages with different number of ordered rules would differ in the number of their representational levels. As a consequence DP has developed systems that only admit a small number of levels that have a psychological parallel in representation, and no intermediate representations. Cognitive phonology assumes that three levels are necessary and sufficient.

Another reason for eliminating rule ordering is that linear application of rules would result in an absurdly long computation in our brains: cognitive mechanisms occur much faster than sequential rule operation would permit. Rules and especially rule ordering are computationally expensive; a declarative approach has to offer a simpler set of representational devices. The theory was termed cognitive phonology by George Lakoff, who emphasised that "Cognitive phonology is to be seen as part of cognitive grammar. As such, it assumes that phonology, like the rest of language, makes use of general cognitive mechanisms, such as cross-dimensional correlations". The attribute 'cognitive' reveals two features in point: psychological plausibility and a mechanism to replace rules, 'cross-dimensional correlations'.

Declarative approaches are constraint-based grammars. Rules, if they are considered processes of symbol-manipulation the application of which is a derivation, find no place in declarative phonology. DP makes reference only to rules in the sense of generalisations, 'in the form of a filter, implication or positive template' (Scobbie, 1993, p. 161). The theory expresses generalisations about the phonology in the form of well-formedness constraints. Three-level phonology in its specific formalism uses *constructions* that are either correlations between levels or well-formedness constraints concerning a particular level. A possible construction scheme is given in (1). Letters from the end of the alphabet denote levels; letters from the beginning of the alphabet stand for features or bundles of features. The construction reads as follows: An X-level A corresponds to B at the Y-level, if it precedes a C at the X-level:

```
(1)   X:  A  C
          |
      Y:  B
```

Identity across levels is the default, which is outweighed by constructions. Constraints, unlike rules, do not apply sequentially but are satisfied simultaneously and are in continuous operation: they are invoked as long as their conditions are met, without extrinsic ranking. They are, in contrast with the constraints of Optimality Theory, inviolable or 'hard' constraints which have to be compatible with each other: any true conflict (the Elsewhere Condition[1] does not lead to a true conflict) would mean that the set of constraints describing the language is inconsistent (Scobbie, 1993).

Direction-neutrality is an additional characteristic feature of constraints which makes the theory simpler: what in generative theory had been an outcome of a rule identified as left- or right-iterative, are now results of the combination of continuous constraint-satisfaction and application from left-to-right, proceeding in the representation in accordance with the natural flow of time. Direction neutrality is also a property of correspondences across levels: in contrast with generative rules, any two-level constraint, describing a correlation rather than a change, is unspecific and neutral to its direction: an X-Y constraint holds in both directions, and is thus a mechanism of both production and perception.

The term three-level phonology implies the theory's concern with *levels* of representation instead of *rules* and *representations*, which were in the forefront of previous scrutiny. Dispensing with rule ordering and striving for psychological plausibility, the number of levels has to be small, as we have pointed out earlier. The three levels are (Goldsmith, 1995a, p. 32)

i. M-level: a morphophonemic level, the level at which morphemes are phonologically specified;

ii. W-level: word-level, the level at which expressions are structured into well-formed syllables and well-formed words, but with a minimum of redundant phonological information,

iii. P-level: phonetic level, a level of broad phonetic description that is the interface with the peripheral articulatory and acoustic devices.

---

[1]Kiparsky's Elsewhere Condition is cited in Kenstowicz (1994, p. 216):
Rules A and B in the same component apply disjunctively to form Θ if and only if

a. The structural description of A (the special rule) properly includes the structural description of B (the general rule).

b. The result of applying A to Θ is distinct from the result of applying B to Θ.

In that case A is applied first, and if it takes effect, B is not applied.

These levels appear psychologically real if one conceives of the morphophonemic level as the one at which morphemes are stored in the mental lexicon of the mind. The word level is phonology, as declarative phonologists think of it, i.e. a set of well-formedness constraints – most of which operates at this level. The phonetic level stores instructions sent to the articulatory organs and information of the inputs decoded by the acoustic ones. The characteristics of the three levels are independent of each other.

2.3 Why three levels?

This conception of phonology aims at a maximally constrained grammar. Other more computationally oriented branches of the theory restrict the number of levels to just two (Karttunen, 1995) and declarative phonology in its more restricted form aims at monostratality (e. g., Bird and Ellison (1994)). Cognitive phonology works with three levels, which according to both Goldsmith and Lakoff are necessary and sufficient. Goldsmith puts down the need for three levels to specific types of orderings (cf. (2)). For feeding relations, he says, two levels would suffice, but bleeding and counterfeeding relations require three. As we will see, none of the four types of conjunctive order of *two* rules demands more than two levels. Yet a two-level constructionist phonology would not be able to accommodate certain phenomena: languages with more than two rules in strict (and specific types of) order cannot be adequately described by constraints making reference to only two levels.

Analysis of any kind of conjunctive order – in which the ordered rules apply subsequently so that "if rule A applies to derive a representation x, a subsequently ordered rule B must apply to x if satisfies the structural description of rule B, the final output is thus the conjunction of the application of the two rules" (Kenstowicz, 1994, p. 216) – is fairly straightforward in cognitive phonology, since we can formulate a systematic way of representing each kind of order as a two-level construction.

The definitions of the four types of extrinsic orders (feeding, counter-feeding, bleeding and counter-bleeding) are given in (2), taken from Kenstowicz (1994, p. 94) for recapitulation.

(2) a. Two rules A and B stand in a potentially feeding relation if the application of A creates new input to B. If B applies, then A is said to *feed* B, if B does not apply, then A and B stand in a *counterfeeding* relation.

b. Two rules A and B stand in a potentially bleeding relation if the application of A removes inputs to B. If B does not apply, then A is said to *bleed* B, if B does apply, then A and B stand in a *counterbleeding* relation.

It is important to note that though the order and the relation of rules is conceived of as being subject to language-specific parameter-setting, sometimes the relation of rules cannot be classified as belonging exclusively to one of the four. A feeding rule can also bleed the consequent rule depending on the specific representation, and in like manner a counterfeeding relation can also be counterbleeding applied to a specific class of inputs. The relatedness of feeding and bleeding on the one hand, and counterfeeding and counterbleeding on the other, is not realised by the representational inventory of generative phonology; we will see that the constructions corresponding to the ordering relations describe this phenomenon adequately by using just two kinds of interlevel correlations, simplifying the grammar one step further.

In a feeding or bleeding relation it generally holds that the rule ordered earlier does not create or take away the *target* of the second rule; its effect is producing or erasing potential environments of the second rule[2]

i. $A \rightarrow B / \_\_ C$[3]

[2]We can conceive of an ordering in which the first rule creates targets for the second rule, but these rules, if they apply in the same environment can be just as well stated as one rule, dismissing the intermediate step; if the environment of one is the proper subset of the other, the two rules can be reformulated as two rules standing in a disjunctive order (i. e. one is the Elsewhere Condition of the other).

[3]All generative rules and constructions are given with a right-environment. The environment of both the generative rules and the construcions could be written as left- or two-sided environment without any significant change to the generality of the the statements.

ii. feeding: D → E/ __ B bleeding: D → E/ __ A

The constructions have to be of the following kind (the symbols are to be interpreted as above)

```
i.  X:  A   C
        |
    Y:  B

ii. feeding             bleeding
    X:  D               X:  D
        |                   |
    Y:  E   B           Y:  E   A
```

Thus, two level constructions take the place of three-step derivations. A crucial feature of both representations is that the environment of the second rule is stated on the second level, where every first level A corresponds to B, which, in a feeding order, is the required environment; the same arrangement blocks the application of the second rule in a bleeding order. The environment of the first rule is not necessarily stated at the first level (if there is no rule that feeds or bleeds it). As a consequence, any number of rules which stand in a strict feeding or bleeding order, or any combination of orders of these kinds can be represented by making reference to just two levels in cognitive phonology.

Patterns of counterfeeding and counterbleeding rules are illustrated with the same three rules, with reversed order.

i. counterfeeding: D → E/ __ B counterbleeding D → E/ __ A

ii. A →B/ __ C

The corresponding constructions are shown below:

```
i.  counterfeeding          counterbleeding
    X:  D     B             X:  D     A
        |                       |
    Y:  E                   Y:  E

ii. X:  A   C
        |
    Y:  B   (C)
```

In counterfeeding and counterbleeding orders, rules do not affect each other's application. The necessity of ordering in these cases is created by the need to block the application of the rules in the reverse arrangement, in which they would stand in a feeding and bleeding relation, respectively. We have shown constructions representing these orders by stating the environment of the two rules on the same level. In counter-relations, though, it is only important that the environment of the first rule be stated in the construction on an earlier level than the output of the second rule. Since the environment of the second rule can be given on either levels, any number of counterfeeding or counterbleeding relations or any combination of just these two kinds can be described on two levels, provided all the environments are stated on the first level. In this way, whatever the effect of one construction is, it does not concern the other construction. The reader can verify for himself the validity of the constructions for different empirical cases of orderings. Some of the types will be shown in the examples. Of primary interest here is the systematic way of rewriting different orders (which in any case made reference to three levels of representation) and the conclusion that all four relations can be described with the help of correlations between two levels. These achievements also allow (and force) us to use the same kind of constructions for related types of conjunctive order and to organise them into two natural groups, feeding and bleeding orders constituting one group, whereas counterfeeding and counterbleeding orders form the other.

Up to this point we have not verified the need for three levels. All we have shown is that no type of conjunctive ordering between two rules necessitates a three-level representation. Now imagine a language with 3 rules, standing in an order so that rule 1 feeds rule 2, which in turn counterfeeds rule 3.

The rule schemes could be written as

i. A →B/ __ C

ii. D → E/ __ B

iii. F → B/ __ G

Let us look at the potential constructions

i. X: A C
   |
Y: B

ii. X: D
   |
Y: E B

We have already shown that in a feeding relation the second construction has to have a second-level environment, so that the output of the first rule, which is also on the second level can trigger the application of the second. We have also seen that in any kind of counter-relation the environment of the first construction should be at an earlier level than the output of the second. Since crucially the environment of the construction for rule 2 is already on the second level, we need a third level for the output of the construction for rule 3. So the third construction would be:

3. X: F G
   |
Y: B

For any strict ordering of three rules in which the first two rules stand in a feeding or bleeding relation (where the environment of the construction needs to be stated at the second level) and the second two rules stand in a counterfeeding or counterbleeding relation (where the environment of the second rule (already at the second level) has to be stated at an earlier level than the output of the second) requires three levels in a constructionist approach. It is also clear that a strict order of three rules where the first two stand in a counterfeeding .or counterbleeding relation (where both constructions can have their environment described on the first level) and the second rule either feeds or bleeds the third (so that the same two levels again will do, since it is only needed that the second level output of the second rule be on the same level as the environment of the third rule) two levels exhaust the set of necessary strata.

Another argument for three levels might be the psychological plausibility of representations. All three levels of cognitive phonology can claim the status of being real, as we explained in the paragraph introducing levels. The characteristics of the levels are independent of each other: all three have different phonotactics and their representations are constructed from different sets of components. A grammar then has the task to assign the adequate representations to each level. If there are three levels of representation, phonology should reflect this property of human cognition.

John Goldsmith's harmonic phonology underlines another important feature of the theory: the focus on phonotactics, stated in the forms of *intralevel* rules or constraints. One of the theory's metaconstraints is that these constraint (of which there are three possible types, corresponding to the three levels) can only operate and they operate as long as they make the representation more *harmonic*, i. e. closer to satisfying the phonotactics of that level (the significance of continuous application lies in describing phenomena that were describable only by iterative rules and cyclic application before). This metaconstraint does not hold for interlevel constructions, though – correspondences between levels do not necessarily improve the phonotactics of levels. The only limitation concerning such constructions is that they cannot lead to derivations with intermediate levels.

2.4 Examples

Let us look at some three-level reanalyses of well known phonological phenomena. All three accounts were chosen to show that cognitive phonology as a highly restricted theory is able to handle processes that were counted among the most problematic cases in generative phonology. To show that cognitive phonology has advantages besides the theoretical considerations as well and to demonstrate that constructions are more than just a rewrite-formalism of previous accounts, we have to provide spectacular examples and refer to linguistic phenomena that required extrinsic rule-ordering in the generative framework. All three examples are taken from Lakoff (1995).

2.4.1 Canadian dialect variation

This simple example is a good start to show what three-level phonology can do. It also illustrates how the generalisations about the systematic correspondence between specific kinds of orders and their constructionist account work. Classically, this example is invoked to show that in some cases minimal variation between dialects which otherwise apparently have the same set of rules can only be explained by the different orders imposed upon their rules. The difference between two Canadian dialects is reflected in the way their speakers pronounce the words writing vs. riding and clouted vs. clouded. Dialect A has the following surface pairs: r[ayD]ing[4](*riding*) vs. r[ʌyD]ing (*writing*), and cl[awD]ed (*clouded*) vs. cl[ʌwD]ed (*clouted*). In Dialect B the differences are neutralised: both *writing* and *riding* are pronounced r[ayD]ing, both *clouted* and *clouded* are pronounced cl[awD]ed. The two rules offered by generative theory are given as (3) and (4):

(3) /ay/, /aw/ → /ʌy/, /ʌw/ / __ voiceless consonant — Vowel Raising

(4) [−cont, +cor] → [+voice] / [−cons, +stress] __ [−cons, -stress] — Flapping Rule

The differences in surface forms are differences of ordering. In Dialect A Vowel raising applies before Flapping (resulting in a counterbleeding order), while in Dialect B the two rules apply in the reverse order, and thus Flapping bleeds Vowel Raising (*na*=not applicable):

(5)

Dialect A				
write	ride	writing	riding	
rʌyt	rayd	wr[ʌyt]ing	r[ayd]ing	Vowel Raising
na	*na*	wr[ʌyD]ing	r[ayD]ing	Flapping
Dialect B				
write	ride	writing	riding	
rayt	rayd	r[ayD]ing	r[ayD]ing	Flapping
rʌyt	*na*	*na*	*na*	Vowel Raising

This seems to be convincing evidence of a grammar's need for rule ordering. Yet, interlevel constructions of three-level phonology are able to describe the dialectal differences without having to reconcile to the arbitrary and computationally expensive device of rule ordering, and actually making reference to only two levels of the three (the word- and the phonetic levels):

(6) *The Flap construction (common to both dialects)*

W: [−cont, +cor]
|
P: [−cons, +stress] [+voice] [−cons, −stress]

The two dialects have then different Raising constructions, the difference being on the level at which the environment of raising is stated:

(7) *The raising construction – Dialect A*

W: [−cons, +stress] [−voice]
|
P: -low

[4]The small capital D denotes a flap.

(8) *The raising construction – dialect B*

W: [−cons, +stress]

P: [−low] [−voice]

In Dialect A, the voicing distinction at the W level where the environment of raising is stated, matches (or unifies with) the raising construction in the case of *writing* and *clouted*, but not in the case of *riding* and *clouded*, so a surface distinction results. In Dialect B, where the environment of vowel raising is given at the P level, neither of the words writing, riding, clouted, clouded unifies with the construction, since all matches the flap construction and has a voiced consonant at the P level, yielding homophonous word pairs.

In Dialect A, Flapping counterbleeds Vowel Raising, so the environment of the raising construction is crucially on the first level. Dialect B's Flapping rule bleeds Vowel Raising, thus the environment of Vowel raising is at the second level. Notice also that although only two levels are mentioned, these are the Word- and Phonetic levels, the reason being that the Flapping construction makes reference to a segment (the flap D) which is not a phoneme of either dialect, but a result of a phonetic neutralisation mechanism.

2.4.2 Icelandic

Armed with the basic devices and mechanisms of the theory, let us look at a more complicated case of rule interaction. Icelandic is a language often cited by Lexical Phonology to prove the need for strict ordering on the one hand and cyclic rule application on the other, making the distinction between lexical and postlexical rules a fundamental component of grammar. Kiparsky also introduces two general stipulative principles of Lexical Phonology, to which the phonology of Icelandic also makes reference: the *Strict Cyclicity Condition* and the *Strong Domain Hypothesis*. Three-level phonology offers a description without rule ordering and cycles; what previously had been lexical rules are now delegated to the M-W either intra- or interlevel correspondences, while postlexical rules reside somewhere on the W-P part in the form of constructions. This example also shows what the previous example did not make clear: that in some cases three levels are necessary to give the proper account

Kiparsky (1984) lists 6 rules, of which the relevant ones are listed below in the order imposed upon them (we will not verify every part of the ordering, the interested reader is referred to Kiparsky (1984)). As Kiparsky states "unless specifically indicated they apply both lexically and postlexically, where permitted by the constraints of the theory" (p. 150).

(9) *Syllabification (lexical and postlexical)*

(10) [+syll, −stress, +lax] → ∅ / __ [−syll, +cor, +lax][5] Syncope (lexical)

(11) a → ö / __ C_0u u-Umlaut (lexical)

(12) ∅ → u / __ r (unsyllabified) u-Epenthesis (lexical and postlexical)

U-Umlaut and Syncope display a strange phenomenon. Let us look at the two forms *böggli* and *bögglu*, derived from the underlying forms *bagg+ul+i* and *bagg+il+u*. The correct surface forms only result if the both rules apply to both underlying forms, though in different orders:

(13) a. bagg+ul+i
bögguli u-Umlaut
böggli Syncope

b. bagg+il+u
bagglu Syncope
bögglu u-Umlaut

Lexical Phonology's way out is to posit cyclic rule application (and specific principles) in the Lexical part of phonology. As can be seen looking at the morphology of the words, both require two cycles of lexical rule application. The list and order of Icelandic rules shows that syncope precedes u-Umlaut (it cannot apply across the board, as forms like *dag+r* →*dagur* show: an epenthetic

[5] l, r, n, d (th), s – a lax dental in onset position

u never triggers u-Umlaut, so u-Umlaut clearly has to be ordered before u-epenthesis.) In the derivational history of *böggli*, the first cycle only makes the process of u-Umlaut available, since the word-final *l*, not being in onset position, does not match the environment of the syncope rule. In the second cycle, prompted by the attachment of the *i* suffix, resyllabification relinks the *l* to the onset of the final syllable, creating an environment for syncope, by which now the *u* is deleted. *Bögglu* has a different derivation, with both syncope and u-Umlaut unavailable on the first cycle, and both applying in the same, second cycle, in the canonical order, syncope applying first and feeding u-Umlaut.

Let us summarise how lexical phonology would depict the representations and changes showing the differences between the two words:

(14)

			Lexical Rules
			Cycle1
bagg+ul	bagg+il	dag+r	Morphology
baggul	baggil	dag(r)	Syllabification
–	–	–	Syncope
böggul	–	–	u-Umlaut
–	–	dag(ur)	u-epenthesis
			Cycle 2
böggul+i	baggil+u	–	Morphology
bögguli	baggilu	–	Syllabification
böggli	bagglu	–	Syncope
–	bögglu	–	u-Umlaut
–	–	–	u-epenthesis
			Postlexical Rules
–	–	dagur	syllabification
–	–	–	u-epenthesis

Lakoff suggests a much simpler solution in the framework of cognitive phonology. Four constructions are needed. A one level well-formedness condition on syllables holding at both the W and P levels states that a syllable consists of an onset cluster, a vowel and an optional coda cluster. Since in Icelandic Cr and Cj clusters are not legal coda clusters, this construction will leave the r and j of any CrC, Cr#, CjC and Cj# clusters unsyllabified.

(15) *The syllable construction*
W, P: $[C_1\ V\ (C_1)]_\sigma$

(16) *The construction for syncope*

$$\begin{array}{lll} \text{M:} & \text{V} & \text{C[}-\text{syll, +cor, +lax] + V} \\ & | & \\ \text{W:} & \emptyset & \end{array}$$

The symbol + appearing outside a feature matrix designates morpheme boundary, thus the final vowel at M-level in this construction is suffixal.

(17) *The construction for u-Umlaut*

$$\begin{array}{llll} \text{M:} & \text{a} & & \\ & | & C_0 & \text{u} \\ \text{W:} & \text{ö} & & \end{array}$$

The statement of the environment of the construction reads: __ C_0u either at the M-level or the W-level

(18) *The construction for u-epenthesis*

$$\begin{array}{lll} \text{W:} & \emptyset & [\text{r}]_{-\sigma} \\ & | & \\ \text{P} & \text{u} & \end{array}$$

These four constructions yield the following representations for the three words discussed above:

(19)

M:	bagg+ul+i	bagg+il+u	dag+r
	|	|	|
W:	böggli	bögglu	dagr
			|
P:			dagur

The word *bagg+ul+i* matches the u-Umlaut construction: its environment condition is met at the M-level, it also unifies with the construction of syncope, which deletes the suffixal u, thus we have the surface form *böggli*; *bagg+il+u* also meets the conditions of the syncope correlation, the suffixal *i* is deleted. As a consequence, the environmental criteria of u-Umlaut are satisfied at the W-level. Finally, *dag+r* does not correspond to the pattern of either syncope or u-Umlaut at the M and W-levels. It matches the construction of u-epenthesis, stated at the W-P levels, yielding a potential condition for u-Umlaut, but a P-level *u* is not among the possible environments of umlaut. The need for three levels here is a consequence of Syncope feeding u-Umlaut, and u-Epenthesis and u-Umlaut standing in a counterbleeding relation.

Several devices of generative phonology have proved to be redundant. Three-level phonology has eliminated rule ordering and cyclic rule application; the need to distinguish lexical and postlexical modules in phonology has been replaced by correspondences between three levels.

2.4.3 Iterativity

A final example is to show an attractive natural consequence of three-level phonology: the account of iterative rules. Lakoff's examples are Slovak and Gidabal iterative shortening, which were described by Kaplan and Kay (1994) as derived by the same generative rule, which being left iterative in Slovak, shortens every long vowel except the first in the words of the language. The same rule operates iteratively from left-to-right in Gidabal, resulting in an alternating sequence of long and short vowels. The generative rule in question was the following:

(20) *Iterative shortening*

$[+\text{syll}, +\text{long}] \rightarrow [-\text{long}] \; / \; [+\text{syll}, +\text{long}]\ C_0$ __

The derivations then would go step by step, sequentially, in the case of a word which underlyingly has four long vowels, requiring two intermediate steps in Slovak and one in Gidabal.

(21)

Slovak

Vː	C	Vː	C	Vː	C	Vː
Vː	C	Vː	C	Vː	C	V
Vː	C	Vː	C	V	C	V
Vː	C	V	C	V	C	V

Gidabal

Vː	C	Vː	C	Vː	C	Vː
Vː	C	V	C	Vː	C	Vː
Vː	C	V	C	Vː	C	V

Two simple constructions without any stipulative iterativeness abolish all intermediate representations. The constructions for the shortening in the two languages are very similar, the only difference being the level at which the environment of the rule is stated. Notice that in Slovak, the rule counterbleeds, in Gidabal it bleeds itself.

(22) *Slovak*

M:	[+syll, +long]	C	[+syll, +long]
			|
W:			[−long]

Resulting in a representation for the above example

(23)

M:	Vː	C	Vː	C	Vː	C	Vː
W:	Vː	C	V	C	V	C	V

(24) *Gidabal*

M:			[+syll, +long]
			|
W:	[+syll, +long]	C	[−long]

And the proper representation is

(25) M: V: C V: C V: C V:
W: V: C V C V: C V

According to Lakoff one advantage of using constructions instead of iterative rules besides not having to stipulate their iterativeness (since not being ordered, constructions apply as long as they are matched) is that we do not have to declare the direction of their application, both can proceed from left to right, the representation matching the real time process. Attractive though this characteristic of the theory may seem, difficulties may arise. Direction-neutral constructions for iterative rules only work for rules in which the environment of the rule is left of the target. In any phenomenon where the environment is on the right side, and the iterative rule is feeding or bleeding (and thus has a second level environment), as in an iterative regressive assimilation process, where it is the last segment which determines the features of a sequence of segments, the construction has to be applied from right to left. Hungarian voice assimilation is a case in point.

2.4.4 Hungarian voicing assimilation

In Hungarian, voicing assimilation is regressive, and affects only obstruents. Adjacent obstruents always agree in voicing determined by the laryngeal specification of the last consonant. The process is left-iterative (see the example *lisztből* below)

(26) liszt+ből 'of flour'
/st/+/b/→/sdb/→/zdb/

The generative rule given in Siptár (1994, p. 204) is (the details and exceptions are beyond the scope of this paper):

(27) obstr → [α voiced] / __ obstr, [α voiced]

The following construction will give the correct results:

(28) M: obstr
|
W: [α voiced] obstr, [α voiced]

The iterative matching of the representation with the construction goes right-to-left. In this specific example the problem can be easily solved if the construction is not a two-level correspondence, but a one-level well-formedness condition, stating that a row of obstruents has to agree in voicing, and it is the last one which determines the value of this feature, the other obstruents get their voicing specification by autosegmental spreading.

2.5 Conclusions

All above analyses proved to have the advantage of eliminating the need for rule ordering. Furthermore, they are formulated in a restricted framework which apparently has greater psychological plausibility than any generative theory of phonology. Reanalyses of phenomena that do not need ordering of rules even in a generative analyses is fairly straightforward, and in these cases a cognitive phonology account might seem nothing more than a simple reformulation of rules in a new but isomorphic language, yet this reformulation, as we have shown in the introductory paragraphs, has serious theoretical and empirical consequences. Of course several questions arise: Do the levels of analysis have any content behind their name? If so, each analysis in this framework should proceed so that representations are only assigned to levels where they are legal, i. e. only segments that are part of the phonemic inventory of the language should appear on the word level, phonetic differences should be assigned to the P-level, and the M–W constructions should only stand for what previously had been called lexical rules, while the postlexical rules should only be reanalysed as W–P constructions. On the other hand, we have also seen that in some specific ordering relations, the pure number of levels is important: if we have three rules ordered strictly so that the first feeds or bleeds the second and the second counterfeeds or counterbleeds the third, three levels are inevitably needed, even if all the three rules belong to the lexical module. These are questions of theoretical kind and can only be resolved by empirical investigation.

References

Bird, Steven and T. Mark Ellison. 1994. One-level phonology: Autosegmental representations and rules as finite automata. *Computational Linguistics*, 20(1):55–90.

Chomsky, Noam and Morris Halle. 1968. *The Sound Pattern of English*. Harper and Row.

Goldsmith, John. 1995a. Harmonic phonology. In *The Last Phonological Rule* (Goldsmith, 1995b), pages 21–61.

Goldsmith, John, editor. 1995b. *The Last Phonological Rule*. The University of Chicago Press, Chicago and London.

Kaplan, Ronald M. and Martin Kay. 1994. Regular models of phonological rule systems. *Computational Linguistics*, 20(1):331–378.

Karttunen, Lauri. 1995. Finite state constraints. In Goldsmith (Goldsmith, 1995b), pages 173–195.

Kenstowicz, Michael. 1994. *Phonology in Generative Grammar*. Blackwell, Oxford.

Kiparsky, Paul. 1984. On the lexical phonology of icelandic. In *Nordic Prosody*, volume 3. Almgist & Wiksell, Stockholm, pages 135–165.

Lakoff, George. 1995. Cognitive phonology. In Goldsmith (Goldsmith, 1995b), pages 117–145.

Scobbie, James. 1993. Constraint violation and conflict from the perspective of declarative phonology. *CJL/RCL*, 38(2):155–167.

Siptár, Péter. 1994. A mássalhangzók [consonants]. In Ferenc Kiefer, editor, *Strukturális magyar nyelvtan. 2. kötet: Fonológia. [A Structural Grammar of Hungarian. Volume 2: Phonology]*. Akadémiai Kiadó, Budapest, pages 183–272.

Chapter 3

HPSG Phonology

3.1 Introduction

In contrast to classical generative phonology, which is derivational and is based on the ordered operation of rules, much recent work has tended towards a new model that can be described in terms of constraints on well-formedness[1]. Though this work has an increasingly declarative flavour, most versions retain procedural devices for repairing representations that fail to meet certain constraints or allow constraints to override each other. This is in contrast with the interpretation of constraints in grammar formalisms like LFG, GPSG and HPSG. In these approaches all constraints must be observed (i.e. no ill-formed intermediate representations can be created). There is a recent school of phonology called Declarative Phonology (DP) that adopts this interpretation of constraints. HPSG has a grammar formalism and an implemented constraint unification mechanism that seems to be suitable for the implementation of declarative phonological models.

On the other hand, HPSG, as it is defined in Pollard and Sag (1987) and Pollard and Sag (1994), deals only with syntax and semantics. The phonological content of words is limited to orthographic strings, supplemented with the concatenation operation. Doing phonology in HPSG is thus motivated by the facts that HPSG lacks phonology and that it seems to be a suitable platform to implement DP descriptions.

How could phonological representations in HPSG be enriched so that a more elaborate treatment of phonology can be accommodated? Many of the current phonological accounts moved closer to a constraint-based perspective, but there are several problems with these accounts. First, they are often too informal, incoherent, inconsistent or indeterminate. Second, when a clear theoretical statement is found, it is usually expressed in procedural terms. Finally, even when explicit and non-procedural generalisations are found, they are stated in a non-linear model. The objective is thus to adopt a formal, non-procedural, non-linear model of phonology and integrate it into HPSG.

The distinctive features of such a model are as follows:

- Phonological representations describe a class of utterances (i.e. linguistic types).
- Lexical representations are partial, and phonological constraints are cast as generalisations in a lexical or prosodic inheritance hierarchy.
- Derivation consists in the gradual refinement of descriptions further constraining the denoted objects (lexical representations are unified with all relevant phonological constraints during the composition of prosodic structure). This is in contrast with the generative tradition, where derivation is not a process of refinement but one of alteration.

[1]This chapter is based on Bird and Klein (1994) and Walther (1993).

3.2 The theoretical framework of HPSG

HPSG is a constraint-based grammar formalism that employs typed feature structures. Every object that can be the value of an attribute has a type[2]. The type system is defined using two kinds of type declarations:

- *a subsumption (type-subtype) ordering* over types that defines a type hierarchy (termed *partitions* in the Appendix of Pollard and Sag (1994))

- *an appropriateness condition* (termed *feature declaration* in Pollard and Sag (1994)): what attributes must be defined for an object of a certain type, and of what type the values of the attributes must be

An example of subsumption ordering is in (1).

(1) $sign \Longrightarrow morph \vee stem \vee word \vee phrase$

As another example, (2) could be the appropriateness condition for objects of type *sign*[3].

(2) $\left[\begin{array}{ll}\textsc{phon} & phon\\ \textsc{synsem} & synsem\end{array}\right]_{sign}$

The feature structure in (2) expresses that objects of type *sign* have a feature PHON (phonology), the value of which is of type *phon*, and a feature SYNSEM (syntax and semantics), of type *synsem*[4]

Appropriateness conditions are inherited by subtypes. For example, since *morph* is a subtype of *sign* according to (1), it inherits all the constraints obeyed by *sign*. Later, we will see more examples of how types are declared and constrained.

3.3 Two Varieties of String-Based Phonology

As we declared in the Introduction, the objective is to adopt a formal, non-procedural, non-linear model of phonology to HPSG. Such a model is Autosegmental Phonology (Goldsmith (1990)). There are at least two ways autosegmental phonological accounts can be formulated in a declarative way. One approach is based on finite-state automata (FSAs) and the other is based on the HPSG list notation.

3.3.1 Finite-State Phonology

Bird and Ellison (1994) describe an algorithm for converting autosegmental representations to FSAs. For a description of their algorithm see Section 1.5 in Chapter 1. Attempts had been made to apply finite-state transducers (FSTs) to nonlinear models (Kay (1987); Kornai (1991)). FSAs have the advantage over FSTs that they form a Boolean lattice under intersection, union and complement i.e. the intersection, union and complement of any two FSAs is an FSA, which is not the case with any two FSTs. Closedness under intersection is crucial, because the intersection operation on FSAs corresponds to the unification of feature structures, which is the only structure building operation in formalisms using feature structures, such as HPSG. Intersection is used in the algorithm of Bird and Ellison to represent links between autosegmental tiers.

An unquestionable benefit of using the FSA encoding and the algorithm of Bird and Ellison is that *any* autosegmental representation and rule-system can be mechanically translated into a declarative (and thus potentially HPSG-compatible) represenation. This is not the case when the

[2]In the terminology of HPSG the term *sort* is used instead of *type*, which is used in the sense that is conventional in linguistics (i.e. to denote linguistic types, as opposed to tokens). In this chapter, however we are going to use the word *type* instead of *sort*.

[3]We adopt the usual convention of writing feature names in (small) CAPITALS and types in *italics*. Structure sharing is denoted by numbers in boxes, as usual in HPSG.

[4]In Pollard and Sag (1994), objects of type (or sort) *sign* also have features like QSTORE and RETRIEVED, which are used to account for quantifier scopes.

list notation is used. Bird and Klein (1990) also claim that using the FSA encoding may also have a computational benefit over other notations (such as using HPSG lists). This claim is rather questionable, since the intersection of FSAs is a very time-consuming complex operation, that also involves the minimization of the resulting FSA (see Chapter 1 or Bird and Ellison (1994)). There is no finite-state transducer that could form the intersection of two regular expressions. The time-complexity of a system featuring Bird and Ellison's One-Level Phonology is thus much higher than that of a real finite-state system, since the on-line application of the intersection operation cannot be avoided because the unification of any two feature structures containing a *phon* substructure (i.e. the unification of any two *signs*) involves taking the intersection of the two FSAs that represent those *phon* substructures. Moreover, if there are any computational benefits of FS encoding, in order to exploit them, the interpretation of regular expressions must be delegated to a specialized engine instead of using HPSG's general constraint solver.

3.3.2 List Notations

Using finie state phonology is thus a possible but not necessaryly feasible extension to HPSG phonology. A more conservative technique is to use HPSG's list notation to represent phonological structure. This approach, however, has the disadvantage that no mechanical procedure is defined to turn autosegmental analyses into a description with list notation. Its applicability seems to be limited to the description of morphological operations (including those in templatic languages) and prosodic structure building operations. In contrast, it is not at all obvious, how an autosegmental analysis featuring e.g. spreading could be represented using the list notation.

In order to be able to use lists to represent phonological structure, the type system must allow parametrised types of the form *list(α)* where α is an atomic type (3).

(3) a. $list(\alpha) \Longrightarrow e\text{-}list(\alpha) \vee ne\text{-}list(\alpha)$

b. $$\left[\begin{array}{ll}\text{FIRST} & \alpha \\ \text{REST} & list(\alpha)\end{array}\right]_{ne\text{-}list(\alpha)}$$

e-list stands for empty list and is an atomic type, while *ne-lists* (non-empty lists) are represented by an embedded structure of lists as in (3-b). The elements of a list are the values of the FIRST attributes of the embedded structures[5]. We can now treat the usual Kleene plus and Kleene star notations (α^+ and α^*) for non-empty and possibly empty sequences of objects of type α as abbreviations for *list*(α) and *ne-list*(α), respectively.

Another useful notation is parenthesised (optional) elements in a list. We could represent $\langle a(b)\rangle \frown L$ (i.e. the concatenation ($\frown$) of the list containing a and an optional b with a list L) as

(4) $$\left[\begin{array}{ll}\text{FIRST} & a \\ \text{REST} & \left[\begin{array}{ll}\text{FIRST} & b \\ \text{REST} & \boxed{1}list\end{array}\right]_{list} \vee \boxed{1}\end{array}\right]_{list}$$

Finally, it is useful to define another notation that recursively assigns a type τ to each position (sublist) in a list. The definition is recursive, as given in (5).

(5) $$map(\tau) \equiv \left[\text{REST} \quad map(\tau)\right]_{\tau} \vee \langle\rangle$$

If a list $\langle abc\rangle$ has the type *map(seg-list)*, it will appear as follows:

---

[5] There is an abbreviated notation for lists in HPSG, in which the elements are listed in a pair of angle brackets(e.g. $\langle xyz\rangle$). $\langle\rangle$ is a shorthand for *e-list*.

(6)
$$\left[\begin{array}{ll} \text{FIRST} & a \\ \text{REST} & \left[\begin{array}{ll} \text{FIRST} & b \\ \text{REST} & \left[\begin{array}{ll} \text{FIRST} & c \\ \text{REST} & \langle\rangle \end{array}\right]_{\textit{seg-list}} \end{array}\right]_{\textit{seg-list}} \end{array}\right]_{\textit{seg-list}}$$

3.3.3 A prosodic type hierarchy

As we have seen, the type system must be declared by defining a subsumption (type-subtype) ordering over types and appropriateness conditions for them. The types involved in phonology thus form a prosodic type hierarchy. The type *phon* can, for example, be defined to have the following immediate subtypes:

(7) $\textit{phon} \Longrightarrow \textit{utterance} \vee \textit{phrase} \vee \textit{pword} \vee \textit{foot} \vee \textit{syl} \vee \textit{segment}$

Each of these types may have further structure. The appropriateness condition for type *segment* could be defined as in (8) assuming a skeletal structure following e.g. Clements (1985, 248). Note that the skeletal structure of segments can be straightforwardly expressed using the feature structure representation.

(8)
$$\left[\begin{array}{ll} \text{LARYNGEAL} & \left[\begin{array}{ll} \text{SPREAD} & \textit{boolean} \\ \text{CONSTRICTED} & \textit{boolean} \\ \text{VOICED} & \textit{boolean} \end{array}\right] \\ \text{SUPRALARYNGEAL} & \left[\begin{array}{ll} \text{MANNER} & \left[\begin{array}{ll} \text{NASAL} & \textit{boolean} \\ \text{CONTINUANT} & \textit{boolean} \\ \text{STRIDENT} & \textit{boolean} \end{array}\right] \\ \text{PLACE} & \left[\begin{array}{ll} \text{CORONAL} & \textit{boolean} \\ \text{ANTERIOR} & \textit{boolean} \\ \text{DISTRIBUTED} & \textit{boolean} \end{array}\right] \end{array}\right] \end{array}\right]_{\textit{segment}}$$

Using this definition for the type *segment*, a constraint like English homorganic nasal assimilation (*hna*) can be straightforwardly defined. This phenomenon does not occur across phrase boundaries and so the constraint will be part of the appropriateness conditions of type *phrase*. Let us assume that *phrase* is a list of *segments*. Using the list notation, SC as an abbreviation for SUPRALARYNGEAL and CONT for CONTINUANT, *hna* can be defined as a negative filter like in (9).

(9)
$$\textit{hna} \equiv \neg \left\langle \cdots \left[\text{SL} \left[\begin{array}{ll} \text{MANNER}\,|\,\text{NASAL} & + \\ \text{PLACE} & \boxed{1} \end{array}\right]\right]_{\textit{segment}} \left[\text{SL} \left[\begin{array}{ll} \text{MANNER}\,|\,\text{CONT} & - \\ \text{PLACE} & \neg\boxed{1} \end{array}\right]\right]_{\textit{segment}} \cdots \right\rangle_{\textit{phrase}}$$

This constraint can also be expressed using the *map* type introduced in (5) and HPSG's FIRST / REST encoding for lists. This makes the constraint look less suspicious.

(10) a. $\textit{hna} \equiv \textit{map}(\textit{hom-nas})$

b.
$$\textit{hom-nas} \equiv \neg \left[\begin{array}{ll} \text{FIRST}\,|\,\text{SL} & \left[\begin{array}{ll} \text{MANNER}\,|\,\text{NASAL} & + \\ \text{PLACE} & \boxed{1} \end{array}\right] \\ \text{REST}\,|\,\text{FIRST}\,|\,\text{SL} & \left[\begin{array}{ll} \text{MANNER}\,|\,\text{CONT} & - \\ \text{PLACE} & \neg\boxed{1} \end{array}\right] \end{array}\right]$$

Standard techniques can now be used to move the negation in (10) inwards.[6]

[6]The following (De Morgan and generalized De Morgan) equivalences can be employed to normalize expressions with disjunction and negation:

3.3.4 Prosodic Constituency

Prosodic constituency could be represented like phrase structure using a non-empty DAUGHTERS list. The DAUGHTERS of a *phrase* would be of type *foot*, while a *foot* would have a non-empty list of DAUGHTERS of type *syl* (syllable). However, there appears to be no linguistic motivation for building such a structure.

An alternative is to define a *phrase* to contain a list of feet, a list of syllables and a list of segments at the same time, properly co-indexed (in HPSG, the term *structure sharing* is used for this), and each list must obey a number of constraints, stating e.g. that the periphery of phrases is exempt from certain sandhi (e.g. assimilation) phenomena, that feet have no more than three syllables, and only certain combinations of heavy and light syllables are permissible etc. The latter two constraints constrain objects of type *foot* while the lack of assimilation between segments in different phrases is the result of making the segments of a *phrase* the domain of assimilation constraints.

3.4 Morphology and the Lexicon

3.4.1 Linguistic Hierarchy

In HPSG, the subsumption ordering over types underlies the organisation of the lexicon as well. The nodes of the inheritance network thus defined contain generalisations about lexical types.

Along the phonological dimension of signs, lexical entries will have to observe any morpheme or word level constraints that apply to the language in question. When words combine as syntactic phrases, they will also have to satisfy all constraints on well-formed phonological phrases.

In some languages there may be a special interaction between the prosodic and the lexical hierarchy. The tongue root harmony in Yoruba is for example restricted to the lexical class of nouns. If *atr* is the type expressing the constraint on harmonic utterances, we can limit this constraint to nouns by defining the following constraint on nouns:

(11)
$$\begin{bmatrix} \text{PHON} & phon \wedge atr \\ \text{SYNSEM}\,|\,\text{LOCAL}\,|\,\text{CAT} & \begin{bmatrix} \text{HEAD} & noun \\ \text{LEX} & + \end{bmatrix} \end{bmatrix}_{noun}$$

3.4.2 Morphological Complexity

There are at least two alternative ways of representing morphologically complex (e.g. affixed) word forms. One alternative is handling morphological complexity in an analogous manner to syntactic complexity, where heads (e.g. affixes in our case) subcategorise for arguments, and morphemes combine in a Word-Grammar Scheme. Such an approach would analyse the English 3rd person singular prefix *-s* in the manner shown in (12).

(12)
$$\begin{bmatrix} \text{PHON} & \text{s} \\ \text{SYNSEM}\,|\,\text{LOCAL}\,|\,\text{CAT} & \begin{bmatrix} \text{HEAD}\,|\,\text{VFORM} & fin \\ \text{SUBJ} & \text{NP}[nom]_{[3rd,sing]} \\ \text{SUBCAT} & verb\text{-}stem \end{bmatrix} \end{bmatrix}_{affix}$$

By applying appropriately modified versions of the Head Feature Principle, Subcategorisation Principle and linear order statements, this would yield a tree-structured sign when combining with a verb stem, such as *walk*.

$$\neg\begin{bmatrix} \text{A} & \phi \\ \text{B} & \psi \end{bmatrix} \equiv \neg\begin{bmatrix} \text{A} & \phi \end{bmatrix} \vee \neg\begin{bmatrix} \text{B} & \psi \end{bmatrix}$$

$$\neg\begin{bmatrix} \text{A} & \phi \end{bmatrix} \equiv \begin{bmatrix} \neg\text{A} & \top \end{bmatrix} \vee \begin{bmatrix} \text{A} & \neg\phi \end{bmatrix}$$

$\neg\text{A}\ \top$ means that the feature A is not appropriate for this feature structure.
This technique eventually yields the disjunctive normal form of expressions.

(13)
$$\underset{verb}{\begin{bmatrix} \text{PHON} & \boxed{1} \frown \boxed{2} \\ \text{DTRS} & \left\langle \underset{verb\text{-}stem}{\begin{bmatrix}\text{PHON} & \boxed{1}\langle \text{wɔːk}\rangle\end{bmatrix}}, \underset{affix}{\begin{bmatrix}\text{PHON} & \boxed{2}\langle \text{s}\rangle\end{bmatrix}} \right\rangle \end{bmatrix}}$$

Another alternative is representing affixes as *partially instantiated word forms*. This can be done by adding a new feature MORPH to the definition of *sign* with a value *morph*. The type *morph* has two subtypes, *affix-morph* and *basic-morph*. The latter is the type of unaffixed stems, the former is that of affixed ones.

(14)
a. *morph* $\Longrightarrow$ *affix-morph* $\vee$ *basic-morph*

b. $\underset{morph}{\begin{bmatrix}\text{STEM} & \textit{stem}\end{bmatrix}}$

c. $\underset{affix\text{-}morph}{\begin{bmatrix}\text{AFFIX} & \textit{affix}\end{bmatrix}}$

d. *affix* $\Longrightarrow$ *prefix* $\vee$ *suffix*

The suffix -s can thus be defined as given in (15).

(15)
$$\underset{3ps}{\begin{bmatrix} \text{PHON} & \boxed{1} \frown \boxed{2} \\ \text{MORPH} & \underset{affix\text{-}morph}{\begin{bmatrix} \text{STEM} & \underset{verb\text{-}stem}{\begin{bmatrix}\text{PHON} & \boxed{1}\end{bmatrix}} \\ \text{AFFIX} & \underset{suffix}{\begin{bmatrix}\text{PHON} & \boxed{2}\langle \text{s}\rangle\end{bmatrix}} \end{bmatrix}} \end{bmatrix}}$$

This description of the suffix -*s* would have to be enriched to allow for the allomorphic alternation -s~-z~-ɪz. The first pair of allomorphs can be handled by underspecifying the suffix for voicing, and defining a voicing assimilation constraint similar to the homorganic nasal assimilation in (10). Moreover, the thus remaining two alternants would have to be properly constrained regarding (complementary) environments in which they can occur.

3.5 Examples

In this section, two examples will be given to illustrate the application of the HPSG list notation to give a declarative description of phonological processes. The first example features nonconcatenative morphology, which was a major empirical motivation for autosegmental phonology. The second example is a declarative analysis of French schwa, which exhibits an alternation that is both prosodically and lexically conditioned.

3.5.1 Sierra Miwok Templatic Morphology

Sierra Miwok has a nonconcatenative templatic morphology that features an intercalation of vowels with consonantal verb roots. This intercalation is similar to the one classical Arabic has, but the data is simpler.

Descriptive Overview

In Sierra Miwok, there are three types of verb stem that differ in the syllable structure of the basic form, which is the form used for present tense. As shown in (16), each type also has three other forms, depending on the morphological and syntactical context (Goldsmith (1990)).

(16)

Gloss	Basic stem	Second stem	Third stem	Fourth stem
Type I				
bleed	kicaaw	kicaww	kiccaw	kicwa
jump	tuyaaŋ	tuyaŋŋ	tuyyaŋ	tuyŋa
take	patɨɨt	patɨtt	pattɨt	pattɨ
roll	huteel	hutell	huttel	hutle
Type II				
quit	celku	celukk	celluk	celku
go home	woʔlu	woʔull	woʔʔul	woʔlu
catch up with	nakpa	nakapp	nakkap	nakpa
spear	wimki	wimikk	wimmik	wimki
Type III				
bury	hamme	hameʔʔ	hammeʔ	hamʔe
dive	ʔuppi	ʔupiʔʔ	ʔuppiʔ	ʔupʔi
speak	liwwa	liwaʔʔ	liwwaʔ	liwʔa
sing	mɨlli	mɨliʔʔ	mɨlliʔ	mɨlʔi

Note that the stem type distinction is only relevant to the Basic stem forms. Note also that the 2nd to 4th stems of Type III verbs feature a default glottal stop as the third consonant of the stem.

Analysis

The association between a segment on the consonantal or the vowel melody tier and a slot on the timing tier (skeleton) can be straightforwardly represented by structure sharing. The PHON attribute of the basic form of the verb *kicaaw* can be defined as given in (17).

(17)
$$_{phon}\begin{bmatrix} \text{CON} & \langle \boxed{1}k\ \boxed{3}c\ \boxed{5}w \rangle \\ \text{VOW} & \langle \boxed{2}i\ \boxed{4}a \rangle \\ \text{SKEL} & \langle \boxed{1}\boxed{2}\boxed{3}\boxed{4}\boxed{4}\boxed{5} \rangle \end{bmatrix}$$

Coindexing adequately encodes association in this case, since it has a slot-filling (rather than the more general temporal) interpretation.[7]

The basic templates themselves can be defined as given in (18).

(18) a.
$$_{basic\text{-}I}\begin{bmatrix} \text{CON} & \langle \boxed{1}\boxed{3}\boxed{5} \rangle \\ \text{VOW} & \langle \boxed{2}\boxed{4} \rangle \\ \text{SKEL} & \langle \boxed{1}\boxed{2}\boxed{3}\boxed{4}\boxed{4}\boxed{5} \rangle \end{bmatrix}$$

b.
$$_{basic\text{-}II}\begin{bmatrix} \text{CON} & \langle \boxed{1}\boxed{3}\boxed{5} \rangle \\ \text{VOW} & \langle \boxed{2}\boxed{4} \rangle \\ \text{SKEL} & \langle \boxed{1}\boxed{2}\boxed{3}\boxed{4}\boxed{5} \rangle \end{bmatrix}$$

c.
$$_{basic\text{-}III}\begin{bmatrix} \text{CON} & \langle \boxed{1}\boxed{3} \rangle \\ \text{VOW} & \langle \boxed{2}\boxed{4} \rangle \\ \text{SKEL} & \langle \boxed{1}\boxed{2}\boxed{3}\boxed{3}\boxed{4} \rangle \end{bmatrix}$$

The PHON attribute of the lexical entries of verbs do not have to contain these templates. Lexical entries contain only unpredictable information like the segmental contents of the consonantal and vowel tiers. The value of the PHON attribute of the verb *kicaaw*, for example, could be defined as (19):

[7] An autosegmental analysis featuring e.g. autosegment spreading would not be so straightforward to represent using the list notation.

(19) $${}_{phon}\left[\begin{array}{ll} \text{CON} & \langle \text{k c w} \rangle \\ \text{VOW} & \langle \text{i a} \rangle \end{array}\right]$$

Lexical entries also have to be inserted into the type hierarchy; e.g. the lexical entry of any verb would be the subtype of one of three verbal root types. We also have to specify how each verbal root type determines the appropriate basic stem form. To do this, we first define *basic* as a subtype of *stem*.

(20) $\textit{stem} \Longrightarrow \textit{affixed} \vee \textit{basic}$

Additionally, *stems* are defined to inherit their SYNSEM value from their ROOT.

(21) $${}_{stem}\left[\begin{array}{ll} \text{SYNSEM} & \boxed{1} \\ \text{MORPH} \mid \text{ROOT} \mid \text{SYNSEM} & \boxed{1} \end{array}\right]$$

(20) ensures that *basic* inherits this constraint from *stem*. Now we define the three verbal root types to be a subtype of *v-root*, which is the type of every verbal root, and to have a phonology that contains the template of the basic stem of the appropriate verbal root type.

(22) a. $\textit{v-root} \Longrightarrow \textit{v-root-I} \vee \textit{v-root-II} \vee \textit{v-root-III}$

b. ${}_{\textit{v-root-I}}\left[\text{PHON} \quad \textit{basic-I}\right]$

c. ${}_{\textit{v-root-II}}\left[\text{PHON} \quad \textit{basic-II}\right]$

d. ${}_{\textit{v-root-III}}\left[\text{PHON} \quad \textit{basic-III}\right]$

We then impose the following constraint on *basic*:

(23) $${}_{basic}\left[\begin{array}{ll} \text{PHON} & \boxed{1} \\ \text{MORPH} \mid \text{ROOT} & {}_{\textit{v-root}}\left[\text{PHON} \quad \boxed{1}\right] \end{array}\right]$$

This ensures that the basic stems assume the form determined by the template that belongs to their root type. The second, third and fourth stem types, on the other hand, themselves define the template that determines the skeletal structure of the 2nd to 4th stem of any verb without regard to the verb's own lexical specification. These issues are discussed in Klein (1993).

3.5.2 French Schwa

This section gives a declarative analysis of the French schwa–zero alternation, which is restricted to lexically specified words and exhibits free variation in prosodic contexts where both alternants result in well-formed syllabic structure, while only one variant can appear in contexts where the other variant would result in ill-formed prosodic structure.

Descriptive Overview

French schwa (unlike English schwa) is a full vowel, which is usually realised as the low-mid front rounded vowel /œ/ (and sometimes as the high-mid front rounded vowel /ø/ in certain predictable environments). Its distinctive characteristics is that in certain environments it *may* fail to be realised phonetically.[8] The alternation is manifested in forms like (24), where dots indicate syllable boundaries.

(24) a. six melons [si.mœ.lɔ̃] / [sim.lɔ̃]
b. sept melons [sɛt.mœ.lɔ̃], *[sɛtmlɔ̃]

[8]We shall not be concerned with another œ∼ ∅ alternation known as elision, which is a phonologically conditioned allomorphy involving alternations such as *le/l'*. Elision, in contrast to the schwa alternation, does not allow free variation in any case.

Note that *sept melons* requires the schwa to break up the *tml* cluster that would be unsyllabifiable. But prosodic constraints do not alone determine the alternation. Instead, they interplay with lexical constraints.

The following data indicates that schwa alternation cannot be treated as a general epenthesis process as (24) would suggest.

(25)

Cluster	Schwa Possible/Obligatory		Schwa Impossible	
rdr	bordereau	[bɔr.dœ.ro]	perdrix	[pɛr.dri]
rʃ	derechef	[dœ.rœ.ʃɛf]	torchon	[tɔr.ʃɔ̃]
skl	squelette	[skœ.lɛt]	sclérose	[skle.roz]
ps	dépecer	[de.pœ.se]	éclipser	[ek.lip.se]

Consequently, we shall assume that schwa must be encoded in lexical representations. The data in (25) may suggest that there is a lexical schwa whenever there is an orthographic *e*. However, the data in (26) indicates that there is no such correspondence.

(26)

Orthography	With Schwa	Without Schwa
bordereau	[bɔr.dœ.ro]	—
fais-le	[fɛ.lø]	—
six melons	[si.mœ.lɔ̃]	[sim.lɔ̃]
pelleterie	—	[pɛl.tri]

Analysis

The declarative analysis we shall give is based on the background assumptions that the alternating schwa is

i. prosodically conditioned,

ii. lexically conditioned,

iii. not in direct correspondence with orthographic *e*.

Since prosodic structure plays an important role in conditioning the alternation, we shall first define a type for syllable structure.

(27)
$$_{syl}\left[\begin{array}{ll}\text{ONS} & \textit{onset}\\ \text{NUC} & \textit{nucleus}\\ \text{CODA} & \textit{coda}\end{array}\right]$$

Syllable structure will be represented by an independent tier encoded as a sequence of such syllables, where segmental constituents of the syllable are coindexed with a separate segmental tier, as defined in (28). The type *phrase* denotes phonological phrase in (28), since that is the domain of syllabification in French. Note that the indices in (28) [9] range over lists of segments that may be empty in the case of onsets and codas.

(28) a.
$$_{phrase}\left[\begin{array}{ll}\text{SYLS} & \left\langle {}_{syl}\left[\begin{array}{ll}\text{ONS} & \boxed{1}\\ \text{NUC} & \boxed{2}\\ \text{CODA} & \boxed{3}\end{array}\right]\right\rangle \frown \boxed{4}\\ \text{SEGS} & \boxed{1}\frown\boxed{2}\frown\boxed{3}\frown\boxed{5}\end{array}\right] \rightarrow {}_{phrase}\left[\begin{array}{ll}\text{SYLS} & \boxed{4}\\ \text{SEGS} & \boxed{5}\end{array}\right]$$

b.
$$_{phrase}\left[\begin{array}{ll}\text{SYLS} & \langle\rangle\\ \text{SEGS} & \langle\rangle\end{array}\right]$$

[9]The $\rightarrow$ in (28-a) stands for material implication. Note that since $\phi \rightarrow \psi$ is equivalrnt to $\neg\phi \vee \psi$, the implication this expression can be turned into a disjunction, which can be further normalized using the De Morgan technique given in footnote (10).

(28) states that the segments of any well-formed phrase must be parsed into syllables.

We have seen that the schwa–zero alternation is also lexically conditioned. We shall represent lexical entries that exhibit the alternating behaviour (with free variation under favourable prosodic conditions) using the optional list element notation defined in (4). The PHON attribute of the lexical entry of *melons* could thus be defined as $_{phon}\left[\text{SEGS} \quad \langle \text{m (œ) l ɔ̃} \rangle\right]$. In contrast to this, lexical entries of words featuring non-alternating schwa contain non-optional schwa segments in their representation.

The syllabification process is constrained by phonotactics, for which a preliminary definition can be given as follows:

(29) a. *onset* ⟹ *word-internal-onset* ∨ *word-initial-onset*
 b. *word-internal-onset* ⟹ ⟨(*cons*) (*glide*)⟩ ∨ ⟨*obstr liquid*⟩
 c. *word-initial-onset* ⟹ *word-internal-onset* ∨ ⟨*obstr sonorant*⟩ ∨ ⟨*s stop liquid*⟩ ∨ ⟨*p n*⟩
 d. *coda* ⟹ *word-internal-coda* ∨ *word-final-coda*
 e. *word-internal-coda* ⟹ ⟨(*cons*)⟩
 f. *word-final-coda* ⟹ *word-internal-coda* ∨ ⟨(*s*) *stop* (*liquid*)⟩

In addition, some more constraints must be given that express the Onset Maximization Principle. The constraint in (30) excludes syllabifications where an empty onset follows a non-empty coda, while (31) prohibits the syllabification of obstruents into the coda if there is no obstruent in the onset (*obstr liquid* sequences are thus syllabified into the onset).

(30) $$onset\text{-}max\text{-}1 \equiv \neg \; _{phrase}\left\langle \ldots \; _{syl}\left[\text{CODA} \quad \textit{ne-list}\right] \; _{nha\text{-}syl}\left[\text{ONS} \quad \textit{e-list}\right] \ldots \right\rangle$$

(31) $$onset\text{-}max\text{-}2 \equiv \neg \; _{phrase}\left\langle \ldots \; _{syl}\left[\text{CODA} \quad \langle \ldots obs \ldots \rangle\right] \; _{nha\text{-}syl}\left[\text{ONS} \quad \neg\langle \ldots obs \ldots \rangle\right] \ldots \right\rangle$$

The second syllable in the above constraints is defined to be of type *nha-syl* (i.e. *non-hache-aspiré-syllable*), because these constraints do not hold for the lexical class of so called hache aspiré words. The phrase *sept haches*, for example, is syllabified as [sɛt.aʃ] and not as *[sɛ.taʃ], and similarly, the only correct syllabification for *quatre haches* is [katr.aʃ], while *[kat.raʃ] and *[ka.traʃ] are ill-formed. The type *syl* thus has two subtypes, and syllables of subtype *ha-syl* (*hache-aspiré-syllable*) are defined to have and empty onset.

(32) a. *syl* ⟹ *ha-syl* ∨ *nha-syl*
 b. $_{ha\text{-}syl}\left[\text{ONSET} \quad \textit{e-list}\right]$

Now, hache aspiré words will be lexically specified as having an initial *ha-syl*[10]. The value of the PHON attribute of the word *haut* would be:

(33) $$_{phon}\begin{bmatrix} \text{SYLS} & \left\langle \; _{ha\text{-}syl}\left[\text{NUC} \quad \langle \boxed{1} \rangle\right] \right\rangle \\ \text{SEGS} & \langle \boxed{1}\ \text{o} \rangle \end{bmatrix}$$

3.6 German Declarative Syllabification

Markus Walther gives a declarative account of syllabification that substantially differs from the account given in the previous section (Walther, 1993). His segmental representations, the representation of syllabic structure and of phonotactic constraints, and the way the Onset Maximization Principle is realised in his account is different from the way these issues have been handled above.

Walther's approach is modular:

- There is a basic sonority-driven syllabification algorithm, which is constrained by

[10]Hache aspiré words could also be stipulated to start a new syllabification domain, but that would mean that they start a new phonological phrase and we have no reason to think that.

- language-particular instantiations of syllable length-restricting constraints and formal phonotactics.
- Phonotactic constraints are part of segmental representations.

The syllabification algorithm:

- avoids extrinsic ordering and structure-changing operations (i.e. it is declarative),
- is based on partially specified subsyllabic functions modelled as types which are assigned on the basis of elementary sonority differences,
- incorporates ambisyllabicity and the domain-bounded nature of syllabification.

For German Final Devoicing, which is a syllable-dependent phonological process, a novel non-feature-changing analysis is proposed. Arithmetic constraints are used to model phonetic interpretation of phonological structure in parallel with the building of phonological representations.

3.6.1 Descriptive Overview

The monomorphemic monosyllables in (34) illustrate the possible complexity of syllables in German:

(34) a. Aal *eel*, oh *oh!*, Maus *mouse*, Luft *air*, grau *grey*, Flug *flight*, Stroh *straw*, Sprung *jump(n.)*
b. Schnee *snow*, blau *blue*, Ball *ball*, rot *red*, Wurm *worm*, Milch *milk*, Faust *fist*, Helm *helmet*, Worms *(town name)*, Mumps *(disease name)*, Herbst *autumn*, Dienst *service*

The data above shows that obstruents precede sonorants in the onset, with matters being reversed in the coda. The extra peripheral consonants may only be coronal obstruents. Onsets allow one (/ʃ/), codas up to two extra segments (e.g. /st/).

The following list contains monomorphemic forms, which are maximal in the sense that only coronal obstruents may follow, e.g. in the form of inflections[11]:

(35) a. viel [fiːl] *much*, doof [doːf] *stupid*, schön [ʃøːn] *beautiful*, Stab [ʃtaːp] *pole*
b. fein [faɪn] *fine*, drauf [dʁauf] *on it*, neun [nɔyn] *nine*, Raub [ʁaup] *robbery*
c. Film [fɪlm] *film*, darf [daʁf] *may*, gern [gɛʁn]] *with pleasure*, gelb [gɛlp] *yellow*

Short vowels tolerate two coda consonants (35-c), long vowels (35-a) and diphthongs (35-b) only allow one coda consonant. No more than two vowel positions are allowed in a syllable:

(36) Theater [teː.aː.tɐ] *theatre*, Oase [oː.aː.zə] *oasis*

As a descriptive statement, the following syllable template can be given:

(37) (C)(C)(C) V X (C)(C)(C)

(37) entails that final full vowels must be *phonologically*[12] long:
Zoo [tsoː] *[tsɔ]

Schwa and syllabic sonorants, which usually result from schwa–zero alternations, do appear in final position. The deviant behaviour of schwa can be attributed to its functioning as a default vowel. Its distribution is predictable, and no schwa-initial words exist in German.

The following monomorphemic polysyllables show that 0-4 medial consonants are possible.[13]

(38) a. Aorta [aː.ɔʁ.taː] *aorta*, Eosin [eː.oː.ziːn] *eosin* $(C_{20}H_8O_5Br_4)$
b. Aroma [aː.ʁoː.maː] *aroma*, Kabine [kaː.biː.nə] *cabin*
c. Alkohol [al.koː.hoːl] *alkohol*, Tablett [taː.blɛt] *tray*

[11]The bracketed transcriptions in the examples below represent the output of phonology before phonetic interpretation. The phonetic interpretation of unstressed long vowels usually yields a short vowel contrary to what the bracketed forms listed here imply.

[12]This does not mean that syllable-final vowels must also be phonetically long. Phonetical length is also influenced by stress. Unstressed vowels in open syllables are short phonetically. For a description of Phonetic Interpretation see the section on arithmetic constraints below.

[13]With regard to the syllabification of medial obstruent clusters only one prominent pattern is shown.

d. Ernte [ɛʁn.tə] *harvest*, Elster [ɛls.tɐ] *magpie*
e. extra [ɛks.tʁaː] *extra*, extrem [eks.tʁeːm] *extreme*, *[eʁks.tʁeːm]

Syllabification in German is a domain-sensitive process; derivational prefixes (39-a), compound members (39-b) and C-initial derivational suffixes (39-c) all form separate syllabification domains. In contrast to this, inflectional and V-initial derivational suffixes (39-d) amalgamate with their respective stems to form larger domains of syllabification:

(39) a. Ur.ur.oma {Ur-}{Ur-}{Oma} *grand-grandmother*, be.vor.mun.den {be-}{vor-}{mund+n}
b. Va.ter.un.ser {Vater}{unser} *Lord's prayer*, Klapp.fahr.rad {Klapp}{fahr}{rad} *foldable bike*, Herbst.anfang {Herbst}{anfang} *beginning of autumn*
c. frag.los {frag}{-los} *unquestionably*, gelb.lich {gelb}{-lich} *yellowish*, Tau.chen {[taʊ.çən]}{Tau}{-chen} *little rope*
d. At.mung {Atm+-ung} *breath (n.)*, far.big {farb+-ig} *coloured*, For.ma.lis.ten {form+-al+-ist+-n} *formalists*

A minimal pair is **Tauchen** [taʊ.çən] *(little rope)* versus **tauchen** [taʊ.xən] *(to dive (inf.))*. The palatal fricative in **-chen** fails to assimilate in backness to the preceding vowel, as predicted by the German **ich/ach** alternation.

The following examples show clear cases of ambisyllabicity:

(40) a. Mitte [mɪtə] *middle*, lachen [laxən] *to laugh*, Hammer [hamɐ] *hammer*
b. trockener [tʁɔkənɐ] *dry (comp.)*
c. Dinge [dɪŋə] *thing (pl.)*, schwanger [ʃvaŋɐ] *pregnant*

Single medial consonants are ambisyllabic after short stressed vowels (40-a) and unstressable schwa (40-b). Segment /ŋ/ is not found in strict onset position, but can be ambisyllabic (40-c).

3.6.2 Declarative Syllabification

Segmental Representation and Sonority

The segmental representations Walther (1993) assumes are based on articulatory gestures which are temporally extended entities as defined by Articulatory Phonology (Browman and Goldstein, 1989). *Segments* are represented as feature bundles consisting of exactly two feature slots, each bearing a gesture. There is a HEAD gesture slot (for primary articulation — the main oral gestures) and a COMP gesture slot (for secondary articulation properties — voicing, nasality, laterality and lip rounding). *Gestures* have a constriction degree, a constriction location and a constriction shape (for some gesture types). Every constriction occupies a temporal interval. A gesture can be *active* or *inactive*. Segments that contain inactive gestures are e.g. /h/: HEAD: *inactive* and /ə/: HEAD: *inactive*, COMP: *inactive*. All this can be represented assuming the types given in (41).

(41) a. $\textit{scalar} \Longrightarrow \textit{temporal\_interval} \vee \textit{constriction}$
b. $\textit{constriction} \Longrightarrow \textit{degree} \vee \textit{location} \vee \textit{shape}$
c. ${}_{\textit{segment}}\left[\begin{array}{ll}\text{HEAD} & \textit{gesture}\\ \text{COMP} & \textit{gesture}\end{array}\right]$
d. $\textit{gesture} \Longrightarrow \textit{active} \vee \textit{inactive}$
e. ${}_{\textit{gesture}}\left[\text{TI} \quad \textit{temporal\_interval}\right]$
f. ${}_{\textit{active}}\left[\text{CD} \quad \textit{degree}\right]$
g. $\textit{active} \Longrightarrow \textit{oral} \vee \textit{velum} \vee \textit{glottis}$
h. ${}_{\textit{oral}}\left[\text{CL} \quad \textit{location}\right]$
i. $\textit{oral} \Longrightarrow \textit{lips} \vee \textit{tongue}$
j. $\textit{tongue} \Longrightarrow \textit{tip} \vee \textit{body}$
k. ${}_{\textit{body}}\left[\text{CS} \quad \textit{shape}\right]$

The representation in (41) assumes that some types (such as *temporal_interval*, and *constriction degree, location* and *shape*) are scalar types. Using *scalar types* along with *arithmetic constraints* to encode phonology-to-phonetics mappings could open up the possibility of such exciting applications as phonetic implementation within HPSG. See the section on phonetic implementation and arithmetic constraints.

Walther's declarative syllabification is sonority-driven. Sonority can be thought of as a segment's 'overall loudness'. It defines a partial ordering that partitions the set of segments into equivalence classes. It is usefulness as an abstraction device for defining syllable contours. Syllables can be defined as instances of a hat-shaped sonority pattern, which has come to be known under the term *sonority sequencing generalization* (Selkirk, 1984). However no satisfactory definition in terms of identifiable phonetic properties has been established so far. This is not surprising, because the sonority scale varies from language to language. On the other hand, language particular sonority scales seem to be learnable. This is demonstrated by Ellison, who gives an algorithm for learning the scale in Ellison (1992). Starting from a random initial sonority scale Ellison's program finds its way to the optimal language-particular scale. For German, the scale in (42) can be assumed:

(42)	obstruents	nasals	/l/	/ʁ/	high vowels	other vowels
	⟶		increasing sonority			⟶

For the purpose of defining syllable structure it is sufficient to think of sonority as a partial relation *more_sonorous(Segment1, Segment2)*.

Sonority-driven declarative syllabification

Walther's syllabification algorithm assigns one of four subsyllabic functions to each segment in a string. The subsyllabic function symbols are: **N**ucleus, **O**nset, **C**oda and **CO** for the ambisyllabic case. A segment can be defined to have one of these subsyllabic functions if it satisfies the conjunction of two basic constraints. There are four such constraints, which we can label C_1, C_2, C_3 and C_4, respectively. A pair of constraints holds for a segment depending on the sonority differences between the segment and its neighbouring segments. The constraints are defined as given in (43). A more expressive name than C_x is also given for each constraint.

(43)	constraint	direction	change	alias
	C_1	next segment is	more sonorous	*in_onset*
	C_2	previous segment is	less sonorous	*not_in_coda*
	C_3	next segment is	less sonorous	*in_rhyme*
	C_4	previous segment is	more sonorous	*in_coda*

The definition of subsyllabic functions then is as follows:

(44)
$$O = in\_onset \wedge not\_in\_coda$$
$$N = not\_in\_coda \wedge in\_rhyme$$
$$C = in\_rhyme \wedge in\_coda$$
$$CO = in\_onset \wedge in\_coda$$

The boundary conditions for the left- and rightmost elements are C_2 *(not_in_coda)* and C_3 *(in_rhyme)*, respectively. These represent zero sonority at the boundaries of syllabification domains. Note that we do not need a formal equivalent of the Maximize Onset Principle.

But note also that the constraints C_1 to C_4 are only defined for segments which are adjacent to segments which are either more or less sonorous than they are. This is not so in the case of *sonority plateaus.* These can be handled by employing *virtual sonority difference* markers *up* and *down*. In the case of real sonority difference they correspond to that difference. In the case of consonantal plateaus, however, virtual sonority continues to rise after an uprise and fall after a downfall, because such plateaus are either in a coda or in an onset. Vocalic plateaus, on the other hand, display an alternating behaviour in virtual sonority (*up, down*), which results in the syllabification of a sequence of long vowels into different syllables.

An application of the revised syllabification algorithm is shown in (45).

(45)

N	C	C	CO	O	N	C	O	N	CO	N	C
ɛ	k	s	t	ʁ	a	ː	t	e	ː	o	ː
up	down	down	down	up	up	down	up	up	down	up	down

Long vowels and diphthongs receive the labelling N-C in final position and N-CO in nonfinal position. The distinction between light and heavy syllables is captured by referring to the absence or presence of a Coda element. In languages where this does not suffice, a more specific instantiation of this criterion would have to state that in order for a syllable to count as heavy, the coda segment must exist and be of type consonantal. The CO label that is assigned to the second element of long vowels by virtue of the syllabification algorithm corresponds to glide formation **Theo** [teːjoː] or glottal stop insertion **Theater** [teːʔaːtɐ].

It is incorrect that the algorithm uniformly assigns ambisyllabic function CO to all (virtual) sonority minima. After twofold sonority downfall (*down down* virtual sonority pattern) no ambisyllabicity is possible. Also in other potentially ambisyllabic circumstances (*up down up* pattern) postvocalic consonants only appear ambisyllabic in certain configurations (46-b). The solution is that the segment occupying a sonority minimum is of type *in_onset* (= $O \vee CO$) iff it can form a well-formed onset with its right-hand neighbour according to phonotactics. Formally, we employ a disjunction for *in_onset* vs. *in_rhyme* constraint placement.

(46) a. Verfil.mung, Ord.nung, Sau.na, Verfeu.rung, Sah.ne, At.mung
filming, order, sauna, burning, cream, breathing
b. Schmu[g]ler, ho[p]la, E[ŋ]e, Samm.lung, Wid.mung, Schaff.ner
smuggler, oops, restriction, collection, dedication, conductor

We can implement the distinction between potentially ambisyllabic configurations and those which are never ambisyllabic by a small finite state automaton with 4 states, which tracks the virtual sonority differences. The basic syllabification algorithm is altered so that the partial subsyllabic function assigned depends on the state of the automaton at a given position.

Length restrictions

The syllable length restrictions that were given as a template in (37) can be defined as a constraint on syllables as given in (47) in the form of regular expressions.

(47)
$$length\_restriction \Rightarrow (ConsClust\ MinSyll)^{+} ConsClust \cap \overline{Schwa\ AnySeg^{*}}$$
$$ConsClust \Rightarrow \epsilon | Cons | Cons\ Cons | Cons\ Cons\ Cons$$
$$MinSyll \Rightarrow FullVowel\ AnySeg | SchwaOrConsNucleus$$

Phonotactics

Walther (1993) uses implicational constraints and cross-subcategorisation between adjacent segments to describe phonotactics. These implications are stored in the segment descriptions themselves. An implicational constraint for *labial obstruents* for example, is like the following:

(48) If the segment (a labial obstruent) is *in_onset* and the next segment is also *in_onset* then the second one must be *non_nasal*.

Morphological and lexical concatenation must connect the right and left contexts of individual segments. Lexicalised exceptions can be handled by type-marking the offending onset segments of items like **Smog, Smoking, Snob** etc. as *lexically_specified*, while normal occurrences of segments would be marked *lexically_unmarked*. The implicational phonotactics is conditioned to constrain only the *lexically_unmarked* segments.

Domain-sensitivity

Syllabification domains in German are identified with the prosodic category phonological word. A distinguished feature WORD (a subtype of *interval*) is used to indicate which phonological word a segment belongs to. Two prosodic words are identical if their temporal intervals are co-indexed. Morphemes like **-lich** which open up a new syllabification domain of their own lexically impose a precedence constraint between their own phonological word and the left-adjacent one. Morphemes like **-ig** specify that the value of their phonological WORD is coindexed with the word to their left.

Arithmetic constraints and phonetic interpretation

Arithmetic constraints can be used to express a precedence relation between temporal intervals (such as gestures) which can also be phonetically interpreted. An interval (sp_1, ep_1)[14] precedes an interval (sp_2, ep_2) if and only if $ep_1 < sp_2$. We might complement the computation of the relative temporal arrangement of individual gestures with an approach to the computation of the phonetic length of these gestural intervals, which recognizes length-determining coefficients throughout the whole prosodic hierarchy. The overall impact of prosodic structure on gestural length can be defined as, say, the product (or some other suitable function) of level-wise length coefficents k_j along a projection line. The phonetic length of phonologically long vowels, for example, is influenced by the stress pattern on the foot level of prosodic structure ($k_{stressed}$ vs. $k_{unstressed}$) as well as be the speech rate at the utterance level ($k_{speechrate}$).[15]

3.6.3 Syllable-dependent phonological processes

Final Devoicing has long been thought of as one of the prime candidates for a feature-changing rule (there is absolute neutralization of voicing contrast for obstruents in final position).

(49) a. Lo[p]–Lo[b]es, Ra[t]–Ra[d]es, Sar[k]–Sar[g]es, akti[f]–akti[v]es, Gra[s]–Gra[z]es
praise–(gen), wheel–(gen), coffin–(gen), active–(neut), grass–(gen)
b. orange *(adj.)* [oʁaŋʃ]–Orange *(n.)* [oʁaŋʒə]
c. Ra[t]–Ra[t]es, Bioto[p]–Bioto[p]e, star[k]–star[k]en
council–(gen), biotope–(pl), strong–(weak form)
d. kin[t]lich kind-lich *childlike*, kin[d]isch kind+-isch *childish*
e. Ja[kt] *hunt (n. sg.)* – ja[g]en *to hunt (inf.)* – Ja[kd]en *hunt (n. pl.)*
f. Ebbe [ɛbə] *low tide*, Bagger [bagɐ], Kladde [ʁladə] vs. Ebbstrom [ɛp][ʃtʁoːm] *stream of low tide*

Final Devoicing also applies to non-native voiced obstruents in loanwords (49-b). There is a multiple final devoicing (49-e). The domain is the coda (absolute syllable-final position – not morpheme-final as is indicated by (49-d)), and (49-f) shows that ambisyllabic consonants are not devoiced.

A declarative (not feature-changing) account can be given by defining initial entries for alternating voiced and non-alternating voiceless obstruents as exemplified for the labial stops in (50).

(50) a. $/\mathrm{b}/ : \exists S.lip\_closure(S) \wedge (in\_onset(S) \Rightarrow voiced(S))$
b. $/\mathrm{p}/ : \exists S.lip_closure(S) \Rightarrow voiceless(S)$

It predicts languages without Final Devoicing, e.g. English, to have different segmental descriptions.

The Final Devoicing constraint:

(51) $\forall S.obstruent(S) \wedge C(S) \Rightarrow voiceless(S)$

(51) is a lexical constraint imposed on the segmental lexicon, it must be conjunctively added to all segment descriptions, following the quantor's prescription. This yields a representation for alternating voiced obstruents like the one in (52).

(52) $/\mathrm{b}/ : \exists S.Slip\_closure(S) \wedge (in\_onset(S) \Rightarrow voiced(S)) \wedge (C(S) \Rightarrow voiceless(S))$

[14] i.e. of starting point sp_1 and of endpoint ep_1

[15] Walther also proposes to treat all geminates as single segments having the subsyllabic function CO and thus achieve geminate inalterability without the need to define any constraints stating it. To cover the length distinction of geminates as compared to 'normal' segments he proposes to interpret CO structures in phonetic implementation. This could even be achieved in a compositional fashion by distinguishing k_{in_onset} and k_{in_coda} length coefficients respectively, whose subsyllabic function counterparts are both contained in CO. Unfortunately, such a treatment can not be trivially carried over to languages like Hungarian, where the distribution of (true) geminates is not as predictable as that of ambisyllabic consonants in German or in Italian.

3.7 Conclusion

In this chapter, we have reviewed some of the possible ways to represent phonology that make it possible to integrate it into HPSG. We have seen that very different solutions to the same problems are possible, yet they all have the common property that only structure building operations are allowed, which are expressed in the form of constraints that are all unified to yield a representation of possible surface strings. We have seen how alternations that were traditionally handled by structure changing rules can be dealt with by underspecifying lexical representations and adding a constraint to the prosodic type where the rule operates. It may depend on the particular representation that we choose what this type is.

References

Bird, Steven and T. Mark Ellison. 1994. One-level phonology: Autosegmental representations and rules as finite automata. *Computational Linguistics*, 20(1):55–90.

Bird, Steven and Ewan Klein. 1990. Phonological events. *Journal of Linguistics*, 26:33–56.

Bird, Steven and Ewan Klein. 1994. Phonological analysis in typed feature systems. *Computational Linguistics*, 20(3):455–491.

Browman, C. P. and L. Goldstein. 1989. Articulatory gestures as phonological units. *Phonology*, 6:201–251.

Clements, George N. 1985. The geometry of phonological features. *Phonology Yearbook*, 2:225–252.

Ellison, T. M. 1992. *Machine Learning of Phonological Structure*. Ph.D. thesis, Dept. of Computer Science, Univ. of Western Australia, Perth.

Goldsmith, John. 1990. *Autosegmental and Metrical Phonology*. Oxford: Blackwell.

Kay, Martin. 1987. Nonconcatenative finite-state morphology. In *Proceedings of the Third European Conference of the Association for Computational Linguistics*, pages 2–10.

Klein, Ewan. 1993. An HPSG approach to Sierra Miwok verb stems. In J. M. Scobbie and T. Mark Ellison, editors, *Phonology and Computation*. University of Edinburgh.

Kornai, András. 1991. *Formal Phonology*. Ph.D. thesis, Stanford University.

Pollard, Carl and Ivan A. Sag. 1987. *Information-Based Syntax and Semantics, Volume 1: Fundamentals*. University of Chicago Press, Chicago and London.

Pollard, Carl and Ivan A. Sag. 1994. *Head-driven Phrase Structure Grammar*. University of Chicago Press, Chicago and London.

Selkirk, E. O. 1984. On the major class features and syllable theory. In M. Aronoff and R. T. Oehrle, editors, *Language Sound Structure*. MIT Press.

Walther, Markus. 1993. Declarative syllabification with applications to german. In T. Mark Ellison and James M. Scobbie, editors, *Computational Phonology — Edinburgh Working Papers in Cognitive Science 8*. Available by anonymous FTP from `scott.cogsci.ed.ac.uk., pages 55--79.`

Part II

Lexical Syntax

Chapter 4

Lexicon in Transformational Grammar

4.1 The Principles and Parameters Approach (Government and Binding Theory)

4.1.1 Meaning and Sound: A Minimalist Flashback To the Principles and Parameters Approach

Before taking a closer look at how a most influential transformational grammatical theory in the 1980s and 1990s gives an account of the lexicon and linguistic description in general, a terminological remark is in order here.

> "The Principles and Parameters Approach is sometimes called Government-Binding (GB) Theory. (....) True, early efforts to synthesize current thinking in these terms happened to concentrate on the theories of government and binding (Chomsky, 1988), but these modules of language stand alongside many others: case theory, thematic theory, and so on. (....) Furthermore, insofar as the theories of government and binding deal with real phenomena, they will appear in some form in every approach to language; this approach has no special claim on them." (Chomsky and Lasnik, 1995, p29)

To achieve descriptive adequacy, it seems necessary to enrich the format of permissible systems, but in doing so the grammar loses the property of feasibility so that the problem of the way of language acquisition is still uncracked. This issue is inherent in the kinds of rule systems that were being mused on in the 60s and partly 70s and it made the theory shift in another direction, namely, **assigning overall principles which underlie rule application generally to universal grammar (UG)**, which is a system of all principles which are common to all human languages and are on hand to each individual prior to experience, that is to say, a biologically given system in a child's brain. It embodies a set of absolute universals, notions, and principles which do not vary from one language to the next. There are language-specific properties that vary cross-linguistically and for these a range of choices is made available by UG.

The actual rules of grammar can then be given in the simplest form, with these principles guaranteeing that they will operate in such a way as to yield the observed phenomena in their full complexity. This grammar can reach a limit when rules are blotted out altogether and the outward rules are deduced from overall principles of UG, in the sense that the interaction of the principles would yield the phenomena which the rules had been thought up to describe. **Language variation is restricted to some options [=parameters] as to how the principles apply.** The Principles and Parameters approach held that languages have no rules in the familiar sense, and no theoretically significant grammatical constructions except as taxonomic artifacts. There are universal principles and a finite array of options as to how they apply, but no language-particular rules and no grammatical constructions of the traditional kind within or across languages.

The Principles and Parameters research has discovered such principles of grammar which cross-cut the traditional division of syntax into constructions named by their key morpheme or semantic content, e.g. passive, A-binding, or wh-question. Some properties of passive are held in common with properties of reflexive anaphora, while others with wh-questions. Syntactic difference among languages are put down to the role of experience in choosing values for certain open parameters in the theory. Since the articulations of the theory largely cross-cut traditional construction boundaries, the setting of a parameter for any articulation can be reckoned to have consequences that cross-cut constructions. The same property of Principles and Parameters research that predicts amazing linkages in cross-linguistic variation yields surprising linkages for the linguist attempting to explain language. A slight change in theory can be expected to have surprising consequences over a range of linguistic constructions.

> " ...Inquiry into generative grammar (....) has pursued the working hypothesis that UG is a simple and elegant theory, with fundamental principles that have an intuitive character and broad generality. By dissolving the notion of construction and moving toward "rule-free" systems, the Principles and Parameters Approach carries this tendency considerably forward. A related assumption is that UG is "nonredundant," in the sense that phenomena are explained by interaction of priciples in one particular way. (....) In its basic structure, the language faculty has properties of simplicity and elegance that are not characteristic of complex organic systems, just as its infinite digital character seems biologically rather isolated." (Chomsky and Lasnik, 1995, p29)

In sum, the Principles and Parameters approach strives to reduce descriptive statements to language-invariant and language-particular ones. The language-invariant statements are principles (including parameters), while the language-particular statements specify particular values of parameters. That is why the notion of construction, in the traditional sense, effectively disappears. For instance, there are no such constructions as verb phrase, or interrogative or relative clauses, or raising constructions. These descriptive objects (which have no theoretical status) are formed by the interaction of overall principles. The parametric options are restricted to properties of the lexicon and the point in derivation from deep structure (D-structure) at which structures are mapped to the phonetic form (see in the next paragraph).

UG specifies some linguistic levels (i.e. representational systems). **Each linguistic level provides systematic information about structural descriptions, and the structural descriptions are the linguistic expressions, and each linguistic expression is a sequence of representations, one at each linguistic level. The structural description may be regarded as a complex of instructions for performance systems, providing relevant information to their functions.** The performance systems fall into two types, (α) articulatory-perceptual, and (β) conceptual-intentional. The two linguistic interface levels providing the instruction for the above two performance systems are usually called the phonetic form for (α), and the logical form for (β) in the Government and Binding theory of the 80s. Each language will determine a set of pairs drawn from these performance systems as its formal representation of sound and meaning.

Before examining the lexicon in the Principles and Parameters approach we might want to show some more concepts of the theory, with special attention to the fact that the lexicon is a module in this framework.

4.1.2 Modular System

By a standard assumption language in the 80s grammar roughly consists of a lexicon and a computational system, the two major components, or "subcomponents of the rule system" (Chomsky, 1988, p5). The lexicon characterizes the lexical items which come up in structural descriptions, whereas the latter uses the lexical items to generate the form of structural descriptions. Deriving a particular linguistic expression involves a choice of items from the lexicon and a computation that constructs the pair of interface representations.

Generative grammars have assumed that there is a system of representation at which lexical items are inserted. This level expresses lexical properties in a form accessible to the computational system. From another aspect, principles at this level specify the way lexical properties are projected in a grammatical structure from the lexicon. Even if a level (the deep structure (DS) in the transformational generative tradition) is not at all indispensable, its role to link up the computational system and a set of lexical items (i.e. the lexicon) must be universally incorporated in any grammar (see the minimalist program).

The principles can be naturally grouped into modules of language. **A characteristic feature of generative grammars is that their model of language is modular, i.e. made up of subsystems.** Why are batches of principles taken to be a subsystem?

Firstly, the subsystems are complete in that they include all the essential rules and principles of their area and they are reckoned to cover the range of data of a particular aspect of grammar. Secondly, the subsystems are coherent in that they shape one particular aspect of grammar. Thirdly, the subsystems are autonomous in that the setup and workings of the principles cannot interfere with anything from outside the subsystem. Of course, an output of each subsystem may serve as an input for another and these intermodular interface levels are potential levels of linguistic representation, that is, representational levels.

The modular grammatical concept, with the exception of the latest minimalist theory, simply took the following hypotheses for granted.

i. There must be a level that serves as **input for semantic interpretation**

ii. There must be an interface level where **the primal structure of a sentence builds up** (thematic structure, X-bar tree graphs)

iii. There must be an interface level where **lexical insertion takes place and idiosyncratic properties of lexical elements enter the grammatical structure**

iv. There must be **a level representing the primal structure which is transformed by some** kind of movement, deletion, or other **transformational operation into a surface structure** (we do not think of **the** surface structure of Chomsky (1965) or Chomsky (1988) but **any** representational level which functions in that way. In other words, there are transformations which take the primal structure of a sentence as their input.

We would like to point out that **the above four hypotheses or the theoretical objects which they refer to are not closely related, let alone, depend on one another by some linguistic necessity, and, therefore, they may have been split into different subsystems of a syntactic theory**, although in fact they were conflated at deep structure, for example, in Chomsky (1965). Also, in the Extended Standard Theory and the Government and Binding Theory out of the above four assumptions (ii), (iii), and (iv) remained inseparable from their conditions or requirements being met at deep structure.

Returning to our starting point, UG consists of interacting subsystems, from one point of view these are the various subcomponents of the rule system of grammar:

α the base: the lexicon and the categorial component

β the transformational component (!syntax = categorial component and transformational component)

γ Phonetic Form component

δ Logical Form component

From another point of view, we can isolate subsystems of principles (bounding theory, government theory, theta-theory, case theory, control theory, binding theory in Chomsky (1988).

Base rules generate deep structure through insertion of lexical items into structures generated by the categorial component in accordance with their feature structure. The categorial component rules meet the conditions of the X-bar theory.

Besides the arbitrariness of the linguistic sign, among restrictions on lexicon properties there are options as to how functional elements are realized, also, variations in global properties of the heads, for instance, verb-complement ordering, specifier-head ordering, etc.), and some limited choice of substantive elements. (The categorial component specifies such parameters). Thus, by choosing an item from the lexicon the deep structure phrase markers will be fully determined for a language with parameters fixed.

How the lexicon is built is largely determined by the overall model of grammar. One takes a set of building blocks from the lexicon and puts them together in a tree in keeping with the principles of the grammar, moving elements from their original position to their final resting place in the course of derivation as needed. As a consequence, the lexicon always serves as the outset for such a derivation.

4.1.3 Lexical Specifications Determine Syntactic Distribution: Semantic and Categorial Selection

An item in the lexicon is either functional or lexical. The latter has substantive content while the former does not. A lexical item has

- categorial features (N,V,Adj,.....)
- grammatical features (gender, number, tense, mood, ...)
- inherent semantic features that determine semantic selection (strict subcategorization (Chomsky, 1965) and thematic roles) and syntactic features that determine categorial selection.
- phonological matrix

It has been an open issue since 1965 **whether categorial selection can be reduced to semantic selection**. Clearly, what complements a lexical head takes is related to its meaning, in other words, lexical specifications determine syntactic distribution, but it is unpredictable from semantic considerations what complements appear in fact (see later in this subsection).

The four basic lexical categories are determined by the following features : +N, -N, +V, -V

We can characterize the basic categories in the following way: adjectives are +V,+N, verbs are +V, -N, nouns are +N, -V, prepositions are -V,-N. The categories with +V features are purely predicative (V and A). The -N categories are potential case-assigners (P and V).

Subcategorizational phenomena are typical examples of syntactic distributions determined by lexical specifications:

(1) a. *Peter saw.
 b. Peter saw a dog.
 c. *Peter saw in the garden
 d. Peter put the man out onto the street.
 e. Peter spoke.
 f. Peter spoke about history.
 g. Peter was on the roof.
 h. *Peter was on.
 i. The TV was on.

Lexical entries include subcategorization frames, the formal realizations of selectional restrictions. For instance, run, V (.....); drive, V, (.....NP); fond, A (....PP).

Adjuncts, or adverbial phrases, are non-subcategorized constituents and thus they can be left out without affecting grammaticality and thus they are never included in subcategorization frames. But complements, or internal arguments (see 12.1.4), which are subcategorized arguments, affect grammaticality when left out. Implicit arguments are indicated in the subcategorization frames in parentheses: eat, V (......(NP)). All representatives of all lexical categories have subcategorization frames.

On the subject of **categorial (c-) selection** linked with **semantic (s-) selection**, the generative tradition has often presupposed more content than the mere pairing of sound and meaning. For one thing, **lexical items impose various requirements on the syntactic structures** in which they find themselves. These requirements shape up part of the speaker's knowledge of those lexical items. For argument-taking items (verbs, adjectives, nominalizations, (e.g. *keenness*, *amazement*)), these include semantic (s-) selection, selection for syntactic categories (c-selection) and assignment of arguments to syntactic positions: **linking** (Pesetsky, 1995). S-selection is simply the consequence of a predicate's lexical semantics for argument structure. This includes selection for relational categories like agent or patient (see 12.1.4) as well as independent categories like proposition or thing. The fact that we utter but do not eat propositions entails that a verb like *tell* allows a propositional argument while a verb like *drink* does not. The semantics of belief and wonderment will entail that *wonder* does take an interrogative argument whereas *believe* does not. Likewise the semantics of all these verbs tells us the semantic relations (agent, experiencer, etc.) borne by their arguments. C-selection in addition to s-selection tells us that *say* but not *drink* c-selects a CP complement, or that *wonder* and *believe* differ in whether the subordinate complementizer triggers *wh*-movement. Furthermore, c-selection tells us that *ask* may take an interrogative argument

DP, while the interrogative argument of *wonder* may not, or, for instance, determines the fact that the semantically close verbs like *enjoy* and *like* differ in their ability to take a gerund or infinitive argument in

(2) a. She liked hearing / to hear the concert.
    b. She enjoyed hearing / *to hear the concert.

Linking involves relations between the semantic categories of a predicate's arguments and their syntactic positions. It tells us that the propositional argument of *believe* is its object and not its subject, that the agent of *throw* is a subject and not an object, and that *like* and *please* differ in their placement of their experiencer in examples (3) and (4):

(3) a. MARY liked the play.
    b. *The play liked Mary.

(4) a. *Mary pleased the play.
    b. The play pleased MARY.

Since s-selection is so rooted in the irreducible pairing of sound and meaning, there have been speculations that the lexical entries of predicates need not specify their c-selection and linking properties directly. Pursuing these speculations, the lexical entries of a predicate will not contain explicit information concerning c-selection and linking, instead, most instances of either must be explained as consequences of s-selection aided by UG principles that map semantic categories onto syntactic categories and syntactic positions. Since s-selection itself is simply an aspect of lexical semantics, we will have a theory of the lexicon in which children learn pairings of meaning and sound, and it is UG that does the rest.

Thus, for example, the fact that *drink* selects things and not propositions along with the fact that UG determines that CPs do not denote things is sufficient (i.e. no particular c-selectional facts need to be learned) to explain the following:

(5) a. Sue said that the Earth is flat.
    b. *Sue drank that the Earth is flat.

The verb *believe* selects a declarative propositional complement but never an interrogative one:

(6) a. Jane believes that Eve is single.
    b. *Jane believes whether Eve is single.

Likewise, if UG requires agents to be linked to subject position, then the linking facts in ((7)) follow from the s-selectional properties of the main verbs and the linking principles supplied by UG:

(7) a. Bill threw the ball.
    b. *The ball threw Bill.

This reasoning should have the consequence that the facts in ((5)), ((6)) and ((7)) have cross-linguistic validity because the agent of *throw* cannot be linked to object positions or we would not expect a verb meaning 'drink' to select a proposition, nor would we expect a verb meaning 'believe' to select an interrogative complement in any language.

Furthermore, the same reasoning can apply to morphologically complex items, once UG is provided with the ability to project the argument structure of complex words from their component parts. Thus, an agent-patient verb (see 12.1.4) like *destroy* will pass on its s-selectional properties to the process reading of *destruction* in accordance with mechanisms of morphological inheritance given by UG plus language-specific facts about the suffix *-ion*, and for any nominalized verb a similar procedure to obtain is expected.

Once we leave c-selection, linking and morphological inheritance, it is clear that language acquisition involves facts that can only be the consequence of a child's experience. Linguistic experience, combined with the properties of the innate linguistic system, yields the native language that the child comes to know and use. Hypotheses about language should put as great a burden as possible on the biologically given system (which has been called UG) and as small a burden as possible on the child's linguistic experience. As a bare minimum, lexical entries consist of arbitrary pairs of meaning and sound and knowledge of these pairings is obviously an outcome of the child's linguistic experience. Of course, no one believes that these pairings exhaust the content of the lexicon. Thus, a child learns agreement-class and conjugation-class membership for nouns and verbs, availability for affixation and quirky case assignment facts.

Our model of the lexicon (cf. Pesetsky (1995)) should allow such facts as what aspects of lexical semantics coupled with what principle of UG could predict ((2)), ((3)) and ((4)) or the fact that the past tense of *go* is *went*, or that *-ion* suffixed to *destroy* yields *destruction*, or that the Russian verb *vladet'* 'command' governs the instrumental case.

On the whole, if we discover some pattern of c-selection, linking or morphological inheritance which neither looks like declensional class or quirky case nor is derivable from lexical semantics and current views of UG, one may make a number of responses. We might abandon the view that the lexicon is maximally simple to accommodate the case at hand, positing some new mechanism employed by the child in acquiring the pattern under discussion. Alternatively, we might identify our characterization of the problem in such a way that the problem disappears and we do this by changing our view of UG.

4.1.4 Semantic Dependencies Split Up Into Linguistically Significant Classes

The roles of the participants of an eventuality are assigned by the predicate. An event is thought of as a stage play, the thematic roles or theta-roles are the stage roles, and the arguments bearing the theta-roles (marked as Θ-role in the literature) are the actors on the stage. Θ-roles are listed in the lexicon, for each predicative head, in the thematic grid of the head.

Some widely acknowledged Θ-roles and their contents are listed below:

(8)	AGENT	PETER called his dog.
	THEME	PETER is in the garden.
	PATIENT	Peter stroked HIS DOG.
	EXPERIENCER	PETER felt warm. The computer pleases PETER.
	GOAL	Peter gave the fish to THE DOG.
	SOURCE	Peter got the fish from A FISHERMAN.
	LOCATION	Peter is in THE PARK.

The thematic/theta-grid of a verb is as follows:

(9) send (1, 2, 3)

The thematic roles listed in a thematic grid will be assigned to arguments in a syntactic structure:

(10) PETER (O1) sent TWO LETTERS (02) to HIS BROTHER (O3).

There is an important difference between subcategorization frames and thematic grids, namely, subcategorization frames deal with the selection of complements only, and external arguments, such as PETER in the earlier examples, are not represented in subcategorization frames, but are listed in thematic grids. Ultimately any variety of theta theory splits up possible semantic dependencies into linguistically significant classes called Θ-roles, and characterizes how each of them is represented in linguistic structure. Thematic roles may be assigned by a lexical head to a complement of that head as defined by X-bar theory, or they may be assigned compositionally by the head and its complements to a subject position; the former are called **internal arguments**, while the latter are called **external arguments**. Several authors claim that all languages canonically assign the agent role to an external argument, and the patient or theme role to an internal argument (Baker, 1988, p48).

The fundamental principle concerning **the biuniqueness condition on theta role assignment** is **the thematic criterion**, informally, every term of LF that requires a Θ-role, that is each argument, must be associated with one and only one position to which Θ-roles are assigned, and each Θ-role determined by the lexical properties of the head is uniquely associated with one and only one argument. Chomsky (1988) states the criterion about **chains** [taken together, a moved category and its trace makes up a chain] of surface structure and LF:

Θ-criterion: Given the structure S, there is a set K of chains, $K = \{C_i\}$, where $C_i = (\alpha_1^i, ..., \alpha_n^i)$, such that:

(i) if α is an argument of S, then there is a $C_i \in K$ such that $\alpha = \alpha^i_j$ and a Θ-role is assigned to C_i by exactly one position P.

(ii) if P is a position of S marked with a Θ-role R, then there is a $C_i \in K$ to which P assigns R , and exactly one α^i_j in C_i is an argument. In case (i) α has the θ-role assigned by P. (Chomsky, 1988, p335))

(Each argument is in a chain which bears exactly one Θ-role and each Θ-position [a syntactic position where a Θ-role must be assigned] is in a chain which contains exactly one argument.)

By this time it is clear to the reader that at D-structure all the phrases must show up in the position to which the theta role which they receive is assigned. As Chomsky (1988) puts it, the D-structure is a pure representation of thematically relevant Grammatical Functions, for example, subject-of, object-of, complement-of, or head-of. Grammatical Functions that are relevant to theta-role [called GF thetas, op.cit.] have an argument that bears a theta-role. Each argument slot determined obligatorily at D-structure must be filled by some argument with the appropriate grammatical function, and each argument must fill exactly one theta role as determined by its grammatical function. As an example, *whose luggage* and *Tom's luggage* must both show up in the position marked by t (trace) at the D-structures:

(11) The airline lost Tom's luggage.

(12) Whose luggage did the airline lose **t**?

(13) Tom's luggage was lost **t** by the airline.

Baker (1988) proposes a strengthening of the notion of D-structure such that it directly represents GF-theta, and in general, it is a direct representation of thematic structure:

> **The Uniformity of Theta Assignment Hypothesis (UTAH)**: Identical thematic relationships between items are represented by identical structural relationships between those items at the level of D-structure.

UTAH supports grammatical function changing processes, or the **unaccusative hypothesis**, according to which certain intransitive verbs with nonagentive subject NPs have that NP as a structural object at D-structure (in details see the next subsection). This NP then becomes the subject at S-structure via move-alpha [=move any constituent]. For instance, the S-structure of the ransitive sentence

(14) Eve melted the ice cream into mush

according to UTAH will imply the following D-structure:

(15) [S Eve [VP melted [the ice cream] into mush]]

The related intransitive structure will imply the S-structure of

(16) The ice melted into mush

according to UTAH will imply the following D-structure:

(17) [S e [VP melted [the ice cream] into mush]]

The same thematic relationship holds between the ice cream and the melting action in both sentences above, which is represented by having the same structural relationship hold between them at D-structure. It is time to turn to an elaborated view of how D-structure is tied up with potential lexicon-internal construction of lexical items.

4.1.5 Argument Structures and Syntactic Structures

In the following we condense a possible approach by Hale and Keyser (1993) that seeks to determine the extent to which one can understand some observed limitations on argument structure in terms of the essential nature of pregrounded principles and elements. **The proper representation**

of predicate argument structure is itself a syntax. That is to say, each lexical head projects its category to a phrasal level and determines within that projection an unambiguous system of structural relations holding between the head, its categorial projections, and its arguments (specifier [=a maximal projections's first branching constituent which is not an X'], of present, and complement). The authors refer to these projections as **lexical argument structures** or **lexical relational structures, LRS** and call on the conventional tree diagrams and labels V, N, V-bar, or VP, and so forth to represent them.

The following ususal definitions of X-bar relations are relevant to further discussion:
m-command: α governs β if and only if

i α does not dominate β and

ii every XP [=maximal projection] that dominates α also dominates β.

government: α governs β if and only if

i α is a governor (governors are heads),

ii α m-commands β,

iii no barrier intervenes between α and β. (Maximal projections are barriers to government).

There are two relevant constraints on argument structure. Firstly, the variety of relations between arguments and the head and its projections is highly restricted, a circumstance that is thrown back in the matchingly restricted range of semantic (thematic) roles, acknowledged in a broad and well-informed linguistic literature, and, secondly, the depth of embedding in lexical structures is uniformly slight, on the whole, permitting one complement VP for a given lexical entry, shunning full use of the recursive capacity inherent in complementation.

Hale and Keyser (1993) are lead to their syntactic view of LRS through investigating denominal verbs, which are derived by a lexical process. Denominal verbs (e.g. *sneeze, shelve, thin*) are derived through the operation of the head movement variant of move alpha, called incorporation, studied by Baker (1988). If denominal verb formation takes place by incorporation it is reckoned to be bound to syntactic principles that govern the application of incorporation.

The first class of denominal verbs, unergative verbs (or true intransitive verbs), for example, *laugh, calve, sneeze, dance, sleep*), has an initial lexical projection of a verb and its nominal complement:

(18) The cow calved.

(19) Peter sneezed.

(20)

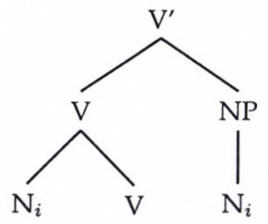

The lexical structure representation (LRS) of an unergative verb involves **incorporation** of the nominal head of an abstract verb's NP complement into the abstract V. **The Head Movement Constraint,** which states that an X head may only move to the Y head that properly governs it, must be respected. The resulting compound, of which only the N component is phonologically realized, corresponds to the denominal verb. Unergatives, thus, have an initial lexical structure of a simple transitive type. Of course, the verbs *make, do, have, take* at first glance project the same structure but they do not involve incorporation. Moreover, unergatives of one language often match a simple transitive VP structure without incorporation in another language (e.g. *sleep* is unergative in English but transitive in Basque), or a transitive VP modified by visible incorporation (e.g. Jemez *-zaae-'a* 'song-do' ='sing'). The relation between the simple transitive structure and the incorporation structure belongs to the class of phenomena known as lexical alternation. A transitive verb that takes an expletive subject (e.g. *It cowed a calf* meaning 'A cow had a calf') is nonexistent in English and several other languages. This hypothetical verb would incorporate its subject and not its complement and it is well-known that a subject that originates as an external argument cannot incorporate into the verb that heads its predicate, and incorporation from the subject position would violate the **Empty Category Principle stating that an empty category must be properly governed.**

Similarly to the unergative class, another N-projected complement is in the LRS of light verbs (e.g. *make, have, take* in *make an agreement, have a look, take a bath* and cognate verbs (e.g. *live a life.*

A second, more complex further verb class includes verbs that have the P-projected complement in the LRS. As a first subtype within these, the location (e.g. *shelve, corral, box* and locatum (e.g. *hobble, saddle*) verbs (as a shorthand, referred to as location verbs), with the same surface representation as the true transitives, share the following LRS:

(21) Peter shelved all the books.

(22) Peter saddled his horse.

(23)

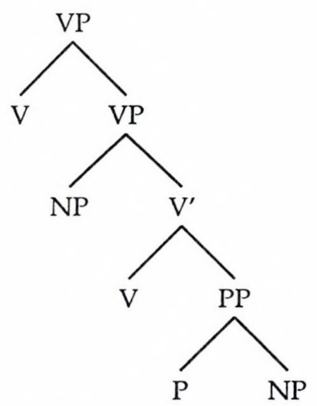

The prepostion P is supposed to be a hidden P, a nonovert variant of the category, the same as in *put a lamp ON the table* or *provide a house WITH paint*. There are two further P-projected subtypes, the *put/smear* type (e.g. *smear, daub, rub*) and the *get/splash* type (e.g. *splash, drip, pour*).

(24) She smeared/put butter on the bread.

(25) He splashed/got mud on the wall.

Without going into the depth, the surface form of these verbs is derived by three applications of head movement, the first of which incorporates the lower N (e.g *shelf, saddle*) into the preposition which governs it. The compound so formed is then moved into the verb that governs it, there forming a compound that makes the final hop to incorporate into the matrix verb. Let us recall that the initial LRS representation of location verbs shares the essential relational structure with the *smear* class, with the exception that the phonologically overt morpheme realized in the matrix verb position is not a verb but a noun, originally heading the complement of the PP in LRS.

Minimality, a syntactic principle, is at issue in explaining why English and many other languages lack hypothetical verbs like *shelve, bush* in the structures *He shelved the books on* or *She bushed a trim* meaning 'He gave the bush a trim'. Although the trace in NP is coindexed with the verb to which its antecedent N is adjoined, this verb cannot govern the trace, the (abstract) preposition being the closest governor, defining PP as the minimal governing domain for the trace, and the PP being a barrier to government from the distant verb outside (Chomsky, 1986).

After the underlyingly N-projected and P-projected classes a final third verb type includes verbs with an underlyingly A-projected complement, the ergative verbs, e.g. *thin, narrow, tighten, loosen, clear* in the structure

(26) The cook thinned the gravy.

with the LRS:

(27)

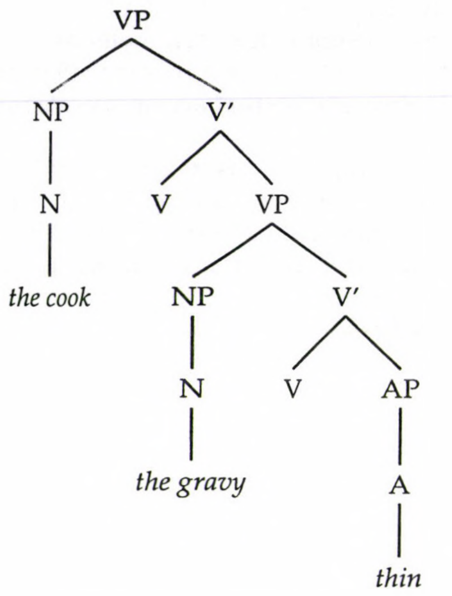

Verbs of this type, like the others we have looked through, are derived by head movement, in this case the incorporated elements are adjectival. The lower verb projects a structure that is parallel to the P-projected verb class, but with the PP of the latter is replaced by AP. It is a fundamental property of an AP that it be attributed of an entity. Thus, just with the PP complements, a subject necessarily shows up in the [Spec,VP] (the NP *the gravy* in the above example). The upper V projects the LRS associated with the clausal relationship.

An interesting contrast between this A-projected group and the P-projected group can shed light on the nature of the LRS. An ergative transitive verb exhibits an unconditional transitivity alternation along the ergative pattern, that is, with the object of the transitive verb and the subject of the intransitive verb. Inchoative use in Hale and Keyser (1993) is when an otherwise transitive ergative verb used intransitively, i.e. non-causatively. The middle construction is well-formed with all non-N-projected verb classes. In the following contrastive examples those marked with *a* are transitive sentences, *b* sentences are middle constructions and *c* sentences are inchoatives.

α) Underlyingly P-projected verbs:

(28) a. Peter shelved the books.
b. These books shelve easily.
c. *These books shelve.

(29) a. We smeared mud on the wall.
b. Mud smeared on the wall easily.
c. *The butter smeared on the wall.

(30) a. The pigs splashed mud on the wall.
b. Mud splashed on the wall easily.
c. Mud splashed on the wall.

β) Underlyingly A-projected verbs

(31) a. Peter cleared the screen.
b. The screen cleared easily.
c. The screen cleared.

Like the inchoative, the middle construction involves s-syntactic movement of an argument bearing the internal subject relations, in the case of the middle, this is an object in s-syntax. Transitive verbs that can undergo middle formation are those whose s-syntactic object is an affected argument, that is, those whose s-syntactic object corresponds to an internal subject in LRS.

Within the underlyingly P-projected class only the subtype *splash/get* shows the same free alternation, the other subtypes, however, show only the middle construction (allowed for all ergative verbs), which is an intransitive alternative (in the transitivity alternation) that is compromised by requirements to be met, such as the use of the generic, a modal, or an adverb like *easily*.

This contrast between the types *splash* and *smear* is explained by supposing internally oriented manner modifier for the former, while externally oriented manner modifiers for the latter. The manner component modifiers of *splash* type verbs are adverbial modifications to the VP and in particular to the event depicted by the verb and its most prominent direct argument. In contrast, transitive verbs of the *smear* type involve a manner component that relates, not internally to the LRS, but to the external argument. Their LRS representations are the following, where the manner component of a verb is marked as a bracketed tag on an appropriate V node, and the internal licensing relation is marked by coindexing the manner component with the internal subject:

(32) *splash*:

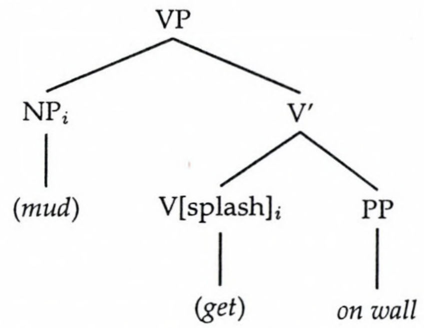

(33) *other subtypes*: *smear, shelve*:

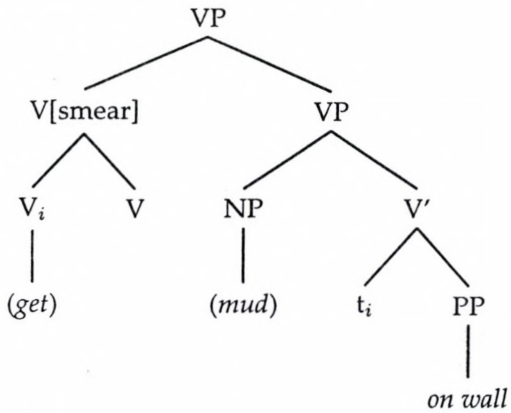

Only at D-structure is the manner component [smear] properly licensed, since it is only at D-structure that the required external argument is visible to the manner tag associated with the verb. The inchoative is impossible for *put* and location verbs as well, because raising the object to the specifier of the inflection phrase (with the conventional notation: [Spec,IP]) blocks licensing of the externally oriented manner modifier. By contrast, a *splash/get* type verb although structurally identical to location verbs in lexical syntax, is devoid of all means or manner modification; and, as expected since no licensing is required, the inchoative is possible for *splash/get*.

Why is the middle construction grammatical for denominal location verbs *shelve/box* if the inchoative blocks licensing of externally oriented manner modifiers and location verbs with regard to middle formation do not differ from ergative verbs? The middle construction is likely to be formed from a (causative) transitive structure, rather than the intransitive structure in which the internal subject is immediately dominated by the uppermost VP node, that is, the structure associated with the inchoative. Let us suppose that internally oriented manner components are associated with the inner verb, whereas externally oriented manner components are associated with the upper (causative verb). Licensing of a manner component, then, is association with a particular verbal element in the LRS representation. Verbs that cannot appear in the inchoative construction (e.g. *smear*, or *shelve*) have their externally oriented manner modifiers associated with the causative verb: they must be transitive. On the other hand, if the middle construction is formed from the transitive, that construction will not interfere with the licensing of externally oriented manner modifiers, since their required locus, the causative verb, is present in the transitive LRS representation.

In sum, **certain gaps in the lexicon can be made to follow if the formation of the lexical items in question is subject to principles known to be operative in syntax**, and, consequently, the structures over which lexical derivations are defined are true syntactic objects, over which syntactic relations and principles are defined. If incorporation involves head movement, it has clear implications, namely, that an argument structure itself is a syntactic object, since it is to be identified with the syntactic structure projected by the lexical heads. Head movement is a process constrained by syntactic principles and is reckoned therefore to limit the range of theoretically possible incorporations. Since lexical processes under examination affect the argument structures of lexical items, we are justified that argument structures themselves are syntactic objects. To put this another way, if denominal verb derivations were not constrained by syntactic processes but they were simple lexical processes of category change, the range of possible denominal verb types would include transitive verbs with an incorporated external argument like *cow* in *It cowed a calf* 'A cow had a calf' (see at the unergative verbs). Hale and Keyser (1993) asserts that the notion of argument structure is to be identified with the notion LRS.

Let us turn to the fact that unergatives (the *laugh/sleep/calve* type verbs) have no subject at all in their LRS representation and ergatives (the *thin/tighten/clear* type verbs) have a intransitive-transitive alternation. In the LRS of change-of-state verbs (equated with ergatives for the present purpose) and location/locatum verbs (*shelf/box*) the appearance of the subject in the inner VP is forced, being required by the complement within that VP. Since the complement in the inner VP is a predicate in the LRS representation of these verbs, full interpretation [a principle which requires that no elements can remain uninterpreted at LF or PF] of the inner verb requires that a subject appear, internal to VP, so that predication can be realized locally, thereby correctly linking the complement of the inner VP to the subject of that VP. The [Spec,VP] position of VP in LRS representation of a lexical verb is filled only when it is forced by some principle. In the case of ergative/change-of-state verbs (*thin/tighten*) and location/locatum verbs (*shelve/box*) just considered, the appearance of subject is forced by **predication**, whereas for the unergative (*laugh/calve*) class nothing forces a subject since the LRS complement of these verbs is a [VP [V, NP]], which is not a predicate. In short, VP complements in LRS representations are not predicates. In truth, the subject is excluded from the LRS representations of unergatives. Full Interpretation will guarantee that ergatives and location verbs have a subject in the inner VP and unergative verbs lack one. For the former group absence of subject would leave the complement of the inner VP uninterpreted, whereas for the latter group a subject present in LRS would be uninterpreted for lack of predicate in the complement position. No stipulation is needed to make this system work.

As far as subject appearance is concerned, firstly, an LRS specifier position can be occupied only if the appearance of an argument there is internally motivated by predication, such arguments are the subjects of VP-internal predicates (the subjects of a PP or AP complement of a head). **Predication requirement** forces the appearance of a subject for those verbs whose complement is inherently predicational, namely, for PP or AP complements.

(34) a. *shelve*:

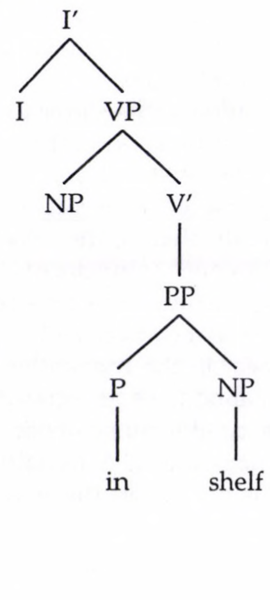

b. *clear*:

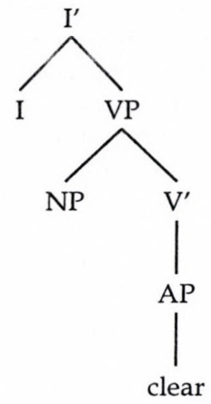

Secondly, there is VP-external motivation for forcing subject; externally forced subjects are products of the construction itself. The appearance of a subject is forced by properties of the matrix, for example, the transitive features of a causative verb or affix, or case/agreement features of inflection:

(35) *Hungarian futtat make (sb) run*

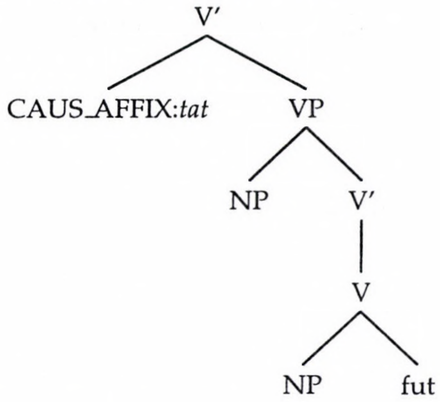

At D-structure, of course, these subjects are identical structurally and both raise to [Spec, IP] but they have distinct interpretation.

As a last issue of this subsection, unergatives lack causative forms in English, while in many languages overt causative morphology has properties that force the appearance of a subject in its immediate complement VP, the English causative does not force a subject and that is why the following unergative construction is ungrammatical: *Sneeze the child 'make the child sneeze'*. Suppose that the English causative is devoid of properties, in particular properties that could force the appearance of an NP in the specifier position of its complement. But the causative could acquire properties through incorporation of an overt head of its complement. This would block lexical insertion of a subject into the lower specifier position, with the result that English causative of the unergative is ungrammatical. On the other hand, this happens exactly in the case of ergatives (*clear, lengthen*), where this incorporation yields the transitive variant. This transitivity conversion simply hinges on whether head movement (incorporation) applies or not. The point is that this unconditional alternation for these ergative verbs is thus not stipulated, since it follows directly from their essential nature, namely, a non-overt causative verb (in English) has no properties of its own and acquires properties through incorporation.

To conclude, lexical categories project **unambiguous syntactic structures**. Crucially, such an unambiguous projection forbids *n*-ary branching, where *n* is greater than two, and it prohibits the projection of a phrasal category through more than one intermediate level, for instance, V projects at most V-bar and VP. This accounts for the first constraint (see at the beginning of this subsection) that the variety of relations between arguments and the head is restricted by restricting arguments to the complement and specifier positions in lexical relational structure representations. These positions match the grammatical relations an argument may bear in lexical relational structures, specifier of VP or complement to V, complement to P, and so forth. And there are no lexically determined roles beyond these. Argument structure representations are also subject to **Full Interpretation** [a principle in the Minimalist approach: no elements can remain uninterpreted at LF or PF]. This is relevant primarily in connection with the specifier position, which can be occupied only if the appearance of an argument there is internally motivated by predication. Since complement APs and PPs but not VPs in the LRS sense are predicates in lexical representation structures, VP recursion is excluded, which accounts, in part at least, for the observed second constraint (at the beginning of this subsection) that embedding permits one (and no more than one) complement VP for a given lexical entry.

Talking of syntactic representations and selectional properties, another crucial principle of government and binding theory until Minimalist Theory had been what is known as the **Projection Principle**, which states that representations at each syntactic level (D-, S-structure, and LF) are projected from the lexicon in that they observe (or, in other words, represent) the lexical selection (i.e. subcategorization) properties of lexical items categorially Chomsky (1988, p29). This principle presupposes the existence of a lexicon listing the idiosyncratic properties, in particular, what thematic relations lexical items may have with other phrases, to wit, what phrases they subcategorize and assign thematic role to. The consequence is that categories moved by move alpha [=move any constituent] will leave phonetically null copies [= traces] behind them to preserve the representation of these selectional properties.

4.2 Lexicon In the 1995 Version Of the Minimalist Program

4.2.1 The Minimalist Theory On the Whole

The *Categories and Transformations (Chomsky, 1995)* came up with a radical version of the transformational-generative grammar in which the semantic and phonetic interface of language are directly linked, and the selected elements build into a structure at once. This yielded derivation branches off to a potential semantic interpretation on the one hand, and a potential phonetic interpretation on the other. The Minimalist model of grammar is a derivational system as its predecessors, but it differs in that the mechanism creating the basic representation and the one performing transformations are the same: the generalized transformations (GT). The model is also unlike previous ones in the assumptions it makes with regard to, firstly, **the formal licensing mechanism** which the grammar makes available, namely, limited to **feature checking** in specifier-head and head-adjunction configurations, and, secondly, **the nature of the constraints on derivations** (economy conditions). As early as the 1993 minimalist approach, the conceptual foundation to be explored (even later, in the 1995 theory) was already conspicous:

> "[in the Principles and Parameters approach] One crucial assumption has to do with the way in which the computational system presents lexical items for further computation. (....) All items that function at LF are drawn from the lexicon before computation proceeds and are represented in the X-bar format. (....) This picture requires conditions to ensure that D-structure has basic properties of LF. (....) If they [the conditions] are not met [at LF], the expression receives some deviant interpretation at the interface and there is nothing else more to say. The Projection Principle and the Theta Criterion have no independent significance at LF. But at D-structure the two principles are needed to make the picture coherent; if the picture is abandoned, they will lose their primary role. These principles are therefore dubious on conceptual grounds, though it remains to account for their empirical consequences, such as the constraint against substitution into a thematic position. If the empirical consequences can be explained some other way and D-structure eliminated, then the Projection Principle and the thematic criterion can be dispensed with." (Chomsky, 1993, p20)

The Minimalist Program pledges itself to back up the hypothesis that there are no conditions linking lexical properties to interface levels, such as a projection principle. Along the lines of section 4.2 of the Minimalist *Categories and Transformations (Chomsky, 1995)* the following two subsections will display the skimpy picture which we can gain of how the lexicon is set out and how lexical items build into structures.

4.2.2 The Content Of the Lexicon

Exceptions are linguistic elements: whatever that does not follow from general principles. The lexicon is exactly the list of these unique elements. The above general principles break down into two subcategories:

α the principles of UG

β specific linguistic principles, i.e. the phonological, morphological bits of information, choice of parametric options, and whatever may enter into language variation. We further call for an optimal coding for idiosyncratic properties (the exceptions) in the unified (see later) lexical entry.

On the simplest assumptions, the lexical entry provides the information needed for further computation. Therefore, a lexical entry will comprise

α the relation of sound and meaning

β the subcategorization of the lexical entry which are sufficient for the logical form and the phonetic form to interpret the incoming derivation.

The lexical entry, however, will not contain the *phi*-features [=information about number, gender, person, tense, or mood] because all these do follow from the lexical categories (N,V,ADJ,.....) by the principles of UG. Furthermore, the entry should not specify phonetic or semantic properties which are universal or specific to that language: the predictable interactions of phonemes, or potential concrete versus abstract usages of nouns, for instance, the fact that the noun *book* can be used to refer to something which is simultaneously abstract and concrete as in the following sentence:

(36) The BOOK that I am writing will weigh five pounds.

The optimal coding includes the list of semantic and phonetic specific features and unpredictable formal features, for example, for the invariant noun forms (*scissors*).

One can choose a noun (or any lexical item) in two steps:

i. we form a numeration (N, I) [**Numeration** (a term of the 1995 Minimalist Program): a set of ⟨lexical item, *i* ⟩ pairs, where *i* is the index of the lexical item, i.e. the number of times we choose that lexical item from the lexicon. We will thus have a batch of lexical items with indices to further connecting them by generalized transformations (operations to build phrases).]

ii. an operation SELECT introduces the N in question into the derivation by adding it to the set of the syntactic derivations and reduces its index by one.

Specific features (plural, nominative) are added to it in the course of these two steps.

As an illustration, let us take the following example. Suppose that *book* is chosen as part of the array from which a derivation proceeds to form PF and LF representations. In the first step we form a numeration which includes (book,*i*), with the index *i* and in the second we introduce *book* into the derivation by the operation Select, which adds book to the set of syntactic objects generated and reduces its index by 1. The optional features of a particular occurrence of *book* (say, accusative and plural) are added by either step 1 or step 2, presumably in the first step, a decision that reduces reference sets and hence computability problems. Then the numeration N will include

(37) (*book*, /accusative/, /plural/, 2)

UG prescribes that such features must exist but UG does not prescribe which features must be chosen concretely. It may well be that a noun is selected together with a greater nominal configuration by SELECT. Such features, as case or thematic features could belong to it.

In this section of the minimalist program you can find the null hypothesis which is essential in some other areas of the minimalist framework. The **null hypothesis** states that case and thematic features are added arbitrarily as a noun is selected for the numeration. In the numeration case and *phi*-features are specified, whether by the lexical entry (intrinsic features) or by the operation that forms the numeration (optional features). Greater structures are relevant only for checking of features of the noun which are already present in the numeration.

Speaking of categories, a lexical entry of a verb also represents the instructions for the phonological component and for the interpretation of the logical form representation: a phonological matrix and some array of semantic properties. The lexical entry must suffice to determine that the particular lexical item has the categorial property V, perhaps by explicit listing. Selectional properties, insofar as they are determined by semantic properties, whether by UG or by specific language-particular rules, will not be listed in the lexicon. In the same way, an individual specification of tense and *phi*-features will not be indicated in the entry because that is determined by its category V (by UG). The case-assigning property of a verb is intrinsic, either determined by the semantic features of the lexical items or listed as idiosyncratic. Further features associsated with the verb which are unpredictable from the lexical entry, for instance, tense or *phi*-features, may stem from

- α choosing them arbitrarily and assigning them to the verb as it enters the numeration, or
- β operations of the overt syntax or the phonological component (including morphology) that form complex words by association with other elements (e.g. adjunction to the functional head T through overt V raising).

Thus in this β case, if overt syntactic operations are involved, the categories involved will be marked in the lexicon or the transition to numeration) as allowing or requiring affixation. The word might reach the phonological component uninflected, the PF form resulting from interaction with functional elements within the phonological component. (The ways can vary across languages or even within one language.)

4.2.3 The Means of Coding

In which form is the information coded in the lexical entry? In the case of a noun (*book*) the optimal representation should include the standard phonological matrix, or some arbitrary coding (say, 23), interpreted within the phonological component as a phonological matrix. Similarly, for PAST TENSE, a phonological matrix DENTAL, or an arbitrary coding, interpreted in the phonological matrix as DENTAL. (Recall that English verbs regularly form their past tense form by [d] or [t].)

In the case of an unpredictable form of a lexical entry, e.g. the English copula *be*, the lexical coding will provide whatever information the phonological rules need to assign a form to the structure [copula, F], where F is some set of formal features (tense, person, etc.). We may present the information as a list of alternants (*am, is, were, ...*), each with its formal features, or by some coding that allows the phonological component to pick the alternant. It would be a methodological mistake to generalize this worst case to all cases, to infer from the existence of the worst case that it holds for all lexical items.

The lexicon contains substantive elements (V,N,ADV,....) with their idiosyncratic properties and some functional elements, among which the lexical representation of agreement or tense may bring about special problems. Postulation of functional categories needs justification, either by phonetic or semantic interpretation or theory-internal arguments. Tense, complementizer, and determiner have semantic properties, while agreement does not. T is finite or nonfinite, with further subdivisions, D is the locus of referentiality, while C is the indicator of mood or force in the Fregean sense, i.e. declarative, interrogative, etc. The choice among the options of a given type is arbitrary, part of the process of forming a numeration from the lexicon, as in the case of *phi*-features of verbs, or case and some *phi*-features of nouns. Functional features have phonological properties, of course. For instance, declarative C is invariantly *THAT* (with a null option).

The lexicon provides optimal coding for the exceptions. Suppose that specific morphological properties of a language constrain the phonetic correlate of formal features, e.g. verbs indicate persons with prefixes, while numbers with suffixes. Then the lexical entry will abstract from these properties, including only the information that they do not determine.

For each lexical item in a given language the idiosyncratic codings are given in a unified lexical entry. One might propose that instructions for phonological rules, instructions for logical form interpretation, and formal features appear in distinct sublexicons, which are accessible at different points in the computational process. Such elaborations might involve new levels and relations among various parts of the derivation.

References

Baker, Mark. 1988. *Incorporation: A Theory of Grammatical Function Changing*. Chicago: Univ. of Chicago Press.

Chomsky, N. 1993. A minimalist program for linguistic theory. In Hale and Keyser (Hale and Keyser, 1993). also in The Minimalist Program : Chapter 3.

Chomsky, N. 1995. *Categories and Transformations*. Cambridge, Mass.: MIT Press. Chapter 4.

Chomsky, N. and Howard Lasnik. 1995. The theory of principles and parameters. In *The Minimalist Program*. MIT Press. Chapter 1.

Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, Mass.: MIT Press.

Chomsky, Noam. 1986. *Barriers*. Linguistic Inquiry Monographs 13. MIT Press.

Chomsky, Noam. 1988. *Lectures on Government and Binding*. Dordrecht: Foris.

Hale, Kenneth and Samuel Jay Keyser. 1993. On argument structure and the lexical expression of syntactic relations. In Kenneth Hale and Samuel Jay Keyser, editors, *The View from Building 20*, Current Studies in Linguistics 24. Cambridge, Mass. MIT Press.

Pesetsky, D. 1995. *Zero Syntax*. Current Studies in Linguisitcs 27. Cambridge, Mass.: MIT Press.

Chapter 5

Argument Structure and grammatical relations in HPSG

This chapter is a survey of the HPSG treatment of argument structure and grammatical relations.

The first section gives an introduction to feature structures as used in HPSG and also explains some of the principles responsible for driving complementation procedures by constraining phrase structure. It also explains in what ways a conception of argument structure is realized in standard HPSG.

The second section survey recent trends in (extended) HPSG and discusses the treatment of subcategorization through valence features.

The third section is an investigation into the constraint-based account of certain complementation patterns, among others, raising, control and passive constructions.

Throughout this chapter we assume a working familiarity with feature structures and the attribute-value matrix representation entertained by HPSG. In the feature structures we only indicate the relevant substructures and, for reasons of easier readability, often abbreviate long feature paths.

5.1 Standard HPSG treatment of argument structure and grammatical relations

This section explores in what ways complementation and subcategorization is treated in standard HPSG (Pollard and Sag, 1994). The evolution of the ideas presented here can be traced through earlier works, especially Pollard and Sag (1987), Sag and Pollard (1989) and Hinrichs and Nakazawa (1989b). These earlier conceptions are not crucially different from the one described in this section though they are outdated in some aspects, therefore we omit their detailed exposition.

5.1.1 Introduction to HPSG features

HPSG uses a typed attribute value structure formalism to state constraints about feature structure that model linguistic objects. One of the most important types that really stand for linguistic signs in a traditional sense are shown in the partial hierarchy below:

(1) *HPSG types of linguistic signs*

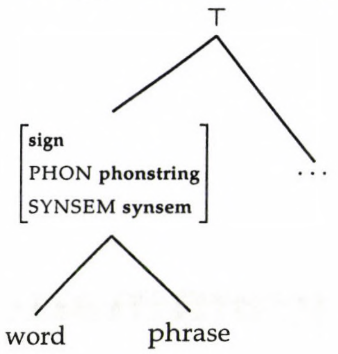

For a quick start, we present a typical AVM feature structure for a non-phrasal sign with values of a rather generic type and explain briefly their intended import.

(2) *a typical feature structure of a word*

$$\begin{bmatrix} \textbf{word}\ \text{PHON}\ \textbf{phon} & \\ \text{SYNSEM} & \begin{bmatrix} \textbf{synsem} & \\ \text{LOC} & \begin{bmatrix} \textbf{loc} & \\ \text{CAT} & \begin{bmatrix} \textbf{cat} \\ \text{HEAD}\ \textbf{substantive} \\ \text{SUBCAT}\ \textbf{list} \end{bmatrix} \\ \text{CONTENT}\ \textbf{content} & \\ \text{CONTEXT}\ \textbf{context} & \end{bmatrix} \\ \text{NON-LOC}\ \textbf{non-loc} & \end{bmatrix} \end{bmatrix}$$

The SYNSEM basically encodes all sorts of categorial, syntactic, semantic and contextual information. Besides this, a sign has the PHON attribute, the value of which is sometimes indicated above the AVS for easier readability. Phrasal signs also have the DAUGHTERS substructure encoding constituent structure information (see section 5.1.3). The division between local and non-local features, appearing in the distinct LOC and NON-LOCAL substructures within the *synsem* domain, is relevant only in the context of the lexical treatment of extraction (unbounded dependency constructions).

Within the local structure we find syntactic, semantic and contextual information featuring in the CAT, CONTENT and CONTEXT values, respectively.

The CAT substructure contains specifications that is traditionally treated under the rubric of categorial, subcategorization and selectional information. This is the substructure that mainly concerns us in relation to the treatment of subcategorization and argument structure.

The CAT value has two attributes, namely HEAD and SUBCAT. The substructure defined by the value of the former is basically the residence of features that are universally assumed to percolate up to the phrasal projections of the word in question. The sort of the HEAD value is more or less the equivalent of the X'-theory category information without the projection level or of other traditional notions of part-of-speech information. In the general case it carries not much more information than syntactic category, but in the case of lexical categories, it may specify formal features that are specific to the category in question and that do not have a semantic impact, e.g., case of nouns. The partial hierarchy below shows the HPSG conception about parts-of-speech.

(3) *Type-hierarchy fragment for categorial information*

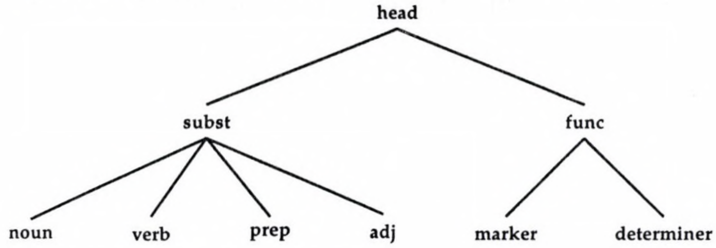

Let us conclude our brief introduction to features with an example that illustrates a partial representation of the syntactic substructure of the English personal pronoun *him*:

(4) *partial representation of* him

$$\left[\text{SYNSEM}\mid\text{LOC}\left[\begin{array}{l}\text{CAT}\left[\begin{array}{l}\text{HEAD}\left[\begin{array}{l}\textbf{noun}\\ \text{CASE } \textbf{acc}\\ \text{MOD } \bot\end{array}\right]\\ \text{SUBCAT } \langle\rangle\end{array}\right]\\ \text{CONTENT}\mid\text{INDEX}\left[\begin{array}{l}\text{PERSON } \textbf{3rd}\\ \text{NUMBER } \textbf{sing}\end{array}\right]\end{array}\right]\right]$$

The substructure taken as the value of the SUBCAT attribute — the other feature of the CAT structure — encodes the subcategorization requirements of the word including all sorts of categorial and semantic selectional information. We devote the next section to the detailed discussion of this feature.

5.1.2 Subcategorization

The SUBCAT value of a sign is in fact its valence, it specifies what sort of other signs the item should be combined with in order to be **saturated**. This is also referred to as the sign's **grammatical arguments**. Formally speaking, the value of SUBCAT is a list (ordered set, indicated with angle brackets) of partially specified **synsem** objects restricting the range of signs that combine with the item. In HPSG the term grammatical argument is much broader than in other syntactic theories in as much as they also include dependent elements usually classified as subjects or specifiers. It is important then that subjects in HPSG are not selected via the intercession of an INFL like in GB theory in clauses containing simple finite verbs.

Note that the list of **synsem** objects corresponding to the complements (in the broad sense) does not constitute an unstructured set, but is strictly ordered. Its order, however, does not reflect the surface order of the constituents corresponding to **synsem** objects, rather a sort of **obliqueness hierarchy**[1]. The example below shows the subcategorization frame of the English verb *give*, the symbols serve only explanatory purposes for the time being, the labels of the items in the subcat list are meant to encode categorial restrictions with the semantic contribution indicated in brackets.

(5) *subcategorization of the verb* give

$$\left[\text{SUBCAT}\left\langle \text{NP}_1(\text{AGENT}),\text{NP}_2(\text{THEME}),\text{PP}_3(\text{GOAL})\right\rangle\right]$$

In HPSG, there is no separate theory for case, case of nominal arguments is treated as part of the subcategorization information included in entries that select them (see a discussion in 5.3.1). In particular finite verb forms assign nominative case to their subjects (initial elements on the subcat list), while other verb forms or predicative non-verbal categories do not contain case specification for their arguments. It is apparent then that the distinction between inherent (lexical) and structural case in HPSG is not stated explicitly.

[1]This notion is very closely related to the obliqueness hierarchies familiar from other theories, e.g., the SUBJ-OBJ-OBJ2 hierarchy of Functional Control in Lexical-Functional Grammar.

In the following, we present the substructure of the lexical representation of an English verb, containing local syntactic (subcategorizational and categorial) and semantic features of the element.

(6) SYNSEM| LOC *substructure of walks*

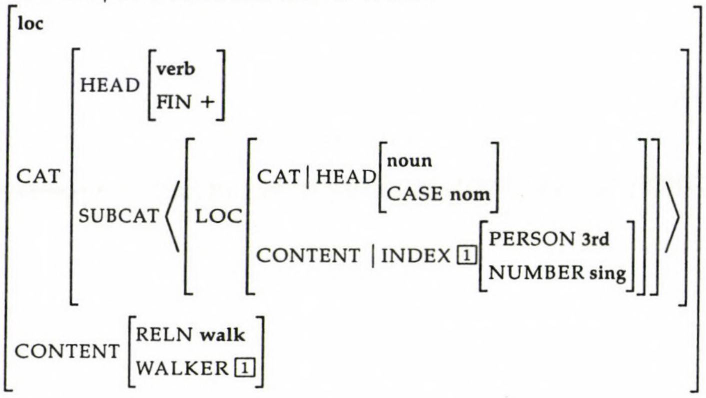

Note that the meaning of this word represented in the same formalism of attribute value matrices as the other features. Nothing really hinges on the type of semantic representation above, therefore we choose not to go into details about this. The main point is that the semantic substructures of the arguments on the SUBCAT list are linked into the representation of the item's CONTENT structure. Therefore, in standard HPSG, linking of semantic arguments to syntactic arguments (elements on the subcat list) is expressed via simple structure sharing. As this is thought to be part of a word's lexical representation, any kind of generalisations about linking can only be stated in the form of lexical (redundancy) rules[2].

5.1.3 Principles constraining phrase structure

As HPSG is not a derivational but a constraint based theory, any generalisation on phrase-structure or phrasal projections should be driven by constraints. A necessary consequence of such an approach is that phrase structure should be encoded in the phrases' representation and any restrictions on phrase structure is a constraint on (a substructure of) phrasal sorts. The constituent (immediate dominance) structure of the item is encoded in the DAUGHTERS substructure of phrases, the value of this feature is of sort varying according to the type of construction in question, i.e. the type of the immediate dominance construction, e.g., head-complement construction or head-adjunct construction. Below, we indicate the constituent structure types in standard HPSG.

(7) *Constituent structure types in standard HPSG*

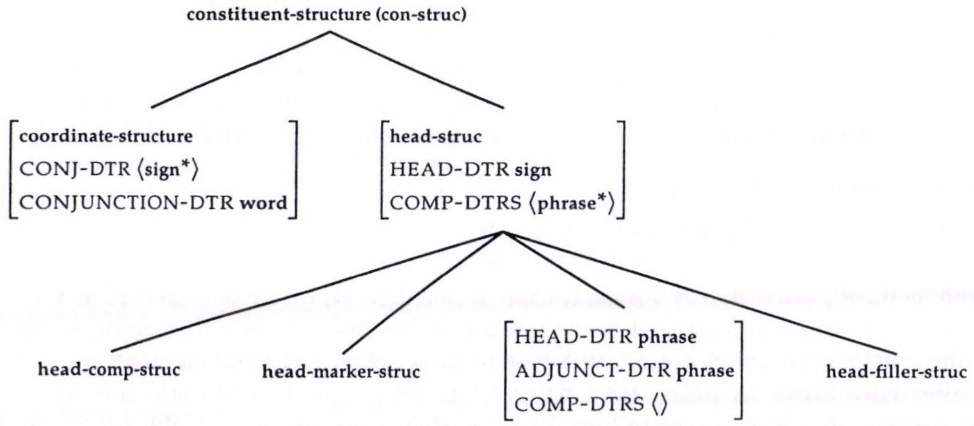

For reasons of perspicuity, however, in the examples, we indicate constituency information in a usual tree structure representation format, while keeping to the feature structure encoding

[2]Thematic roles and any generalisation on the linking of particular types of arguments to particular types of subcategorized elements can hardly be expressed in this framework.

when stating the constraints. The edges in the tree stand for attribute paths, the labels on these edges disambiguate the type of the constituent, i.e. adjunct, head, etc., corresponding to the ADJ-DAUGHTERS HEAD-DAUGHTER, etc. attributes in the DAUGHTERS substructure (these are declared in the hierarchy fragment (7)). The following two representations then are equivalent.[3]

(8) *equivalent representations of constituent structure*

a. *AVS encoding of constituent structure*

$$\left[\boxed{0}\ \text{DAUGHTERS}\begin{bmatrix}\textbf{head-comp-struc}\\ \text{HEAD-DAUGHTER}\ \boxed{1}[\ldots]\\ \text{COMP-DAUGHTERS}\left\langle \boxed{2}[\ldots], \boxed{3}[\ldots]\right\rangle\end{bmatrix}\right]$$

b. *tree-structure encoding of constituent structure*

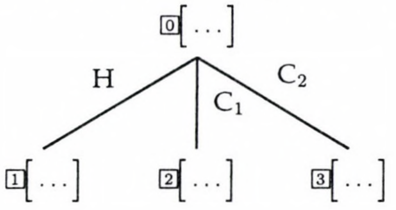

According to the ontology, then, phrases have the following general characterization:

(9) *phrases*

$$\begin{bmatrix}\textbf{phrase}\\ \text{PHON}\ \textbf{phonstring}\\ \text{SYNSEM}\begin{bmatrix}\textbf{synsem}\\ \ldots\ \text{CAT}\begin{bmatrix}\text{HEAD}\ \textbf{head}\\ \text{SUBCAT}\ \langle\textbf{synsem}^{*}\rangle\end{bmatrix}\\ \ldots\ \text{CONTENT}\ \textbf{content}\end{bmatrix}\\ \text{DTRS}\ \textbf{con-struc}\end{bmatrix}$$

HPSG's most important universal principles comprise the four principles constraining phrase structure. These are discussed in detail in the rest of this section.

The first of these four principles is a constraint that specifies which are the possible phrasal constructions available in a language. The **Immediate Dominance Principle** (IDP) says that each and every phrase must satisfy one of the Immediate Dominance Schemata. The relevant three schemata driving the projection of lexical categories are explicitly stated below:

(10) *Immediate Dominance Schemata*

a. HEAD-SUBJECT SCHEMA
The SYNSEM| LOC| CAT| SUBCAT value is ⟨⟩, and the DAUGHTERS value is of sort **head-comp-struc** whose HEAD-DAUGHTER value is a phrase whose COMP-DAUGHTERS value is a list of length one.

b. HEAD-COMPLEMENT SCHEMA
The SYNSEM| LOC | CAT| SUBCAT value is a list of length one and the DAUGHTERS value is an object of sort **head-comp-struc** whose HEAD-DAUGHTER value is of type **word**.

c. *Head-Adjunct Schema*
The phrase's DTRS value is of type **head-adj-struc** and the HEAD-DTR value is token-identical to the ADJ-DTR|SYNSEM|LOC|CAT|HEAD|MOD value.

As the first item on the subcategorization list corresponds to the subject while the others are the complements of the item, the principle says that a headed phrase is such that it has an almost saturated (a list containing one item) or a totally saturated (i.e. empty) subcategorization list. The main prediction of this restriction is that any syntactic operation that refers to an item with a different ex-

[3] Note that the sort of the DAUGHTERS value can be recovered from the labels of the edges.

tent of saturation is bound to failure. A phrase satisfying the HEAD-COMPLEMENT SCHEMA is the analogue of a bar-level projection in GB, while one satisfying the HEAD-SUBJECT SCHEMA is a the analogue of a maximal phrasal projection. Note that when embedded as complements, items are *not necessarily* subcategorized for with respect to their phrasal status — whether the sign is of type **phrase** is unrecoverable from its SYNSEM substructure —, but rather with respect to their extent of saturation. As a consequence, a word having no complements qualifies as a totally saturated sign just as a maximal phrasal projection, which entails that there is no need to postulate vacuous phrasal projections for these words unlike in X'-bar theory. Analogously, words having a subcat list of length one are free to get involved in subject-head or adjunct-head constructions without turning into a phrase beforehand. Unfortunately, this advantage is not realised in standard HPSG, where the requirement that complements should be phrases are already stated in the feature declaration of immediate constituent types (cf. (7)), and therefore one needs another ID schema to allow for (among others) vacuous phrasal projections. Note that in any case a complement is in principle allowed to be not or partially saturated, unlike in GB theory.

Now we can introduce some useful abbreviations for substructures often referred to. Note also that, from now on, in the examples we will indicate only the relevant substructure of the AVSs.

(11) *some abbreviations*

a. $\mathrm{NP}_{\boxed{1}}$ $$\begin{bmatrix} \textbf{synsem} \\ \text{LOC}\begin{bmatrix} \text{CAT}\begin{bmatrix} \text{HEAD } \textbf{noun} \\ \text{SUBCAT } \langle\rangle \end{bmatrix} \\ \text{CONTENT} \mid \text{INDEX } \boxed{1} \end{bmatrix} \end{bmatrix}$$

b. VP:$\boxed{1}$ $$\begin{bmatrix} \textbf{synsem} \\ \text{LOC}\begin{bmatrix} \text{CAT}\begin{bmatrix} \text{HEAD } \textbf{verb} \\ \text{SUBCAT } \langle \textbf{synsem} \rangle \end{bmatrix} \\ \text{CONTENT } \boxed{1} \end{bmatrix} \end{bmatrix}$$

c. S:$\boxed{1}$ $$\begin{bmatrix} \textbf{synsem} \\ \text{LOC}\begin{bmatrix} \text{CAT}\begin{bmatrix} \text{HEAD } \textbf{verb} \\ \text{SUBCAT } \langle\rangle \end{bmatrix} \\ \text{CONTENT } \boxed{1} \end{bmatrix} \end{bmatrix}$$

The ID schemata basically encode in what ways lexical items are allowed to project phrases. In the first phase (Schema 2) a lexical item is combined with a collection of other phrases yielding a phrase that has only one missing argument. In the next step the phrase can project further into a maximal projection, i.e., a category with an empty subcategorization list. Obviously one needs other principles to guarantee that arguments are saturated not only in fixed steps but also under the right conditions, that is satisfying the subcategorization requirements of th elexical heads. This is guaranteed by the **Subcategorization Principle** (SP):

(12) *The Subcategorization Principle*
In a headed phrase, the list value of the DAUGHTERS| HEAD-DAUGHTER| SYNSEM| LOC |CAT | SUBCAT (a list) is the concatenation of the list value of the SYNSEM| LOC| CAT| SUBCAT with the respectively ordered list of the SYNSEM values of the elements in the DAUGHTERS| COMPLEMENT-DAUGHTERS list.

On the one hand this principle ensures that the non-head immediate constituents of the phrase (i.e. complements) have properties complying with the subcategorization requirements of the head daughter. On the other hand, it also declares that the phrase inherits only the unsaturated part of the subcat list from the head. In other words, the complements present in the phrase are cleared off from the subcat list of the head (get saturated) and only the unsaturated part is inherited by the mother from the head. The figure below illustrates this effect:

(13) *Constraints of the SP*

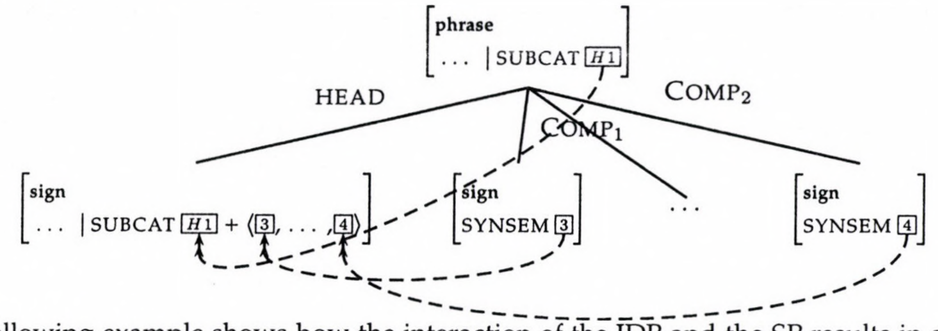

The following example shows how the interaction of the IDP and the SP results in a unique syntactic structure for the maximal phrasal projection of *give*.

(14) *Phrasal projection of* give

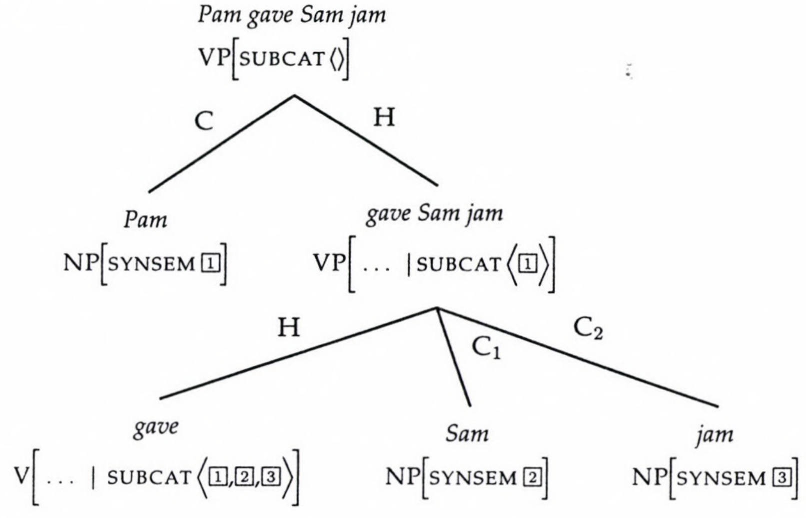

It is apparent that the constructions combining the subject and the intermediate phrasal projection, on the one hand, and the complements with the head, on the other, are generalized under the same construction type namely a DAUGHTERS value of sort **head-comp-struc**.

As the saturated complements of a head are not present on the subcat list of the phrase, and as subcategorization is necessarily strictly local, it is ensured that no item subcategorizing for a phrase can "look into its phrase structure", i.e., no subcategorization can depend on any feature of the already saturated complements of a subcategorized head. The projection level of the phrase, that is encoded in the length of the SUBCAT value, can in fact be subcategorized for as discussed below in sections 5.3.1 and 5.3.3 in connection with raising predicates.

Note also that, since, as we mentioned earlier, subcategorization is restricted to the SYNSEM substructure of arguments, a head cannot subcategorize for phonetic or constituent-structure properties.[4]

The phonetic form of phrases is controlled by the so-called **LP Principles**. Linear precedence relations (at least in English in the general case) are thought to mirror the order of members on the list value of the DAUGHTERS| COMPLEMENT-DAUGHTER attribute of the phrase, that, in turn, mirrors the order of the elements specified in the SUBCAT list of the head.[5] For the discussion of the linear order of constituents and word order variation as well as of problems concerning discontinuous constituency, we refer the reader to Uszkoreit (1987),Pollard (1991) and Reape (1994).

We noted earlier that the other feature within the categorial information is the collection of features that gets inherited by the phrasal projections. The universal principle ensuring HEAD feature percolation is referred to as the **Head Feature Principle**:

---

[4]This also entails that the linear order of the constituents of a phrase cannot in any way be altered or disrupted by the selecting head when it is embedded as its complement.

[5]Note the rather bizarre situation that the substructure encoding immediate dominance relations and not intended to encode linear precedence information still does exploit a list structure representation.

(15) *Head Feature Principle*
In a headed phrase the values of SYNSEM| LOC| CAT| HEAD and DAUGHTERS| HEAD-DAUGHTER| SYNSEM| LOC| CAT| HEAD are token-identical.

We mention here that there are other ID schemata — not discussed in this introductory survey — including, among others, one that is responsible for generating inverted structures such as English interrogative sentences. Finally we present the fourth universal principle that controls the way semantics of phrases are composed.

(16) *Semantics Principle*
The SYNSEM| LOC| CONTENT value is token-identical with that of DAUGHTERS| ADJ-DAUGHTER| LOC | CONTENT if the DAUGHTERS value is **head-adj-struc** and with that of DAUGHTERS| HEAD-DAUGHTER| LOC| CONTENT otherwise.

This principle basically ensures that the CONTENT value of phrases is inherited from the semantic head of the phrase (the semantically active functor), i.e., the adjunct-daughter in the case of an adjunct-head construction and the head-daughter otherwise.

(17) *semantic composition of adjunct structures*

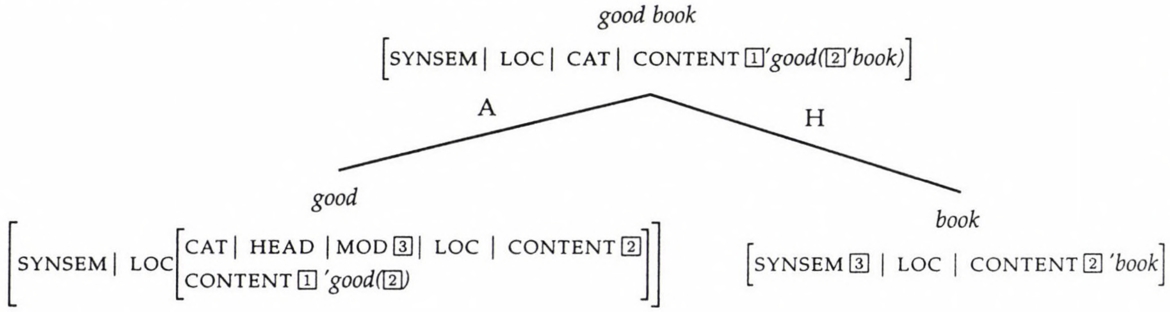

(18) *semantic composition of head-complement structures*

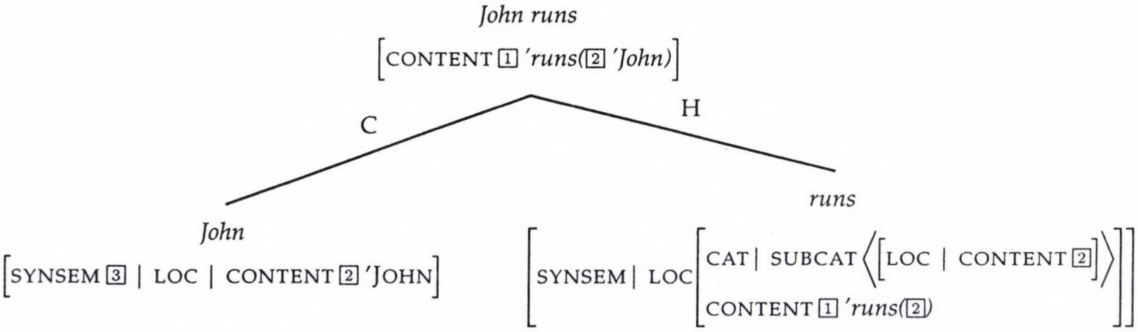

As noted earlier, subject/head and complement/head relations are treated under the same DAUGHTERS value. Since the Semantics Principle is only sensitive to this value, one consequence of this is that the subject has a semantic contribution that is by no means more than its contribution to the semantics of the verb phrase. This, in general, means a shared value between the arguments CONTENT attribute and some feature of the verb's CONTENT.

Pollard and Sag (1994, pp342-43) gives a brief indication of what an HPSG linking theory would look like. For a more comprehensive attempt, we refer the interested reader to Wechsler (1991).

5.2 Subcategorization and valence in Modern HPSG

5.2.1 Valence Features and argument structure

Pollard and Sag (1994, chapter 9) argues for an extended version of HPSG, in which subcategorization features are treated differently from the "standard" approach characterized in the first section.

The essence of the new idea (coming from R. Borsley (1987; 1989a) is to split subcategorization features into three separate features called the **valence features.**

The separation of the subcat list into distinct SUBJ and COMPS lists is motivated by the clumsy treatment of non-predicative prepositions in standard HPSG. These prepositions, if they subcategorize for a complement, always realize it after the head. If, however, there is only one list containing all arguments of a head, then this complement is basically indistinguishable from the subject like in other categories. In order to disallow the wrong order of constituents, prepositions were lexically specified as inverting and formed phrases with the help of the ID schema for inverted subject constructions. With a separate feature for subjects, it is possible to lexically specify non-predicative prepositions with an empty SUBJ list, and hence ensure a true complement status for their arguments.

With the distinction of subjects and complements, it is possible to realize the apparent diversities in the behavior of complements and specifiers. This prompts Pollard and Sag to introduce a third valence feature called SPR for specifiers. Note that, though formally the SUBJ and SPR values are of type **list** the theory constrains them to be lists containing at most one element.

Obviously, for these new features to work, it is necessary to introduce new ID schemas as well as new construction types (i.e., new types of DAUGHTERS values, namely **head-subj-struc** and **head-spec-struc**). The Subcategorization Principle also has to be replaced with the **Valence Principle**.

(19) *Valence Principle*
In a headed phrase, for each valence feature F, the F value of the head daughter is the concatenation of the phrase's F value with the list of SYNSEM values of the F-DAUGHTERS value.

The above analysis is also motivated by the fact that Welsh (cf. Borsley (1989b)) finite clauses exhibiting a systematic VSO pattern are argued to be best analyzed with ID Schema 2 (or actually the new version requiring the phrase to have an empty COMPS list). This requires, however, the "subject" (least oblique argument) to be specified as the first element on the COMPS list as well as the SUBJ value to be the empty list in the lexical representation of finite verb forms.

This split view on valence has been taken over in Modern HPSG and much recent work. Counter to Borsley's original intention, however, the valence features did not replace the old SUBCAT list but came as additional features of signs. It is only the valence features that get saturated in the course of the projection of the phrase, the subcat list remains the same (actually due to structure sharing its members get specified). Thereby the meaning of subcat list is much closer to conventional conceptions of argument structure, which is reflected by its new name ARG-S.

5.2.2 Mapping argument lists to valence values

Although in Pollard and Sag (1994, chapter 9.) valence features an the subcat list are present in a lexical representation, there, subcat lists are always thought to be the concatenation of the list values of the valence features. Despite apparent redundancy, the presence of both valence and argument-structure specification can be defended on grounds of the existence of alternative mappings from ARG-S (SUBCAT) to valence features (cf. Bouma (1997?) for possible generalizations on these mappings). One such alternative, in fact incomplete, mapping can be the case of pro-drop as argued in Manning and Sag (1995). The main argument in favor of the duplicate representation, however, is that the Binding Theory (cf. Pollard and Sag (1994, chapter 6)) is defined with the help of the **obliqueness hierarchy**, i.e. the order of elements on the ARG-S list. Manning and Sag (1995) argues that binding relations remain constant as defined by ARG-S irrespective of the different mappings to valence values and surface order. They also argue that assuming different structures makes it possible to state parametric variation of causatives across languages.

5.3 Complementation Patterns

5.3.1 Unsaturated complements

In the first section we saw that subcategorization can refer to the level of saturation of the selected complement. The full power of this possibility is seen in the HPSG treatment of raising and control predicates. In GB theory usual Subject-to-subject raising predicates (SSRPs) are thought to select for an infinitival IP (a clause) the subject of which moves out of the embedded clause (e.g., to matrix subject position) to receive case. The analogue of this movement is the structure sharing of the first elements on the subcat list of the predicate and its selected VP complement. The range, however, of raising in HPSG is broader than in GB in as much as it includes predicates enabling exceptional case marking constructions, i.e., allow the subject of their infinitival clause complement to be assigned accusative case. Pollard and Sag (1994, chapter 3.) lists a great deal of arguments in favor of a Subject-to-object raising analysis of these predicates (SORPs).

(20) *Raising predicates*

a. *Lexical representation of a Subject-to-subject raising predicate*

seem

$$\begin{bmatrix} \text{CAT} & \begin{bmatrix}\text{HEAD } \textbf{verb} \\ \text{SUBCAT } \langle \boxed{1}, \text{VP}[\text{SUBCAT } \langle \boxed{1}\text{NP}\rangle]:\boxed{2}\rangle \end{bmatrix} \\ \text{CONTENT} & \begin{bmatrix} \text{RELN } \textbf{seem} \\ \text{ARG1 } \boxed{2}\end{bmatrix}\end{bmatrix}$$

b. *Lexical representation of a Subject-to-object raising predicate*

believe

$$\begin{bmatrix} \text{CAT} & \begin{bmatrix}\text{HEAD } \textbf{verb} \\ \text{SUBCAT } \langle \text{NP}_{\boxed{1}}, \boxed{2}, \text{VP}[\text{SUBCAT } \langle \boxed{2}\text{NP}\rangle]:\boxed{3}\rangle \end{bmatrix} \\ \text{CONTENT} & \begin{bmatrix} \text{RELN } \textbf{believe} \\ \text{ARG1 } \boxed{1} \\ \text{ARG2 } \boxed{3}\end{bmatrix}\end{bmatrix}$$

Note that subject and object control (or equi) predicates only differ from SSRPs and SORPs, respectively in that, unlike the latter, they assign a semantic role to their own arguments.

(21) *Representation of control predicates*

a. *A subject-control predicate*

try

$$\begin{bmatrix} \text{CAT} & \begin{bmatrix}\text{HEAD } \textbf{verb} \\ \text{SUBCAT } \langle \text{NP}_{\boxed{1}}, \text{VP}[\text{SUBCAT } \langle \text{NP}_{\boxed{1}}\rangle] : \boxed{2}\rangle \end{bmatrix} \\ \text{CONTENT} & \begin{bmatrix} \text{RELN } \textbf{try} \\ \text{ARG1 } \boxed{1} \\ \text{ARG2 } \boxed{2}\end{bmatrix}\end{bmatrix}$$

b. *An object-control predicate*

persuade

$$\begin{bmatrix} \text{CAT} & \begin{bmatrix}\text{HEAD } \textbf{verb} \\ \text{SUBCAT } \langle \text{NP}_{\boxed{1}}, \text{NP}_{\boxed{2}}, \text{VP}[\text{SUBCAT } \langle \text{NP}_{\boxed{2}}\rangle] :\boxed{3}\rangle \end{bmatrix} \\ \text{CONTENT} & \begin{bmatrix} \text{RELN } \textbf{persuade} \\ \text{ARG1 } \boxed{1} \\ \text{ARG2 } \boxed{2} \\ \text{ARG3 } \boxed{3}\end{bmatrix}\end{bmatrix}$$

Raising and control do not only differ with respect to whether the embedded subject's index is reentrant with a matrix thematic role, but also in that there is a different kind of structure sharing in the two cases. In particular, raising predicates exhibit **strong connectivity**, i.e. the whole SYNSEM substructure is shared, while in the case of SORPs, only the indices are shared between the arguments (**weak connectivity**). Note that in the case of weak connectivity information outside the referential index of the NP need not be the same or should be duplicated in the lexical representation by pure stipulation. Such a distinction could be motivation if sharing of the whole synsem structure would result in a type mismatch. And indeed, object control cases with strong connectivity would yield a case mismatch between the matrix object and the embedded subject, since case is represented as a feature within the synsem substructure (in particular as the value of SYNSEM|LOC|CAT|HEAD|CASE) . This is a consequence of the lack of distinction between inherent and structural case, i.e. the structural case of Noun phrases is not determined relative to their constituent structure position but assigned lexically. (See Heinz and Matiasek (1994) for a comprehensive discussion of case assignment in HPSG.)

A more serious problem is the fact that based on the above considerations the weak and strong connectivity representation of shared arguments would motivate a parallel treatment of subject to subject raising and subjects control (the case when there can be no case mismatch) on the one hand, and subject to object raising and object control on the other, as opposed to the above.

Control predicates selecting infinitival complements typically include many that allow *for to* infinitive clauses as well as subjectless infinitives. Unlike GB theory, it is not true that one has to assume a null subject in the subjectless case in order to regard these as a natural class. Pollard and Sag (1994, pp125-127) stipulates that these predicates subcategorize for a *marked* infinitival VP or S, and that *for* marks saturated complements while a zero marker marks VPs. The problem is that one does not gain much from this stipulation given that a disjunctive specification is needed for these entries, anyway. Firstly, the categories VP and S cannot be generalized over, secondly the semantics in the two cases is different in terms of the sentential complement does not trigger structure sharing of indices. we refer the reader to Pollard and Sag (1994, chapter 8) for a detailed discussion and a different treatment of complement control.

The former problem is immediately solved in extended HPSG, where one can easily characterize S and VP as a natural class with the feature specification [COMPS ⟨⟩]. The advantage of this seems to be rather tiny in the light of the fact that the VFORM value is idiosyncratically constrained in all instances of VP selection, and will hardly form a natural class with the finite VFORM value.

It is also important to emphasize that HPSG, following Gazdar, K.Pullum, and Sag (1982), treats auxiliaries as verbs. In particular they are typical instances of subject raising predicates. The infinitival marker *to* is also treated the same way, only it totally inherits the CONTENT value of the subcategorized VP. In this respect, it parallels non-predicative prepositions that are treated to denote only an index that they inherit from their NP complements.

(22) a. *Infinitive particle as an underspecified auxiliary*

to

$$\begin{bmatrix} \text{CAT} \begin{bmatrix} \text{HEAD} \begin{bmatrix} \text{VFORM } \textbf{inf} \\ \text{AUX} + \end{bmatrix} \\ \text{SUBCAT } \langle \boxed{1}, \text{VP}[\textbf{base}, \text{SUBCAT}\langle \boxed{1} \rangle]:\boxed{2} \rangle \end{bmatrix} \\ \text{CONTENT } \boxed{2} \end{bmatrix}$$

b. *Non-predicative preposition to*

to

$$\begin{bmatrix} \text{CAT} \mid \text{HEAD } \textbf{prep} \\ \text{SUBCAT } \langle \text{NP}[\textbf{acc}, \text{CONTENT } \boxed{1}] \rangle \\ \text{CONTENT } \boxed{1} \end{bmatrix}$$

5.3.2 Passive

Passive as a Lexical Rule

Following Pollard and Sag (1987), HPSG treats passive as relational. Given the main principles of subcategorization explained above, one can easily characterize the subcategorization patterns of some English verbs and their passive counterparts.

(23) *Subcategorization of some verbs and their passive counterparts*

verb	*active*	*passive*
read	SUBCAT $\langle NP_1, NP_2 \rangle$	SUBCAT $\langle NP_2, PP[by]_1 \rangle$
give	SUBCAT $\langle NP_1, NP_2, PP[to]_3 \rangle$	SUBCAT $\langle NP_2, PP[to]_3, PP[by]_1 \rangle$
give	SUBCAT $\langle NP_1, NP_2, NP_3 \rangle$	SUBCAT $\langle NP_2, NP_3, PP[by]_1 \rangle$
promise	SUBCAT$\langle NP_1, NP_2, VP[inf] \rangle$	SUBCAT $\langle NP_2, VP[inf]_3, PP[by]_1 \rangle$

The parallel between the two can be captured with a lexical rule, which is still the standard HPSG treatment of voice alternation.

(24) *The Passive Lexical Rule*

$$\begin{bmatrix} \textbf{word} \\ \text{SYNSEM}\,|\,\text{LOC} \begin{bmatrix} \text{CAT} \begin{bmatrix} \text{HEAD} \begin{bmatrix} \textbf{verb} \\ \text{VFORM } \textbf{active} \end{bmatrix} \\ \text{SUBCAT } \langle NP_{\boxed{1}}, NP_{\boxed{2}} \rangle \oplus \boxed{3} \end{bmatrix} \\ \text{CONTENT } \boxed{3} \end{bmatrix} \end{bmatrix}$$

$$\downarrow$$

$$\begin{bmatrix} \textbf{word} \\ \text{SYNSEM}\,|\,\text{LOC} \begin{bmatrix} \text{CAT} \begin{bmatrix} \text{HEAD} \begin{bmatrix} \textbf{verb} \\ \text{VFORM } \textbf{passive} \end{bmatrix} \\ \text{SUBCAT } \langle NP_{\boxed{2}} \rangle \oplus \boxed{3} \oplus \langle PP[\text{BY}]_{\boxed{1}} \rangle \end{bmatrix} \\ \text{CONTENT } \boxed{3} \end{bmatrix} \end{bmatrix}$$

Surprisingly enough, though alternative mapping strategies between argument structure and valence *can* be expressed in current HPSG, passive is not treated this way. Keeping to the unchanged idea of the relational passive, in modern HPSG, a passive verb receives the standard representation of a derived entry (cf. Meurers (1995)). The nested complex representation and the presence of the old argument structure is thought of as a virtue rather than a vice. As reported by Manning and Sag (1995), in a lot of languages, such as Russian, binding relations are different in a passive and active sentence, in as much as in the passive an oblique complement can be bound by either the surface or the logical subject. This is treated with a disjunctive interpretation of the binding theory, i.e. its conditions can apply to either the derived or the embedded argument structure. Here STEM represents the stem of a derived word (complex predicate).

(25) *Manning and Sag's Universal characterization of passive with "nested"* ARG-S

$$\begin{bmatrix} \ldots \begin{bmatrix} \text{HEAD}\,|\,\text{VFORM } \textbf{passive} \\ \text{ARG-S } \langle \boxed{3} \rangle \oplus \boxed{4} \; (\oplus \langle \boxed{6}_{\boxed{1}} \rangle) \end{bmatrix} \\ \ldots \text{ CONTENT } \boxed{5} \\ \text{STEM} \begin{bmatrix} \ldots \begin{bmatrix} \text{HEAD}\,|\,\text{VFORM } \textbf{active} \\ \text{ARG-S } \langle \boxed{2}_{\boxed{1}}, \boxed{3} \rangle \oplus \boxed{4} \end{bmatrix} \\ \ldots \text{ CONTENT } \boxed{5} \end{bmatrix} \end{bmatrix}$$

Auxiliary Passive as Raising

An alternative treatment of passive is found in Kathol (1994), where auxiliary passives are treated on a par with other raising constructions. There, it is assumed that the passive auxiliary subcategorizes for an unsaturated verb, whose arguments are raised and rearranged in the auxiliary entry's argument structure. Note that German dative or impersonal passives can be treated similarly (for an alternative unified treatment of German passives see Pollard (1994)).

(26) *Kathol's representation of the German passive auxiliary*

werden

$$\left[\begin{array}{l} \text{SUBJ}\ \langle \text{NP[nom]}_{\boxed{2}}\rangle \\ \\ \text{COMPS}\ \boxed{1} \oplus \left\langle \text{V}\left[\begin{array}{l}\text{VFORM}\ \textbf{part ii}\\ \text{SUBJ}\ \langle\text{NP}\rangle \\ \text{COMPS}\ \langle \text{NP[acc]}_{\boxed{2}}\rangle \oplus \boxed{1}\end{array}\right]\right\rangle \end{array}\right]$$

In fact, binding conditions in English and German seem to question the universal treatment of passive as in (24), given that there is no alternative for the passive agent to bind an oblique complement. If, however, one adopts Kathol's treatment of the passive in terms of inheriting valence features but still keep the ARG-S feature, one is forced to say that it mirrors the surface relations, i.e., it is the concatenation of the valence lists.

Note that, just like in the original treatment of Subject-to-Object raising and control predicates, there is only weak connectivity between the inherited NPs (see section 5.2.2).

5.3.3 Subcategorization Inheritance

Except for the above treatment of passive, so far, the discussion of unsaturated complements was restricted to intermediate projections, i.e., phrases with a subcat list of length one. The possibility of generalizing over raising is explored in recent work on **subcategorization inheritance**. Gerdeman (1994) examines the idea of function composition in categorial grammar (cf. Moortgat (1985)) and shows in what ways an HPSG treatment of subcat list inheritance can accommodate better to the requirements of morphology and syntax.

The idea of subcategorization inheritance originates from Hinrichs and Nakazawa (1989a) and Hinrichs and Nakazawa (1994), who apply it to German auxiliaries. In morphology a similar problem arises with the subcategorization of derived words (for this treatment, see Krieger and Nerbonne (1991)). It is seemingly a conflict with HPSG principles, that, while one would choose to treat the affix as the head of a derived entry, subcategorization information clearly parallels that of the stem. Here it would be rather problematic to say that the affix subcategorizes for a saturated phrase because of apparent case mismatches (i.e., English nominalized verbs, for instance, typically realize their argument in genitive case). With general subcategorization inheritance, however, the affix can without problem be treated as the head (morphological and semantic) that subcategorizes for an unsaturated complement and the rest of its subcat list is reentrant with that of the complement.

(27) *Parallel representations for affix and auxiliary*

a. *representation of German the modal auxiliary* kann

kann

$$\left[\begin{array}{l}\text{HEAD}\ \textbf{verb}\\ \text{SUBCAT}\ \langle \text{VP[SUBCAT}\ \boxed{1}\text{]},\boxed{1}\rangle\end{array}\right]$$

b. *representation of the English productive affix* -ness

-ness

$$\left[\begin{array}{l}\text{HEAD}\ \textbf{noun}\\ \text{SUBCAT}\ \langle \text{AP[SUBCAT}\ \boxed{1}\text{]},\boxed{1}\rangle\end{array}\right]$$

This is rather problematic if we persue the old conception of subcat list here. As, in the course of phrasal projection the subcat list gets saturated and as there is no way to select for a word not a phrase, one can never guarantee that *all* the complements are inherited. This is only possible if one has a feature that preserves lexical arguments throughout the projection. But this is still not sufficient, one also needs another feature that gets saturated (valence) and the structure sharing of which with the SUBCAT or ARG-S can guarantee that all complements are there. If, however, we want to enable alternative mappings between argument structure and valence, the checking of word status — besides being rather complicated if not impossible — requires that the inheriting lexical item knows about all available mappings of the word subcategorized for. Thus we must see that enabling checking for exhaustive subcat lists of subcategorized elements has serious consequences in as much as we loose the nice restrictedness of subcategorization explained above.[6]

A straightforward solution to this problem is to introduce distinct rules combining words and phrases — though we remark that the obvious generalization of the two type of rules is not easily expressible in the current framework of HPSG.

5.4 Further reading

An alternative version of HPSG theory is developped to handle complex predicates in the treatment of verbal constructions, such as causatives, passive, verbal particles, tense and aspect constructions, in Ackerman and Webelhuth (1998).

Throughout our survey, we concentrated only on local features of items though we argued that subcategorization is stated for synsem objects, which include non-local features. For a justification of lexicalised non-local feature selection see the treatment of *easy/tough* constructions in (Flickinger and Nerbonne, 1992) and (Pollard and Sag, 1994, chapter 4.,pp.166-171).

On the phenomenon of complement extraposition, we refer the reader to (van Eynde, 1996; Bouma, 1996; Keller, 1995)

Krenn and Erbach (1994) develops an HPSG account of idioms and some support verb constructions.

Problems with the treatment of adjuncts and complements in HPSG are discussed in (van Noord and Bouma, 1994; Verspoor, 1996; Kasper, 1994)

References

Ackerman, Farrell and Gert Webelhuth. 1998. *A Theory of Predicates*. CSLI Publications, Stanford.

Borsley, Robert. 1987. Subjects and complements in HPSG. Technical report, Technical Report no. CSLI-107-87, Stanford,CSLI.

Borsley, Robert. 1989a. Phrase-structure grammar and the *barriers* conception of clause structure. *Linguistics*, pages 27:843–863.

Borsley, Robert D. 1989b. An HPSG approach to Welsh. *Journal of Linguistics*, pages 25:333–354.

Bouma, Gosse. 1996. Complement clauses and expletives. *Complutational Linguistics in the Netherlands*.

Bouma, Gosse. 1997? Valence alternation without lexical rules. unpublished manuscript.

Flickinger, Daniel and John Nerbonne. 1992. Inheritance and complementation: a case study of *easy* adjectives and realted nouns. *Computational Linguistics*, 19(3):260–309.

Gazdar, Gerald, Geoffrey K.Pullum, and Ivan A. Sag. 1982. Auxiliaries and related phenomena in a restricted theory of grammar. *Language*, pages 58:591–638.

Gerdeman, Dale. 1994. Complement inheritance as subcategorisation inheritance. In Nerbonne et al. (Nerbonne, Pollard, and Netter, 1994), pages 341–364.

[6]It might suffice to guarantee that the subcat list of the item has a fixed number of elements. This, however, would imply that if an item can take a word that is inherently specified for the empty subcat-list, then it will also take a totally saturated phrase with similar other features, which is clearly not viable in the case of derivational suffices.

Heinz, Wolfgang and Johannes Matiasek. 1994. Argument structure and case assignment in German. In Nerbonne et al. (Nerbonne, Pollard, and Netter, 1994), pages 199–236.

Hinrichs, Erhard and Tsuneko Nakazawa. 1989a. Flipped out: AUX in German. In *Papers from the 25th Annual Regional Meeting of the Chicago Linguistic Society*. Chicago Linguistic Society, pages 187–202.

Hinrichs, Erhard and Tsuneko Nakazawa. 1989b. Subcategorization and VP structure in German. In *Proccedings of the Third Symposium on Germanic Linguistics*.

Hinrichs, Erhard and Tsuneko Nakazawa. 1994. Linearizing AUXs in German verbal complexes. In Nerbonne et al. (Nerbonne, Pollard, and Netter, 1994), pages 11–38.

Kasper, Robert. 1994. Adjuncts in the mittelfeld. In Nerbonne et al. (Nerbonne, Pollard, and Netter, 1994), pages 39–70.

Kathol, Andreas. 1994. Passive without lexical rules. In John Nerbonne, Carl Pollard, and Klaus Netter, editors, *German in HPSG*. CSLI Publications, pages 237–272.

Keller, Frank. 1995. Towards an account of extraposition in HPSG. In *Proceedings of the EACL, Dublin*.

Krenn, Brigitte and Grefor Erbach. 1994. Idioms and support verb constructions. In Nerbonne et al. (Nerbonne, Pollard, and Netter, 1994), pages 365–396.

Krieger, Hans-Ulrixh and John Nerbonne. 1991. Feature-based inheritance networks for computational lexicons. Technical report, Technical Report RR-91-31, DFKI.

Manning, Christopher and Ivan Sag. 1995. Dissociations between argument structure and grammatical relations. Working Draft.

Meurers, Walt Detmar. 1995. Towards a semantics for lexical rules in HPSG. In *Proceedinngs of the Conference on Formal Grammar*. Barcelona.

Moortgat, Michael. 1985. Functional compostition and complement inheritance. In Geer Hoppenbrouwer, editor, *Meaning and the Lexicon*. Foris, Dordrecht.

Nerbonne, John, Carl Pollard, and Klaus Netter, editors. 1994. *German in HPSG*. CSLI Publications.

Pollard, Carl. 1991. On head non-movement. In *Tilburg Symposium on Discontinuous Constituency*. University of Tilburg. also in: Arthur Horck and Wietske Sijtsma eds., *Discontinuous Constituency*. The Hague: Mouton.

Pollard, Carl. 1994. Toward a unified account of passive in German. In Nerbonne et al. (Nerbonne, Pollard, and Netter, 1994), pages 237–272.

Pollard, Carl and Ivan Sag. 1987. *Information-based Syntax and Semantics, Volume 1: Fundamentals*. CSLI Lecture Notes no. 13. CSLI Publications,distributed by the University of Chicago Press.

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

Reape, Mike. 1994. Domain union and word order variation in German. In Nerbonne et al. (Nerbonne, Pollard, and Netter, 1994), pages 151–198.

Sag, Ivan and Carl Pollard. 1989. Subcategorization and HPSG. In Mark R. Baltin and Anthony S. Kroch, editors, *Alternative Conceptions of Phrase Structure*. University of Chicago Press, pages 139–181.

Uszkoreit, Hans. 1987. *Word Order and Discountinuous Constituents in German*. CSLI Lecture Notes No. 8. CSLI Publications.

van Eynde, Frank. 1996. An hpsg treatment of "it" extrapositions without lexical rules. In *EACL???*

van Noord, Gert and Gosse Bouma. 1994. Adjuncts and the processing of lexical rules. In *Proceeding of the 15th International conference on Computational Linguistics (COLING)*, pages 250–256.

Verspoor, Cornelia M. 1996. A perspective on PPs. In C Grover and E. Vallduvi, editors, *Edinburgh Working Papers in Cognitive Science, Studies in HPSG*. University of Edinburgh, pages 229–271.

Wechsler, S. 1991. *Argument Structure and Linking*. Ph.D. thesis, Stanford University.

Chapter 6

Argument Structure in Construction Grammar

6.1 Introduction

Construction Grammars — reviving ideas implicit in traditional linguistics — were born as a reaction to the work of linguists following the Chomskian tradition. The main point of divergence is the fact that people working in the former framework do not except the clearcut division between lexical, syntactical and (suprasegmental) phonological information that is of central importance in the latter theories. Their main claim is that all these types of information have the same structure and therefor any division is an unwarranted theoretical construct. In what follows we will outline the consequences of this approach for the treatment of the argument structure of verbs.

Section 6.2 is a brief introduction to Construction Grammars in general and it will present some of the arguments brought in favour of this approach. Then in section 6.3 we turn to the problem of the argument structure of verbs within these theories. We conclude this paper with an example from English in section 6.4.

6.2 Constructions

To highlight the problems that led to the birth of **Construction Grammars** (usually associated with (Fillmore and Kay, 1993)) let us first consider the case of the lexicon as it was conceived of in post-Chomskian linguistics. It was assumed that the lexicon consists of a list of words associated with all the information necessary, i.e., syntactic, semantic and phonological information which was considered relevant for the working of the latter components of grammar. One such type of information was the argument structure especially in the case of verbs.

But it was further generally assumed that the lexicon is also the depository of idiomatic information. But since the lexicon consisted of a list of words and the Lexical Integrity Hypothesis determined what it means to be a word, namely something that enters syntax as a unit, there was a problem. Under these assumptions idiomatic constructions such as *kick the bucket* — which consists clearly of several words —had to be introduced into the lexicon as a special case of the lexeme *kick* (say $\texttt{kick}_7$) which has all the properties of its base verb except that its argument is fixed (namely it must be *the bucket* and its meaning is a one place predicate (*to die*) rather than a relation. Along similar lines we can describe the construction *bite to dust* which has the further peculiarity that even the argument structure of the hypothetical lexeme that is responsible for the construction is different from the argument structure of the verb of which it is supposed to be a special case, i.e., there is very little they have in common.

But this does not only apply to real idioms. Consider the case of a semi-idiomatic construction like the one in the following sentence:

(1) I have been to London before.

It would be very difficult to find the word that is to "blame" for this sentence, although the productive rules of the grammar would produce the sentence with the preposition *in* which is clearly ungrammatical. To account for this phenomenon of blocking we must assume that the relevant information is present somewhere in the grammar. We will return to further such phenomena as we describe to problem of argument structure later on.

It is only natural that these kinds of phenomena lead to the assumption that the lexicon contains not only words but also (semi-)idiomatic constructions, i.e., entities that consist of possibly several words but their meaning is idiomatic. But what is special about idioms? The fact that either their form or their meaning is unpredictable from the form or the meaning of their component parts. But we can take this argument further and we can say that the same applies to syntactic rules: they put certain formal requirements on their arguments and the resulting structure and they are associated with a semantics of their own, i.e., they are also only a case of a form–meaning pair, and so they are not different from lexical constructions. The same applies to suprasegmental phonological patterns which are of the same kind: a form associated with a meaning.

This reasoning lead to the following definition of constructions ((Goldberg, 1995, p4)):

(2) **Constructions**
C is a CONSTRUCTION iff$_{def}$ C is a form-meaning pair $< F_i, S_i >$ such that some aspect of F_i or some aspect of S_i is not strictly predictable from C's component parts or some other previously established constructions.

This helps to solve — or rather to dissolve — problems that pose a serious challenge to component grammarians, such as the problem of preverbs in Hungarian . The problem is the following: how can we treat verb–preverb combinations in Hungarian when the preverb can move quite freely in syntax yet semantically it forms often a unit with its verb. They often behave like an adverbial: *be*(`in`) + *megy*(`go`) = `go-in`; but very often they form an idiomatic unit with the verb: *be* + *rúg*(`kick`) = `get drunk`. Although the preverbs show a peculiar behaviour that seems to point to the fact that they are not mere adverbials (e.g., they cannot act as topics in neutral sentences), these properties are shared by certain adverbials that were traditionally classified with preverbs in traditional accounts. But now these problems can readily be solved if we accept that preverbs are syntactically in a special class of adverbials, and idiomatic verb–preverb constructions are construction that consist of two words. This is in accordance with the nature of constructions, namely that everything not compositional on its parts forms a construction. The question whether they are lexical does simply make no sense in this framework.

As a consequence we loose the clearcut division between different components of the grammar on the one hand, and on the other hand generation and parsing become similar tasks, namely identifying the constructions that play a role in the input representation and finding the structure(s) that satisfy the constraints imposed by the other component of the relevant construction(s).

A further advantage is the fact that the same element can now be part of several different constructions. To see the importance of this fact consider the following Hungarian examples:

(3) a. ˈMari ˈlátta ˈJánost.
Mary saw John-ACC
'Mary saw John'
b. ˈLátta Jánost Mari.
ˈLátta Mari Jánost.
Mary did see John.
c. ˈMari látta Jánost.
It was Mary who saw John.
d. ˈJánost látta Mari.
It was John who Mary saw.
e. ˈJánost ˈMari látta.
As for John, it was Mary who saw him.
f. ˈJánost ˈlátta Mari.
As for John, Mary did see him.

All the above sentences describe the same situation, namely the fact that *John saw Mary*. But they differ in their topic-focus articulation, i.e., in the distribution of old and new information in the sentence with respect to the discourse. Instead of postulating movements of elements into hypothesized Topic and Focus positions which are supposed to take care of the word order variation we can say that there are two types of constructions at work in these Hungarian examples, namely the Subject-Predicate- and the Transitive-VP–Constructions, which take care of the predicate–argument relations in the sentence, and the Topic– and Focus–Constructions, which take care of word order in Hungarian. Since constructions express constraints on syntactic and semantic components, it is no wonder that a word can act in several constructions.

The above example shows two further features of Construction Grammars to which we want to call attention. First, as shown by examples c. and d., certain sentences have two forms with the same meaning. This is due to the fact that the constraints posed by the relevant constructions may not fully determine the word order, they can result in underspecification. The relevant constructions only require that the sentence be headed by the verb and thus they do not constrain the order of the subject and the object. This fact is responsible for the free variation rather than some "stylistic movement" or "scrambling".

Comparing the two examples there is one more thing to note, namely the fact that the two sets of sentences are the same except for their intonation pattern. This shows that the suprasegmental phonology, i.e., intonation pattern are parts of construction or even constructions on their own thus blurring a further distinction that is unnecessary in Construction Grammar. It is not a special phonological component that is responsible for those pattern, they are just identified as any other construction in generation and parsing.

A further important methodological principle of Construction Grammar is the Principle of No Synonymy. This principle states that there can be no to constructions in the language that are completely synonymous. I.e., no two construction can carry the same syntactic and semantic information (constraints) about their constituents; as we will see later, even in the case of constructions that result from the application of "transformations" of argument structures of different constructions, and which thus may end up being F(ormally)-synonymous and S(emantically)-synonymous, there must be at least a difference in their pragmatics — technically speaking, they must not be P(ragmatically)-synonymous.

6.3 Argument Structure

6.3.1 Problems

The following section is an introduction to the treatment of the argument structure of verbs in the framework of Construction Grammar. Before we turn to the actual problem we must make one important remark. As we said in the preceding chapter, Construction Grammar formalisms are non–transformational and it makes not really sense to speak of components of the grammar. But as in the preceding section we used the pretheoretical notion of syntactic positions, similarly in what follows we will use the terms lexicon and lexical in some intuitive sense which has however no theoretical significance. The phenomena described below are more thoroughly discussed in (Goldberg, 1995).

Whatever is the representation of verbs in any grammar formalism, it certainly must make reference to the argument structure of the verb. But what and how should be represented? Should we represent both the syntactic and the semantic argument structure of the verb or is it enough to specify the latter since it determines the former in accordance with universal mapping rules? And is it enough to represent the cases of the arguments — in the sense of (Fillmore, 1968) — or do we need a more alaborate representation? In what follows we will present the answers Construction Grammars gave to the above questions.

First let us start with some observations. Let us take a look at the following pair of sentences:

(4) a. John eats.
  b. *John devours.

The fact that verbs with roughly the same semantics come with a different syntactic argument structure seems to indicate that we must also explicitly specify the syntactic argument structure.

On the other hand there are verbs which can occur with several argument structures such as the following pairs of sentences:

(5) a. I brought Pat a glass of water.
b. I brought a glass of water to Pat.

But the two argument structures differ in their semantics in that the first requires its object to be animate whereas the second doesn't as shown by the following pair of sentences[1]:

(6) a. *I brought the table a glass of water.
b. I brought a glass of water to the table.

This is a very common phenomenon with verbs that have several but related argument structures. To give just one further example let us take a look at the following pair of sentences:

(7) a. I loaded the truck with hay.
b. I loaded hay onto the truck.

The first sentence suggests that the truck is full of hay as an effect of my action whereas the second means only that there is some hay on the truck.

Now the question arises whether we should list all argument structures for a verb in the lexicon and how we should capture the fact that they are related. In traditional approaches all forms had to be listed and people introduced productive lexical rules to produce one form from another, i.e., one form was considered to be basic and the others were considered to be derivative on it. Or they just introduced Lexical Redundancy Rules which specified the relations among lexical items.

This seems to be quite a natural way to cope with the above problems but there are certain arguments that seem to show that they sought the solution in the wrong direction. Let us take a look at the following list of sentences:

(8) a. Pat kicked the wall.
b. Pat kicked Bob black and blue.
c. Pat kicked the football into the stadium.
d. Pat kicked at the football.
e. Pat kicked his foot against the chair.
f. Pat kicked Bob the football.
g. The horse kicks.
h. Pat kicked his way out of the operating room.

In the above case people accepting the lexical solution are compelled to list all the argument structures — syntactic and semantic — as separate lexical entries for the verb *kick* and they must possibly postulate several lexical rules to account for their relation. But since they are postulated on exactly the basis of the fact that the verb can show up in these types of sentences they constitute no explanation of the facts.

A more serious threat to this approach is posed by the following sentences:

(9) a. John sneezed the napkin off the table.
b. John talked himself blue in the face.

The problem is the following: we cannot really claim that these types of verbs come with these argument structures since the two cases — not unlike in the case of sentence h. in example (8) — have something idiomatic about them.

But these examples show us the way to the solution of all the above mentioned problems. At least in the above cases it is natural to assume that what we have here are not separate lexical entries for the verbs in question but we have certain — more or less idiomatic — constructions

[1]The examples (5) and (6) are taken from (Goldberg, 1995) but there is a remark to be made about them. The constituent *to Pat* in (5).b. is ambiguous between a locative and a benefactive interpretation and it is the second interpretation that (5).a. is a paraphrase of. The example in (6).a. is ungrammatical exactly because the Ditransitive Construction requires its second argument to be animate, but this applies to the metaphorical interpretation of the Caused-Motion Construction (cf. below) — which licences the second sentences in the pairs — as well, as shown by the following pair of sentences:

a. *I gave the table a glass of water.

b. *I gave a glass of water to the table.

The second sentence is only grammatical because it has a genuine locative interpretation.

that can be applied to the verbs due to their semantic peculiarities. Now we can see how the representations for the lexical entries and constructions proposed by construction grammarians solve the above mentionad problems along these lines.

6.3.2 Arguments and Construction Grammar

The lexical constructions for the verbs themselves — traditionally called the lexical entries — must at least contain some indication of the semantic argument structure of the verb since this will determine the constructions that the verb can participate in. We use a semantics that determines the so called **participant roles** that are associated with a verb. A lexical entry thus looks as follows:

(10) HAND: < **hander, handed, handee** >

Thus the construction for *hand* determines that the verb has three semantic arguments, namely the *agent*, i.e., the person who does the handing, the *patient*, i.e., the object that is handed over and the *recipient*, i.e., the person who receives the object[2] In what follows we will claim that this information is enough to determine the syntactic possibilities of the verb.

There is one more thing to be said about this representation. Consider the following examples:

(11) a. Jesse robbed the rich (of all their money).
b. *Jesse robbed a million dollars (from the rich).

(12) a. *Jesse stole the rich (of their money).
b. Jesse stole money (from the rich).

If you use our heuristics to determine the participant roles for the two verbs, it will be apparent that the frames do not differ, yet the two verbs can show up with different syntactic argument frames as shown by the examples above. So there must be some difference in their semantics. This difference is captured by the notion of **lexical profiling** of participants. Profiled participants are those that normally show up obligatorily in finite clauses.[3] In the lexical representation we use **boldface** to highlight the profiled arguments thus the representations for *rob* and *steal* look as follows[4]:

(13) a. ROB: < **thief, target**, goods >
b. STEAL: < **thief**, target, **goods** >

As we have said above this information is enough to determine which construction a lexeme can appear in and the fusion of the lexical entry with specific constructions will will produce the traditional argument structures, which as we see is only an epiphenomenon.

Before we can turn to this problem however we first have to show what forms constructions take in this framework. Let us look at the following example[5] :

---

[2]Goldberg gives the following heuristics for determining the basic meaning and the relevant participant roles for a verb: interpret the verb in gerundial form in the following frame:

No __ing occured.

The number and types of participants understood to be involved in the situation determine the semantics associated with the verb.

[3]There are contexts which can override this obligatoriness and certain constructions — such as the passive or the middle constructions — serve explicitly to suppress arguments.

[4]Goldberg argues that the profiling of arguments can be motivated independently, since the semantic constraints associated with the profiled arguments indicate that they play a prominent role in the way the situation is viewed. A similar contrast motivates the difference in profiling and thus in syntactic behaviour between *eat* and *devour* (cf. example (4)).

[5]We will present the constructions in what follows in the form given in (Goldberg, 1995) although sometimes the thematic roles that she assigns to specific arguments are at least questionable and depend very much on the definition of the specific role which in most cases varies from author to author. In the following example it might be argued that the third argument should rather be assigned the thematic role THEME — since it is normally not affected by the action, which again is sometimes taken to be the definition of a PATIENT.

(14) **Ditransitive Construction**

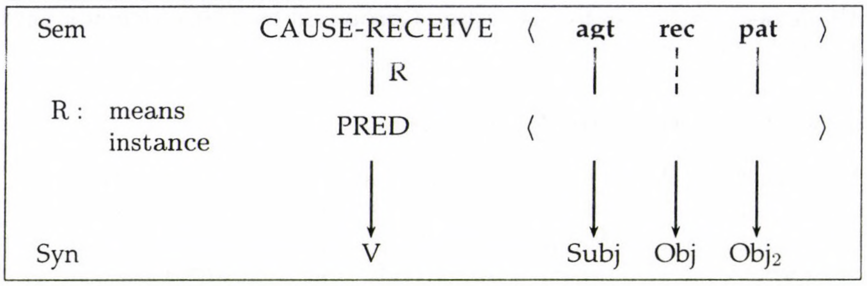

The obvious part of the construction specifies the syntax and the semantics of the construction (first and third row). The semantics specifies that the basic meaning of the ditransitive construction is that the agent causes the recepient to receive the patient (X CAUSES Y TO RECEIVE Z) and the syntactic argument structure consist of a subject, an object and a secondary object under the appropriate mapping indicated by the arrows. The boldface indicates the so called (**constructionally**) **profiled arguments** which are just by definition the arguments that are linked to direct grammatical relations (i.e., SUBJ, OBJ or OBJ_2). Now we can turn to the middle row that is the place to be filled by a lexical construction. Thus PRED is a variable that will be instantiated by the fusion of the two constructions. Similarly, R is a variable for the relation between the basic meaning and the verb meaning. The above construction constrains this relation to be either *means* or *instance*. The type of lines between the first and the second row serves to indicate a relation between the participant (lexical) and argument (constructional) roles: solid lines indicate obligatory fusion, whereas dashed lines show that these roles can be contributed by the construction. We will see an example of the role of this difference later on.

Now we have all that we need to account for the ditransitive use of *hand*, which results from the fusion of the relevant constructions given in examples (14) and (10) above. The resulting construction looks as follows:

(15) **Ditransitive HAND**

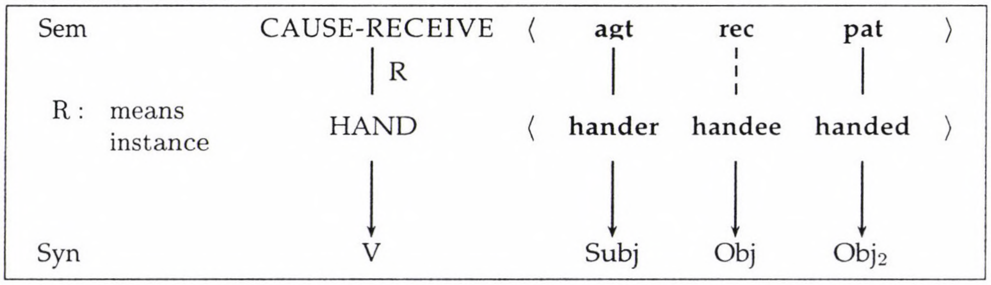

What we end up with is an instantiation of the variables and the empty places of the constructional frame with information originating from the lexical construction that contained enough material to constrain the way semantic arguments (participant roles) can be linked to syntactic arguments via linking with the argument roles of the construction frame. The resulting construction corresponds to what is traditionally thought of as a lexical entry of the verb.

To turn back to the role played by the dashed line in the construction let us see how this tool helps us to account for an argument structure of the verb *kick*. Using our test we can establish that the basic meaning of the verb and the participant roles look as follows:

(16) KICK: < **kicker, kicked** >

But this construction meets the requirements of the CAUSE–RECEIVE construction since it contains the obligatory agent and patient and as we said before the dashed line indicates exactly that the relevant role can be contributed by the construction itself. Thus the combination of the two constructions of examples (16) and (14) will result in the following construction:

(17) **Ditransitive KICK**

Sem	CAUSE-RECEIVE	⟨	**agt**	**rec**	**pat**	⟩
	R					
R: means	KICK	⟨	**kicker**		**kicked**	⟩
Syn	V		Subj	Obj	Obj_2	

This structure will be responsible for the following type of sentence mentioned in example (8) repeated here for convenience:

f. Pat kicked Bob the football.

The same mechanism accounts for those semi–idiomatic constructions that were mentioned in example (9). Since the first example will play a central part in section 6.4 we will only present the proposed solution for the second sentence:

b. John talked himself blue in the face.

The construction associated wih the verb will obviously look as follows in our framework:

(18) TALK: < **talker** >

The above sentence is a case of the English Resultative Construction whose frame looks as follows:

(19) **Resultative Construction**

Sem	CAUSE-BECOME	⟨	**agt**	**pat**	result–goal	⟩
	R					
R : means instance	PRED	⟨				⟩
Syn	V		Subj	Obj	$Obl_{AP/PP}$	

Accordingly, the basic meaning of the construction is identified as **X causes Y to become Z** and shows the appropriate argument roles and the links to the respective surface constituents. Since the above construction indicates that it can supply both the patient and the result–goal arguments it is not surprising that it can combine with the construction for *talk* thus resulting in the above composite structure:

(20) **Resultative TALK**

Sem	CAUSE-BECOME	⟨	**agt**	**pat**	result–goal	⟩
	R					
R : means instance	TALK	⟨	**talker**			⟩
Syn	V		Subj	Obj	$Obl_{AP/PP}$	

This construction then accounts for the example mentioned above[6].

Along similar lines we can account for a form associated with *kick* mentioned in example (8):

[6] There would be more to say about this construction. While the fact that the lexical entry supplies only one participant may explain why the OBJ is preferably pronominal (reflexive), this does not account for cases of fake objects — arguments coreferential with the subject — and cases of real objects — , the grammaticality of which is questionable for many speakers of English anyway.

b. Pat kicked Bob black and blue.

It is again the construction that is responsible for the extra argument since the lexical construction of the verb contains only two arguments — an agent and a patient (cf. (16)).

6.4 The Caused Motion Construction

6.4.1 The base case

We conclude this part with a concrete example, the English Caused Motion construction that will serve to highlight some further features charateristic of the constructional approach to grammar. The most important property that deserves mention is that constructions come in families whose members are related through different kinds of relations (so called **links**) to a basic construction — capturing the basic meaning of the construction.

In the case of the Caused Motion Construction the base case is captured by the following construction:

(21) **Caused–Motion Construction**

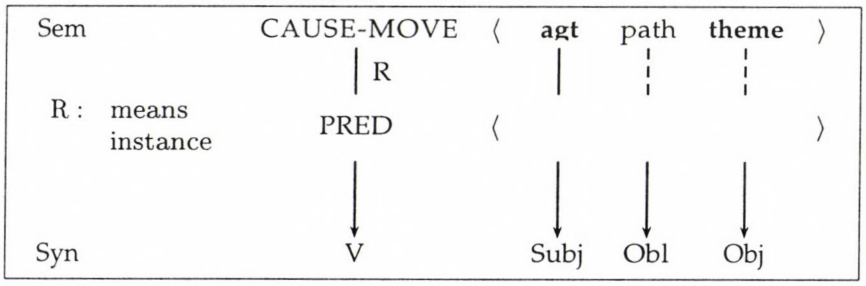

Thus we have a basic meaning, namely **X causes Y to move Z**. Let us take a verb that belongs naturally into this class, say *put* with its lexical roles distributed as follows:

(22) PUT: < **putter, put.place, puttee** >

By fusing the participant roles with the appropriate argument roles we arrive at the representation of a basic case in this family:

(23) **Caused–Motion Construction** + *put*

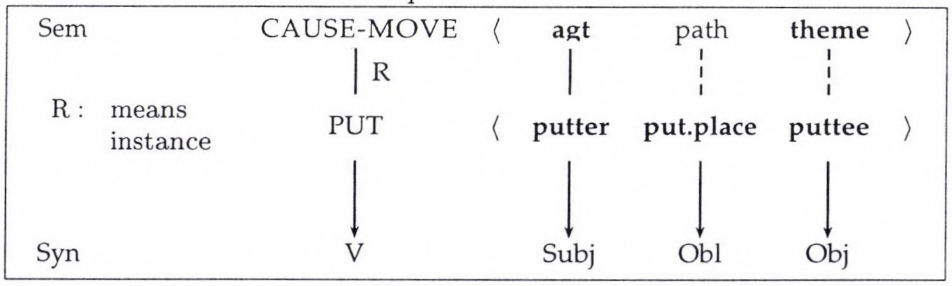

Why should we assume that this is a construction on its own rather than just an epiphenomenon on the basic meaning of *put*? One good reason is that this construction contains twodashed lines, which means that in this construction verbs can show up whose participant role frame is poorer yet matches the above construction. A case at hand is again the verb *kick* whose lexical representation combined with the construction (cf. (16) and (21)) results in the following frame:

(24) **Caused–Motion Construction** + *kick*

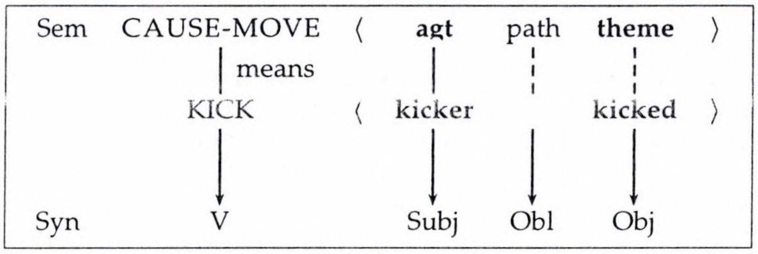

This construction again accounts for a further argument frame of the verb given in example (8):

c. Pat kicked the ball into the stadium.

There is one further point that is worth noting in this case, namely that the second participant role of the verb (**kicked**) matched the patient argument role in the Ditransitive Construction (cf. example (17)) whereas here it matches the theme role. This on the one hand shows that it is reasonable to assume some underspecified representation for the participant roles, on the other hand, it shows that it is the lexeme and the construction that raelises it together that determines the exact angle fron which we view a situation. Take the case of the following sentence:

(25) Pat kicked Bob into the bathroom.

Since the construction determines that *Bob* is the theme of the sentence, the focus is put on the fact that he is in the bathroom now, rather than on his being affected by the act of kicking.

Thus the fact that in certain cases it is this construction that furnishes a role for the argument structure of the verb seems to show that what we have here is a genuin construction. But it is not only this feature that is furnished by the above construction. Compare the above sentences with the following sentence (cf. (8)):

a. Pat kicked the wall.

This indicates that the basic meaning of *kick* is not causative at all. It is the Caused Motion Construction that is responsible for the causativity of the verb in the sentence in question.

A more clearcut demonstration of what we have shown above is furnished by one of our semi–idiomatic cases (cf. example (9)) repeated below:

a. John sneezed the napkin off the table.

This sentence is only possible because the verb *sneeze* can also show up in this construction due to its paticipant role frame:

(26) SNEEZE: < **sneezer** >

The fusion of the two constructions results in a composite that licences the above sentence:

(27) **Caused–Motion Construction** + *sneeze*

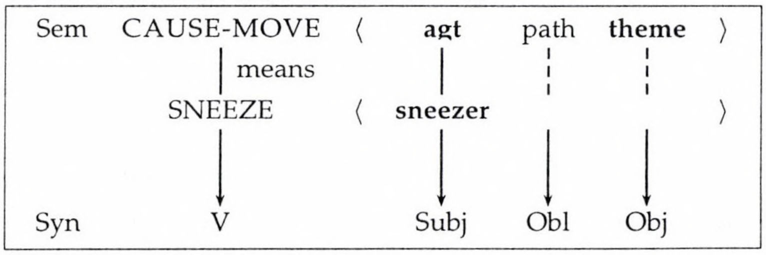

In this case it is impossible to claim that it is the meaning of the verb or the latter combined with the meaning of the preposition that lets us derive the correct meaning compositionally — whose components are almost exceptionlessly furnished by the Caused Motion Construction — as it was sometimes claimed for the cases of *kick* discussed above.

6.4.2 Links

A further argument for the existence of this construction is the fact that it can serve as the base case for a family of constructions that are related to it through links that can be attested in the case of other constructions as well. We now turn to the discussion of some of these relations.

One of these relations is the so called **Subpart Link**. This is one of the relations that belong to the class that serves to suppress argument places — to this class belong among others the English passive and middle constructions. Thus the relation produces the following frame on the Caused Motion Construction:

(28) **Intransitive Motion Construction**

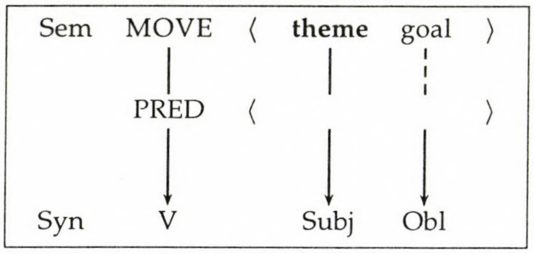

We see that in this construction we supressed the agent argument of the Caused Motion Construction. It is typically this argument that is supressed by the relation and since this argument is usually associated with the subject position that must be filled by some entity, a relinking mechanism causes some other argument to appear as subject of the sentence[7]. This construction is responsible for examples like the following:

(29) The boat floated into the cave.

Another family of links is the family of **Polysemy Links**. These links usually do not change the argument structure (neither the semantic nor the syntactic) but they cause changes in the basic meaning of the verb. One typical example is the prevent relation. When applied to the Caused Motion Construction it results in the following construction:

(30) **Caused–Prevent Construction**

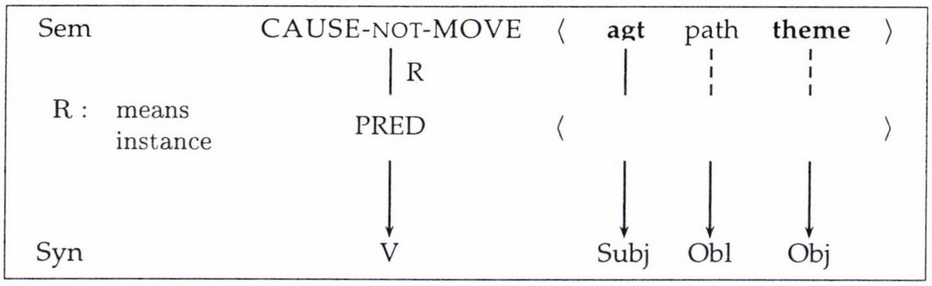

The meaning of this construction can be paraphrased as: **X prevents Y from moving Comp(Z)**. This construction accounts for the following examples[8]:

(31) a. Joe locked Mary into the bathroom.
b. I kept her at arm's length.

This is not the only type of link in this family. To indicate some further examples we specify

[7] Another instance of this link connects the Resultative Construction (cf. (19)) with its intransitive version relating such pairs of sentences as:

a. John hammered the metal flat.
b. The metal became flat.

[8] The application of this type of link to the Ditransitive Construction ((14)) results in a construction with the meaning X CAUSED Y NOT TO RECEIVE Z which accounts for the following example:

Joe refused bill a cookie.

the resulting constructions and give an example for the sentences licensed by them for both the Caused Motion (cf. (21)) and the Ditransitive (cf. (14)) Constructions:

(32) a. X ENABLES Y TO __ Z
(i) Joe permitted Chris an apple.
(ii) Joe permitted Chris into the room.
b. Conditions of satisfaction imply X CAUSES Y TO __ Z
(i) Pat promised Bob a car.
(ii) Pat ordered Bob into the room.

A further important type of link is the so called **Metaphorical Extension Links**. It is argued by Goldberg that the Resultative Construction is metaphorically linked to the Caused Motion Construction under the following mapping:

motion ⟶ change
location ⟶ state

There are certain English expressions that reflect this metaphor directly:

(33) a. The jello goes from solid to liquid in minutes.
b. She slid into madness.

We regard the fact that a family of constructions is related to a basic meaning through the well attested links as a certain indication of the existence of the construction[9].

A last indication of the presence of an explicit construction is that there are explicit semantic constraints associated with it. These constraints can be attested in the case of the Cause Motion Construction. To give just to simple examples: first, the cause argument can only be an agent or a natural force but it cannot be an instrument:

(34) a. *Chris* pushed the piano up the stairs.
b. *The wind* blew the ring into the gutter.
c. **The hammer* broke the vase onto the floor.

Another constraint involves the fact that there can be no mediating cognitive decision by the entity denoted by the direct object:

(35) a. Pat lured Bob into the room.
b. *Pat persuaded Bob into the room.

Although *lured* involves a psychological state it does not imply the existence of a cognitive decision as shown by the following sentence:

(36) Sam lured the mouse out of its hiding place.

These are constraints that are better associated with the constructions than with the semantics of each verb that can appear in them.

6.5 Conclusion

The preceding discussion served to illustrate that all these features are most naturally explained by positing a distinct construction that these properties are associated with. We can similarly argue for the existence of other constructions in English — of which we have seen but a few — that make it evident that in the majority of cases what we held for distinct lexical entries, i.e., the distinct argument structures of a verb — are in many cases mere epiphenomena due to the interplay of 'lexical' and 'syntactical' constructions.

[9] A further type of link is the **Instance Links** which relate constructions to their idiomatic uses as in the following case:

John drives Mary crazy/nuts/bananas/over the edge.

This is an idiomatic case of the Resultative Construction in as far as it requires both the presence of the verb *drive* and an argument that means *crazy*.

References

Fillmore, Ch. J. 1968. The case for case. In E. Bach and R. T. Harms, editors, *Universals in Linguistic Theory*. Holt, Rinehart and Winston, New York, pages 1–88.

Fillmore, Ch.J. and P. Kay. 1993. *Construction Grammar Coursebook*. Reading materials for Linguistics X20, University of California, Berkeley CA.

Goldberg, A.E. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. The University of Chicago Press, Chicago.

Part III

Lexical Semantics

Chapter 7

Introduction

This chapter is a summary of some current issues in **lexical semantics**, a branch of linguistics that is gaining ever more importance as grammars become more and more lexicalized (cf. Section 00). The aim of lexical semantics is to provide a framework, i.e., a logical language for the representation of meanings, which is suitable for characterizing the meanings of lexical entries. Such representations are successful if the structures that the other components of grammar (in particular, syntax) posit and the meaning representations together determine the meanings of complex expressions, i.e., if they favour the **compositional** interpretation of linguistic expressions. So lexical semantics assumes that the linguistic structures (and the semantic operations associated with them) will take care of what meanings have to be combined and how, while it assumes the responsibility of determining what the meanings that are to be combined are like. As a consequence, lexical semantics can hardly be considered a fully autonomous component — neither can the other components, especially morphology and syntax, be considered independent from it.

We will try and review the most important problems and their possible solutions in a systematic way in what follows. Any systematization of this sort is disputable and reflects the particular position the authors take about the issues at hand. Our general approach is the following. We believe that the meanings of natural-language predicates are not arbitrary. For example, the denotation of a one-place predicate (i.e., a set of individuals or a property) is not just any subset of the universe, but its members must have something in common other than having the given property.[1] Therefore, extensions can be given in terms of their relations to other extensions. As a matter of course, circularity is to be avoided; how the entire system of denotations is grounded is a philosophical problem, which should not bother us here. (For example, we can think of the final anchors that help us avoid circularity as directly measurable quantities if we take an empiricist's stance.) Whatever the philosophically correct approach is, we assume that lexical meanings can be **decomposed** in terms of other extensions that characterize what the members of the extension in question have in common. In the case of natural language, the similarities between the members of an extension are chosen in a largely arbitrary manner, either because of their relevance for biological and social reasons, or just owing to historical accidents. At any rate, those similarities have a **conceptual** rather than model theoretic relevance. The system of extensions that arises in this way is called a (natural-language) **ontology** in the literature. They will be the topic of Chapter 2. Ontologies play a role in all those mechanisms in which conceptual similarity between extensions matters, in particular, in relating meanings metaphorically. Should we consider the operations that combine lexical meanings to be purely functional operations (such as application or composition), ontologies and decomposition would not be required. But we believe more mechanisms than just the functional ones are at play in meaning composition.

The body of this chapter is about problems directly related to the interaction of lexical meanings in linguistic structures. The first problem, of a rather general character, is treated in Section 00: if a given lexical item seems to behave differently in different constructs, i.e., it apparently carries different meanings when combined with different lexical items, when is it legitimate to assume that the phenomenon is due to a genuine lexical **ambiguity** (i.e., the surface coincidence of two underlying lexical items)? When do we have to explain the different meanings using one and

[1] This is related to the philosophical problem of inductive reasoning, cf. Goodman (19??). Induction would not be possible at all if we allowed concepts to correspond to arbitrary sets of entities.

the same underlying lexical meaning? This question is to be answered first, because it arises in connection with each of the other problems that we will deal with in the remaining part of the chapter. The tests that have been used to decide whether a lexical item is truly ambiguous and the phenomena related to ambiguity and vagueness will be treated in Section 00.

The remaining part of this chapter is divided into three sections according to the three main types of approach that people have taken towards the problem of systematic polysemy (i.e., polysemy based on one and the same lexical entry rather than due to ambiguities) in the literature. First, we will deal with the decomposition-based approach to metonymical polysemy proposed by Pustejovsky (1995b) and various papers in the same spirit. The most important type of explanation in such frameworks is that the types of the expressions that we can combine in natural language often do not match exactly, and certain distinguished features within the decomposition of the meaning of the lexical entries can help solve the type conflict in such cases. For example, consider the expression *finish the beer*. The verb *finish* semantically selects an argument denoting an activity, yet *beer* has a different type of meaning. But *beer* is associated with various types of activity, in particular 'drinking', 'brewing' etc., which are accessible to the semantic combination operation. The accessible features of a lexical meaning representations are (i) the purpose/goal/typical use of the entity in question (this slot is associated with 'drinking' in the case of *beer*); (ii) the origin/coming into existence of the entity in question (this corresponds to 'brewing' in the case of *beer*); and (iii) the typical parts of the entity or those entities that the given entity is typically part of. Various different types of phenomenon are treated in this spirit, which will be presented in Section 00.

A slightly different approach, first proposed by Nunberg (1995), is explained in Secion 00. This approach relies on the concept of **transfers of meaning**, which consist in the modification of predicate extensions in terms of 'noteworthy' properties of their arguments. For example, in *I'm parked in the backyard* is interpreted as 'my car is parked in the backyard', even though *I* is not polysemous. The explanation of this phenomenon is that the meaning of *park* can be shifted in such a way that it is true when the original predicate 'park' applies to the car of its argument rather than its argument itself, provided that which car belongs to a given individual is one of its 'noteworthy' properties. That this phenomenon is associated with the predicate itself is shown by the fact that #*I consume two gallons of petrol* cannot be interpreted as 'my car consumes two gallons of petrol'.

Finally, the third type of approach that we will examine considers metonymy as an instance of **ellipsis**. This approach is based on the similarity between 'missing links' in discourse (e.g., between two consecutive sentences, or between a definite description and its antecedent), on the one hand, and meaningful connections to be established between lexical items when combined in a particular way (e.g., a verb and one of its arguments). The idea is that, since lexical entries are combined in the particular ways that their syntax legitimates, the differences between the 'gapping' processes in discourse and the systematic polysemy of lexical entries are to be attributed to the **constructions** that the lexical entries participate in. This type of approach explains both the varying degrees of productivity of the lexical processes in question and the relevance of the conceptual dimensions captured by Pustejovsky's (1995b) 'qualia structure'.

Meaning and grammatical category

In principle, a lexical entry consists of a surface form, a meaning representation and a grammatical categorization. However, it is a commonplace that much of the grammatical behaviour of a lexical item is predictable from its meaning. For example, if the meaning of a lexical item is a two-place predicate expressing, say, a process located in time and space, then its grammatical category can hardly be anything else than a transitive verb; if its meaning is an abstract concept corresponding to the same process without reference to time and space, then its grammatical category is most probably an abstract noun; and so on. It is a debated issue just how much of the grammatical behaviour of a lexical item is predictable from its meaning. No doubt, there are cases when meaning alone is not sufficient for explaining the surface behaviour of a lexical item. For example, if the meaning of a lexical item is a sentential connective, it may behave as a conjunction to be placed in between clauses or, for example, as a second-position clitic within the second clause. Or a verb denoting a process may or may not take an argument expressing an obligatory feature of the process:

say, a verb meaning 'lose weight' may or may not have an optional or even obligatory argument expressing how much weight the subject loses.

On the other hand, the danger of circularity often arises when relating meanings to grammatical behaviour. For example, one could say that a verb meaning 'lose weight' and taking two obligatory arguments (subject and weight lost) denotes a two-place relation, whereas another verb, also meaning 'lose weight' but taking just one argument (the subject) denotes a one-place predicate. This would correctly predict their different syntactic behaviour but, at the same time, one could argue that such an explanation is *ad hoc*, because their different semantics is motivated uniquely by their grammatical behaviour. In addition, if one has a verb meaning 'lose weight' with an optional 'weight lost' argument, one should say that it is ambiguous depending on whether it is used transitively or not, which is semantically awkward.

Another type of case can be illustrated with the English preposition *on*, which expresses a temporal relationship in front of the days of the week (*on Monday, on Tuesday* etc.), whereas the same relationship seems to be expressed by the preoposition *in* in connection with names of months (*in January, in February* etc.) and by *at* in connection with periods of the day (*at night, at noon* etc.). It is largely arbitrary whether we say that the grammatical (sub)categories of the prepositions *on, in* and *at* are responsible for their different behaviours or that their meanings are different in such a way that they express the given temporal relation just with the appropriate lexical classes of nouns, or else that their meaning is uniform (largely underspecified), and different nouns select different special senses for them.

Similar examples are legion, not only in connection with more or less lexicalized combinations of lexical items (like the combinations of temporal prepositions and certain types of nouns above), but all the phenomena traditionally referred to as **collocations** or **production idioms** belong here. These are constructs that have compositional meanings, yet their form is impossible to predict on the basis of their meanings. Because of this lack of predictability, we must consider them separate lexical items, although they often do not figure as such in written dictionaries (if they appear at all in them).

Collocations are just one example of **partially predictable** lexical items. Once we face the fact that every (formal or semantic) feature of a lexical item need not be arbitrary,[2] the problem of ambiguity versus polysemy gets somewhat easier to solve, because we can account for certain cases of productive polysemy by postulating very general collocations (if we can argue that they are partially arbitrary).

For example, take the verb *bake*: this verb is usually taken to have (at least) the following two meanings:

(1) a. 'submit something to dry heat' (as in *I baked the potato*);
b. 'create something by baking' (as in *I baked a cake*).

These two meanings are obviously related, so dictionaries do not consider the word *bake* homonymous. This type of productive polysemy can be explained by postulating a collocation corresponding to (1-b), occurring with verbs that express physical influence on some object, which potentially produces a different type of object (especially by transforming the input object), and expresses the production of the output using a form that is partially arbitrary (namely, the noun phrase corresponding to the output is expressed as the direct object, and there is no conventionalized way of expressing the input object). This generalized concept of collocation is usually referred to as a **construction** (cf. Fillmore and Kay (1993)). Constructions, i.e., partially arbitrary associations of form types with meaning types include every single lexical item as well as collocations and, if we want to be absolutely consequent, even syntactic rules.

We want to claim that **all** systematic polysemies are linked to particular constructions rather than just to lexical items. Constructions are partially productive by definition, they may impose largely arbitrary limitations on the lexical items that can participate in them, and they can refer to conceptual categories that are made explicit by lexical decomposition. For example, the **dativus eticus** that can be embodied by the second object of ditransitive constructions in English is linked to a conceptual category that we could call 'making available', which is a rather arbitrary feature

[2]Note that even idioms some of the semantic and syntactic features of the items that occur in them as a rule. For example, an idiom containing an event verb, such as *kick the bucket*, could hardly refer to a state, although the event it refers to has nothing to do with either *kick* or *bucket*.

of certain verbs, but which can stem from the context itself. The verb phrase *wash myself a car* may be appropriate in a context in which washing a car makes it available for some purpose (like taking it from the garage), whereas in other cases the lexical entry of the verb specifies that the verb has the feature 'making available' (e.g., *build myself a house, sing myself a song*).

References

Fillmore, Ch.J. and P. Kay. 1993. *Construction Grammar Coursebook.* Reading materials for Linguistics X20, University of California, Berkeley CA.

Nunberg, Geoffrey. 1995. Transfers of meaning. *Journal of Semantics*, 12:109–132.

Pustejovsky, J. 1995. *The Generative Lexicon.* MIT Press, Cambridge MA.

Chapter 8

Cognitive Semantics

8.1 Introduction

In this section we take a look at what kinds of ontological assumptions lie behind the most influential approaches to the study of meaning. In particular, we examine two concurrent paradigms in contemporary semantics: Truth-conditional Semantics and Cognitive Semantics. However, our treatment will not be completely balanced. As the main topic of this volume is Lexical Semantics, and, as we will see later on, words, rather than sentences, are considered to be the *locus of meaning* according to the cognitive approach (as opposed to the truth-conditional one), the cognitive view will be given somewhat more room in our discussion. The linguist we will take a closer look at to see the charasteristics of cognitive semantics is *Ray Jackendoff* and his *Conceptual Semantics.*

8.2 Two Paradigms of Semantics

8.2.1 Truth-conditional Semantics

A very influential branch of contemporary semantics is *truth-conditional* (or *model-theoretic*) semantics. It is rooted in mathematics and can be traced back at least to the works of Gottlob Frege (Frege (1972)) and, most prominently, to those of *Alfred Tarski* (see, for example, Tarski (1983)). Tarski developed his ideas in order to be able to characterize the concept of truth in formal (that is to say, artificial) languages, but his theory have gained great popularity among linguists dealing with natural language since then. Although Tarski himself expressed his skepticism about the possibility of applying his theory of truth to natural language, subsequent philosophers, *Richard Montague* in the first place (Montague (1970a) ,Montague (1970b), Montague (1972)) have shown that Tarski might have been mistaken. "But what does Truth have to do with Meaning?", one might ask. According to this tradition, very much, and this answer can in fact be taken as the hallmark of this particular attitude to the question "What is Meaning?" And this is completely understandable if we consider the fact that the cardinal motivation behind the enterprise was to create exact means to handle logical relations between propositions[1]. Logic is interested in the valid patterns of inference, that is, those patterns of inference that *preserve truth.* What about meaning? To put it tersely, meaning is the property by virtue of which a sentence[2] can participate in various patterns of inference. In other words, *the meaning of a sentence can be characterized by the set of sentences it entails.* This implies, incidentally, that the basic level of truth-conditional semantics is that of the sentence; all other linguistic entities (for example, words) are assigned meaning according to the roles they

[1]Frege's ultimate goal was to prove that arithmetics (and thus mathematics itself) was a part of logic, but, in order to attain this goal (which was proved unattainable by Russell through discovering paradoxes in naive set theory), he had to develop suitable tools to formalize the informal mathematical reasoning that had been used in mathematical proofs.

[2]Sentences express propositions, and it is propositions that can be said to be truth-bearing entities in the real sense. Although the relationship between a sentence and a proposition it expresses is far from being simple, it is usually possible to identify the proposition on the basis of the sentence by making explicit all implicit information related to the circumstances of the utterance of the sentence in question. That is why it is not always necessary to make a distinction between these philosophically very different entities in a painstakingly consistent manner.

can play in sentences.

We can see that, on this view, semantics is most intimately connected with (and dependent on) truth. Put it more precisely, the objective of this kind of semantics is to identify the *truth conditions* of the sentences in an (artificial or natural) language. This is usually accomplished by invoking a mathematical construct called *a model*, which consists of a non-empty set (the *universe*), and a special function (the *interpretation function*) that maps the signs of the language (that is, individual names, predicate names, names of relations, etc.) onto mathematical constructions defined over the universe (particular elements of the universe, subsets of the universe, ordered pairs of the elements of the universe, etc.) These mappings are then combined in such a way ("compositionally") that the sentences of the language are mapped onto two abstract values (not included in the universe), **True** and **False**.

Originally, this approach was *extensional*, which means that it could not handle several phenomena of natural language (such as modality, or belief-contexts, etc.) This situation changed, however, when *Saul Kripke* succeeded in giving exact semantics to modal logic by formalizing the concept of "possible worlds"(see Kripke (1959)). His framework was developed by Richard Montague into *intensional logic*, a very powerful instrument for modelling various kinds of natural language phenomena. Recent developments in formal semantics can be seen as efforts to overcome the shortcomings of Montague's system. Two such developments should be mentioned. The first is *Dynamic Predicate Logic* by *J. Groenendijk* and *M. Stokhof* (Groenendijk and Stokhof (1991)). This system was made to cope with cross-sentential anaphora, and introduced a new concept of meaning as a potential to change the information state of the hearer.
The second important development is the flourishing of various *algebraic semantics* , which were invented to fill the gap concerning plural nouns and mass nouns in Montague's original system (Montague himself largely ignored these questions). Systems of algebraic semantics (see for example Landman (1991) for a comprehensive textbook, or *Link* in Link (1983), which was one of the earliest developments in the field) widely use constructs from general algebra such as semi-lattices, lattices, etc. and they seem to be somewhat closer to algebra than to logic. This attitude, if successful, can lead semantics further away from classical logic than was thought to be possible at the dawning of formal semantics.

Finally, a remark concerning the ontological status of model-theoretic semantics is in order. It is very important to see that the ontological commitments of model-theoretic semantics are not more than those of mathematical set theory, coupled with the commitment to there being two abstract (Platonic) truth values. This means, in effect, that model-theoretic semantics is practically as neutral from an ontological point of view as mathematics. As in other fields of mathematics, real ontological questions emerge when it comes to *applying* the piece of mathematics in question, that is, when the components of the theory become interpreted by the "human user" having a concrete domain of possible application in mind. To put it briefly, model-theoretic semantics does not involve any realist commitment whatsoever. Such commitments belong to the person applying the mathematical means rather than to the means themselves. As *Joost Zwarts* and *Henk Verkuyl* put it in Zwarts and Verkuyl (1994, p. 2.):

> [T]he use of model theory in the analysis of natural language does not imply a realist commitment. The mathematical tools employed in a model-theoretic approach to natural language are [. . .] neutral with respect to philosophical or epistemological positions. It is hard to see why mentalists should refrain from using set theory, Boolean algebras, lattices, Tarskian assignment functions, possible worlds and other modern tools of analysis . . .

8.2.2 Cognitivist Semantics

One of the most significant features of cognitivist semantics is that it postulates an autonomous level of *mental representations*. Expressions of natural language have their interpretation on this level. In other words, outside reality does not have to play any role in interpreting the expressions of a language; as far as the linguist is concerned it is sufficent for him or her to be able to describe the relationship between the expressions of the language and the corresponding mental entities. Truth and falsity then are notions that can be situated *not* between language and reality but between the level of mental representations and reality. However, *this* relationship should not concern the cognitive semanticist because, to cite a famous slogan, "Meanings are in the head."

8.3 Jackendoff's Conceptual Semantics

Ray Jackendoff is one of the most important authors in the field of cognitive semantics. His works include *Semantics and Cognition* (1983), *Consciousness and the Computational Mind* (1987), and *Semantic Structures* (1990). Since it is his Semantic Structures that deals with language in an extensive fashion, we will mainly draw on that work. However, it is of course not possible to present Jackendoff's views (that range from linguistics to cognitive psychology to philosophy) in an all-embracing manner (because that would be a task to which a separate volume should be devoted), so we can only focus attention on the broad outlines of his work.

8.3.1 Representational Modularity

According to Jackendoff (Jackendoff (1996, pp. 1–5.), Jackendoff (1990, pp. 285ff.)) the mind/brain encodes information in a modular fashion. Jackendoff calls this *Representational Modularity* , because it is the way of representing information that makes different modules (different *representation modules*) differ from each other. There are, for example, a linguistic module, a visual module, an acoustic module, etc., that work with information that is encoded in the particular format they need each (they are different "languages of the mind"). The modules can contain further submodules. For example, the linguistic module contains a phonology module, a syntax module, and a module interfacing with several other faculties, the module of *Conceptual Structure.* This module is where meanings are to be found in Jackendoff's theory.
Further, there are various *interface modules* mediating the flow of information between different representation modules. The task of these interface modules is to "translate " one "language of mind" used by one module into another one used by another module. According to Jackendoff, this process is not a simple code-switch, but more like "a partial homomorphism" (Jackendoff (1996, p.4.)) between two formats of representation: part of the information present in the source representation might have no corresponding piece in the target representation.
According to this picture of grammar, a lexical item is essentially a *correspondence* between phonological, syntactic and conceptual structures.

8.3.2 Lexical Decomposition

It is a very important component of Jackendoff's theory that he accepts the existence of semantic primitives out of which all meanings ("thoughts") are composed, in other words, he believes in *lexical decomposition*[3]. This ontological commitment can be defended on the grounds of the apparent creativity of natural language (Jackendoff (1990, pp. 8–11., 37–41.)). Presupposing that there are strong parallels between syntax and semantics (which, we will see, Jackendoff does presuppose), the arguments from productivity in syntax can be applied to semantics as well: just as there must be primitives in syntax (because not all syntactic structures can be stored in memory, since there are potentially infinitely many of them), there must also be primitives in semantics (because not all thoughts can be stored in the finite brain).

8.3.3 The Conceptual Structure

Conceptual Structure (CS) has the following characteristics(Jackendoff (1990, pp.7–43.), Jackendoff (1996, pp. 5–13)).

[3]On the topic of lexical decomposition see *Wunderlich* in Wunderlich (1994) where he compares three different approaches to lexical decomposition.

- It is built up out of discrete primitive features and functions.
- Its expressions refer to the world *as we conceptualize it.*
- It must contain all the nonsensory distinctions of meaning made by natural languages (e.g. the type–token distinction).
- It is relational rather than linear.
- It does not have to be entirely digital so that it can permit stereotype and family resemblance effects to be formulated.
- It is universal.
- Languages can differ in respect to what elements of CS will surface in their syntax/morphology.
- Syntax only "sees" the argument structure in CS.
- The interface from syntax to CS preserves embedding (i.e. if X is a constituent of Y, then the CS of X will be contained by the CS of Y).

What kind of "nonsensory distinctions" does CS have to encode? At the very least, the following items seem to be necessary:

- CS must contain pointers to all the sensory modalities.
- CS must contain the distinction between tokens and types.
- CS must contain the encoding of quantification and quantification scope.
- CS must be able to abstract actions from the individual actually performing the actions.
- CS must encode taxonomic relations.
- CS must encode some social predicates (such as "is friend of," "is fair," "is obliged to," etc.)
- CS must encode modal predicates (such as the distinction between "is flying," and "can fly," etc.)

8.3.4 The Basic Ontology of Conceptual Structure

The essential units of conceptual structure are *conceptual constituents,* each of which belongs to one of a small set of major ontological categories (or "conceptual parts of speech"). These are THING, EVENT, STATE, ACTION, PLACE, PATH, PROPERTY and AMOUNT[4]. Although they pick out different entities as referents, Jackendoff lists six points of similarity between them (Jackendoff (1990, pp.22–25.)).

(i) Each major syntactic constituent of a sentence maps into a conceptual constituent in the CS of the sentence. For example, in *John ran toward the house, John* and *the house* map to Thing-constituents, the PP *toward the house* corresponds to a Path-constituent, while the entire sentence corresponds to an Event-constituent. The matching is by *constituents,* not by *categories*: the same syntactic category can express many conceptual categories (for example, an NP can express a Thing (*the dog*), or an Event (*the war*), or a Property (*redness*)).

(ii) Each conceptual category supports the encoding of units not only on the basis of linguistic input but also on the basis of the visual (or other sensory) environment (see also footnote (4)).

(1) a. *That* is a robin.
    b. *There* is your hat.

[4]In Jackendoff (1987, pp. 148–152.) Jackendoff argues that *pragmatic anaphora* (i.e., using a demonstrative pronoun to refer back to some entity that is perceptually accessible to the hearer) and the use of phrases such as *same* and *different* (as in the sentence *Bill ate at the same place as Jack did*) are the basic clues to the implicit ontology in natural language.

c. Can you do *this*?
d. The fish was *this* long.

Accompanied by a suitable gesture, the italicised phrases identify a Thing, a Place, an Action and an Amount, respectively.

(iii) Many of the categories (except perhaps Property and Amount) support a type–token distinction. For example, the Place-type expressed by *over your head* can express many different Place-tokens.

(iv) Many categories support quantification.

(2) a. Every dinosaur had a brain. (Things)
b. Everything you can do, I can do better. (Actions)
c. Anyplace you can go, I can go too. (Places)

(v) Each conceptual category has some realizations in which it is decomposed into a function–argument structure; each argument is in turn a conceptual constituent of some major category.

(3) a. John is tall.
b. president of the republic
c. from under the table

In (8) the arguments are *John* (Thing) and *tall* (Property), whereas the whole sentence expresses a State. (9) is an example of a Thing that has another Thing as an argument; and in (10) a Path has a Place as its argument.

(vi) The conceptual structure of a lexical item is an entity with zero or more open argument places. The meanings of the syntactic complements of the lexical item fill in the values of the item's argument places in the meaning of the sentence. For example, *be* in (8) expresses a State-function the arguments of which are found in subject and object positions.

The above observations lead to the basic formation rules of **X-Bar Semantics**. $\bar{X}$-semantics parallels $\bar{X}$-syntax in capturing basic structural symmetries in the field of meanings.
Here are the basic formation rules of $\bar{X}$-semantics (Jackendoff (1990, pp.22–32)).

(R1) $$[\text{Entity}] \rightarrow \begin{bmatrix} \text{Event/Thing/Place/} \cdots \\ \text{Token/Type} \\ \text{F}(\langle \text{Entity}_1, \langle \text{Entity}_2, \langle \text{Entity}_3 \rangle\rangle\rangle) \end{bmatrix}$$

(R2) XP corresponds to [Entity].

(R3) $$\begin{bmatrix} \text{X}^0 \\ \_\_\langle \text{YP} \langle \text{ZP} \rangle\rangle \end{bmatrix} \text{ corresponds to } \begin{bmatrix} \text{Entity} \\ \text{F}(\langle \text{E}_1, \langle \text{E}_2, \langle \text{E}_3 \rangle\rangle\rangle) \end{bmatrix}$$

where YP corresponds to E_2, ZP corresponds to E_3, and the subject (if there is one) corresponds to E_1.

8.3.5 Further Ontological Refinements

A very attractive feature of Jackendoff's theory is that various cross-categorial generalizations can be stated in it in a relatively simple way. For example, verbs and prepositions related to the spatial semantic domain can very often be transferred to non-spatial domains as well (Jackendoff (1990, pp. 25–27.), Jackendoff (1987, pp. 152–158)).

(4) a. The bird went from the ground to the tree.
b. The inheritance went to Philip.
c. Harry went form elated to depressed.

These sentences are intuitively related in the sense that their common structure can be described as follows.

(5)

$$[_{\text{Event}}\ \text{GO}([\ \ ], \begin{bmatrix}_{\text{Path}} & \text{FROM}([\ \ ]) \\ & \text{TO}([\ \ ])\end{bmatrix})]$$

The difference in the meanings of sentences (11–13) is accounted by introducing different GO-functions for each semantic field (ontological domain); for example, in (11) the conceptual structure of the sentence will contain $\text{GO}_{\text{Spatial}}$ (physical space), in (12) it will contain GO_{Poss} ("space of possession"), and in (13) it will contain GO_{Ident} ("space of identification").

A further interesting point concerns inference patterns. The following sentences can be inferred from the sentences (6)a-c, respectively:

(6) a. The bird is in the tree.
b. The money is Philip's.
c. Harry is depressed.

It is possible to account for the above inferences by the following schema:

(7) *At the termination of* $[_{\text{Event}}\ \text{GO}([X], [_{\text{Path}}\ \text{TO}([Y])])]$, *it is the case that* $[_{\text{State}}\ \text{BE}([X], [_{\text{Place}}\ \text{AT}([Y])])]$.

8.3.6 Elaborating Function–Argument Structure

The basic function–argument structure introduced in (**R3**) can be further elaborated depending on which particular category is being specified. For example, the ontological categories EVENT and STATE have the following specialized formation rules (Jackendoff (1990, pp. 33, 87–99.)).

(**R3a**) $$[\text{EVENT}] \rightarrow \left\{ \begin{array}{c} [_{\text{Event}}\ \text{GO}([\text{THING}], [\text{PATH}])] \\ [_{\text{Event}}\ \text{STAY}([\text{THING}], [\text{PLACE}])] \end{array} \right\}$$

(**R3b**) $$[\text{EVENT}] \rightarrow \left[_{\text{Event}}\ \text{CAUSE}\left(\left[\left\{ \begin{array}{c} \text{THING} \\ \text{EVENT} \end{array} \right\}\right], [\text{EVENT}]\right)\right]$$

(**R3c**) $$[\text{STATE}] \rightarrow \left\{ \begin{array}{c} [_{\text{State}}\ \text{BE}([\text{THING}], [\text{PLACE}])] \\ [_{\text{State}}\ \text{ORIENT}([\text{THING}], [\text{PATH}])] \\ [_{\text{State}}\ \text{EXTEND}([\text{THING}], [\text{PATH}])] \end{array} \right\}$$

Further conceptual functions introduced by Jackendoff include:

(**R3d**) $[\text{EVENT}] \rightarrow [_{\text{Event}}\ \text{MOVE}([_{\text{Thing}}\ \ ])]$

Instances of this function can be found in sentences such as

(8) Debbie danced.

and (in the non-spatial domain)

(9) Lila laughed.

(**R3e**) $[\text{STATE}] \rightarrow [_{\text{State}}\ \text{CONF}([_{\text{Thing}}\ \ ])]$

"CONF" stands for *configuration*, and is related to the internal spatial configuration of its Theme. As an example, consider

(10) Sally stood/sat for hours on end.

(11) (**R3f**) $[\text{EVENT}] \rightarrow [_{\text{Event}}\ \text{INCH}([_{\text{State}}\ \ ])]$

This function is responsible for the *inchoative reading* of the following sentence:

(12) Bill stood on the table.

8.3.7 The Architecture of Lexical Entries

Lexical entries contain the following pieces of information:

- phonological form
- syntactic category
- subcategorization frame
- conceptual structure

As an example, let us consider an item having a fairly complicated lexical entry, the verb *drink* (Jackendoff (1990, p. 53)).

(13)
$$\begin{bmatrix} \text{DRINK} \\ \text{V} \\ \_\_\ \langle NP_j \rangle \\ [_{\text{EVENT}}\ \text{CAUSE}([_{\text{THING}}\ \ ]_i, [_{\text{EVENT}}\ \text{GO}([_{\text{THING}}\ \text{LIQUID}]_j, \\ [_{\text{PATH}}\ \text{TO}([_{\text{PLACE}}\ \text{IN}([_{\text{THING}}\ \text{MOUTH OF}([_{\text{THING}}\ \ ]_i)])])])])] \end{bmatrix}$$

Here angle brackets denote the fact that the complement NP is not obligatory but optional. Indexing conventions ensure that the CS of the Subject is substituted into the argument place indexed with i; coindexing between other syntactic and conceptual constituents is always made explicit in the lexical entry (see the coindexing of the direct object NP and the argument of GO in the above lexical entry). CAUSE is a primitive function that takes a THING and an EVENT as an argument, and combines them into an EVENT. GO is also a primitive function mapping a THING and a PATH onto an EVENT, while TO maps a PLACE onto a PATH. Finally, IN takes a THING and yields a PLACE (intuitively, the place occupied by the thing), and MOUTH OF takes a THING and yields "the mouth of" that thing (presumably, a living thing). LIQUID in $[_{\text{Thing}}$ LIQUID$]$ is the *selectional restriction* that one can only drink liquids. This information is already encoded in the lexical conceptual structure (LCS) of the item *drink*, and it constraints the possible substitution of arguments into the argument place coindexed with the direct object. Jackendoff calls this kind of argument substitution *fusing* or *merging*[5].
To sum up, here is the rule for *Argument Fusion* (Jackendoff (1990, p. 53ff.)).

(**R4**) To form the conceptual structure for a syntactic phrase XP headed by a lexical item H:

(a) Into each indexed constituent in H's LCS, fuse the conceptual structure of that phrase YP that satisfies the coindexed position in H's subcategorization feature.

(b) If H is a verb, fuse the conceptual structure of the Subject into the constituent indexed i in H's LCS.

The treatment of selectional restrictions parallels that of totally incorporated arguments, such as the Goal in the following sentence,

(14) Joe pocketed the money.

The verb *pocket* has the following conceptual structure:

(15) $$[_{\text{Event}}\ \text{CAUSE}([_{\text{Thing}}\ \ ]_i, [_{\text{Event}}\ \text{GO}([_{\text{Thing}}\ \ ]_j, [_{\text{Path}}\ \text{TO}([_{\text{Place}}\ \text{IN}([_{\text{Thing}}\ \text{POCKET}])])])])]$$

Here the Path bears no index and thus receives its interpretation "into a pocket" entirely from the verb.

Besides those conceptual structures that involve function–argument organization, Jackendoff considers *restrictive modification*. In this case one constituent is the *modifier* of another. Jackendoff's rule for Restrictive Modification is as follows(Jackendoff (1990, p. 56ff.)):

[5]Fusion of conceptual structures is in effect a kind of *unification*, in that fusion is only possible if no contradiction arises because of feature-clash.

(**R5**) If YP is a daughter of $\bar{\bar{X}}$ in XP,
and the conceptual structure of YP is $[C_y]$,
then the conceptual structure of XP is of the form

$$\begin{bmatrix} \dots \\ [C_y] \end{bmatrix}.$$

An Example

As an example of the above principles in work, Jackendoff provides the derivation of the conceptual structure of the sentence

(16) John went home at 6:00.

as follows (Jackendoff (1990, p. 57)).

The head of the phrase is *go*, the lexical entry of which is

$$(17) \quad \begin{bmatrix} \text{GO} \\ \text{V} \\ \underline{\quad} \text{ PP}_j \\ \left[_{\text{EVENT}} \text{ GO} \left([_{\text{THING}} \quad]_i , [_{\text{PATH}} \quad]_j \right) \right] \end{bmatrix}$$

John and *home* have the following (obviously, simplified) respective entries.

(1) $[_{\text{Thing}} \text{ JOHN}]$

and

(2) $[_{\text{Path}} \text{ TO}([_{\text{Place}} \text{ HOME}])]$.

By Argument Fusion, the entries for *John* and *home* become fused into the entry of the head *go*, yielding

(3) $[_{\text{Event}} \text{ GO} ([_{\text{Thing}} \text{ JOHN}], [_{\text{Path}} \text{ TO}([_{\text{Place}} \text{ HOME}])])]$

The conceptual structure of the PP *at 6:00* is also constructed by Argument Fusion:

(4) $[_{\text{Place}} \text{ AT}_{\text{Temp}}([_{\text{Time}} 6{:}00])]$

If we assume that the PP is a daughter of $\bar{\bar{V}}$, then the Restrictive Modifier Rule will apply, and specify the frame of the CS of the entire sentence as

$$(5) \quad \begin{bmatrix} \dots \\ [_{\text{Place}} \text{ AT}_{\text{Temp}}([_{\text{Time}} 6{:}00])] \end{bmatrix}.$$

The CS of the entire sentence is the fusion of (3) and (5):

$$(6) \quad \begin{bmatrix}_{\text{Event}} \text{GO} ([_{\text{Thing}} \text{ JOHN}], [_{\text{Path}} \text{ TO}([_{\text{Place}} \text{HOME}])]) \\ [_{\text{Place}} \text{ AT}_{\text{Temp}}([_{\text{Time}} 6{:}00])] \end{bmatrix}.$$

8.3.8 Thematic Roles

According to Jackendoff, thematic roles (θ-roles) are part of the level of conceptual structure, not of syntax. Thematic roles are nothing but particular structural configurations in conceptual structure. The following list summarizes these configurations (Jackendoff (1990, pp. 46–50)).

- *Theme* = the first argument of GO, STAY, BE, ORIENT or EXTEND
- *Source* = the argument of the Path-function FROM
- *Goal* = the argument of the Path-function TO
- *Agent* = the first argument of the Event-function CAUSE

- *Experiencer* = an argument of a State-fuction having to do with mental states (see Jackendoff (1990, p. 140)).

By giving such a structural explanation of traditional thematic roles, Jackendoff is able to extend the notion of thematic roles to embrace conceptual roles having no traditionally acknowledged label. Also, in his theory not only NPs receive thematic roles. For example, in the sentences

(18) a. The light changed from red to green.
b. Bill talked Harry into shutting up.

green receives the thematic role of Goal in (a), and *PRO shut up* is also a Goal in (b).

8.3.9 The Action Tier

One of the most interesting points in Jackendoff's theory is his drawing parallels between semantic and phonological theory (Jackendoff (1990, pp. 125–151.)). Jackendoff uses two tiers (familiar from phonological theory such as Autosegmental Phonology), the *Thematic Tier* and the *Action Tier*. Paralleling autosegmental phonology, Jackendoff allows there to be syntactic units sharing a thematic and/or action role as well as units having more than one thematic role. This obviously violates θ-criterion. However, Jackendoff has arguments showing that it is possible to encounter well-formed sentences containing NPs that share the same thematic role (or have more than one). Consider the following sentence as an example:

(19) The box has books in it.

Here, *the box* and *it* do not have distinct θ-roles. For example, it is not possible to question the object of the preposition

(20) *What does the box have books in?

and (22) can be paraphrased so that it will not contain the participant in question twice:

(21) There are books in the box.

Also, the object of the preposition cannot be made reflexive (even though it ought to be):

(22) *The box has books in itself.

To see that there are NPs that have more than one θ-role, consider the verb *buy*. *Buy* involves (at least) the following two components:

X buy Y from Z

(a) Y changes possesion from Z to X
(b) money changes possession from X to Z

Here X and Z have two semantic roles apiece. And it is not possible to find a way out by saying that the roles in the countertransfer (b) do not count, because precisely the presence of the countertransfer distinguishes *buy* from (for example) *obtain.*

What is the reason why Jackendoff introduces the two tiers, one for such thematic roles as Source, Goal, Theme, etc., and one for the couple Actor–Patient? If we use the frame

(23) What NP did was ...

as a test frame for Actors, then in the following sentences we can identify Actors as in Source, Theme and Goal, respectively (Jackendoff (1990, p. 126.)).

(24) a. The sodium emitted electrons. (What the sodium did was emit electrons.)
b. Bill roll down the hill. (What Bill did was roll down the hill.)
c. The sponge absorbed the water. (What the sponge did was absorb the water.)

Thus it is reasonable to suppose that conceptual roles fall into two tiers: a thematic tier dealing with motion and location, and an action tier dealing with Actor–Patient relations[6]. As the intuitive definition of Actor could be "the doer of the action," while that of Patient could be "the affected entity," Jackendoff introduces the formal elaboration of Events in the action tier as follows:

$$\textbf{(R3g)} \qquad [\text{EVENT}] \rightarrow \left[\begin{array}{c} \ldots \\ \text{AFF}(\langle[\text{THING}]\rangle, \langle[\text{THING}]\rangle) \end{array} \right]$$

Here we can give a structural account of Actor and Patient as Actor being the first argument and Patient being the second.

Jackendoff uses the action tier to formalize Talmy's account of causation. Talmy in Talmy (1988) claims that causation is one instance of a broad system of concepts he calls *force-dynamics*. Force dynamics involves two characters one of which, the *agonist*, has a tendency toward performing (or not performing) some action, whereas the other character, the *antagonist*, tries to oppose the agonist's tendency. The agonist–antagonist dyad shows up on the action tier: the agonist as Patient and the antagonist as Actor. For example, the conceptual structure of

(25) Harry forced Sam to go away.

will be

$$\text{(a)} \qquad \left[\begin{array}{c} \text{CAUSE}([\text{HARRY}], \left[\begin{array}{c} \text{GO}([\text{SAM}], [\text{AWAY}]) \\ \text{AFF}([\text{SAM}], \quad) \end{array} \right] \\ \text{AFF}([\text{HARRY}], [\text{SAM}])) \end{array} \right]$$

Similarly, the conceptual structure of

(26) Harry prevented Sam from going away.

will differ from (a) in that Harry's efforts are directed toward Sam's *not* leaving:

$$\text{(b)} \qquad \left[\begin{array}{c} \text{CAUSE}([\text{HARRY}], \left[\text{NOT} \left[\begin{array}{c} \text{GO}([\text{SAM}], [\text{AWAY}]) \\ \text{AFF}([\text{SAM}], \quad) \end{array} \right] \right] \\ \text{AFF}([\text{HARRY}], [\text{SAM}])) \end{array} \right]$$

8.4 Conclusion

In the foregoing sections we took a brief look at the outlines of Jackendoff's semantic theory. We saw that the basic presuppositions of cognitive semantics and truth-conditional semantics are thought to be very different. This might be due to the fact that model-theoretic semantics is often mistaken to hold *realist* (or even *physicalist*) commitments, which, as a matter of fact, it does not imply at all. On the other hand, cognitive semantics (such as Jackendoff's) is often criticised because it does not lend itself very easily for treating those aspects of natural language that have been handled successfully by model-theoretic semantics, such as, for example, quantification. However, there are encouraging signs that the two paradigms might get closer in the future (see for example Piñon (1993) and Zwarts and Verkuyl (1994); the pronounced goal of the latter is to provide a model-theoretic interpretation for Jackendoff's Conceptual Structures). Obviously, both traditions would be able to profit from joining forces with each other.

References

Frege, Gottlob. 1972. *Conceptual Notation and Related Articles.* Oxford University Press.

Groenendijk, J. and M. Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy*, pages 14:39–100.

[6]Jackendoff cites Gruber (Gruber (1965)) who defines Theme as "the object in motion or being located, " Source as "the object from which motion proceeds," and Goal as "the object to which motion proceeds."

Gruber, J. S. 1965. *Studies in Lexical Relations*. Ph.D. thesis, MIT, Cambridge,.

Jackendoff, Ray. 1987. *Consciousness and the Computational Mind*. Cambridge, MA: MIT Press.

Jackendoff, Ray. 1990. *Semantic Structures*. Cambridge, MA: MIT Press.

Jackendoff, Ray. 1996. The architecture of the linguistic–spatial interface. In Bloom, Peterson, Nadel, and Garrett, editors, *Language and Space*. MIT Press, Cambridge, MA, chapter 1.

Kripke, Saul A. 1959. A completeness theorem in modal logic. *The Journal of Symbolic Logic*, pages 24:1–14.

Landman, Fred. 1991. *Structure for Semantics*. Kluwer Academic Publishers.

Link, Godehard. 1983. The logical analysis of plurals and mass terms: A lattice-theoretic approach. In Bauerle, Schwarze, and von Stechow, editors, *Meaning, Use, and the Interpretation of Language*. de Gruyter, Berlin.

Montague, Richard. 1970a. English as a formal language. In B. Visentini et al., editor, *Linguaggi nella societa e nella tecnica*. Milan, pages 189–224.

Montague, Richard. 1970b. Universal grammar. *Theoria*, pages 36:373–398.

Montague, Richard. 1972. The proper treatment of quantification in ordinary english. In Hintikka, Moravcsik, and Suppes, editors, *Approaches to Natural Language*. Reidel, Dordrecht.

Piñon, J. Christopher. 1993. Paths and their names. In *Proceedings of the Chicago Linguist Society 29*. Chicago.

Talmy, L. 1988. Force dynamics in language and thought. *Cognitive Science*, pages 12:49–100.

Tarski, Alfred. 1983. *Logic, Semantics, Metamathematics. Papers from 1923 to 1938*. Clarendon Press, Oxford.

Wunderlich, Dieter. 1994. 'models of lexical decomposition. In E. Weigand and F. Hundsnurscher, editors, *Conference on 'Lexicology and Lexical Semantics'*. Münster.

Zwarts, Joost and Henk Verkuyl. 1994. An algebra of conceptual structure. *Linguistics and Philosophy*, pages 17:1–28.

Chapter 9

Ambiguity and Vagueness

9.1 Introduction

A linguistic utterance can often be interpreted in more than one way. An important distinction can be made on the basis of what the source of the variability of possible interpretations is. The two kinds of sources to be distinguished are *lexical ambiguity* and *vagueness* or generality of meaning. The distinction is also an important lexical issue because it is reflected in the lexical representation of the entries that can have various interpretations.

In this chapter, besides making clear what the difference between ambiguity and vagueness is, some tests will be introduced that have been proposed for deciding between ambiguity and vagueness. It is also going to be discussed how those tests really work and whether they can really be used to for the purpose of identifying how a word should be lexically represented and if yes, then how they should be used.[1]

9.2 Ambiguity tests

9.2.1 A basic test

In order to clarify the notion of vagueness as opposed to ambiguity, a very basic 'test' can be introduced. In fact, this can not usually be used as a real test, because it involves prompting the utterer, which device is not usually available. The 'test' is the following:

An utterance is said to be vague (as opposed to being ambiguous) if the utterer himself can be uncertain about the details of the situation described by the utterance that are not clear from the utterance. This means that when he is prompted to clarify the vague details, he may say 'I do not know'.

The sentence in (1), for example, is vague concerning who ate how much:

(1) The four kids ate three pizzas.

Vagueness is often the characteristics of sentences involving collectives (such as the one in (1)). But the source of vagueness is often lexical. Some tests have been proposed specifically for the detection of lexical ambiguity vs. vagueness.

9.2.2 The absence of crossed readings

In sentence (2), the word *teachers* can refer to a mixture of male and female teachers, i.e. the word is lexically vague concerning sex (gender).

(2) John and Bill both like their teachers.

On the other hand, the word *file* in sentences (3) and (4) is not vague (but ambiguous) in the sense that we do not interpret (3) and (4) as John having a tool and Bill a dossier.

[1] Most of the examples and the explanations in this chapter are based on Lascarides, Copestake, and Briscoe (1996).

(3) John and Bill each have a file.

(4) John has a file and Bill does too.

(3) and (4) are examples of a test for lexical ambiguity. This test seems to be much more usable than the basic ambiguity test since it does not involve prompting the utterer. Instead, the possible interpretations are considered. Here the lexical ambiguity of the word *file* leads to the impossibility of crossed readings. Note that in sentence (2) we used basically the same construction as in (3) and thus the possibility of crossed readings in (2) (i.e. John and Bill liking any combination of male and female teachers) and its impossibility in (3) seems to reflect a difference between the lexical representation of *teacher* (underspecified sex) and *file* (ambiguous, unrelated senses).

9.2.3 Zeugma

The absence of crossed readings along with a figure of speach known as zeugma are often taken to be the most reliable tests for lexical ambiguity as oppsed to vagueness (Zwicky and Sadock (1975), Cruse (1986)). Zeugma is a word punning effect, in which two terms are inappropriately linked together. It can occur with a variety of syntactic constructions, as the examples (5) to (8) indicate:

(5) Some dam busters blew up banks and so did some bank robbers.

(6) I tried to take the plane to Chicago, but it was too heavy.

(7) Heseltine left with a smile, a wave and his wife.

(8) Mr. Pickwick took his hat and his leave.

Zeugma in all of these cases arises from the use of an ambiguous lexical item: *bank* in (5) *take* in (6) and (8) and *with* in (7).

The zeugma test gives the same result for the word *teacher* as the crossed readings test, in that the lack of zeugmatic effect in (9) indicates that the word *teacher* is vague (i.e. underspecified) concerning gender.

(9) Teachers may take maternity or paternity leave.

The phenomena of zeugma and the absence of crossed readings are closely related: intuitively zeugma occurs when incompatible crossed readings are strongly suggested by other information in the sentence.

9.3 How the tests work

9.3.1 The case of multiple lexical entries

The traditional picture of lexical organisation is that distinct word senses should be represented as discreet units, i.e. multiple lexical entries, or MLEs (e.g. Kempson (1977, 81f)). Let's look at how MLEs could account for the absence of crossed readings in e.g. (10).

(10) Texas and Alabama have preservation orders for their most beautiful banks.

First, there must be (at least) two lexical entries for *bank*: *bank_ground* (the mound sense) and *bank_org* (the finantial organisation sense). Moreover, there is no entry for *bank* that is general between these two senses. Then two logical forms are created for (10): one containing the predicate *bank_ground* and the other the predicate *bank_org*. So (10) cannot refer to a mixture of earth mounds and finantial institutions, because neither logical form would be satisfied in this case.

(11) Bank robbers and dam busters blow up banks.

In order to explain zeugma featuring the same lexical items such as in sentence (11), more than just the existence of multiple lexical entries must be assumed. It is necessary to link the compositional semantics to a model of pragmatic reasoning (which is open-ended and involves arbitrary knowledge and interactions as opposed to the relatively simple taxonomic lexical semantic knowledge utilised in linguistic processing), because as with lexical ambiguity in general, the interpreter uses pragmatic information to decide which interpretation of *bank* should be used. Using pragmatic reasoning, a semantically ambiguous sentence like *A bank robber blew up a bank* can be disambiguated. However, in the case of (11) the interpreter cannot use *both* interpretations of *bank* at the same time, even in the case of the distribute interpretation of (11). Knowledge about bank robbers favours the *bank_org* interpretation, while knowledge about dam busters favours *bank_ground*. The pragmatic clues conflict and since neither is more specific than the other, the interpreter cannot decide which interpretation should take precedence. Because of the irresolvable conflict in pragmatic reasoning, the semantic ambiguity cannot be resolved, which produces a zeugmatic effect.

The previous account for zeugma and non-crossed readings is adequate for homonymous words such as *bank*, where senses are unrelated (synchonically). But the cases where the senses are related (such as the physical object and the contents sense of the word *book*) are more frequent. Unfortunately, zeugma and non-crossed readings do not provide reliable tests in these cases; nor is the previous explanation applicable to them.

Another class of examples which cannot be explained the way sketched above involve the repetition of the ambiguous word form, such as (12) and (13).

(12) John banked the money and then he banked the plane.

(13) John had a file and Bill had a file.

(12) is zeugmatic, though the effect is not as pronounced as in (11). Also the crossed readings of (13) are at least strongly dispreferred in spite of there being a possible logical form for them.

9.3.2 The case of constructional polysemy

It can be argued that in many cases when senses are related, multiple aspects of a word's meaning should be encoded in a single entry (see e.g. Copestake and Briscoe (1995)). This is referred to as *constructional polysemy* (in contrast to *sense extension*, where the senses are related but have distinct lexical entries, see Copestake and Briscoe (1995)). Without assuming a single lexical entry, it would be very difficult to account for sentences like those in (14), where the two senses of the word *book* are involved simultaneously.

(14) a. The books on the top shelf are about syntax.
b. That thesis has thousands of pages and is quite unreadable.

On the other hand, sentences featuring the word *book* can also be zeugmatic, as in example (15).

(15) That thesis is orange and unreadable.

Moreover, sentence (16) cannot refor to a situation where Kim wrote and sold two novels to a publishing house while Sandy sold two books in a shop as a saleswoman, that is a crossed reading is not available.

(16) Kim and Sandy both sold two books.

The account of zeugma and the unavailability of crossed readings in Section 9.3.1 depended on there not being a lexical entry which is general between the senses involved. The grammaticality of ((14)a,b), on the other hand, indicates that there must be such an entry. In fact, the zeugmatic effects and non-crossed readings can be explained *on pragmatic grounds* even when there is a single lexical entry.

Note that the non-zeugmatic (14-b) basically differs from the zeugmatic (15) only in that in the latter the predicate *orange* is used instead of the predicate *has thousands of pages*. The two related senses thus either yield a zeugma or not, depending on the particular assertions being made. Moreover, the zeugmatic effect depends also on how those assertions are presented. In contrast to (15), the sentence (17), in which *orange* is not used predicatively, is perfectly acceptable.

(17) That orange thesis is unreadable.

It seems that whether coordinating constituents is zeugmatic or not does not only depend on whether the lexical structure permits co-predication on different aspects of *thesis*, but also on how the assertions rhetorically link together, to contribute to the coherence of discourse. In (14-b) the rhetorical link is strong since the size of books is clearly relevant to their readability, so a causal link is inferrable between the two, and the strong rhetorical link makes the discourse coherent, in the case of (15) on the other hand, the discourse is incoherent since the colour of the cover of a book is not relevant to its readability. In (17) there is no conjunction of predicates, so there is no need to find a rhetorical link.

9.4 Pragmatic reasoning

As we have seen in the previous section, in the case of single lexical entries zeugmatic effects are explainable by the lack of coherence of discourse. We now sketch a model of pragmatic reasoning that makes possible to define how the coherence of discourse affects word interpretation.

9.4.1 Rhetorical relations

As it was mentioned, the coherence of discourse is to a great degree determined by the rhetorical relations that link the assertions that make up the text. Some of the relations are listed in Table 9.1.

Relation	Example
Narration	Max stood up. John greeted him.
Result	John pushed Max. Max fell.
Elaboration	Max painted a picture. He used oils.
Contrast	Max has black hair, but Mary has brown hair.
Parallel	Max has black hair, and Mary has black hair too.

Table 9.1: Rhetorical relations

The different rhetorical relations impose different constraints on the semantics of the assertions they link together. These constraints are summarised in Table 9.2.

Relation	coherence constraints	
Narration	entails that order of events = temporal order	requires distinct common topic
Result	requires causal relation between events	
Elaboration	entails that the event elaborated is the topic	subtype relation between events
Contrast	partial structural isomorphism required	contrasting theme required
Parallel	partial structural isomorphism required	common theme required

Table 9.2: Coherence constraints on rhetorical relations

9.4.2 Interpretation Constraint and Weak Coherence

Lexical sense disambiguation in a discourse context is subject to a constraint that captures the fact that *people try not to infer propositions that lead to weakly coherent or incoherent discourse.* Specifically, interpretations of words that give discourses which are only weakly coherent or incoherent are avoided. We are going to term this constraint the *Interpretation Constraint*.

Many of the the rhetorical relations listed in Table 9.2 require that a topic be identified. In most of the zeugmatic cases, the weak coherence of discourse follows from the fact that the topic of discourse that can be identified is strange. A strange topic results in weak discourse coherence if no explanation for that being the topic of conversion can be nonmonotonically deduced from the interpreter's real world knowledge and the information gathered in the course of the discourse.[2]

[2] The formal pragmatic theory assumed in Lascarides, Copestake, and Briscoe (1996) is DICE (Discourse in Commonsense Entailment). The non-monotonic logic DICE is based on is Asher and Morreau's Commonsense Entailment (CE). DICE and CE are not described here, see Lascarides, Copestake, and Briscoe (1996), Lascarides and Asher (1991) and Asher and Morreau (1991) for details.

9.5 More on how the tests work

9.5.1 Another case of constructional polysemy

In Section 9.3.2, the case of the word *book* was discussed, which we assumed to have a single lexical entry that encodes the *physical object* and the *information container* senses of the word simultaneously. Another type of constructional polysemy is involved in the case of the word *brush* in the sentences in (18).

(18) a. Nylon bristles and plastic handles are used to make cheap brushes.
b. Rembrandt used a brush and so did our janitor.
c. Rembrandt and our janitor used a brush.

The fact that (18-a) is a perfect sentence indicates that there is a vague sense of *brush* the telic ('purpose') attribute of which is underspecified (e.g. Pustejovsky (1991)). On the other hand, (18-b) and (18-c) are zeugmatic. The MLE account of zeugma and the unavailability of crossed readings in Section 9.3.1 cannot explain zeugma in this case, because there is only one logical form for the zeugmatic sentences. An alternative account can be formulated that explains zeugma again on pragmatic grounds. The explanation is based on the way pragmatics affects interpretations. Pragmatic inference involves default reasoning based on real world knowledge and assumptions about the speaker's intentions. (19-a) is thus interpreted as *Rembrandt used a paint brush* by default, while (19-b) as *the janitor used a cleaning brush*. Pragmatic reasoning is needed in the case of (18-b) and (18-c), because the word *brush* is the object of the very general verb *use* in those sentences, so its telic attribute must be specified in order to get an interpretation for the sentences that is not completely devoid of meaning.

Sentence (18-c) is equivalent to the conjunction of (19-a) and (19-b), since applying the generalized quantifiers $\lambda Q \exists x(\text{Rembrandt}(x) \wedge Q(x))$ and $\lambda Q \exists x(\text{janitor}(x) \wedge Q(x))$ of the coordinate NP to the meaning of the predicate $\lambda x(\text{used\_a\_brush}(x))$ yields the same conjunctive formula as the sequence of (19-a) and (19-b).

(19) a. Rembrandt used a brush.
b. Our janitor used a brush.

In order to make the sequence of (19-a) and (19-b) coherent, a rhetorical link that could connect the two sentences must be found. In the case of (19-a) and (19-b), the candidate link would be the *Parallel* relation. The sequence of (19-a) and (19-b) is by default interpreted as (20). Thus the coherence of *Parallel* must be checked in (20).

(20) Rembrandt used a paint brush. Our janitor used a cleaning brush.

But, although it is logically consistent, (20) is at best weakly coherent. *Parallel* requires a common theme. The best theme we could find in this case could be glossed as *people doing something with a brush*. This is not a very good theme, and strong coherence would require that an explanation can be deduced from the information gathered in the course of the discourse why the speaker would want to talk about people using brushes for very different activities. Context could be such that it would make possible such a deduction, but we assume that (18-c) is not uttered in such a context. Since the word interpretations *brush* $\rightarrow$ *paint brush* and *brush* $\rightarrow$ *cleaning brush* result in weak coherence, one or both of these interpretations are invalidated by the Interpretation Constraint, which states that such interpretations must be avoided. But the interpreter cannot decide which interpreatation to drop. So although (18-c) is only pragmatically ambiguous, this ambiguity cannot be resolved and this produces the zeugmatic effect.

9.5.2 Incremental processing

It is established that interpretation of language proceeds incrementally (e.g. Marslen-Wilson and Welsh (1978), Frazier (1979), Crain and Steedman (1985)). The interpreter calculates the semantic content *and* the pragmatic implicature as soon as enough information is available to make an initial decision. This is a necessary strategy, because of limitations of memory, the utility of rapid comprehension and so forth. But this sometimes makes reinterpretation necessary which leads to zeugma. In the case of (18-b) the pragmatic implicatures of *Rembrandt used a brush* are calculated before beginning to parse *and so did our janitor*. This includes specializing *brush* to *paint brush*, which must be retracted when processing *and so did our janitor*, for the same reason as in the case of (18-c) and the same unresolvable conflict occurs. The zeugmatic effect is more pronounced in (18-b) as in the case of (18-c), because it is amplified by the enforced reinterpretation. In the case of (6), here repeated as (21), the zeugmatic effect is caused exclusively by reinterpretation.

(21) I tried to take the plane to Chicago, but it was too heavy.

9.5.3 Zeugma and non-crossed readings with word repetition

Zeugma is created in the case of (22), which features the polysemuous word *bank* having multiple lexical entries, by the need of reinterpretation and weak coherence.

(22) John banked the money and then he banked the plane.

The similarity in the syntactic structure and the word forms implies a *Parallel* relation (besides *Narration*) between the propositions in (22). But as the two occurences of *banked* are pragmatically disambiguated (the meanings of the object arguments provide the clue for the disambiguation), the meaning of the sentence amounts to (23).

(23) John deposited the money. He turned the plane.

(23) is at best weakly coherent. There is in fact no *Parallel* and not even an acceptable common topic for *Narration* can be found. So the interpretations of *banked* are blocked again by the Interpretation Constraint. Moreover, due to incremental processing, the reinterpretation of the word *bank* is forced when processing the second part of the sentence. (24) is zeugmatic for exactly the same reason.

(24) Rembrandt used a brush and our janitor used a brush.

9.6 Conclusion

In this chapter we discussed what ambiguity and vagueness are, and we reviewed some tests that can be used to decide whether an utterance or a lexical item is ambiguous or vague. The tests that were proposed for lexical ambiguity are the absence of crossed readings and zeugma. We have seen how these tests work, and that they have to be used with caution in the case of words that feature constructional polysemy. In the case of such words the zeugmatic effects and the absence of crossed readings are exclusively created during the disambiguation process involving pragmatic reasoning by the lack or weakness of coherence in the discourse.

In the case of words with multiple lexical entries, pragmatic reasoning also plays a role in the creation of zeugma, but unless word repetition is involved, it is not weak coherence that results in zeugma but the fact that a single occurence of a word cannot have two unrelated interpretations. The zeugmatic effect is often amplified (and in some cases exclusively caused) by the fact that due to incremental processing, reinterpretation is enforced at some point of processing the sentence.

Finally, Table 9.3 summarizes the examples, their interpretations and the mechanisms involved.

entry type	word	example	effect	occurence	mechanism
MLE[a]	file	John and Bill each have a file.	NCR[b]	single	no crossed logical form
MLE	bank	Bank robbers and dam busters blow up banks.	zeugma	single	multiple pragmatic interpretations of a single word are impossible
MLE	file	John has a file and Bill does too.	NCR	anaphor	weak coherence and the need for reinterpretation exclude crossed readings
MLE	file	John has a file and Bill has a file.	NCR	repeated	weak coherence and the need for reinterpretation exclude crossed readings
MLE	take	I tried to take the plane to Chicago, but it was too heavy.	zeugma	single	reinterpretation is enforced
SLE/MA[c]	book	The books on the top shelf are about syntax.	no zeugma	single	one predicate is attributive, thus there is no rhetorical linking
SLE/MA	book	Kim and Sandy both sold two books.	NCR	single	weak coherence excludes crossed readings
SLE/MA	thesis	That thesis has thousands of pages and is quite unreadable.	no zeugma	single	strong coherence (causal link)
SLE/MA	thesis	That thesis is orange and is quite unreadable.	zeugma	single	weak coherence (strange topic for *Parallel*)
SLE/VF[d]	teacher	Teachers may take maternity or paternity leave.	no zeugma	single	strong coherence (world knowledge)
SLE/VF	teacher	John and Bill both like their teachers.	vagueness/CR[e]	single	vague feature may remain vague (disambiguation not needed for interpretation)
SLE/VF	brush	Nylon bristles and plastic handles are used to make cheap brushes.	vagueness	single	vague feature may remain vague
SLE/VF	brush	Rembrandt and our janitor used a brush.	zeugma	single	weak coherence blocks disambiguation of vague feature
SLE/VF	brush	Rembrandt used a brush and so did our janitor.	zeugma	anaphor	reinterpretation and weak coherence
SLE/VF	brush	Rembrandt used a brush and our janitor used a brush.	zeugma	repeated	reinterpretation and weak coherence

[a]multiple lexical entries
[b]non-crossed readings
[c]single lexical entry with multiple aspects
[d]single lexical entry with vague (underspecified) feature
[e]crossed reading

Table 9.3: Summary of examples

References

Asher, N. and M. Morreau. 1991. Commonsense reasoning – a modal theory of non-monotonic reasoning. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, pages 387–392, Sydney, Australia.

Copestake, A. and E. J. Briscoe. 1995. Semi-productive polysemy and sense extension. *Journal of Semantics*, 12:15–67.

Crain, S. and M. J. S. Steedman. 1985. On not being led up the garden path: The use of context by the physiological parser. In L. Karttunen D. Dowty and A. M. Zwicky, editors, *Natural Language Parsing*. Cambridge University Press.

Cruse, D. A. 1986. *Lexical Semantics*. Cambridge University Press, Cambridge, England.

Frazier, L. 1979. *On Comprehending Sentences: Syntactic Parsing Strategies*. Indiana University Linguistics Club, Bloomington, Indiana.

Kempson, R. M. 1977. *Semantic Theory*. Cambridge University Press.

Lascarides, A. and N. Asher. 1991. Discourse relations and defeasible knowledge. In *Proceedings of the 29th Annual Meeting of the Association of Computational Linguistics*, pages 55–63, Berkeley, CA.

Lascarides, A., A. Copestake, and E. J. Briscoe. 1996. Ambiguity and coherence. In A. Copestake and A. Lascarides, editors, *Lexicon and Pragmatics. Reader.* ESLLI, Prague.

Marslen-Wilson, W. and A. Welsh. 1978. Parsing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology*, 10:29–63.

Pustejovsky, J. 1991. The generative lexicon. *Computational Linguistics*, 17(4):409–441.

Zwicky, A. M. and J. M. Sadock. 1975. Ambiguity tests and how to fail them. In J.P Kimball, editor, *Syntax and Semantics IV*. Academic Press, New York, pages 1–36.

Chapter 10

Metonymy

This chapter mainly focuses on a phenomenon called **constructional polysemy**. This kind of polysemy, called constructional because the items involved in it are disambiguated on the phrase level (i.e., when put into a **construction**). The type of metonymy which, at the same time, qualifies as constructional polysemy is called **logical metonymy**.

The first section gives the reader a general idea of what we call logical metonymy and explains whether and what kind of a lexical semantic treatment of the phenomena is necessary.

The second section begins with a survey of **lexical decomposition** and the conception of **Generative Lexicon Theory** (Pustejovsky (1995; Pustejovsky (1996b; Pustejovsky and Boguraev (1995)) about it. One of the central issues in GL is the treatment of **constructional polysemy**, the formal treatment of which is discussed afterwards.

The third chapter discusses to what extent and in what ways metonymic constructions are constrained in the lexicon.

The chapter is concluded with an exposition of further issues related to constructional polysemy and logical metonymy.

10.1 Logical metonymy as constructional polysemy

10.1.1 Logical metonymy and the lexicon

Metonymy is essentially a figure of speech, i.e., a rhetorical rather than linguistic device. However, there is one type of metonymy, called **logical metonymy**, which seems more linguistically anchored than the general case. We will only deal with this type in what follows.

The verb phrases in (1) are well known examples of **logical metonymy**.

(1) a. finish the book/beer/sandwich/cigarette
    b. begin the book/beer/sandwich/cigarette
    c. enjoy the book/beer/sandwich/cigarette
    d. want a book/beer/sandwich/cigarette

Logical metonymy is not restricted to verb + direct object structures.

(2) a. The book took two days.
    b. The movie annoyed me a lot.
    c. During my hourly cigarette, someone came to me.

Observe that, for instance, (1-a) can be paraphrased as "finished reading the book" in which the verb's argument is of semantic type *event* which is not normally the type associated with *book*.

The question immediately arises whether this is a pragmatic or lexical problem. After all, knowing in what ways one's finishing a book differs from one's finishing a beer is nothing more than knowing in what ways a book differs from a beer, in particular what one normally does to them. The use of these constructions, however, is much more restricted than what simply the possibility of drawing colloquial inferences based on chunks of world knowledge would alone justify. Note also that, though the natural interpretation of logical metonymies can be overridden given strong contextual information (see (3)) (see Copestake and Briscoe (1995) for a discussion), these constructions doubtlessly have default and conventionalized readings (see Viegas (1995) on lexicalisations).

(3) a. My goat enjoyed the book.
    b. Book-sellers prefer postcards around Christmas.

An adequate explanation of this phenomenon therefore is only feasible if one imagines the distinctions are a result of the different lexical (semantic) features of the items involved. Therefore GL, in accordance with many other theories, chooses to give an account of logical metonymy — at least partly — within the lexicon.

10.1.2 Metonymy is not verbal polysemy

Note that in order to be able to draw the desired inferences from the various uses of say *finish*, one may introduce **meaning postulates** stating, for instance, that every time you *finish a book*, it is implied that you `finish reading it`. This, however, is not necessarily true, as in an appropriate context *finish the book* can fairly mean `finish writing the book`, which leads to either exclusively disjunctive meaning postulates or to the explicit introduction of two senses for *finish*.

One could choose the latter solution anyway without using meaning postulates. This is what traditional **sense enumeration lexicons** do when they enumerate the possible senses of *finish*, one of which would be `finish reading` and which selects for an argument denoting "something that one usually reads".

Even though sense enumeration leads to an immense proliferation of lexical entries, it is still unable to cover each sense, given the very general and creative use of metonymic constructions.

An even more serious problem with sense enumeration of the verb senses is that the polysemous behavior is independent of the actual verb. The default interpretations of the phrases below are a uniform indication that the resolution of the elliptical metonymic link is determined by the argument, and often remains constant with different verbs.

(4) a. finish/begin/enjoy/want a book
    b. finish/begin/enjoy/want a sandwich

Note that in (4-a), the missing link is *reading*, in that it is really *reading the book* that one finishes, begins, enjoys or wants. Enumeration of verbal lexical entries corresponding to different phrasal senses misses this generalization. (See Pustejovsky (1995, chapter 4) for further discussion.)

10.1.3 Metnonymy is not nominal polysemy

Based on the above observations, one could alternatively say that *book* is metonymically interpreted as `reading the book`. If again we wish to enumerate the metonymic senses of words, the most serious problem is that there will be a symmetry between the available senses in the lexicon. As a matter of fact, a very strong asymmetry is apparent between a word's basic or primary sense and its metonymic interpretations. This intuitively very clear statement is confirmed by the fact that the use of metonymic senses is much more restricted than that of basic senses in contexts where the semantic type would otherwise be adequate; (see section 10.3 for a discussion).

The enumeration of the various senses is also problematic in as much as it is hard to find a conceptual or formal difference between the treatment of **contrastively ambiguous** words (words with unrelated multiple meanings, eg. *bank* `institution` or `mount of earth`) and constructionally polysemous ones (Pustejovsky (1995, chapter 2)).

Their distinctness, however, is strongly suggested by the fact that in the case of constructional polysemy one and the same lexical item can be metonymically interpreted with respect to one predicate while having its original sense with respect to the other. In other words, the different senses (see (5)) can be copredicated. (See Copestake and Briscoe (1995) and section 10.3 and 10.4 for more about copredication.)

(5) a. Sally finished and threw away her book.
    b. Orsi began the sandwich but then dropped it.

This is even more apparent in (6). Logical metonymy is not restricted to verb phrases at all, modifiers + noun phrases also exhibit this phenomenon.

(6) a. fast typist/car/runner/motor-way/coffee
b. long record/book/dinner/road/trip/film

What was said about verb phrases carries on to this kind of construction (Boullion and Viegas (1994)), with the adjective being the semantic head like the verb in Verb Phrases. Observe here that the different senses of modification involve a metonymic interpretation though the whole Noun Phrase retain the semantic type of the Noun's primary sense, eg. a *fast typist* is `typist who (usually) types fast`.

This urges one to conclude that handling logical metonymy with sense enumeration of either the semantically active elements or the arguments would be an approach on the wrong track. Instead one should provide an account in which the specific senses only emerge in constructions. The lexical item itself is represented in an underspecified manner, giving rise to a metonymic interpretation in the right context. Obviously as one and the same lexical item can be involved in several constructions at the same time, it may contribute to the phrase- or sentence-level meaning in diverse ways, ie. with different senses.

10.2 Generative Lexicon Theory and Metonymies

Generative Lexicon Theory (GL) (Pustejovsky (1995; Pustejovsky (1996b; Pustejovsky and Boguraev (1993)) provides a theory and representational formalism which integrates lexical decomposition and underspecification as well as provides formal devices to treat constructional sense specification. These are explained in detail in this section. (See Pustejovsky (1995) for detailed exposition, Pustejovsky (1996b) for a quick introduction and J. Fodor (1997) for a critical discussion and T. Briscoe and Copestake (1993) for more formal issues.)

10.2.1 Lexical decomposition

According to many (Boguraev and Pustejovsky (1990; Boguraev and Pustejovsky (1994; Pustejovsky and Boguraev (1993; A. Copestake and de Paiva (1993; Anick and Bergler (1992)) working on natural language lexicology, one of the most pressing problems of semantics — namely the compositional interpretation of sentences — defies any kind of solution without **lexical decomposition**. Lexical decomposition here is meant to suggest that in an ideal lexicon much more information about a word's meaning has to be contained than is the simplistic assumption in many theories (e.g., syntactic or semantic type information).

It can immediately be recognized that even if one wishes to capture some generalizations on selectional restrictions, one usually turns to features which are semi-syntactic in the sense that in a lot of cases they cannot be read off the unanalysed semantic characterization of the item in question. These features, though, in the majority of cases, very strongly related to semantic properties of the item. A plausible relation between semantic and selectional features even in straightforward cases is, however, only feasible with the introduction of meaning postulates, since non-decomposed (atomic) meanings assigned to lexical items deny us the chance of drawing inferences from them. In conclusion if one wants to be able to build semantic classes, — the semantic and syntactic behavior of the members of which is more similar than one could afford to miss a generalization, — one has to rely on a plausible and restricted method of decomposition, ie. decompose meanings into a small and manageable set of **primitives**.

10.2.2 Generative Lexicon Theory

The main objective of GL is to give an account of the different kinds of polysemies as well as to attempt at explaining the creative use of words in context. GL argues that by developing an ontologically plausible and formally rather restrictive framework of lexical semantic representation, operations of which correspond to legitime semantic operations such as word disambiguation in constructions or metonymic interpretation, one could also get closer to an explanation of various aspects of syntactic behavior on grounds of well-motivated semantic classes.

Another very important goal is to establish clearly the notion of semantic well-formedness. This would stand to provide for a clear-cut characterization of possible word meanings. It is, however, very hard to limit an adequate domain for semantic well-formedness, as the highly flexible structure of discourse meaning is best thought to be composed from the relatively independent interpretations of various semantic levels including lexical, discourse or temporal structure. Given this view it is obvious that lexical semantic representation is only one of the many levels which are at play in determining discourse interpretation.

In GL the actual implementation of lexical decomposition is driven by the intuition that natural language semantics is the image of nonlinguistic conceptual principles. This is to say that in many cases the dimensions of decomposition rhyme with the dimensions of our cognitive operations. (For a detailed discussion of a semantic theory built upon this principle see 8.) The latter are said to be well approximated by Aristotle's Modes of Explanation (Moravcsik (1975)), which in turn serves as an ontological basis for Pustejovsky's **qualia structure** discussed in detail in section 10.2.3 (see Pustejovsky (1995, chapter 5,6)).

10.2.3 Qualia structure

GL takes the view that the semantic properties of lexical items can be defined with the help of a small set of "basic substances" and some universal features of these (Pustejovsky and Boguraev (1993; Pustejovsky (1995)). A rather workable implementation of the latter in GL is **qualia structure** that is inspired by and named after the **Aristotelian causal categories**.

Dimensions of metonymic interpretation

One of the most straightforward advantages of lexical decomposition along the categories of qualia structure is that it makes it possible to capture some generalizations on the nature of metonymies.

Consider again some of the examples in (1) repeated here as (7) for convenience.

(7) a. finish the book (reading)
b. begin the beer (drinking)
c. enjoy the sandwich (eating)
d. want a cigarette (smoking)

(8) a. good knife (good for cutting)
b. bright bulb (shining brightly)
c. fast motor-way (can drive fast on it)
d. long record (playing it takes long time)

The events in brackets are the events of the most likely metonymic interpretations of the phrases. We argued earlier that these events are very strongly determined by the arguments' semantics. If one doubts that they are just arbitrary events associated with each nominal object that are available in the metonymy, one has to find in what ways the respective relations between the events and the objects are similar. Well, in (7), it is apparent that the events name the most common event associated with the object, ie. its **usual purpose**. Similarly in the examples of (8), the relevant event linking the adjective to the modified noun is the event associated with the usual purpose of the object.

In the examples in (9), we encounter a different dimension of metonymy. Here, the linking event can be the process that is responsible for the creation of the object or contributed to the **coming into being** of the nominal in question.

(9) a. finish the dinner (cooking)
b. forget the letter (to write)
c. begin a new film (making)
d. regret an article (writing)

In general then we might say that certain aspects of meaning are semantically relevant properties that, in order to capture generalizations on the possible dimensions of metonymic interpretation, have to be included in the lexicon as lexical features.

In GL, it is assumed that these dimensions are relevant in the case of each and every lexical entry and qualia structure is thought to be a necessary component of the lexical representation of words. It is obviously not meant by this that a lexical entry has to have a fully specified value for each and every qualia attribute, only that these are appropriate features for all lexical items — though probably with a slightly diverse semantics for individual- or event-denoting words.

In the following, we give the four essential aspects of meaning assumed in GL and summarize what they are meant to describe (Pustejovsky (1995, 5.4)).

- CONSTITUTIVE: the relation between an object and its constituent parts
- FORMAL: that which distinguishes it within a larger domain
- TELIC: its purpose and function
- AGENTIVE: factors involved in its origin or "bringing it about".

In (10), you see the tentative representation of the qualia structure for *novel* in an Attribute-Value Structure (AVS) representation. (See T. Briscoe and Copestake (1993) for representational issues.)

(10)

novel	
CONSTITUTIVE	*narrative*
FORMAL	*book*
TELIC	*reading*
AGENTIVE	*writing*

Further reading on the application of qualia

It is thoroughly reported in Johnston and Busa (1996) that generalizations about the usual interpretation of a great deal of nominal compounds is very easily captured with a reasonable application of qualia roles.

In Busa (1996b) the semantics of agentive nominals is based on GL's conception of qualia structure.

Pustejovsky (1996b) describes, among many other things, in what ways qualia roles can help in explaining privative metonymic reconstruction of the arguments of verbs like *risk*.

In GL, qualia structure is not restricted to the representation of nominals (Pustejovsky and Anick (1988; Busa (1996a),Busa (1996b)) but also relevant in the case of verbal or adjectival entries. The dual behavior of complex events, for instance, also urges one to analyze transitions as involving a process that is followed by a resulting state. These two events are dimensions of the predicate's meaning and are directly encoded as qualia values of the FORMAL and AGENTIVE aattributes, respectively as argued in Pustejovsky (1995, chapter 9). (See also Pustejovsky (1988), Pustejovsky (1991), Pustejovsky and Boullion (1995), Pustejovsky and Busa (1995) for more on events in GL).

In (Pustejovsky, 1995, 10.2), Pustejovsky argues that the distinction between stage-level and individual-level predicates is best explained with a difference in the qualia structure of predicates rather than not an event type distinction.

10.2.4 Problems with qualia

Metaphysics and the lexicon

The apparently plausible inventory of qualia roles as well as their formal treatment, though it seems convincing, is not without problems. As it is only our intuition about meaning that restricts what can and what can not be the qualia role of a given item, the restrictedness of the theory is highly questionable. Only when one can give a foolproof definition of what it means to be the TELIC role of a word, can GL be viable. Obviously a correct characterization can only be achieved if we have the intuitive metaphysical categories encoded in the structure of the world model used to interpret representational formalisms of our semantics. Until that time GL can only attain explanatory adequacy with respect to a rather tentative model of our socio-culturally determined naive metaphysics. (See Pustejovsky (1995a), Copestake and Briscoe (1995), Busa (1996a) for further discussion.)

Problems, however, already occur if one chooses to ignore the above considerations and is happy with relying on one's intuitions.

Qualia and flexibility

According to an intuitive interpretation of the TELIC qualia role, the usual purpose of *perfume* is to make its wearer smell better. The metonymic construction *enjoy the perfume* can hardly be properly interpreted with assuming that *enjoy* selects the TELIC role of its argument. Rather some default event of perception (see section (30)) is understood to fill in the interpretational gap in the construction. This is also the case with a lot of arguments whose semantic type clearly have a proper TELIC role at least in commonsense metaphysics (see (11)).

(11) a. *enjoy the key = enjoy opening the door with it
b. *enjoy the chair = enjoy sitting on it
c. *enjoy the window = enjoy opening it

In fact, it is rather a challenge to figure out whether the Aristotelian categories are too rigid or not fine grained enough to enable the linguist to give a purely semantic account of metonymy or the latter is just out of the question and one always has to turn to blind listing of unexplainable or exceptional use. In section 10.3, we attempt at characterizing some of the very intricate lexical constraints on out-of-context metonymic interpretation mainly by providing examples.

10.2.5 Means of Semantic Composition

As we saw in the previous section, the regularities of an elliptical structure such as a logical metonymy can be captured if one allows for the semantic argument in the compositional structure to influence the sense of the phrase. This view necessarily implies the the need to reject the traditional "active functor / passive argument" view on composition. In order to allow for this kind of semantic composition, one has to provide an alternative formal treatment of new compositional rules that, however, should be fairly general and restricted for the calculus to attain explanatory adequacy while at the same time should preserve compositionality.

Such means of composition are achieved through some **generative devices** in GL. The reason for this somewhat clumsy name lies in the idea that generative devices — like lexical rules in other theories — can be thought to expand the lexicon by "productively" generating the available senses out of **underspecified** lexical entries. These generative devices include **type coercion** and **selective binding** among others not discussed here. (See also section 10.4.2. We refer the interested reader to Pustejovsky (1995, chapter 7).)

As shown below in section 10.3, some lexical items seem to be sensitive to whether their argument is a "generated" or an "original" sense. Though one can think of the operations of generative devices as productive, — since generated elements are not fully interchangeable with words of the same type — one has to keep track of how the senses came into being. For this reason as well as because the formal status of generative devices within the lexicon is not sufficiently clarified anyway, we can fairly assume that these are really compositional operations the availability of which can in some sense be restricted by the semantic head.

Type coercion (Godard and Jayez (1993)) stands for the situation in which the predicate coerces its argument into the semantic type required when the former would otherwise fail to combine with the argument's root type. This type-mismatch resolution can be thought of as a special semantic operation in the course of which the predicate functor applies not to the root type of the argument but to a special function of the root, say a qualia value.

Interestingly a number of logical metonymies can be explained as type coercions.

(12) a. enjoy the novel = enjoyed reading the novel (TELIC)
b. enjoy the novel = enjoyed writing the novel (AGENTIVE)

(13) a. finish the beer = finish drinking the beer (TELIC)
b. want a cigarette = want to smoke a cigarette (TELIC)

As the name *coercion* suggests type coercion is a marked operation. Its alternative, when the semantic head is applied to the argument's root type (primary sense), reduces to a very straightforward function application, an unmarked operation. This parallels the markedness properties of these constructions as well as the asymmetry mirrored in the feature structure representation.

The operation of **selective binding** strongly resembles type coercion with the difference that here the head semantic type gets the coerced interpretation in the context of a modofier. In the case of *fast typist* the head semantic type is a `typist`, but the modification translates as `fast typing`.

10.3 Productivity of metonymies

It is obvious that the unrestricted use of type coercion together with lexical items having a richly specified qualia structure leads to a serious degree of overgeneration. In this section we briefly survey the constraints on qualia coercion (Pustejovsky (1993; Pustejovsky (1995c)).

10.3.1 Passive and active selection

The grammaticality of coercive constructions is not simply the function of the selectional restrictions of the coerced event predicate. It is not the case in the `finish reading` sense of *finish* that the verb's selectional criteria parallel those of *read* as the below example illustrates, ie. you can not finish everything you can read.

(14) a. read the lines/advertisement/instructions/
b. ??finish the lines/advertisement/instructions/

The generative devices are supposed to explain in a way the seeming polysemies of lexical items which dissolve in a more or less predictable way when in composition. Allowing this way of composition might well exempt us from having to specify a potentially infinite number of senses for the lexical items. Coercion, however, is really not an unrestricted operation. The view is taken in GL that the coercion operations available in a composition are lexically restricted in the predicates' representation. That the applicability of the specific coercion operations is idiosyncratically specified in the case of a particular predicate is thought to be motivated by the contrast between **active-** and **passive selection predicates**. The next example show the diverse coercive behavior of some verbal predicates.

(15) a. finish/*stop a book
b. begin/*start a book

10.3.2 Qualia selection

Note that in most cases there can be several patterns of coercion associated with a predicate. The choice is determined by the selectional restrictions imposed on the thematic argument. Cases arise, however, when there are special constraints on the accessibility of a particular quale value. This refers to situations where a predicate known to be able to coerce its argument refuses to coerce along a certain qualia path though the thematic argument would be of the appropriate type. This can be called **qualia selection**, which refers to the property of the predicate in that it selects for particular qualia roles to bind. The adjective *long*, for instance, selects clearly for TELIC and does not coerce the other qualia values in its coercive sense.

(16) a. long book/record = it takes long to read it/listen to it (TELIC)
b. *long book/record = it takes long to write it/record it (TELIC)

10.3.3 Further idiosyncrasies

The problem is that qualia selection is not argument-insensitive. This is to say, — as the examples in (17) and (18) show, — that a normally accessible qualia value (b) is not available for semantic binding (a) for a predicate that is normally coercive with respect to a certain qualia role (c).

(17) a. *?enjoy the chair = enjoyed sitting on the chair

b. use the chair = sit on the chair
c. enjoy the movie

(18) a. *begin the road = begin driving on the road
b. use the road = drive on the road
c. begin the book

Observe that the following example show how utterly problematic it is to assume that restrictions on logical metonymies derive merely from the selectional properties of the active predicate and semantic properties of the argument.

(19) a. ?*fast liquid
b. fast water
c. fast beer

Note that *fast water* is best interpreted as water that runs (flows) fast, the metonymic linking event being a property (a quale value, eg. FORMAL) that is very likely to be shared by the supertype *liquid*. Still the semantic binding of this qualia value is impossible for one and the same (strongly coercive) adjective *fast*. This would imply that active and passive selection is not only dependent on the arguments semantic properties but might have to refer to arbitrary individual types — or even worse: to arbitrary individual lexical items.

10.4 Non standard composition in other constructions

In this section we give further examples of contextual disambiguation of certain polysemies. Given a more flexible theory of lexical decomposition, similar kinds of non-standard compositional devices are also likely to explain these instances of constructional polysemy.

10.4.1 Non-qualia coercion: arbitrary metonymies?

The case of *Remember*

The different patterns for *remember* taking Verb Phrase or clause arguments clearly exhibit at least three different senses (non-factive (nonF), factive (F) and embedded question (EQ)).

(20) a. I will remember to phone you. (F)
b. I remember phoning you. (nonF)
c. I don't remember what the name of his father was. (EQ)

As a matter of fact, *remember* licenses an NP as an argument. One should note that a very special type of NP coercion is available in the case of the embedded question reading. Here, the explicit reconstruction of an embedded question from the NP — though not very easy — is almost unambiguous but still shows a striking diversity across different NP types.

(21) a. remember the book/piece/house/coffee = remember what it is like/is about (EQ1)
b. remember the name/address/vocation/number = remember what it is (EQ2)

Interestingly, if the NP's denotation is an event the interpretation is ambiguous between the non-factive and the factive reading:

(22) You always make me remember the phone-call/party. (nonF[E]/F[E])

The next example shows that a usual coercive interpretation, that is qualia value coercion, is not available for *remember* with an NP argument.

(23) a. *?remember the book = remember to read/write the book (nonF)
b. *remember the book = remember reading/writing the book (F)

Examples (23-a) and (23-b) show that either in a non-factive or in a factive sense *remember* is reluctant to coerce its argument's qualia roles.

Reference transfer

One of the commonest types of classical rhetorical metonymy, namely, when 'the name of something is used instead of the name of something else', is also the least determined lexically. In particular, it is a type of non-qualia coercion. Nunberg (1995) observes that metonymy is not only possible with contentful words, but also with pronouns:

(24) a. *This is parked out back.* (pointing to the keys of a car)
b. *I am parked out back.*

This phenomenon can be called **reference transfer**. It seems as if the pronoun referred to a car in (24-a), as shown by co-ordination tests:

(25) a. *This is parked out back and does not want to start.*
b. *#This is parked out back and fits only the left door.*

On the other hand, the pronoun in (24-b) seems to refer to the speaker:

(26) a. *I am parked out back and have been waiting for half an hour.*
b. *#I am parked out back and may not start.*

We will not dwell on the question what makes it possible for the pronoun *this* in (24-a) to refer to the car rather than the key; Nunberg calls this a **deferred indexical reference**, whereby a demonstrative or indexical expression refers to an entity that stands in a salient relationship to the actual object. One way of telling the two types of metonymies apart is that deferred indexical reference is incompatible with definite descriptions, whereas the type that we are concentrating on (i.e., the one in (24-b)) is not:

(27) a. *#The key I'm holding is parked out back.*
b. $^{\text{OK}}$*The man with the cigar/Mr. McDowell/... is parked out back.*

The phenomenon illustrated in (24-b) and (27-b) must be attributed to the reference of the **predicate** *park* rather than the noun phrase subject (*I, the man with the cigar* etc.) or the construction in which they occur. This is why Nunberg calls this phenomenon **predicate transfer**. We have seen one argument to this effect, namely, that the reference of the noun phrase corresponds to its literal meaning (which is further supported by the fact that it does not matter whether the noun phrase is a pronoun, a description or a proper name). Another argument comes from the fact that the identity of the predicate matters. Consider (26-b): from the two co-ordinated predicates, *park out back* can co-occur with *I* in the sense 'my car is parked out back', whereas *may not start* cannot (at least not in the sense 'my car may not start'). Similarly, one cannot felicitously say *#I got broken yesterday* in the sense 'my car got broken yesterday', and so on.

Nunberg's analysis also makes interesting predictions on predicate transfer. He formulates the following condition on predicate transfer:

(28) Let $\mathcal{P}$ and $\mathcal{P}'$ sets of properties that are related by a salient transfer function $g_t\colon \mathcal{P} \to \mathcal{P}'$. Then if F is a predicate that denotes a property $P \in \mathcal{P}$, there is also a predicate F', spelt like F, that denotes the property P', where $P' = g_t(P)$.

Clearly, the type of predicate transfer that we have seen in (24-b) and (27-b) is much more special than this formulation, namely, it involves a salient function h from the bearers of property P to the bearers of property P' (e.g., a function from cars to their owners). The reason why Nunberg opts for this rather general formulation is that, in this way, the type of phenomenon at hand can be treated on a par with others, such as metaphors, where no such direct association between the bearers of the properties is involved. (For example, the association between the symbol ∪ and the word *cup* does not involve a mapping from particular instances of ∪ to particular cups.)

Now, what can be the relation between the functions g_t and the functions h? There are two typical cases, according to Nunberg, namely:

(29) a. $g_t(P) = \lambda y \forall x((h(x) = y) \to P(x));$
b. $g_t(P) = \lambda y \exists x(h(x) = y \wedge P(x)).$

If we replace P with 'be parked out back' and h with 'owner of [a car]', we get $g_t(P)$ meaning 'people all cars of whom are parked out back' in (29-a), whereas 'people some car of whom is parked out back' in (29-b). Since we normally assume people have just one relevant car when something like (24-b) is uttered (hence, the two readings coincide), we have to pick different examples to see that predicate transfer is in general ambiguous between these two readings:

(30) a. *We are in Chicago* (says the accountant of his firm);
b. *I am in the Whitney* (says the painter).

These examples are truly ambiguous between the universal and the existential readings, although (30-a) is preferably interpreted existentially because of the fact that, according to our knowledge of the world, it is not costumary to have all one's paintings in musea.

The interesting aspect of the above treatment based on predicate transfer is that it can be easily extended to other cases of metonymy, namely, the cases of systematic metonymy studied by Pustejovsky (1995b). For example, *finish a beer* may be analysed as involving the transfer of the meaning of *finish* using a function h assigning quantities of beer to the event of their drinking. As a matter of course, the existence of such functions h may depend on the lexical content of the predicate as well as on the meaning of its relevant argument. This means that the approach based on predicate transfer is a possible formalization of, rather than an alternative to, Pustejovsky's solution. Still, it may have interesting consequences for semantic analysis.

Lexically fixed coercions

It seems, however, that *remember* can actually coerce its argument into an underspecified event type that denotes some sort of "bringing along". What is more interesting is that this is only available for the non-factive use.

(31) a. remember the book = remember to bring/have the book (nonF)
b. *remember the book = remember bringing/having the book (F)

Such underspecified event coercions (see (32)) are very common, though in most cases it is specified by an available qualia role of the same type.

Some orientation verbs (eg. (32-a)) generally coerce into an event of an underspecified **perception** type, which gets specified by the arguments properties, whereas some of them (eg. (32-b)) coerce into events of type **possession**. Aspectual verbs prefer events of type **creation** as their default coerced event type (eg. (32-c)).

(32) a. like the perfume/picture/play/music
b. want a key/bird/child/bike
c. finish the chair/sculpture/bike/road

The following phrases in (33) show how very strange the interpretation of metonymic constructions is. Note that exclamation marks indicate that the a default coercive reading is unavailable or at least a coerced specific qualia value reading is preferred.

(33) a. enjoy this ?road/movie/book/cigarette/apple/beer
b. want a road/movie/!?book/!cigarette/!apple/!beer
c. finish/begin the road/!?movie/!book/!cigarette/!apple/!beer

Different senses?

The reason we called the default coerced sense of *want* coerced at all was simply because we wanted to generalize over its semantics in order to further reduce the set of different senses. This is also what we had to do with words enabling more than one type of true coercion. The tests, however, show that diffenent types of true coercions go together in one coordinate argument of the verb while the verb's lexical default event coordinated with a true coerced argument triggers a zeugmatic effect in the case of *want*:

(34) a. ?#I want a bike and a cigarette. (i.e. having/smoking)
b. I want a beer and a cigarette.(i.e. drinking/smoking)

Note also that the different NP-complemented uses of *forget* are also mutually exclusive in as much as they sound funny when coordinated:

(35) a. #I forgot the keys and the house. (nonF + EQ1)
b. #I forgot her name and my phonebook. (EQ2 + nonF)
c. #I forgot the phone-call and therefore the party. (F[E] + nonF[E])
d. #I forgot the birthday and the presents. (nonF[E] + nonF)
e. #I forgot the journey and his name. (F[E] + EQ2)
f. #I forgot the visit and the faces. (F[E] + EQ1)

10.4.2 Qualia non-coercion: light verb specification

There are actually other phenomena where qualia values seem to be important in the explanation of polysemous behavior. The following example shows that light verb *use* when combined with its NP argument, undergoes sense modification. Needless to say that these senses are disambiguated in the local context and systematically depend on the NP argument's semantic structure.

(36) a. use the knife/contact lenses/pills/chair/road
b. have a shower/meal/party/class

We might say that here also semantic composition provides access to the argument NP's qualia value for semantic binding. Here the non-root (coerced) type will not be the argument of the uniform event-taking predicate, but actually replaces the underspecified predicate itself. As it needs distinct formal treatment — namely unification (specification) of a general event type with a qualia value of the argument —, in GL this phenomenon is treated with another distinct device called **light verb specification**.

10.4.3 Non-qualia non-coercion: further instances of constructional polysemy

In what follows we shall show some examples of polysemous adjectives, which by virtue of clever decomposition allow treatment that shows the explanatory power of underspecification and non-standard composition.

(37) a. noisy dog/room
b. sad event/man

Note the different kinds of modifications in the case of the adjectives in (37). Totally diverse kinds of inferences can be drawn from one phase than the other, and nevertheless one has the intuition that *sad* just means the same.

The idea would be to turn this into a structural ambiguity. If one would decompose *sad* into a "causation" and a "state", one would end up with a scene where semantically there are two arguments. If the adjective/noun (adjectival modification) construction in syntax is stated generally enough and the adjective's restrictions enable the binding of more than one semantic argument, the polysemy is simply a result of underspecification. In our case *man* would be a different thematic argument in the semantic representation of *sad* than the word *event*. The unambiguity of which results from obvious type restrictions on thematic argumetns. Before existential closure of the unexpressable implicit argument, the phrases' semantic representation would look something like (38).

(38) a. a sad man = $lPle\underline{x}(P(x) \wedge MADE(e, SAD(x)))$
b. a sad event = $lPlx\underline{e}(P(e) \wedge MADE(e, SAD(x)))$

And finally observe that there are several other (constructionally) polysemous adjectives that would allow a similar underspecification/decomposition treatment and that also have their own arbitrary restrictions saying which arguments are available for semantic binding.

(39) a. sad/happy dog/event/day/?room
b. frightened/nervous dog/*event/*day/*room
c. sunny *dog/*?event/day/room lively dog/*?event/*?day/?room

d. noisy/silent dog/*?event/*?day/room
e. quiet dog/event/day/room

10.5 The need for generalizing constructional polysemy

Finally, we would like to argue that the types of metonymy discussed so far could be viewed as special cases of a much more general mechanism, at work in a much larger domain of language, which is usually referred to as **bridging**.

As we have pointed out, logical metonymy is characterized by a great deal of **context dependence**. We have said that, presumably, the identity of the event left implicit in such cases can always be overridden by external context (cf. (3)). The question arises, then, whether logical metonymy can be seen as analogous to other phenomena related to implicit information, such as the retrieval of antecedents of anaphors or the retrieval of rhetorical relations between consecutive sentences.

The fact that there are severe semantic restrictions on what the missing event can be does not falsify the claim that metonymy proper is analogous to other cases of retrieval of implicit information. Similar restrictions apply in those cases as well. As a matter of fact, the severeness of the restrictions on the type of metonymy illustrated in (1) can be explained by the fact that a grammatical relation (namely, the direct object relation) holds between the verb (i.e., *finish/begin/* ...) and the noun phrase. The grammatical relation between the verb and the noun phrase may allow us to argue that the restriction on what the missing event can be is due to a separate **construction**, associated with verbs like *begin*, which prescribes a particular type of metonymical interpretation. Note that the same assumption has been argued independently by Michaelis and Lambrecht (1994) for the metonymical interpretation of noun phrases in the 'nominal extraposition' construction:

(40) It's amazing, the people my sister knows.
'The types/number of people my sister knows are/is amazing'
*'The people my sister knows are amazing'

Michaelis and Lambrecht (1994) argue that it is the nominal extraposition construction that is responsible for the obligatorily metonymical interpretation of the noun phrase *the people*; the same explanation can be applied to the type of metonymy related to verbs like *begin* in (1).

Once we recognize that metonymy is analogous to other instances of 'bridging', i.e., of establishing missing links, what remains to be seen is whether the restrictions on metonymy, as explained by Pustejovsky (1995b), for example, are the same in all these cases (disregarding, of course, the construction specific constraints mentioned above).

For example, consider:

(41) Joe got married yesterday. The minister spoke very harshly.

The accommodation of the definite noun phrase *the minister* in the second sentence makes it necessary to establish a link with entities present in the context of utterance. Obviously, the lexical/encyclopedic knowledge associated with *marry* will legitimate such a link. However, a simple qualia value coercion will not take care of this, because the path connecting the event of marrying someone with the existence of a minister leads through a wedding ceremony (which is part of a marriage) to the minister (who takes part in the ceremony). As far as we can see, such paths are always composed of relations corresponding to the qualia features. This corroborates the idea that logical metonymy should be seen as a special case of bridging.

The same holds for rhetorical relations between sentences of a discourse seen as missing links: the possible implicit relationships between the situations described by two consecutive sentences again correspond to the qualia features, yet there is usually no single lexical source for the resolution:

(42) Joe went home earlier. He met his boss.

With a little effort, we can get three different rhetorical relations between the two sentences in (42). In general the cause ('he went home because he had met his boss'), the goal ('he went home to meet his boss') and the elaboration ('he met his boss while going home') interpretations are parallel to the AGENTIVE, TELIC and CONS quale roles, but there is no lexical entry in these sentences to which we could attribute them.

References

A. Copestake, A. Sanfilippo, T. Briscoe and V. de Paiva. 1993. The acquilex lkb: An introduction. In T. Briscoe and Copestake (T. Briscoe and Copestake, 1993).

Anick, P. and S Bergler. 1992. Lexical structures for linguistic inference. In J. Pustejovsky and S. Bergler, editors, *Lexical Semantics and Knowledge Representation*. Springer Verlag, Berlin.

Boguraev, B. and J. Pustejovsky. 1990. Lexical ambiguity and the role of knowledge representation in lexicon design. In *Proceedings of COLING-1990, Helsinki.*

Boguraev, B. and J. Pustejovsky. 1994. A richer characterisation of dictionary entries. In A. Zampolli B. Atkins, editor, *Automating the lexicon*. Oxford University Press.

Boullion, P. and E. Viegas. 1994. A semi-polymorphic approach to the interpretation of adjectival constructions: a cross-linguistic perspective. In *Proceedings of Euralex 94, Amsterdam.*

Busa, F. 1996a. *Compositionality and the semantics of Nominals.* Ph.D. thesis, Brandeis University.

Busa, F. 1996b. The semantics of agentive nominals in the generative lexicon. In Pustejovsky (Pustejovsky, 1996a).

Copestake, A. and T. Briscoe. 1995. Semi-productive plysemy and sense extension. *Journal of Semantics 12.*

Godard, D. and J. Jayez. 1993. Towards a proper treatment of coercion phenomena. In *Proceeding of the 1993 European ACL.*

J. Fodor, E. Lepore. 1997. The emptiness of the lexicon:critical remarks on j. pustejovsky: "the generative lexicon". *Linguistic Enquiry.*

Johnston, M. and F. Busa. 1996. Qualia structure and the compositional interpretation of compounds. In Pustejovsky (Pustejovsky, 1996a).

Michaelis, L.A. and K. Lambrecht. 1994. On nominal extraposition: A constructional analysis. *BLS*, 20:363–373.

Moravcsik, J. M. 1975. Aitia as generative factor in aristotle's philpsophy. *Dialogue*, 14:622–36.

Nunberg, Geoffrey. 1995. Transfers of meaning. *Journal of Semantics*, 12:109–132.

Pustejovsky, J. 1988. The geometry of events. In C. Tenny, editor, *Studies in Generative Approaches to Aspect*, Lexicon Project Working Papers. MIT.

Pustejovsky, J. 1991. The syntax of event structure. *Cognition.*

Pustejovsky, J. 1993. Type coercion and lexical selection. In J. Pustejovsky, editor, *Semantics and the Lexicon.* Kluwer Academic Publishers, Dordrecht, The Netherlands.

Pustejovsky, J. 1995a. Extentions to qualia structure. Master's thesis, Computer Science Department, Brandeis University.

Pustejovsky, J. 1995b. *The Generative Lexicon.* MIT Press, Cambridge MA.

Pustejovsky, J. 1995c. Linguistic constraints on type coercion. In Saint-Dizier and Viegas (Saint-Dizier and Viegas, 1995).

Pustejovsky, J., editor. 1996a. *Generative Lexicon Theory and The Computational Lexicon.* ESSLLI Reader, Prague.

Pustejovsky, J. 1996b. Lexical shadowing and argument closure. In *Generative Lexicon Theory and The Computational Lexicon* (Pustejovsky, 1996a).

Pustejovsky, J. and P. Anick. 1988. The semantic interpretation of nominals. In *Proceedings of COLING-1988*, Budapest.

Pustejovsky, J. and B. Boguraev. 1993. Lexical knowledge representation and natural language processing. *Artificial Intelligence*, 63:193–223.

Pustejovsky, J. and B. Boguraev. 1995. Lexical semantcs in context. *Journal of Semantics*, 12:1–14.

Pustejovsky, J. and P. Boullion. 1995. Logical polysemy and aspectual coercion. *Journal of Semantics*, 12:133–162.

Pustejovsky, J. and F. Busa. 1995. Unaccusativity and event composition. In J. Higginbotham P.M. Bertinetto, V. Binachi and M. Squartini, editors, *Temporal Reference: Aspect and Actionality*. Rosenberg and Sellier, Turin.

Pustejovsky, James. 1995. *Generative Lexicon Theory*. MIT Press.

Saint-Dizier, P. and E. Viegas, editors. 1995. *Computational Lexical Semantics*. Cambridge University Press.

T. Briscoe, V. de Paiva and A. Copestake, editors. 1993. *Inheritance, Defaults and teh Lexicon*. Cambridge University Press.

Viegas, E. 1995. The semantics-pragmatics of lexicalisation. In Saint-Dizier and Viegas (Saint-Dizier and Viegas, 1995).

Contents