A Mathematical Formalism for Linguistic Theories with an Application in Head-Driven Phrase Structure Grammar

Dissertation submitted to the University of Tübingen

Please cite only with the exact date

of this version!

Frank Richter

April 28, 2000

Contents

guage are the means by which a linguist characterizes a natural language in HPSG 94, the description language is of particular interest in a formalization of object-based HPSG. To make the intuitions that underly the architecture of HPSG 94 mathematically precise, it is thus necessary to define a class of formal languages in which the principles of Pollard and Sag 1994 can be faithfully expressed, and to state a model theory for these languages that makes explicit which collection of feature structures is licensed by a set of descriptions of each formal language.

2.2.2 Formal Reconstruction and Different Perspectives

I begin my formal reconstruction of HPSG 94 with the definition of the formal languages of King's SRL (Speciate re-entrant logic) in Section 2.2.2.1. SRL has been used as the common basis of three different ways of characterizing the meaning of an HPSG 94 grammar. The first one that I will present (Section 2.2.2.2) is the account of King 1999, which holds that the task of a grammar of a natural language is to determine a certain subclass of models of the entire class of models of the grammar. King calls the class of models that the grammar determines the class of *exhaustive models* of the grammar. According to King, the class of exhaustive models comprises one particular exhaustive model, $\mathcal{L}$, that contains exactly the system of all possible entities in the natural language. The actual tokens of the natural language that we observe in utterances in the real world are in $\mathcal{L}$, as are the non-actual tokens that are licensed by the grammatical system but never realized as actual utterance events in our world. The other models in the class of exhaustive models are indistinguishable from $\mathcal{L}$. In Section 2.2.2.3, I show how King's exhaustive models are related to the object types of Pollard and Sag 1994. The object types of Pollard and Sag will be formalized as totally well-typed, sort resolved and potentially infinite abstract feature structures that stand in an abstraction relation with the entities in any given exhaustive model of a grammar. Both the abstraction relation and the abstract feature structures are closely related to the constructs of the same name that we have seen in HPSG 87. In Section 2.2.2.4, I will turn to Pollard's (1999) notion of the *strong generative capacity* of an HPSG 94 grammar. The strong generative capacity of a grammar is a set of pairwise non-isomorphic mathematical representations of the possible linguistic tokens of the language. Each representation is isomorphic to each possible token in an equivalence class of indistinguishable possible tokens of the language, and for each equivalence class of indistinguishable possible tokens, the strong generative capacity of the grammar contains a representation of the tokens in the class. One of the three definitions of the strong generative capacity of a grammar of Pollard 1999 will establish a direct connection to the exhaustive models of a grammar, thus completing the picture of how the three explanations of the meaning of a grammar are related.

2.2.2.1 SRL

My presentation of SRL follows the definitions of King 1999,[34] which generalizes some of the original definitions of King 1989. For ease of comparison with corresponding notions in HPSG 87 and with the extended formalism of HPSG 94 of Chapter 3, I adopt the notational conventions that I use there.

The fundamental intuition about formal languages that underlies SRL is that each expression of its formal languages is true or false *of an entity* in an interpretation. Given an interpretation, an expression thus denotes a set of entities. It is because of this property that logics like SRL are sometimes referred to as logics of descriptions. The idea is that the languages are used to describe entities. As formal languages of HPSG, they are used to describe the entities that constitute a natural language. SRL provides a class of formal languages, each of which consists of a signature and a class of interpretations of the signature. Each signature provides the non-logical symbols from which the formal language is constructed.

Definition 15 Σ *is an* ***SRL signature*** *iff*

> Σ *is a triple* $\langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$,
>
> $\mathcal{S}$ *is a set,*
>
> $\mathcal{A}$ *is a set, and*
>
> $\mathcal{F}$ *is a total function from the Cartesian product of* $\mathcal{S}$ *and* $\mathcal{A}$ *to the power set of* $\mathcal{S}$.

I call each element of $\mathcal{S}$ a *species*, each element of $\mathcal{A}$ an *attribute*, and $\mathcal{F}$ the *appropriateness function*. The symbols :, $\sim$, $\approx$, $\neg$, [,], $\wedge$, $\vee$, and $\rightarrow$ are reserved symbols, and I will henceforth assume that none of them are a species or an attribute symbol. In contrast to King 1989, infinite sets of

[34] The definitions of King 1999 are consistent with the ones presented in Pollard 1999.

species are admitted. The original definition was more restrictive, because King 1989 was interested in logical aspects of SRL that will not play a role below.

It is useful to have a compact terminology to talk about the appropriateness of attributes and species. I say that an attribute α is appropriate to a species σ if $\mathcal{F}\langle\sigma,\alpha\rangle$ is a nonempty set of species. A species σ' is appropriate for an attribute α at some species σ if σ' is an element of $\mathcal{F}\langle\sigma,\alpha\rangle$.

A striking difference between an SRL signature and an 87 signature is that an SRL signature does not include a sort hierarchy. Instead, it simply provides a set of species. This simplification also means that the appropriateness function is much simpler, because it no longer has to enforce the inheritance of appropriate attributes and attribute values in accordance with the ordering of the sorts in the sort hierarchy. In Pollard and Sag 1994 the appropriateness function is given in so-called "feature declarations" that obey the 87 appropriateness conditions that we have seen in 87 signatures: If an attribute α is 87 appropriate to some sort σ, then it is also 87 appropriate to all of its subsorts; and the sort that is the value of the 87 appropriateness function at the subsorts of σ and α is at least as specific as the value of $\mathcal{F}$ at σ and α. At first blush, the lack of an explicit sort hierarchy seems to be a problem for using SRL in formalizing HPSG 94 grammars, because, just as HPSG 87 grammars, they make heavy use of sort hierarchies and of attribute inheritance. However, King 1999, pp. 329–331, shows that the inclusion of a sort hierarchy in the signature of the formal languages for HPSG 94 is mathematically superfluous and that the essential information that is encoded in sort hierarchies can be expressed without them. Before I can explain why this is the case, I need to introduce interpretations of signatures:

Definition 16 *For each SRL signature* $\Sigma = \langle\mathcal{S},\mathcal{A},\mathcal{F}\rangle$, I *is a* Σ ***interpretation*** *iff*

> I *is a triple* $\langle\mathsf{U},\mathsf{S},\mathsf{A}\rangle$,
>
> U *is a set,*
>
> S *is a total function from* U *to* $\mathcal{S}$,
>
> A *is a total function from* $\mathcal{A}$ *to the set of partial functions from* U *to* U,
>
> *for each* $\alpha \in \mathcal{A}$ *and each* $u \in \mathsf{U}$,

$$\mathsf{A}(\alpha)(u) \textit{ is defined iff } \mathcal{F}\langle \mathsf{S}(u), \alpha\rangle \neq \emptyset, \textit{ and}$$
$$\textit{if } \mathsf{A}(\alpha)(u) \textit{ is defined then } \mathsf{S}(\mathsf{A}(\alpha)(u)) \in \mathcal{F}\langle \mathsf{S}(u), \alpha\rangle.$$

U is the *set of entities in the universe*, S is the *species assignment function*, and A is the *attribute interpretation function*.

S partitions the universe of entities by assigning each entity exactly one species. I sometimes simply say that an entity has a species. Informally, I also say that a species denotes a set of entities, which is the set of entities that have that species. The attribute interpretation function provides a denotation for the attribute symbols. Each attribute denotes a partial function from entities to entities. The attribute interpretation function must obey appropriateness. That means that an attribute is defined on an entity exactly if the attribute is appropriate to the species of the entity; and if an attribute is defined on an entity, u, then the species of the entity u' which is the result of interpreting α on u, is appropriate for the species of u and α.[35]

With the definition of Σ interpretations in hand, I can now return to the question of how SRL can express HPSG 94 sort hierarchies without representing them in its signatures. The reason that an SRL signature suffices to express the sort hierarchies of HPSG 94 grammars as characterized in Pollard and Sag 1994, pp. 395–396, is the intended interpretation of sort hierarchies. Algebraically, the sort hierarchies of HPSG 94 are finite partial orders.[36] The terminology that Pollard and Sag 1994 uses to describe the direction of the ordering relation is not entirely consistent. Pollard and Sag assume a sort, *object*, which subsumes all other sorts and denotes all entities, and they call all sorts subsorts of *object*. Sorts that do not have a proper subsort are called maximal or maximally specific. This terminology agrees with the sort hierarchy of 87 signatures (DEFINITION 1, page 30). On the other hand, Pollard and Sag 1994, pp. 17–18, assumes a reversed ordering that agrees with the one in Carpenter 1992, where the more inclusive sorts are lower in the ordering. For example, according to that convention *sign* is immediately below

[35] The reader might notice that the behavior of attribute interpretation with respect to appropriateness corresponds to the total well-typedness of concrete feature structures.

[36] Pollard and Sag 1994, p. 395, also postulates that, for each sort σ_1, for each sort σ_2, for each sort σ_3, if σ_1 is a subsort of σ_2 and of σ_3, then either σ_2 is a subsort of σ_3 or *vice versa*. In other words, no sort is an immediate proper subsort of two distinct sorts. While that condition holds for the grammar of Pollard and Sag 1994, it is not met by grammars with so-called multiple inheritance hierarchies (see Kathol 1995; Sag 1997; Przepiórkowski 1999a, among many others). Since I am interested in a formalism that captures as much of the literature of HPSG 94 as possible, I will ignore Pollard and Sag's additional restriction.

the sorts *word* and *phrase*. Since the direction of the ordering is merely a matter of convention, and no mathematical properties depend on which one we choose, I choose to retain the direction of ordering of 87 signatures to keep my terminology coherent across the different formalisms, and because that ordering is consistent with the intuitive understanding of the terms subsort, supersort, and maximally specific sort.

With respect to the denotation of sorts, Pollard and Sag assume that if σ' is a subsort of σ, then σ' denotes a subset of σ. The root sort, *object*, denotes all entities in the interpretation. Each entity is in the denotation of exactly one maximally specific sort. Given the assumptions about the sort hierarchy, the denotation of each of its sorts is thus fully determined by the denotation of the maximally specific sorts that it subsumes. Given the assumptions about attribute inheritance of Pollard and Sag 1994, the domain and range of the function in the denotation of each attribute is also fully determined by the appropriateness of the attribute to maximally specific sorts. For each signature with a sort hierarchy of the kind described above and with the respective attribute inheritance, we can thus give a function that maps it to an SRL signature such that each interpretation of the original signature is also an interpretation of the derived SRL signature. In addition, as King 1999, p. 330, shows, the use of nonmaximal sorts in descriptions of the formal language can easily be understood as metasyntactic notation for bigger, disjunctive descriptions that only use maximal sorts. It follows immediately that the formal languages of SRL are capable of expressing all aspects of HPSG 94 that are related to sort hierarchies.

The SRL signature in Figure 2.4 illustrates how a sort hierarchy is expressed in an SRL signature without having it explicitly included. It corresponds directly to the 87 signature with sort hierarchy in Figure 2.2. $\mathcal{S}$ is the set of maximally specific sorts of the 87 signature. The set of attributes, $\mathcal{A}$, does not change. The appropriateness function, $\mathcal{F}$, is now a total function, and it states for every species and every attribute which species are appropriate for them. If no species is appropriate for a given pair then the value of $\mathcal{F}$ at that pair is the empty set. For example, DRIVER is not appropriate to *man*, thus $\mathcal{F}\langle man, \text{DRIVER}\rangle = \emptyset$. For every species, σ, in the 87 signature for which an attribute, α, is 87 appropriate, the value of the corresponding appropriateness function in our SRL signature without sort hierarchy, $\mathcal{F}\langle \sigma, \alpha\rangle$, is the set of the maximally specific sorts subsumed by its sort value in the 87 signature. For example, since *person* is appropriate to the pair '$\langle vw, \text{OWNER}\rangle$' in the 87 signature of Figure 2.2, the value of $\mathcal{F}\langle vw, \text{OWNER}\rangle$

is now the set $\{man, woman\}$.

Figure 2.4: An example of an SRL signature

$\mathcal{S} = \{vw, bmw, man, woman\}$,

$\mathcal{A} = \{\text{DRIVER}, \text{OWNER}, \text{LIKES-BEST}\}$, and

$$\mathcal{F} = \left\{ \begin{array}{l} \langle\langle vw, \text{OWNER}\rangle, \{man, woman\}\rangle, \langle\langle vw, \text{DRIVER}\rangle, \{man, woman\}\rangle, \\ \langle\langle vw, \text{LIKES-BEST}\rangle, \{\}\rangle, \langle\langle bmw, \text{OWNER}\rangle, \{man, woman\}\rangle, \\ \langle\langle bmw, \text{DRIVER}\rangle, \{man, woman\}\rangle, \langle\langle bmw, \text{LIKES-BEST}\rangle, \{\}\rangle, \\ \langle\langle man, \text{LIKES-BEST}\rangle, \{vw, bmw, man, woman\}\rangle, \\ \langle\langle man, \text{DRIVER}\rangle, \{\}\rangle, \langle\langle man, \text{OWNER}\rangle, \{\}\rangle, \\ \langle\langle woman, \text{LIKES-BEST}\rangle, \{vw, bmw, man, woman\}\rangle, \\ \langle\langle woman, \text{DRIVER}\rangle, \{\}\rangle, \langle\langle woman, \text{OWNER}\rangle, \{\}\rangle \end{array} \right\}$$

From a more general perspective, the fact that an explicit sort hierarchy is mathematically superfluous in HPSG 94 but not in HPSG 87 is a consequence of the fact that signatures no longer form the basis of a subsumption ordering of feature structures that is meant to capture their relative degree of informativeness. The lack of a subsumption ordering on the level of entities is directly reflected by the fact that there are no entities of nonmaximal sorts in the interpretations of SRL signatures. The entities are no longer thought of as information bearing objects, but as models of linguistic entities of some kind.

The fact that each 87 signature with a sort hierarchy can be translated into a corresponding SRL signature without a sort hierarchy that has the same interpretations does not mean that sort hierarchies are not linguistically relevant. In principle, it is conceivable that psycholinguistic evidence which supports the representational reality of a sort hierarchy can be found. One research area where this question could be investigated is language acquisition.[37] Green 2000 programmatically explores the idea that first language acquisition proceeds along incremental and largely monotonic extensions of a sort hierarchy that is extended as more and more types of linguistic en-

[37]To the best of my knowledge, ideas about the psychological reality of sort hierarchies, although alluded to occasionally, are mere speculations that have not been tested in scientific experiments yet. Moreover, the computational results about the description languages of HPSG 94 of Section 3.3 caution against far-reaching assumptions about their psychological reality.

tities are distinguished by the child.[38] From the practical perspective of a grammar writer, sort hierarchies provide the means for a more compact notation of principles and for a useful structuring of the empirical domain in the eyes of the linguist. Even if finite sort hierarchies are superfluous from a mathematical perspective, they can still be an expedient syntactic construct for writing grammars.

The logical language of SRL is designed to describe (sets of) entities. Given an interpretation, essentially, we can say that entities have a certain sort ("sort assignment"); or that a certain attribute is defined on an entity, whose interpretation then leads to another entity; or that one entity equals another entity. To achieve this, SRL provides a set of terms, which are composed of the reserved symbol, ':', and the attributes of a given SRL signature, Σ, and a set of descriptions, which are built from the terms, the species of Σ, and reserved symbols for expressing sort assignment and equality, '$\sim$' and '$\approx$'. The descriptions are closed under negation, conjunction, disjunction, and implication. Other logical connectives can be obtained by combinations of these. I define Σ terms and Σ descriptions simultaneously:

Definition 17 *For each SRL signature* $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, $\mathcal{T}^\Sigma$ *and* $\mathcal{D}^\Sigma$ *are the smallest sets such that*

$: \in \mathcal{T}^\Sigma$,

for each $\alpha \in \mathcal{A}$ *and each* $\tau \in \mathcal{T}^\Sigma$, $\tau\alpha \in \mathcal{T}^\Sigma$,

for each $\sigma \in \mathcal{S}$, *for each* $\tau \in \mathcal{T}^\Sigma$, $\tau \sim \sigma \in \mathcal{D}^\Sigma$,

for each $\tau_1 \in \mathcal{T}^\Sigma$, *for each* $\tau_2 \in \mathcal{T}^\Sigma$, $\tau_1 \approx \tau_2 \in \mathcal{D}^\Sigma$,

for each $\delta \in \mathcal{D}^\Sigma$, $\neg\delta \in \mathcal{D}^\Sigma$,

for each $\delta_1 \in \mathcal{D}^\Sigma$, *for each* $\delta_2 \in \mathcal{D}^\Sigma$, $[\delta_1 \wedge \delta_2] \in \mathcal{D}^\Sigma$,

for each $\delta_1 \in \mathcal{D}^\Sigma$, *for each* $\delta_2 \in \mathcal{D}^\Sigma$, $[\delta_1 \vee \delta_2] \in \mathcal{D}^\Sigma$, *and*

for each $\delta_1 \in \mathcal{D}^\Sigma$, *for each* $\delta_2 \in \mathcal{D}^\Sigma$, $[\delta_1 \rightarrow \delta_2] \in \mathcal{D}^\Sigma$.

For each SRL signature Σ, I call each element of the set $\mathcal{T}^\Sigma$ a Σ *term*, and each element of $\mathcal{D}^\Sigma$ a Σ *description*.

[38]The formal foundations underlying Green's account are, however, not entirely clear to me, as her terminology fluctuates between HPSG 87, HPSG 94 and terms that apparently stem from other sources not directly related to HPSG formalisms. It remains to be seen if her research program can be rigorously rendered in a formalism for HPSG.

A Σ term consists of the reserved symbol, ':', followed by a (possibly empty) finite string of Σ attributes. Just as in Chapter 2, I will use the term "paths" to refer to finite strings of attributes. The symbol ':' can be viewed as the single variable of the languages of SRL. When we define denotations for our expressions, it will denote the entities being described. I call Σ descriptions of the form $\tau \sim \sigma$ *sort assignments*; Σ descriptions of the form $\tau_1 \approx \tau_2$ are *path equations*.

Superficially, the Σ descriptions of SRL do not bear any apparent resemblance to the conventional, loosely specified syntax of AVM diagrams that linguists use. King 1989, pp. 106–135, addresses this issue, and provides a translation from the expressions of formally defined languages of attribute value matrices to SRL's descriptions that, under certain conditions, preserves the intuitively intended meaning of the attribute value matrices. For my discussion of the SRL-based formalisms of HPSG 94, it suffices to know that a translation is possible in principle, and to appeal to an informal understanding of how AVM diagrams and Σ descriptions correspond. Once I have presented my extension of SRL in Section 3.1, I will return to the issue of linguistic notation in Section 3.2 and define an AVM syntax for the extended formalism that is modeled after the notational conventions of the HPSG literature.

Before discussing a few examples of descriptions, I want to define the denotation of descriptions in interpretations. To understand the denotation of terms, it might initially be helpful to compare the term interpretation function, T_{I}, of a given Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, to the iterated transition function, Δ^*, of a given concrete feature structure under Σ with transition function Δ. The denotation of the term consisting of the colon corresponds to the transition from one node to another by interpreting the empty path, and the denotation of a term of the form $:\pi$, where π is a nonempty path, corresponds to the transition from one node to another by interpreting π. The elements of U correspond to the nodes of the concrete feature structure.[39] That means that by interpreting a term at an element, u, of U, we transit from u to another element of U. In the special cases of a cyclic transition or a term consisting only of the colon, the element that we reach is identical

[39] There is, of course, no connectivity condition on interpretations that stipulates that each element of U be reachable from some root element in U. The Σ interpretations of SRL are more general structures than (totally well-typed and sort resolved) concrete feature structures. I will discuss the relationship between concrete feature structures and interpretations of SRL signatures in more detail in Sections 2.2.2.3 and 2.2.2.4.

to u. Since descriptions are built from terms, their interpretation depends directly on the interpretation of terms, and I define the two interpretation functions simultaneously:

Definition 18 *For each SRL signature $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, T_I is the total function from $\mathcal{T}^\Sigma$ to the set of partial functions from U to U, and D_I is the total function from $\mathcal{D}^\Sigma$ to the power set of U such that for each $u \in \mathsf{U}$,*

$T_\mathsf{I}(:)(u)$ is defined and $T_\mathsf{I}(:)(u) = u$,

for each $\tau \in \mathcal{T}^\Sigma$, for each $\alpha \in \mathcal{A}$,

$T_\mathsf{I}(\tau\alpha)(u)$ is defined
iff $T_\mathsf{I}(\tau)(u)$ is defined and $\mathsf{A}(\alpha)(T_\mathsf{I}(\tau)(u))$ is defined, and

if $T_\mathsf{I}(\tau\alpha)(u)$ is defined

then $T_\mathsf{I}(\tau\alpha)(u) = \mathsf{A}(\alpha)(T_\mathsf{I}(\tau)(u))$,

for each $\tau \in \mathcal{T}^\Sigma$, for each $\sigma \in \mathcal{S}$,

$$D_\mathsf{I}(\tau \sim \sigma) = \left\{ u \in \mathsf{U} \;\middle|\; \begin{array}{l} T_\mathsf{I}(\tau)(u) \textit{ is defined, and} \\ \mathsf{S}(T_\mathsf{I}(\tau)(u)) = \sigma \end{array} \right\},$$

for each $\tau_1 \in \mathcal{T}^\Sigma$, for each $\tau_2 \in \mathcal{T}^\Sigma$,

$$D_\mathsf{I}(\tau_1 \approx \tau_2) = \left\{ u \in \mathsf{U} \;\middle|\; \begin{array}{l} T_\mathsf{I}(\tau_1)(u) \textit{ is defined,} \\ T_\mathsf{I}(\tau_2)(u) \textit{ is defined, and} \\ T_\mathsf{I}(\tau_1)(u) = T_\mathsf{I}(\tau_2)(u) \end{array} \right\},$$

for each $\delta \in \mathcal{D}^\Sigma$, $D_\mathsf{I}(\neg\delta) = \mathsf{U} \backslash D_\mathsf{I}(\delta)$,
for each $\delta_1 \in \mathcal{D}^\Sigma$, for each $\delta_2 \in \mathcal{D}^\Sigma$, $D_\mathsf{I}([\delta_1 \wedge \delta_2]) = D_\mathsf{I}(\delta_1) \cap D_\mathsf{I}(\delta_2)$,
for each $\delta_1 \in \mathcal{D}^\Sigma$, for each $\delta_2 \in \mathcal{D}^\Sigma$, $D_\mathsf{I}([\delta_1 \vee \delta_2]) = D_\mathsf{I}(\delta_1) \cup D_\mathsf{I}(\delta_2)$, and
for each $\delta_1 \in \mathcal{D}^\Sigma$, for each $\delta_2 \in \mathcal{D}^\Sigma$, $D_\mathsf{I}([\delta_1 \rightarrow \delta_2]) = (\mathsf{U} \backslash D_\mathsf{I}(\delta)) \cup D_\mathsf{I}(\delta_2)$.

For each SRL signature Σ, I call T_{I} the Σ *term interpretation function with respect to* I, and D_{I} the Σ *description interpretation function with respect to* I. As discussed already, a Σ term interpretation function with respect to $\langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$ is a partial function from U to U. More specifically, the colon denotes the identity function, and we obtain the denotation of each term in which a nonempty string of attributes succeeds the colon by functional composition of the denotation of the colon and the attribute interpretation of each of the attribute symbols in the term, with the interpretation of the last attribute coming first.

The denotation of Σ descriptions follows naturally from the denotation of Σ terms: Each Σ description $\tau \sim \sigma$ denotes the subset of entities in U on whose members the Σ term interpretation function (with respect to $\langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$) of τ is defined and yields an entity of species σ. Each Σ description of the form $\tau_1 \approx \tau_2$ denotes the set of those entities in U on which the Σ term interpretation function (with respect to $\langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$) of both τ_1 and τ_2 is defined and yields the same entity in both cases. Negation, conjunction, disjunction and implication of Σ descriptions are interpreted classically as set complement of the denotation of the description, as the intersection and union of the denotations of the two descriptions of the complex description, and as the union of the set complement of the denotation of the first description in the implication with the denotation of the second description, respectively.

A signature expresses how we conceive of the world that we want to talk about as being structured. Assume the particular SRL signature given in Figure 2.4. For the purposes of the following examples, I will take it as fixed and leave it implicit when talking about terms and descriptions. According to the signature, the entities in our universe are of sort *man*, *woman*, *bmw* and *vw*; and there are no other entities. A *bmw* and a *vw* have a DRIVER and an OWNER, who are either a *woman* or a *man*.[40] But neither a *vw* nor a *bmw* have the property of liking somebody or something best. Each *man* and each *woman*, on the other hand, likes some entity of the universe best, and the entity that they like best can be of any sort. Neither a *woman* nor a *man* has a DRIVER or an OWNER. Let us now construct one of the many possible interpretations of this SRL signature.

Suppose the following scenario involving two cars on the parking lot in front of the Seminar für Sprachwissenschaft in Tübingen: The first car, which

[40] Note that these appropriateness conditions exclude, among other conceivable scenarios, multiple owners of a single car.

belongs to Anke, is a Volkswagen. Anke is in the driver's seat. We know that Anke is married, and she likes her husband best. Her husband, of course, likes her best. The second car is a Volkswagen, too. It is owned by Kordula, but since she is on vacation, Detmar is driving it. Kordula and Detmar like each other best. We can capture some facts of this little scenario in an interpretation of our signature. The two cars, Anke, her husband, Detmar and Kordula are the entities in the interpretation. They are assigned the obvious species given the intuitive meaning of the species symbols. The attributes receive a denotation that obeys appropriateness and gives the attributes their natural denotations with respect to the above scenario. Figure 2.5 shows this interpretation, which I will call $\mathsf{I}_{2.5}$.

Figure 2.5: An interpretation of the SRL signature of Figure 2.4

Let $\mathsf{I}_{2.5} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, with:

$\mathsf{U} = \{$Anke, Anke's husband, Detmar, Kordula, first car, second car$\}$

S(Anke) = *woman*,
S(Kordula) = *woman*,
S(Anke's husband) = *man*,
S(Detmar) = *man*,
S(first car) = *vw*, and
S(second car) = *vw*

A(LIKES-BEST)(Anke) = Anke's husband,
A(LIKES-BEST)(Anke's husband) = Anke,
A(LIKES-BEST)(Detmar) = Kordula,
A(LIKES-BEST)(Kordula) = Detmar,
A(OWNER)(first car) = Kordula,
A(OWNER)(second car) = Anke,
A(DRIVER)(first car) = Detmar,
A(DRIVER)(second car) = Anke,
A(LIKES-BEST) is undefined on the first car and on the second car,
A(OWNER) is undefined on Anke, Anke's husband, Detmar and Kordula, and
A(DRIVER) is likewise undefined on Anke, Anke's husband, Detmar and Kordula

We can now inspect the denotation of descriptions in the interpretation $\mathsf{I}_{2.5}$. In (15a) we see that the set of women in $\mathsf{I}_{2.5}$ consists of Anke and Kordula. There is no BMW in $\mathsf{I}_{2.5}$ (15b). Finally, only the second car is a Volkswagen whose owner is also its driver (15c), because Anke drives her own car, whereas Detmar drives Kordula's car; and there is no other car in $\mathsf{I}_{2.5}$.

(15) a. $D_{\mathsf{I}_{2.5}}(:\sim \mathit{woman}) = \{\text{Anke, Kordula}\}$

b. $D_{\mathsf{I}_{2.5}}(:\sim \mathit{bmw}) = \emptyset$

c. $D_{\mathsf{I}_{2.5}}([:\sim \mathit{vw} \wedge :\text{OWNER} \approx : \text{DRIVER}]) = \{\text{second car}\}$

Note that in a different interpretation of the same signature, the descriptions in (15) might denote different sets of entities. As soon as there are BMWs in the interpretation, '$:\sim \mathit{bmw}$' denotes that set of cars. Moreover, it is easy to see how the symbols of the sort hierarchy of the 87 signature of Figure 2.2 that are not present in our corresponding SRL signature can be understood as an abbreviation of disjunctive descriptions. For example, '$:\sim \mathit{person}$' can be defined as a metasyntactic notation for '$[:\sim \mathit{man} \vee :\sim \mathit{woman}]$', which denotes the set comprising Anke, Anke's husband, Detmar and Kordula. In other words, a sort assignment with a sort σ that is not a species is interpreted as the disjunction of the sort assignments with the species that σ subsumes in the envisaged sort hierarchy of an 87 signature.

An HPSG 94 grammar is a pair consisting of a signature and a set of expressions of the description language, the principles of grammar. DEFINITION 19 captures this conception of a grammar with the notion of an SRL grammar:

Definition 19 Γ *is an* ***SRL grammar*** *iff*

Γ *is a pair* $\langle \Sigma, \theta \rangle$,

Σ *is an SRL signature, and*

$\theta \subseteq \mathcal{D}^{\Sigma}$.

Since an SRL grammar always includes a fixed signature, Σ, I will simply talk about descriptions of SRL grammars below instead of Σ descriptions of SRL grammars.

The purpose of a grammar is to describe a natural language. In HPSG 94, the most important means to delineate a natural language is clearly the set of principles. It is, therefore, a prerequisite for determining what an SRL grammar means to define what a set of descriptions means. The denotation of a set of descriptions, θ, in a given interpretation, I, is the set of all entities in I of which each description in θ is true:

Definition 20 *For each SRL signature Σ, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, Θ_{I} is the total function from the power set of $\mathcal{D}^{\Sigma}$ to the power set of U such that for each $\theta \subseteq \mathcal{D}^{\Sigma}$,*

$$\Theta_{\mathsf{I}}(\theta) = \left\{ u \in \mathsf{U} \,\middle|\, \begin{array}{l} \textit{for each } \delta \in \theta, \\ u \in D_{\mathsf{I}}(\delta) \end{array} \right\}.$$

I call Θ_{I} the *theory denotation function with respect to* I. For example, the theory $\theta_1 = \{[: \sim bmw \rightarrow \neg : \approx :], [: \sim vw \rightarrow : \text{OWNER} \sim woman]\}$ describes every entity in the interpretation $\mathsf{I}_{2.5}$, i.e., $\Theta_{\mathsf{I}_{2.5}}(\theta_1) = \{$Anke, Anke's husband, Detmar, Kordula, first car, second car$\}$. This is so, because each of the descriptions of θ_1 describes every entity in $\mathsf{I}_{2.5}$: The first, '$[: \sim bmw \rightarrow \neg : \approx :]$,' describes all entities that are not BMWs or not identical with themselves. This is true of all entities that are not BMWs and false of all BMWs. Since $\mathsf{I}_{2.5}$ does not contain any entity of sort *bmw*, the description is true of all entities in $\mathsf{I}_{2.5}$. The second description in θ_1 describes entities that are not VWs or whose owner is a woman. That is true of all entities that are not VWs and of all VWs whose owner is a woman. In other words, it is true of all entities in $\mathsf{I}_{2.5}$, because Anke, Anke's husband, Detmar and Kordula are men and women, and the two VWs are owned by Anke and Kordula, respectively.

As the example shows, the theory denotation of a finite set of descriptions equals the description denotation of the conjunction of the descriptions that the set contains. A finite theory can thus be expressed by the conjunction of the descriptions in it, and for the finite case, theories and theory denotation functions are not strictly necessary. However, the inductive definition of the languages of SRL does not permit infinite conjunctions. The theory denotation functions with respect to an interpretation thus add the possibility of interpreting infinite theories.

Denotations of theories in interpretations do not yet determine the meaning of a grammar by themselves. The theory denotation function only tells us of which entities in an interpretation a set of descriptions is true. This property can, however, be exploited in a first approximation of the intended interpretations of a grammar. Clearly, the linguist is only interested in those interpretations of a grammar in which every description is true of every entity: Every description of the grammar describes every entity in the interpretations, or, equivalently, no description is false of any entity in the intended interpretations. Interpretations that have that property with respect to a given grammar are called *models* of the grammar:

Definition 21 *For each SRL grammar* $\Gamma = \langle \Sigma, \theta \rangle$, *for each* Σ *interpretation* $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$,

$$\mathsf{I} \textit{ is a } \Gamma \textit{ \textbf{model} iff } \Theta_{\mathsf{I}}(\theta) = \mathsf{U}.$$

The theory, θ, of an SRL grammar, Γ, denotes the entire universe of entities, U, in each Γ model, $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$. Since every entity in the interpretation $\mathsf{I}_{2.5}$ is in the denotation of each element of θ_1, $\mathsf{I}_{2.5}$ is a Γ model of the grammar Γ that consists of the signature of Figure 2.4 and θ_1. Intuitively speaking, no entity in $\mathsf{I}_{2.5}$ violates the conditions of not being a BMW and of being owned by a woman if it is a Volkswagen. Not every theory has a nonempty model. Consider the theory, θ_2, that contains the single description '$: \sim vw$.' A Γ model with the theory θ_2 may only contain entities of sort *vw*, because '$: \sim vw$' is only true of VWs and false of all other entities. But by virtue of the signature, each Volkswagen must have an owner who is either a woman or a man. Therefore, no Γ model with the theory θ_2 can contain a Volkswagen; each Γ model with the theory θ_2 must have an empty universe.

Considering the notion of a model of a grammar from a linguistic perspective, it restricts the candidates for the meaning of a grammar to those interpretations whose entities do not violate any principles. For example, assuming that a grammar contains a correct specification of subject verb agreement in English, the sentence *John love Mary* cannot be in any model of the grammar, because it violates subject verb agreement. The sentence contains at least one entity that is not described by at least one description in the theory of the grammar. The models of SRL grammars will be the starting point for all three explications of the meaning of HPSG 94 grammars in the following three sections.

My introduction to SRL deliberately omitted the logic of SRL and results pertaining to it, because the logic of SRL is of no immediate relevance for the characterization of natural languages in the SRL formalism. For completeness, I mention some of the logical aspects of SRL here.[41] For each SRL signature Σ with a finite set of species, King 1989 provides a Hilbert and Ackermann style calculus, and shows that the inference relation that the calculus determines is sound and complete with respect to entailment, i.e., for each set of Σ descriptions θ and for each Σ description δ, θ infers δ if and only if θ entails δ. Kepser 1994 then proceeds to show that for each SRL signature

[41]See King 1999, pp. 328–329, and the literature cited there for a more comprehensive overview over the logical results about SRL.

with a finite set of species and a recursive set of attributes, the satisfiability problem is decidable: There is an effective algorithm that decides for each SRL signature of the kind indicated, and for each description δ generated by that signature, whether there is an interpretation of the signature in which δ has a nonempty denotation. In Section 3.3, I will return to the question of the significance of these and related results to a theoretical linguist who uses formal languages of the kind that HPSG 94 envisages in order to describe natural languages.

2.2.2.2 Exhaustive Models

In a substantial reformulation of an earlier attempt to characterize the meaning of HPSG 94 grammars in King 1994, King 1999 investigates the question of when an SRL grammar is true of a natural language. King (1999) formulates three necessary conditions for an SRL grammar to be true of a natural language. These conditions are met if a natural language belongs to a certain class of models of a given SRL grammar. King calls that class of models the class of *exhaustive models* of a grammar. The meaning of an SRL grammar is thus determined as delineating the class of its exhaustive models, and the grammar is true of a language, $\mathcal{L}$, if $\mathcal{L}$ is an exhaustive model of the grammar.

In this section, I present the motivation behind King's notion of exhaustive models, and their definition. As the introductory, short description in the preceding paragraph of King's approach reveals, the perspective of King 1999 is realistic in the sense that the task of SRL grammars is to directly characterize natural languages without the intervention of some modeling mathematical structure.[42] The natural languages themselves are the intended models of grammars. The realistic approach to the meaning of scientific theories is the most important feature that distinguishes King's explanation of the meaning of HPSG 94 grammars from the explanations of Pollard and Sag 1994 and of Pollard 1999, which choose a representational approach. Despite their fundamentally different assumptions about the ontological status of the structures that grammars characterize, the three formalisms are very similar and mathematically closely related. Understanding the technical side to the

[42] The term *realistic* is meant in a pre-theoretical sense here; it is not meant to be a technical characterization of King's philosophy of science. In particular, King does not explicitly call himself a realist. I use the term to refer to his view because it emphasizes his predominant interest in describing linguistic behavior directly, and stresses the contrast to the representational view of Pollard and Sag.

idea of exhaustive models will thus also illuminate many crucial aspects of the other two approaches, and exhaustive models will become the mathematical link between them. For expository reasons, I adopt the realistic terminology of King for the rest of the present section.

At first it might be surprising that the notion of the models of an SRL grammar defined in Definition 21 of the previous section is insufficient to determine the meaning of a grammar. The basic problem is, however, very easy to see: A non-contradictory grammar, i.e., a grammar that has a nonempty model, does not have a unique model; and the models of an interesting grammar differ in linguistically significant ways. A trivial example of an obviously unintended model is the model with the empty universe. A less trivial example of an unintended model is a model of Pollard and Sag's (1994) grammar which contains exactly one token of the sentence *Kim likes bagels* and nothing else. Presupposing that the grammar contains the necessary lexical entries and that there are no hidden mistakes in the grammar of Pollard and Sag, every entity in the sentence *Kim likes bagels* is described by every principle in the theory of their grammar. But the grammar is supposed to license many more sentences that do not occur in this particular small model. For example, Pollard and Sag's fragment of English also contains the sentence *We believe there to have been many solutions to that problem*, and this sentence must therefore also be in the intended model of their grammar. Intuitively, every model that is missing an expression of English that is licensed by a given grammar of English is in a relevant sense too small to be the intended model of that grammar. The intended model should contain instances of all expressions of English that are permitted by the theory.

What is missing so far is a model theory that tells us unambiguously which models are the intended models of an SRL grammar. If we knew which models are the "right" models of a grammar, then these models would also provide the intended meaning of the grammar. In the terminology of generative grammar, the question of the right models of a given grammar is connected to the question which linguistic structures are *generated* by the grammar. Adapting the traditional terminology of generative grammar to the constraint-based framework of HPSG 94, we can slightly rephrase the question: In a constraint-based framework, we ask which linguistic structures are *licensed* by a grammar. This question is of utmost importance to generative grammarians. Given a natural language or a certain fragment thereof, linguists want to check whether their grammar generates exactly the sentences of the language or of the language fragment.[43] If the grammar generates sentences that are not in the language, generative grammarians say that the grammar *overgenerates*. If the grammar does not generate sentences that are in the language, generative grammarians say that the grammar *undergenerates*. Correspondingly, we can say that an HPSG 94 grammar *overlicenses* if it licenses linguistic structures that are not in the language; and it *underlicenses* if it does not license linguistic structures that are in the language. With an arbitrary model of a grammar, we cannot reliably determine whether the grammar overlicenses or underlicenses the target language. If the chosen model contains all the linguistic structures that the grammar is supposed to license, there might still be some "bigger" model of the grammar that contains structures that are judged ungrammatical, and we would then want to say that the grammar overlicenses. On the other hand, the existence of a model that is "too small" does not mean that a grammar underlicenses. This is so because there might be another, "bigger" model of the grammar that contains the structures that are missing in the smaller model. The theory of the intended models of a grammar is supposed to give us models that tell us if a grammar overlicenses or underlicenses. The exhaustive models of a grammar are these intended models. If an exhaustive model of a grammar contains structures that are not in the language, the grammar overlicenses. If an exhaustive model of a grammar does not contain structures that are in the language, the grammar underlicenses.[44] Of course, it is conceivable that a grammar simultaneously overlicenses and underlicenses a language.

The situation can be illustrated with our little car scenario of interpretation $\mathsf{I}_{2.5}$ from the previous section. A linguist who wants to write a grammar of (a fragment of) a natural language faces a given intended model, the natural language, and the task of the linguist is to write a grammar that neither overlicenses nor underlicenses the language. As a simple, non-linguistic illustration of the problem of the linguist, consider the little scenario $\mathsf{I}_{2.5}$. In that scenario, we find the following observationally given facts: Kordula is the owner of the first car, Detmar drives the first car, and Kordula and Detmar

[43] In this context, a *sentence* means a phonological string together with (at least) a syntactic and a semantic structure. Phonological strings together with unintended syntactic or semantic structures are not sentences of the language in that sense, even if the phonological string in question belongs to a sentence of the language.

[44] For the sake of giving the reader a first, simple intuition about the idea of exhaustive models, this characterization of overlicensing and underlicensing of a grammar relative to an exhaustive model is oversimplifying a bit. It will be refined presently.

like each other best. Anke is the owner and the driver of the second car, and she and her husband like each other best. Kordula and Anke are women, and Detmar and Anke's husband are men. Recall that $\mathsf{I}_{2.5}$ is an interpretation of the SRL signature of Figure 2.4. I call that signature $\Sigma_{2.4}$. The sort symbols of $\Sigma_{2.4}$ allow us to talk about entities of sort *man, woman, vw* and *bmw*, and their attributes are DRIVER, OWNER and LIKES-BEST. The appropriateness function restricts the attributes to entities where they intuitively make sense. Note that with this signature, we cannot talk about particular men and women such as Anke and Detmar, and their cars. We can only talk about men, women, Volkswagens and BMWs, and three kinds of relationships in which they stand to each other. With our grammar, we thus necessarily abstract away from specific entities, and focus on the configurations in which entities of a certain kind stand. Informally, the interpretation $\mathsf{I}_{2.5}$ contains two kinds of configurations. In the first configuration, a Volkswagen is owned by a woman who likes a man best that drives her car and likes the woman best. In the second configuration, a Volkswagen is owned and driven by a woman who likes a man best who likes the woman best. A grammar whose intended model is $\mathsf{I}_{2.5}$ in the sense that the grammar neither overlicenses nor underlicenses $\mathsf{I}_{2.5}$ should thus license models containing configurations of cars, women and men that are isomorphic to the two configurations in $\mathsf{I}_{2.5}$.

As a first attempt at a grammar of $\mathsf{I}_{2.5}$, consider the $\Sigma_{2.4}$ theory $\theta_1 = \{[:\sim bmw \rightarrow \neg :\approx:], [:\sim vw \rightarrow :\text{OWNER} \sim woman]\}$ of Section 2.2.2.1. Given an interpretation, θ_1 is true of all entities that are not a BMW, of all Volkswagens that belong to a woman, and of all entities that are neither a Volkswagen nor a BMW. In Section 2.2.2.1, we already saw that $\mathsf{I}_{2.5}$ is a model of the grammar $\langle \Sigma_{2.4}, \theta_1 \rangle$, because $\mathsf{I}_{2.5}$ does not contain a BMW, the first Volkswagen belongs to Kordula, the second Volkswagen belongs to Anke, and there are no other cars. However, there are configurations in other models of $\langle \Sigma_{2.4}, \theta_1 \rangle$ that are not isomorphic to the ones in $\mathsf{I}_{2.5}$. For example, $\mathsf{I}_{2.5}$ does not contain the legitimate configuration in which a Volkswagen is owned by a woman who likes a man best who likes her best, but the driver of the Volkswagen is a second man who likes a second woman best who likes the second man best. Figure 2.6 shows a model of $\langle \Sigma_{2.4}, \theta_1 \rangle$ that contains exactly such a scenario: The third car is a Volkswagen that is owned by Anne, while Mike drives it. Anne and Manfred and Mike and Maria like each other best.

I conclude that $\langle \Sigma_{2.4}, \theta_1 \rangle$ overlicenses $\mathsf{I}_{2.5}$, because there are models of $\langle \Sigma_{2.4}, \theta_1 \rangle$ that contain configurations of entities that are not isomorphic to any configurations in the target model, $\mathsf{I}_{2.5}$. In fact, it is easy to see that $\langle \Sigma_{2.4}, \theta_1 \rangle$

Figure 2.6: A second model of $\langle \Sigma_{2.4}, \theta_1 \rangle$

Let $\mathsf{I}_{2.6} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, with:

$\mathsf{U} = \{$Anne, Manfred, Maria, Mike, third car$\}$

S(Anne) = *woman*,
S(Maria) = *woman*,
S(Manfred) = *man*,
S(Mike) = *man*,
S(third car) = *vw*, and

A(LIKES-BEST)(Anne) = Manfred,
A(LIKES-BEST)(Manfred) = Anne,
A(LIKES-BEST)(Maria) = Mike,
A(LIKES-BEST)(Mike) = Maria,
A(OWNER)(third car) = Anne,
A(DRIVER)(third car) = Mike,
A(LIKES-BEST) is undefined on the third car,
A(OWNER) is undefined on Anne, Manfred, Maria and Mike, and
A(DRIVER) is likewise undefined on Anne, Manfred, Maria and Mike

licenses infinitely many pairwise non-isomorphic configurations of entities that are not isomorphic to a configuration of entities in $\mathsf{I}_{2.5}$. The exhaustive models of $\langle \Sigma_{2.4}, \theta_1 \rangle$ must contain infinitely many different configurations, because the theory θ_1 is very unrestrictive. For example, imagine a scenario that only contains a man who likes another man best who likes himself best. But there is also a scenario with three men in which the first likes the second and the second likes the third best; and the third likes himself best. This series of scenarios can be continued with four men, with five men, and so on *ad infinitum*, and they are all models of $\langle \Sigma_{2.4}, \theta_1 \rangle$. Therefore, they must all be in a model of $\langle \Sigma_{2.4}, \theta_1 \rangle$ that is exhaustive. Of course, there are infinitely many other possible scenarios involving men, women, and Volkswagens that are possible models of $\langle \Sigma_{2.4}, \theta_1 \rangle$, which I will not consider here.

Clearly, we need a stronger, more restrictive theory about the configurations of entities in $\mathsf{I}_{2.5}$ in order to obtain a grammar of which $\mathsf{I}_{2.5}$ is an exhaustive model. The theory θ_2 is a new attempt at constructing a grammar of $\mathsf{I}_{2.5}$:

(16) $\theta_2 =$

$$\left\{ \begin{array}{l} [: \sim woman \leftrightarrow [: \text{LIKES-BEST} \sim man \wedge : \text{LIKES-BEST LIKES-BEST} \approx :]]\,, \\ [: \sim man \leftrightarrow [: \text{LIKES-BEST} \sim woman \wedge : \text{LIKES-BEST LIKES-BEST} \approx :]]\,, \\ [: \sim bmw \rightarrow \neg : \approx :]\,, \\ [: \sim vw \rightarrow [: \text{OWNER} \sim woman \wedge : \text{OWNER} \approx : \text{DRIVER}]] \end{array} \right\}$$

The theory θ_2 enforces the relationship between men and women that we observe in $\mathsf{I}_{2.5}$ in all models of the grammar $\langle \Sigma_{2.4}, \theta_2 \rangle$. Every woman likes a man best who likes her best, and every man likes a woman best who likes him best. No BMWs are admitted. The last description in θ_2 excludes the $\Sigma_{2.4}$ interpretation of Figure 2.6 from being a $\langle \Sigma_{2.4}, \theta_2 \rangle$ model, because it is only true of those Volkswagens whose owner is also the driver, and the owner of the third car is Anne, whereas Mike is the driver. So far, it looks as if $\langle \Sigma_{2.4}, \theta_2 \rangle$ were the grammar that we are after, because it does not overlicense $\mathsf{I}_{2.5}$. Unfortunately, the new grammar, $\langle \Sigma_{2.4}, \theta_2 \rangle$, underlicenses $\mathsf{I}_{2.5}$. It requires that the owner of each Volkswagen be also its driver, but in $\mathsf{I}_{2.5}$, Detmar is the driver of Kordula's car. In other words, $\mathsf{I}_{2.5}$ is not even a model of the grammar $\langle \Sigma_{2.4}, \theta_2 \rangle$, because the last description of θ_2 does not describe Kordula's car.

As a grammar of $\mathsf{I}_{2.5}$, $\langle \Sigma_{2.4}, \theta_1 \rangle$ is too unrestrictive, whereas $\langle \Sigma_{2.4}, \theta_2 \rangle$ is too restrictive. However, based on the experience with these two grammars, it is not difficult to find a grammar that imposes the right conditions. We must be careful not to strengthen θ_1 too much, so $\mathsf{I}_{2.5}$ remains a model of the grammar. At the other end of the spectrum, we must relax θ_2 sufficiently, so every description in the new grammar is still true of every entity in $\mathsf{I}_{2.5}$. Consider the theory θ_3 in (17):

(17) $\theta_3 =$

$$\left\{ \begin{array}{l} [: \sim woman \leftrightarrow [: \text{LIKES-BEST} \sim man \wedge : \text{LIKES-BEST LIKES-BEST} \approx :]]\,, \\ [: \sim man \leftrightarrow [: \text{LIKES-BEST} \sim woman \wedge : \text{LIKES-BEST LIKES-BEST} \approx :]]\,, \\ [: \sim bmw \rightarrow \neg : \approx :]\,, \\ \left[: \sim vw \rightarrow \left[\begin{array}{l} : \text{OWNER} \sim woman \wedge \\ [: \text{OWNER} \approx : \text{DRIVER} \vee : \text{DRIVER} \approx : \text{OWNER LIKES-BEST}] \end{array} \right] \right] \end{array} \right\}$$

Let Γ_3 be $\langle \Sigma_{2.4}, \theta_3 \rangle$. In addition to the descriptions of θ_1, θ_3 contains descriptions that are true of all entities in a $\Sigma_{2.4}$ interpretation if every woman and every man likes a person of the opposite sex best, and the other person returns their feelings. Moreover, every Volkswagen in a model of Γ_3 is owned by a woman who is either also the driver, or the man she likes best is the driver. This is where θ_3 is not as restrictive as θ_2. BMWs remain excluded from the models. It is straightforward to check that $\mathsf{I}_{2.5}$ is an exhaustive model of Γ_3, because it contains all kinds of configurations of entities that can occur in models of Γ_3. In particular, the configuration under the third car in the model $\mathsf{I}_{2.6}$ has no isomorphically configured counterpart in an exhaustive model of Γ_3, because the last description in θ_3 does not describe the third car: Neither is its driver the woman who owns it (Anne), nor the man she likes best (Manfred).

It is now time to define a terminology that allows me to be more precise about what I mean by configurations of entities and isomorphically configured entities in an interpretation. To simplify some formulations, I will from now on say that an entity, u, is an element of a Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$ if $u \in \mathsf{U}$, and I write $u \in \mathsf{I}$ to indicate that u is a member of the universe of I. The first step towards defining a notion of configurations of entities in an interpretation is the definition of componenthood. An entity, u', of a Σ interpretation I is a *component* of an entity, u, of I, if there is a Σ term τ such that the interpretation of τ at u is defined and yields u'. Informally, there is a path (preceded by the colon) whose interpretation leads from u to u', i.e., u' is in the subalgebra of I generated by u. Definition 22 captures the notion of componenthood:

Definition 22 *For each SRL signature* $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, *and for each* Σ *interpretation* $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, Co_I *is the total function from* U *to the power set of* U *such that for each* $u \in \mathsf{U}$,

$$\mathsf{Co}_\mathsf{I}(u) = \left\{ u' \in \mathsf{U} \;\middle|\; \begin{array}{l} \textit{for some } \tau \in \mathcal{T}^\Sigma \\ T_\mathsf{I}(\tau)(u) \textit{ is defined, and } T_\mathsf{I}(\tau)(u) = u' \end{array} \right\}.$$

$\mathsf{Co}_\mathsf{I}(u)$ is the *set of components of* u *in* I. Conversely, I say that an entity, u, of a Σ interpretation I is a *common ancestor* of two entities u_1 and u_2 in I if u_1 and u_2 are components of u in I.

As configurations of entities I want to single out parts of Σ interpretations that comprise exactly one freely chosen entity with the set of its components and the relationships that exist between all these entities by virtue of the attribute interpretation function, but that do not contain any other, unrelated entities. For example, I want to have a simple terminology that allows me to talk about the second car and everything that we can see about the second car in the interpretation $\mathsf{I}_{2.5}$ of Figure 2.5: that Anke owns it, that Anke drives it, and that Anke likes her husband best, who in turn likes her best. What I am after is clearly a subinterpretation of the interpretation $\mathsf{I}_{2.5}$. The subinterpretation comprises an entity, the second car, of $\mathsf{I}_{2.5}$, all its components in $\mathsf{I}_{2.5}$, and the sort assignment and the attribute interpretation function restricted to these entities. DEFINITION 23 formalizes subinterpretations of that kind of a given interpretation. I_u is the subalgebra generated by u in I:

Definition 23 *For each SRL signature $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, for each $u \in \mathsf{I}$,*

$\mathsf{I}_u = \langle \mathsf{U}_u, \mathsf{S}_u, \mathsf{A}_u \rangle$ *is the Σ* ***interpretation under** u **in** I iff*

$$\begin{aligned}
&\mathsf{U}_u = \mathsf{Co}_\mathsf{I}(u),\\
&\mathsf{S}_u = \mathsf{S} \cap (\mathsf{U}_u \times \mathcal{S}),\\
&\mathsf{A}_u = \mathsf{A} \cap (\mathcal{A} \times \{\mathsf{U}_u \times \mathsf{U}_u\}), \textit{ and}\\
&\mathsf{I}_u = \langle \mathsf{U}_u, \mathsf{S}_u, \mathsf{A}_u \rangle.
\end{aligned}$$

By DEFINITION 23, the interpretation under Kordula in $\mathsf{I}_{2.5}$ is the interpretation $\mathsf{I}_{\text{Kordula}} = \langle \mathsf{U}_{\text{Kordula}}, \mathsf{S}_{\text{Kordula}}, \mathsf{A}_{\text{Kordula}} \rangle$, with
$\mathsf{U}_{\text{Kordula}} = \{\text{Kordula, Detmar}\}$,
$\mathsf{S}_{\text{Kordula}} = \{\langle \text{Kordula}, \textit{woman} \rangle, \langle \text{Detmar}, \textit{man} \rangle\}$, and
$\mathsf{A}_{\text{Kordula}} = \{\langle \text{LIKES-BEST}, \{\langle \text{Kordula}, \text{Detmar} \rangle, \langle \text{Detmar}, \text{Kordula} \rangle\} \rangle\}$,
because the only component of Kordula (besides Kordula herself) is Detmar, Kordula is a woman, Detmar is a man, and they like each other best. It is convenient to have a name to refer to an arbitrary entity in an interpretation and the subinterpretation under that entity:

Definition 24 *For each SRL signature $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, for each $u \in \mathsf{I}$,*

$\langle u, \mathsf{I}_u \rangle$ *is the* ***configuration of entities under** u **in** I iff*

I_u *is the Σ interpretation under u in I.*

I will omit the reference to I if an interpretation is fixed. In that case I will simply talk about the configuration (of entities) under some entity, u. Assume the interpretation $\mathsf{I}_{2.5}$. Then $\langle \text{Kordula}, \mathsf{I}_{\text{Kordula}} \rangle$ is the configuration (of entities) under Kordula.

Crucially, we want to be able to compare configurations under different entities in possibly different interpretations of the same SRL signature. This is necessary in order to determine if one interpretation contains a configuration under an entity that is in fact isomorphic to a configuration under a different entity in another interpretation. Alternatively, we might also want to say that, in one and the same interpretation, all components of two distinct entities are equally configured. For that purpose, I define what it means for a pair of an entity and an interpretation of which the entity is a member to be congruent to a second pair of an entity and an interpretation. Two configurations under different entities in $\mathsf{I}_{2.5}$ that are congruent are $\langle \text{Kordula}, \mathsf{I}_{\text{Kordula}} \rangle$ and $\langle \text{Anke}, \mathsf{I}_{\text{Anke}} \rangle$. But $\langle \text{Anke}, \mathsf{I}_{\text{Anke}} \rangle$ and $\langle \text{Detmar}, \mathsf{I}_{\text{Detmar}} \rangle$ are not congruent, although the configuration under Anke and the configuration under Detmar each have two components on which LIKES-BEST is defined; and in both cases LIKES-BEST, when interpreted on one component, denotes the other component. The reason that the two configurations are nevertheless not congruent is that Anke is a woman and Detmar is a man. DEFINITION 25 makes the notion of an SRL congruence that I have used in my examples explicit:

Definition 25 *For each SRL signature* Σ, *for each* Σ *interpretation* $\mathsf{I}_1 = \langle \mathsf{U}_1, \mathsf{S}_1, \mathsf{A}_1 \rangle$, *for each* $u_1 \in \mathsf{I}_1$, *for each* Σ *interpretation* $\mathsf{I}_2 = \langle \mathsf{U}_2, \mathsf{S}_2, \mathsf{A}_2 \rangle$, *for each* $u_2 \in \mathsf{I}_2$,

f *is an* ***SRL congruence from*** $\langle u_1, \mathsf{I}_1 \rangle$ ***to*** $\langle u_2, \mathsf{I}_2 \rangle$ ***in*** Σ

iff f *is a bijection from* $\mathsf{Co}_{\mathsf{I}_1}(u_1)$ *to* $\mathsf{Co}_{\mathsf{I}_2}(u_2)$,

for each $u \in \mathsf{Co}_{\mathsf{I}_1}(u_1)$, $\mathsf{S}_1(u) = \mathsf{S}_2(f(u))$,

for each $\alpha \in \mathcal{A}$, *for each* $u \in \mathsf{Co}_{\mathsf{I}_1}(u_1)$,

$\mathsf{A}_1(\alpha)(u)$ *is defined iff* $\mathsf{A}_2(\alpha)(f(u))$ *is defined, and*
if $\mathsf{A}_1(\alpha)(u)$ *is defined then* $f(\mathsf{A}_1(\alpha)(u)) = \mathsf{A}_2(\alpha)(f(u))$.

In the definition of SRL congruences, the interpretations in the two pairs of an entity and an interpretation need not be interpretations under the respective entities in the pair: u_1 is an element of I_1, but not all elements of I_1 are necessarily components of u_1. However, since I_1 by definition obeys the appropriateness conditions imposed by its SRL signature, it is ensured that the interpretation under u_1 is a subinterpretation of I_1. Each interpretation, I, contains the interpretation under each u that is an element of I as a subinterpretation. In the special case in which I_1 only contains entities that are components of u_1 and I_2 only contains entities that are components of u_2, the existence of an SRL congruence from $\langle u_1, \mathsf{I}_1\rangle$ to $\langle u_2, \mathsf{I}_2\rangle$ implies that I_1 and I_2 are isomorphic. Informally, I say that the configurations under u_1 and u_2 are isomorphic, or the components of u_1 and u_2 are isomorphically configured.

A first entity in a first Σ interpretation and a second entity in a second Σ interpretation are *SRL congruent in* Σ exactly if there is a species and attribute preserving bijection from the components of the first entity in the first interpretation to the components of the second entity in the second interpretation, and the bijection maps the first entity to the second entity:

Definition 26 *For each SRL signature* Σ*, for each* Σ *interpretation* $\mathsf{I}_1 = \langle \mathsf{U}_1, \mathsf{S}_1, \mathsf{A}_1\rangle$*, for each* $u_1 \in \mathsf{I}_1$*, for each* Σ *interpretation* $\mathsf{I}_2 = \langle \mathsf{U}_2, \mathsf{S}_2, \mathsf{A}_2\rangle$*, for each* $u_2 \in \mathsf{I}_2$*,*

$\langle u_1, \mathsf{I}_1\rangle$ *and* $\langle u_2, \mathsf{I}_2\rangle$ *are* ***SRL congruent in*** Σ

iff for some f*,* f *is an SRL congruence from* $\langle u_1, \mathsf{I}_1\rangle$ *to* $\langle u_2, \mathsf{I}_2\rangle$ *in* Σ*.*

The terminology of DEFINITION 22–26 allows us to be precise about an important aspect of King's idea of using exhaustive models to explicate the meaning of HPSG 94 grammars that I have already mentioned in the discussion of the car scenario of Figure 2.5. This aspect is intimately tied to King's realistic approach to scientific theories, which takes the natural language itself to be the target model of the grammar instead of some mathematical representation of the language. When I formulated the grammar of $\mathsf{I}_{2.5}$, I did not say that the envisioned grammar should have $\mathsf{I}_{2.5}$ as its unique exhaustive model. What I said was that the grammar should license configurations of entities that were configured isomorphically to the configurations involving Anke, Kordula, Detmar, Anke's husband, the first car, and the second car. To reformulate this in our new terminology, it is not required that every configuration under some entity in every model of the grammar be in the intended model of the grammar; it is only required that the intended model contain at least one configuration under an entity *that is SRL congruent* to any such configuration. That means that not every possible model is part of the intended model. In the example above, the fact that it is only required that a configuration under an entity that is SRL congruent to the configuration under the third car of Figure 2.6 must be part of an exhaustive model of the SRL grammar $\langle \Sigma_{2.4}, \theta_1\rangle$ means that it is not necessary that the concrete scenario involving the third car, Anne, Manfred, Maria, and Mike is part of the intended model. What must be part of an exhaustive model of $\langle\Sigma_{2.4}, \theta_1\rangle$ is at least one scenario that looks just like the one in Figure 2.6: There must be a Volkswagen that is owned by a woman and driven by a man, and both people like different persons of the opposite sex best. Crucially, the scenario might involve different entities from the ones in $\mathsf{I}_{2.6}$, namely a different Volkswagen and different people.

Similarly, linguistic grammars have a class of exhaustive models which differ from each other with respect to the number of SRL congruent configurations under entities in them, and with respect to the nature of the objects in the universe of the models. The significance of this will become clear when looking at the characterization of King, 1999, p. 343, of the content of an SRL grammar of a natural language. The idea is to formulate "a strong and well-defined necessary condition for an HPSG grammar to be true of a natural language" (King, 1999, p 336): An SRL grammar $\langle\Sigma, \theta\rangle$ is true of a natural language only if

1. the natural language can be construed as an interpretation, I, of the SRL signature Σ, where I is a system of linguistic tokens,

2. each description in θ is true of each entity in the natural language, i.e., the natural language, I, is a $\langle\Sigma, \theta\rangle$ model, and

3. some description in θ is false of some component of each entity, u', in any Σ interpretation, I', for which no entity u in the natural language, I, has isomorphically configured components.

I will discuss the implications of King's three conditions for a mathematical characterization of the meaning of SRL grammars in turn:

A crucial aspect of condition 1 is that natural languages are understood as systems of linguistic tokens. King takes tokens to be objects, events, states, locations, or complexes (King, 1999, p. 316). Most importantly, linguistic tokens are at least partly empirical entities that are accessible to observation and testing. Typical manifestations of linguistic tokens are the utterances of a speaker. In order to account for the fact that almost all of the infinitely many grammatical sentences of a natural language never occur as actual utterances, King concedes that he must subsume non-actual tokens under his notion of linguistic tokens. While the actual linguistic tokens are those that did occur or will occur in our world, non-actual tokens happen not to be manifest in our universe. In the context of an architecture for phonology expressed in HPSG, Höhle 1999, pp. 74–77, elaborates this concept of tokens and refers to the utterance tokens in the denotation of a grammar as *possible utterance events*.[45] Adopting Höhle's terminology, I will usually call King's tokens the *possible tokens* of a language.[46]

The linguistic knowledge of native speakers determines what a legitimate token of their language is, independent of whether a token is actual or non-actual. In order to account for the familiar differences in the grammaticality judgments of native speakers that might depend on the particular situation in which a native speaker is tested, King appeals to Chomsky's notion of an ideal speaker/hearer of a language. The ideal speaker/hearer of a natural language is a fictional and flawless representative of the community of speakers of the natural language. According to King, 1999, p. 317, a linguistic token belongs to a natural language exactly if the linguistic knowledge of the ideal speaker/hearer judges it grammatical.

King holds that we do not know much about the exact nature of linguistic tokens yet. But we know enough about them to propose empirically testable hypotheses about configurations of linguistic tokens. A formalism in which these hypotheses can be meaningfully expressed must be compatible with the assumption that the entities that are subject to description are empirical entities and as such accessible to observation and testing. In a long discussion of the approach of Pollard and Sag 1994, which takes linguistic object types as the structures that are subject to grammatical description, King argues against the widespread representational interpretation of the meaning of grammars and in favor of a realistic interpretation that avoids the additional abstraction step from (at least partly) empirical entities to the theoretical construct of linguistic types. According to King, considerations of ontological parsimony suggest that linguistic types should be avoided in a linguistic formalism, unless to do so would compromise the simplicity of grammars. Condition 1 implies that the entities in the intended interpretation of the SRL grammar of a natural language are possible tokens of the

[45] Höhle sketches an interpretation function that maps phonological linguistic tokens in possible utterance events to physical utterance events. The theory of this function makes clear that King's token objects can be regarded as abstract entities that are not necessarily identical with physical events.

[46] Note that the endorsement of possible tokens calls into question an unqualified technical characterization of King's view as realistic. As mentioned before, I use the term merely as a convenient label that highlights some important aspects of King's theory. I have nothing to say about philosophy of science.

natural language. In Section 2.2.2.3, I will return to the issue of possible tokens as elements of the interpretation of grammars in a mathematical reconstruction of the distinction between linguistic types and linguistic tokens of Pollard and Sag 1994.

Condition 2 narrows down the intended interpretations of linguistic SRL grammars to models. For any linguistic SRL grammar $\langle \Sigma, \theta \rangle$, only those Σ interpretations are candidates for the intended interpretation whose entities are linguistic tokens, and each description in the set of principles, θ, describes every token in the interpretation. While condition 1 ties the intended interpretation to the signature of the grammar, condition 2 links it to the descriptions in θ.

Condition 3, finally, says that the natural language is an exhaustive model of the grammar. More technically, the condition is equivalent to saying that for every entity u' in every model I' of the grammar, there is an entity u in the natural language, I, such that $\langle u, \mathsf{I} \rangle$ and $\langle u', \mathsf{I}' \rangle$ are SRL congruent. Essentially, condition 3 affirms that the intended model of the grammar contains every possible token of the natural language, and nothing else. If there is a model of a hypothetic SRL grammar of a natural language that contains a configuration under an entity for which there is no isomorphically configured possible token of the language, then we must conclude that the grammar is flawed, because it permits possible tokens that cannot exist.

The problem with formulating the requirement that the intended model of a grammar is the model that contains exactly the possible tokens of the language is that there is no way in which the concept of a linguistic token—as opposed to a non-linguistic token—can be made mathematically precise. To get around that problem, the formal characterization of the intended linguistic model must be accomplished without an appeal to the unknown properties that distinguish linguistic entities from other entities in the world. To achieve this, King says that a grammar is only true of a natural language if the natural language is in a certain class of models, namely in the class of exhaustive models of the grammar. The class of exhaustive models of a grammar is determined independently of the question of whether the entities in a particular model in that class are all linguistic, partly linguistic, or not linguistic at all. The intended interpretation of the grammar remains the natural language, but all other models in the class of exhaustive models that the grammar determines are descriptively indistinguishable from the natural language. We then know that any model in the class of exhaustive models equals the natural language that we intend to describe in all descriptively relevant respects, and any model in that class can be used in model theoretic and computational investigations of the natural language. To put it in slightly different words, every model in the class of exhaustive models emulates the natural language model, and the existence of a configuration under an entity in any exhaustive model therefore implies the existence of an SRL congruent configuration under a possible token of the natural language.

King's conditions 1–3 characterize natural languages as being in a class of models of SRL grammars of these natural languages. If it turned out that we cannot conceive of a natural language as a model of an SRL grammar or not even as an interpretation of an SRL signature, this might easily jeopardize the scientific value of writing grammars of natural languages in the SRL formalism, unless there is a systematic relationship between the actual structures of natural languages and the interpretations of SRL signatures. Below, when I criticize the SRL formalism as being insufficient for capturing how HPSG linguists think of natural languages, I will claim that the interpretations of SRL are structures that are not rich enough to capture the structure of natural languages (as conceived by HPSG linguists). However, the richer interpretations with which I will replace the interpretations of SRL are a proper extension of SRL interpretations. If the richer formalism that I will propose is on the right track, the meaning of SRL grammars can still be regarded as a valuable approximation to the structure of natural languages.[47]

In the remainder of this section I will provide a precise definition of the class of exhaustive models of a grammar. There are different equivalent possibilities of defining the class of exhaustive models. I will present two of them. The more intuitive definition is algebraic and uses SRL congruence. It is phrased in parallel to King's condition 3. Alternatively, the class of exhaustive models can be defined on the basis of the denotation of descriptions alone.

I begin with an algebraic definition of exhaustive models. For that purpose, I first define what it means for one Σ interpretation to simulate another

[47] From a type-based, representational perspective, the same problem would look slightly different. In that case, the set of object types that an SRL grammar specifies is understood to stand in a systematic relationship to the phenomena of the empirical domain. If the SRL formalism proved to be inadequate for expressing HPSG 94 grammars, it would mean that the previously assumed systematic relationship between the set of object types and the empirical domain can no longer be maintained. SRL grammars would remain a valuable approximation, if a new and possibly more complicated systematic relationship could be found.

Σ interpretation. The idea is that a first Σ interpretation, I_1, simulates a second Σ interpretation, I_2, if for each configuration under an entity in I_2, there is an SRL congruent counterpart in I_1. As far as configurations under entities in I_2 are concerned, I_1 can therefore be regarded as a perfect substitute of I_2: It can differ from I_2 because it might contain configurations under entities for which there is no SRL congruent counterpart in I_2; but it contains a counterpart to every configuration under an entity that is in I_2.

Definition 27 *For each SRL signature Σ, for each Σ interpretation $\mathsf{I}_1 = \langle \mathsf{U}_1, \mathsf{S}_1, \mathsf{A}_1 \rangle$, for each Σ interpretation $\mathsf{I}_2 = \langle \mathsf{U}_2, \mathsf{S}_2, \mathsf{A}_2 \rangle$,*

> I_1 ***simulates*** I_2 ***in*** Σ
>
> *iff for each $u_2 \in \mathsf{U}_2$, for some $u_1 \in \mathsf{U}_1$, $\langle u_1, \mathsf{I}_1 \rangle$ and $\langle u_2, \mathsf{I}_2 \rangle$ are SRL congruent in Σ.*

The notion of one interpretation simulating another interpretation leads directly to the definition of exhaustive models:

Definition 28 *For each SRL signature Σ, for each $\theta \subseteq \mathcal{D}^{\Sigma}$, for each Σ interpretation I,*

> I *is an* ***exhaustive*** $\langle \Sigma, \theta \rangle$ ***model*** *iff*
>
> > I *is a $\langle \Sigma, \theta \rangle$ model, and*
> >
> > *for each Σ interpretation I',*
> >
> > > *if I' is a $\langle \Sigma, \theta \rangle$ model then I simulates I' in Σ.*

A model, I, of a grammar Γ is an exhaustive model of Γ exactly if for every configuration under an entity in any other model of Γ, we can find an SRL congruent counterpart in I. Loosely speaking, every configuration of entities in any model has a mirror image in an exhaustive model. In terms of linguistic grammars, the intended model in the class of exhaustive models of a grammar contains every possible grammatical token that is licensed by the grammar; and since it is a model, it does not contain any ungrammatical configurations under a token.

The relationship between an algebraic definition of exhaustive models as in DEFINITION 28 and a characterization that is based on description denotations becomes clear when looking at the notion of two entities in two interpretations of the same SRL signature being indiscernible and comparing it to SRL congruence. Two entities in two Σ interpretations are indiscernible if it is impossible to tell the two entities in their interpretations apart by any Σ description that the formal language provides:

Definition 29 *For each SRL signature Σ, for each Σ interpretation $\mathsf{I}_1 = \langle \mathsf{U}_1, \mathsf{S}_1, \mathsf{A}_1 \rangle$, for each $u_1 \in \mathsf{U}_1$, for each Σ interpretation $\mathsf{I}_2 = \langle \mathsf{U}_2, \mathsf{S}_2, \mathsf{A}_2 \rangle$, for each $u_2 \in \mathsf{U}_2$,*

> $\langle u_1, \mathsf{I}_1 \rangle$ *and* $\langle u_2, \mathsf{I}_2 \rangle$ *are* ***indiscernible in*** Σ
> *iff for each $\delta \in \mathcal{D}^{\Sigma}$, $u_1 \in D_{\mathsf{I}_1}(\delta)$ iff $u_2 \in D_{\mathsf{I}_2}(\delta)$.*

Proposition 1 asserts that indiscernibility and congruence are in fact equivalent. If it is impossible to distinguish an entity, u_1, in a Σ interpretation, I_1, from another entity, u_2, in another Σ interpretation, I_2, by any description of the formal language, then $\langle u_1, \mathsf{I}_1 \rangle$ and $\langle u_2, \mathsf{I}_2 \rangle$ are congruent, and *vice versa*:[48]

Proposition 1 *For each SRL signature Σ, for each Σ interpretation $\mathsf{I}_1 = \langle \mathsf{U}_1, \mathsf{S}_1, \mathsf{A}_1 \rangle$, for each $o_1 \in \mathsf{U}_1$, for each Σ interpretation $\mathsf{I}_2 = \langle \mathsf{U}_2, \mathsf{S}_2, \mathsf{A}_2 \rangle$, for each $o_2 \in \mathsf{U}_2$,*

> $\langle o_1, \mathsf{I}_1 \rangle$ *and* $\langle o_2, \mathsf{I}_2 \rangle$ *are SRL congruent in* Σ *iff* $\langle o_1, \mathsf{I}_1 \rangle$ *and* $\langle o_2, \mathsf{I}_2 \rangle$ *are indiscernible in* Σ.

The equivalence of indiscernibility and congruence leads to the possibility of a characterization of the simulation of models in terms of the denotation of descriptions. For any given SRL signature Σ, for each set of Σ descriptions θ, and for each Σ interpretation I, to say that if a Σ interpretation I' is a $\langle \Sigma, \theta \rangle$ model then I simulates I' is equivalent to saying that if I' is a $\langle \Sigma, \theta \rangle$ model and an arbitrary theory θ' describes something in I' then θ' also describes something in I. In effect, as far as description denotation is concerned, I is a perfect substitute for the model I'. Proposition 2 summarizes this result:

Proposition 2 *For each SRL signature Σ, for each $\theta \subseteq \mathcal{D}^{\Sigma}$, for each Σ interpretation I,*

[48] I omit the proof of Proposition 1 and Proposition 2 at this point. Both propositions are special cases of the corresponding propositions, Proposition 10 and Proposition 11, of the extended formalism of Section 3.1, which I prove in Appendix A on page 376 and on page 380, respectively.

for each Σ *interpretation* I',

if I' *is a* $\langle\Sigma,\theta\rangle$ *model then* I *simulates* I' *in* Σ

iff for each $\theta' \subseteq \mathcal{D}^{\Sigma}$, *for each* Σ *interpretation* I',

if I' *is a* $\langle\Sigma,\theta\rangle$ *model and* $\Theta_{\mathsf{I}'}(\theta') \neq \emptyset$ *then* $\Theta_{\mathsf{I}}(\theta') \neq \emptyset$.

Since the simulation of models can alternatively be expressed in terms of the denotation of theories as made precise in the preceding proposition, we can replace the simulation condition in the definition of exhaustive models (Definition 28) by the equivalent characterization of the relationship between two interpretations in terms of nonempty denotations of sets of descriptions in them, and we get an alternative way of describing the class of exhaustive models of a grammar:

Corollary 1 *For each SRL grammar* $\langle\Sigma,\theta\rangle$, *for each* Σ *interpretation* I,

I *is an exhaustive* $\langle\Sigma,\theta\rangle$ *model iff*

I *is a* $\langle\Sigma,\theta\rangle$ *model, and*
for each $\theta' \subseteq \mathcal{D}^{\Sigma}$, *for each* Σ *interpretation* I',
if I' *is a* $\langle\Sigma,\theta\rangle$ *model and* $\Theta_{\mathsf{I}'}(\theta') \neq \emptyset$ *then* $\Theta_{\mathsf{I}}(\theta') \neq \emptyset$.

Intuitively, Corollary 1 asserts that an interpretation, I, of an SRL grammar Γ is an exhaustive model exactly if any arbitrary theory that denotes something in some model of Γ also denotes something in I: Whatever we can describe in any Γ model thus finds a counterpart in I. This corresponds to the requirement of the algebraic definition that every configuration under an entity in any Γ model find an SRL congruent counterpart in I.

The advantage of the non-algebraic characterization of the exhaustive models of a grammar in Corollary 1 over the algebraic characterization of Definition 28 is that it does not need any additional algebraic notions such as congruence and simulation. The definitions of the formal language of SRL and the definition of the model of a grammar are sufficient to say what the class of exhaustive models of a grammar is.[49]

[49]Richter et al. 1999, p. 293, and Richter 1999, p. 90, choose this characterization of exhaustive models exactly because it needs fewer preliminary definitions.

THEOREM 1 provides a final important mathematical result about exhaustive models. It assures that each SRL grammar has an exhaustive model. Since there are exhaustive models for each SRL grammar, it makes sense to explain the meaning of an arbitrary SRL grammar in terms of the class of its exhaustive models.

Theorem 1 (King 1999) *For each SRL signature Σ, for each Σ theory θ, there exists a Σ interpretation I such that I is an exhaustive $\langle \Sigma, \theta \rangle$ model.*

It follows immediately that each SRL grammar that has a nonempty model, i.e., a model that contains at least one entity, also has a nonempty exhaustive model.

2.2.2.3 Types and Tokens

At first blush, King's explanation of the meaning of SRL grammars of Section 2.2.2.2 appears to be so different from the view of Pollard and Sag 1994—summarized in Section 2.2.1 above—that it is hard to see how they are related. Although the philosophical differences in their understanding of the meaning of grammars are indeed considerable, it is nevertheless possible to make both approaches to language and grammar mathematically explicit within the single formal framework of SRL, and doing so provides interesting insights into the relationship between the two approaches. Pollard and Sag's account of the entities that are specified by a grammar will also allow us to see HPSG 94 under the perspective of some of the technical apparatus of HPSG 87, thus making some aspects of the relationship between the two types of formalisms clearer on an intuitive and on a technical level.

Firstly, Pollard and Sag (1994) hold that an HPSG 94 grammar is about the object types of a language and not about possible tokens. A grammar admits a collection of object types. This view is motivated by their assumption that the knowledge of the speakers of a natural language is not about individual tokens of sentences like *John loves Mary*, but rather about the unique object type *John loves Mary*. Pollard and Sag do not specify the relationship between the object type of a sentence and the observable tokens of it in the world, but it seems natural to assume that there must be some kind of abstraction function that maps the tokens of a sentence to their common, unique object type. Secondly, Pollard and Sag assume that object types are modeled by some kind of totally well-typed and sort resolved feature structures. By contrast, the entities in the exhaustive models of King are not necessarily feature structures. The entities in interpretations of SRL signatures are simply given as a set of entities in an interpretation, without any commitment to their internal structure.[50] Moreover, the intended linguistic model in the class of exhaustive models that is the natural language is thought of as a system of possible tokens. The possible tokens of a language are entities in an interpretation. Linguistic tokens in an interpretation are not mathematical entities like feature structures; they are the empirical entities that constitute a natural language.

There are numerous ways to reconstruct the relationship between the views of King on the one hand and Pollard and Sag on the other. The most comprehensive mathematical presentation of this issue is put forth by King (1996, pp. 33–36), who provides a translation from Pollard and Sag's signatures with sort hierarchies to SRL signatures with the same interpretations, and constructs Pollard and Sag's object types as equivalence classes of indiscernible pairs of an entity and an interpretation, $\langle u, \mathsf{I}\rangle$, where $u \in \mathsf{I}$. By choosing equivalence classes of entities in interpretations as the objects of interest, the distinction between indiscernible (or SRL congruent) linguistic entities in interpretations is lost. Loosing the distinction between indiscernible possible tokens in interpretations captures the central idea of Pollard and Sag's object types. King proceeds to define feature structures under an SRL signature. *King feature structures* are essentially potentially infinite, totally well-typed and sort resolved abstract feature structures.[51] A definition of *King feature structure admission* says what it means for a set of Σ descriptions to admit a set of King feature structures under Σ. The basic idea is that a King feature structure is admitted by a set of descriptions iff each (abstract) node of the King feature structure satisfies each description in the set. Feature structure admission can thus be seen as saying for each SRL grammar Γ which set of King feature structures is admitted by Γ.[52]

Abstraction functions finally establish the connection between interpre-

[50]As I will show later, for each entity in an interpretation, a uniquely determined concrete feature structure can be constructed from the subalgebra generated by the entity. In that sense, the subalgebra generated by an entity in an interpretation can be understood as indicating an internal structure of the entity. In principle, however, the internal structure of an entity in an interpretation is unrelated to the structure of the subalgebra that it generates.

[51]The definition of *94 abstract feature structures* in DEFINITION 32 below is very similar to the definition of King feature structures.

[52]The definition of *morph admission*, DEFINITION 77, on page 195 is essentially an extension of King's feature structure admission.

tations and King feature structures: The abstraction functions that King defines are total functions from the set of entities in a Σ interpretation to the set of King feature structures under Σ. King shows that for each signature Σ, and for each Σ interpretation I, I is an exhaustive Γ model if and only if the abstraction function from the set of entities in the Σ interpretation I to the set of King feature structures under Σ determines a bijection from the set of object types determined by I to the set of King feature structures under Σ admitted by Γ (King, 1996, p. 36). In other words, the object types of any exhaustive model of a grammar stand in a one-to-one correspondence to the King feature structures admitted by that grammar. It is therefore justified to regard the set of King feature structures admitted by the grammar as the set of object types that Pollard and Sag originally had in mind. From King's point of view, the set of King feature structures that are admitted by an HPSG 94 grammar are structured mathematical objects that represent the object types of a natural language; and the object types of a natural language that they represent are equivalence classes of indiscernible possible tokens in the natural language.

In the present section I will investigate the same general picture from a slightly different perspective. I am interested in clarifying how elements of the HPSG 87 formalism relate to the distinction between possible tokens and types in HPSG 94. By relating interpretations of SRL grammars to sets of concrete feature structures, I obtain an alternative way of deriving Pollard and Sag's object types from the exhaustive models of an SRL grammar. What is interesting about choosing this route is that it establishes the link between exhaustive models and feature structure representations of object types by using definitions and techniques that we know from HPSG 87. Studying the distinction of possible tokens and object types in terms of two kinds of feature structures of HPSG 87 and their relationship sheds light on the change of perspective from HPSG 87 to HPSG 94 regarding the task of a grammar.

Pollard and Sag 1994 envisions an architecture of grammar in which a grammar determines a set of abstract feature structures by an admission function that assigns each grammar a set of abstract feature structures. The summary of the mathematical reconstruction in the SRL formalism of Pollard and Sag's object types has revealed their relationship to the exhaustive models of a grammar. Since the members of the set, S, of abstract feature structures admitted by a grammar stand in a one-to-one correspondence to the indiscernibility classes of possible tokens in an exhaustive model of the grammar, the answer that S gives to the question which linguistic data a grammar predicts is empirically equivalent to the answer that the exhaustive models of the grammar give. The admitted abstract feature structures represent the object types of a natural language; and the set of admitted abstract feature structures is the content of an HPSG 94 grammar. According to Pollard (personal communication), this is the standard view of generative grammarians specialized to the case where the structural representations of linguistic objects are abstract feature structures. Given the properties of abstract feature structures that we have seen in Section 2.1.2, it is an obvious thought to assume that if abstract feature structures correspond to the object types of a natural language, then the system of possible tokens of a natural language corresponds to a collection of concrete feature structures. In order to spell out this idea, I first need a definition of the collection of *94 concrete feature structures.*

Just as 87 concrete feature structures, 94 concrete feature structures are concrete feature structures in the sense of the general DEFINITION 2 with a few additional restrictions. The restrictions of 87 concrete feature structures were finiteness, acyclicity and the appropriateness of their attributes. The restrictions of 94 concrete feature structures are total well-typedness and sort resolvedness. A minor complication with defining 94 concrete feature structures as a special case of concrete feature structures is that in the current SRL-based framework, signatures do not contain a sort hierarchy, whereas concrete feature structures were defined as concrete feature structures under an 87 signature with a sort hierarchy. In addition, the appropriateness function, $\mathcal{F}$, of 87 signatures is a partial function from the Cartesian product of the set of sorts and the set of attributes to the set of sorts, whereas in SRL signatures, $\mathcal{F}$ is a total function from the Cartesian product of the set of species and the set of attributes to the power set of the set of species. To accommodate these differences, I need to introduce an auxiliary 87 signature, Σ', in the definition of 94 concrete feature structures under the SRL signature Σ:

Definition 30 *For each SRL signature* $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, $\mathbb{C}$ *is a* ***94 concrete feature structure under*** Σ *iff*

$\sqsubseteq = \left\{ \langle \sigma, \sigma \rangle \in \mathcal{S} \times \mathcal{S} \,\middle|\, \sigma \in \mathcal{S} \right\}$, $\mathcal{F}' = \{\}$, $\Sigma' = \langle \mathcal{S}, \sqsubseteq, \mathcal{A}, \mathcal{F}' \rangle$,

$\mathbb{C} = \langle Q, \hat{q}, \Delta, \Lambda \rangle \in \mathbb{CFS}^{\Sigma'}$,

for each $q_1 \in Q$, *for each* $q_2 \in Q$, *for each* $\alpha \in \mathcal{A}$,

$\Delta \langle q_1, \alpha \rangle$ *is defined iff* $\mathcal{F} \langle \Lambda(q_1), \alpha \rangle \neq \emptyset$, *and*
if $\Delta \langle q_1, \alpha \rangle$ *is defined then* $\Lambda(\Delta \langle q_1, \alpha \rangle) \in \mathcal{F} \langle \Lambda(q_1), \alpha \rangle$.

If Σ is an SRL signature, I write $\mathbb{CFS}^{\Sigma}_{94}$ for the class of 94 concrete feature structures under Σ. Note that $\langle \mathcal{S}, \sqsubseteq, \mathcal{A}, \mathcal{F}' \rangle$ is an 87 signature. Its sort hierarchy, $\langle \mathcal{S}, \sqsubseteq \rangle$, is the Noetherian partial order in which each species only subsumes itself, and the partial appropriateness function, $\mathcal{F}'$, is undefined on each pair of a species and an attribute. The signature Σ' is needed, because I defined concrete feature structures with reference to 87 signatures instead of SRL signatures. However, neither the partial order of the sort hierarchy nor the appropriateness function are actually referred to in the general definition of concrete feature structures, DEFINITION 2. Therefore, I can fix them in any way that satisfies the requirements of the definition of 87 signatures without consequences. Note also that the sort resolvedness of 94 concrete feature structures follows immediately from the fact that they are defined under SRL signatures, which only contain species and leave the sort hierarchy implicit.

With a precise notion of 94 concrete feature structure in hand, I can now pursue the question of how 94 concrete feature structures are related to the entities in interpretations of SRL signatures. DEFINITION 31 implies a straightforward answer:

Definition 31 *For each SRL signature* $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, *for each* Σ *interpretation* $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$,

$$\mathbb{CFS}^{\mathsf{I}}_{94} = \left\{ \begin{array}{l} \langle Q, \hat{q}, \Delta, \Lambda \rangle \\ \in \mathbb{CFS}^{\Sigma} \end{array} \middle| \begin{array}{l} \textit{for some } u \in \mathsf{I}, \\ Q = \mathsf{Co}_{\mathsf{I}}(u), \\ \hat{q} = u, \\ \Delta = \left\{ \begin{array}{l} \langle \langle u_1, \alpha \rangle, u_2 \rangle \in \\ (Q \times \mathcal{A}) \times Q \end{array} \middle| \begin{array}{l} \textit{for some } \alpha \in \mathcal{A}, \\ \textit{for some } u_1 \in \mathsf{Co}_{\mathsf{I}}(u), \textit{ and} \\ \textit{for some } u_2 \in \mathsf{Co}_{\mathsf{I}}(u), \\ \langle \alpha, \langle u_1, u_2 \rangle \rangle \in \\ \mathsf{A} \cap (\mathcal{A} \times (\mathsf{Co}_{\mathsf{I}}(u) \times \mathsf{Co}_{\mathsf{I}}(u))) \end{array} \right\}, \\ \Lambda = \mathsf{S} \cap (\mathsf{Co}_{\mathsf{I}}(u) \times \mathcal{S}) \end{array} \right\}.$$

It can easily be checked that for each SRL signature Σ, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, and for each $u \in \mathsf{I}$, there is a concrete feature structure in $\mathbb{CFS}^{\mathsf{I}}_{94}$ with root node u. For each $u \in \mathsf{I}$, Q is a set, because the collection of components of each entity in an interpretation is a set. Under each root node $\hat{q} = u$, the connectivity condition on concrete feature structures is satisfied, because the transition function Δ is derived from the attribute interpretation function restricted to the components of u: The transition from $u_1 \in Q$ to $u_2 \in Q$ with attribute α is defined iff the attribute interpretation of α at the component u_1 of u is defined and yields the component u_2 of u. But being a component of some u in an interpretation is defined as being reachable by finitely many interpretations of attributes, and connectivity in concrete feature structures means that every node in the concrete feature structure can be reached from the root node in finitely many transitions that consume finitely many attributes. Finally, for each $u \in \mathsf{I}$, Λ is a total function from Q to $\mathcal{S}$, because Q is the set of components of u, and Λ is the species interpretation function restricted to the components of u.

Since $\mathbb{CFS}^{\mathsf{I}}_{94}$ is defined for interpretations of SRL signatures, which must obey the appropriateness function, and since the requirements of SRL's appropriateness function correspond to the total well-typedness of feature structures, the elements of $\mathbb{CFS}^{\mathsf{I}}_{94}$ can be characterized even more specifically as 94 concrete feature structures:

Proposition 3 *For each SRL signature* Σ*, for each* Σ *interpretation* I*,*

$\mathbb{CFS}^{\mathsf{I}}_{94}$ *is a set of 94 concrete feature structures.*

For each SRL signature Σ and for each Σ interpretation I, I call $\mathbb{CFS}^{\mathsf{I}}_{94}$ the *set of 94 concrete feature structures determined by* I. For each entity $u \in \mathsf{I}$, there is exactly one 94 concrete feature structure with root node u in $\mathbb{CFS}^{\mathsf{I}}_{94}$. If we look at an arbitrary element $\mathbb{C} = \langle Q, \hat{q}, \Delta, \Lambda \rangle$ with root node $\hat{q} = u$ of $\mathbb{CFS}^{\mathsf{I}}_{94}$, we can see an interesting correspondence of $\mathbb{C}$ to the configuration of entities under u in I, $\langle u, \langle \mathsf{U}_u, \mathsf{S}_u, \mathsf{A}_u \rangle \rangle$: Q equals U_u, Λ equals S_u, and there is a bijection from A_u to Δ. It follows that for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, there is a bijection from the set of configurations of entities under some u in I to $\mathbb{CFS}^{\mathsf{I}}_{94}$ such that for each $u \in \mathsf{I}$, $\langle u, \langle \mathsf{U}_u, \mathsf{S}_u, \mathsf{A}_u \rangle \rangle$ is mapped to the 94 concrete feature structure in $\mathbb{CFS}^{\mathsf{I}}_{94}$ with the root node u. A 94 concrete feature structure determined by I with root node u is essentially another way of looking at the configuration under u in I.

Exhaustive models of an SRL grammar Γ contain SRL congruent counterparts to every configuration under an entity that can be found in some Γ model. For an exhaustive Γ model I, $\mathbb{CFS}^{\mathsf{I}}_{94}$ thus contains at least one 94 concrete feature structure that corresponds to an SRL congruent counterpart of any configuration under an entity that can occur in some Γ model.

Intuitively speaking, we might say that it contains at least one 94 concrete feature structure of every shape that is licensed by Γ. But just as in the case of exhaustive models where multiple occurrences of isomorphic configurations under an entity in an exhaustive model are possible, $\mathbb{CFS}^{\mathsf{I}}_{94}$ might contain multiple 94 concrete feature structures of the same shape. To be more precise, whenever an exhaustive Γ model I contains n configurations under an entity that are SRL congruent to $\langle u, \mathsf{I}_u \rangle$ for some u in I, then $\mathbb{CFS}^{\mathsf{I}}_{94}$ contains n 94 concrete feature structures with different root nodes such that they are pairwise isomorphic with the 94 concrete feature structure with the root node u. To put it slightly differently, these 94 concrete feature structures are identical except for their nodes. Assuming King's interpretation of the meaning of SRL grammars, if I is the exhaustive Γ model that is the natural language, then isomorphic 94 concrete feature structures in $\mathbb{CFS}^{\mathsf{I}}_{94}$ are mathematical representations of different tokens in the natural language of the same kind of entity, for example of different tokens of the sentence *Kim likes bagels.*

In HPSG 87, we encountered a corresponding phenomenon with 87 concrete feature structures. The fact that two distinct 87 concrete feature structures could potentially subsume each other prevented them from forming a partial order under subsumption, and we introduced 87 abstract feature structures in order to abstract away from distinct entities that carried the same information in the same way. The abstraction function collapsed 87 concrete feature structures that mutually subsumed each other into a single 87 abstract feature structure. Apart from being slightly unintuitive under the view that feature structure subsumption between different feature structures corresponds to the relative informativeness of feature structures, the potential isomorphisms between distinct 87 concrete feature structures prevented them from forming a partial order under subsumption. This was one of the reasons for why they could not form a Heyting algebra under subsumption either. The Heyting algebra of feature structures was an important part of the formalization of the HPSG 87 formalism of Section 2.1, because it provided the algebraic operations on feature structures on which an HPSG 87 grammar relies.

The perspective that HPSG 94 takes regarding concrete and abstract feature structures is very different, but the same technique that is used in HPSG 87 to get from concrete to abstract feature structures can be used in HPSG 94 for a different purpose, namely for capturing the relationship between representations of tokens and representations of types. If I is the natural language, then the elements of $\mathbb{CFS}^{\mathsf{I}}_{94}$ are complete mathematical representations of the possible tokens in the natural language, and it does not make sense to ask about the relative informativeness of the isomorphic mathematical representations of two tokens. However, we might want to collapse them into a single representation of the object type of the two tokens. From HPSG 87 we know that the abstraction relation, Abst_Σ, of DEFINITION 10 does exactly what we are looking for: It abstracts away from the distinction that is caused by isomorphic concrete feature structures having distinct nodes.[53] The domain of the *94 abstraction function under* Σ that can be derived as a special case from Abst_Σ is the collection of 94 concrete feature structures, $\mathbb{CFS}^{\Sigma}_{94}$, and its range the set of 94 abstract feature structures, $\mathbb{AFS}^{\Sigma}_{94}$, which I define as follows:

Definition 32 *For each SRL signature* $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, $\mathbb{A}$ *is a* ***94 abstract feature structure under*** Σ *iff*

$$\sqsubseteq = \left\{ \langle \sigma, \sigma \rangle \in \mathcal{S} \times \mathcal{S} \,\middle|\, \sigma \in \mathcal{S} \right\},\ \mathcal{F}' = \{\},\ \Sigma' = \langle \mathcal{S}, \sqsubseteq, \mathcal{A}, \mathcal{F}' \rangle,$$

$$\mathbb{A} = \langle \beta, \varrho, \lambda \rangle \in \mathbb{AFS}^{\Sigma'},$$

for each $\pi \in \mathcal{A}^*$, *for each* $\alpha \in \mathcal{A}$,

$$\textit{if } \pi\alpha \in \beta \textit{ then } \lambda(\pi\alpha) \in \mathcal{F}\langle \lambda(\pi), \alpha \rangle,$$

for each $\pi \in \mathcal{A}^*$, *for each* $\alpha \in \mathcal{A}$,

$$\textit{if } \pi \in \beta \textit{ and } \mathcal{F}\langle \lambda(\pi), \alpha \rangle \neq \emptyset \textit{ then } \pi\alpha \in \beta.$$

If Σ is an SRL signature, I write $\mathbb{AFS}^{\Sigma}_{94}$ for the set of 94 abstract feature structures under Σ. Due to the restrictions on 94 abstract feature structures under Σ, they are totally well-typed. The sort resolvedness follows immediately from the fact that they are defined under SRL signatures.

All relevant properties of Abst_Σ hold of the 94 abstraction relation between $\mathbb{CFS}^{\Sigma}_{94}$ and $\mathbb{AFS}^{\Sigma}_{94}$. In particular, it is a total function, and it maps isomorphic 94 concrete feature structures to the same 94 abstract feature structure. The nodes of the 94 concrete feature structures are turned into equivalence classes

[53]Again, we need to modify the definition of Abst_Σ minimally in order to account for the fact that Σ is now an SRL signature instead of an 87 signature. As with other definitions before, this modification is unproblematic.

of paths that lead from the root node to the same node. Since an exhaustive Γ model, I, of a grammar Γ always contains SRL congruent counterparts to all existing configurations under an entity in any Γ model, and since it may only differ from other exhaustive Γ models with regard to how many SRL congruent "copies" of configurations under an entity it contains, the set of 94 concrete feature structures determined by I, $\mathbb{CFS}^{\mathsf{I}}_{94}$, only differs from the sets of 94 concrete feature structures determined by other exhaustive Γ models in the number of isomorphic 94 concrete feature structures of a given shape. Suppose that for a given grammar Γ and an exhaustive Γ model I, $\mathbb{AFS}^{\mathsf{I}}_{94}$ is the set of 94 abstract feature structures that we get by collecting the results of applying the 94 abstraction function to the members of $\mathbb{CFS}^{\mathsf{I}}_{94}$. By the properties of the 94 abstraction function and by the properties of the sets of 94 concrete feature structures determined by exhaustive Γ models it follows that for each SRL signature Σ, for each SRL grammar $\Gamma = \langle \Sigma, \theta \rangle$, and for each pair of exhaustive Γ models I and I', $\mathbb{AFS}^{\mathsf{I}}_{94} = \mathbb{AFS}^{\mathsf{I}'}_{94}$. That means that the set of 94 abstract feature structures that we abstract from each $\mathbb{CFS}^{\mathsf{I}}_{94}$ is actually independent of the choice of the exhaustive Γ model, I. It is always the same set of 94 abstract feature structures. Following the terminology of Pollard and Sag, I call it the *set of 94 abstract feature structures admitted by* Γ, and write $\mathbb{AFS}^{\Gamma}_{94}$.

Constructing the set of the 94 abstract feature structures, $\mathbb{AFS}^{\Gamma}_{94}$, admitted by a grammar Γ from a set of 94 concrete feature structures determined by an exhaustive Γ model can be shown to be equivalent to the definition of King feature structure admission of King 1996.[54] The members of $\mathbb{AFS}^{\Gamma}_{94}$ are mathematical representations of the object types of the natural language. In fact, arriving at $\mathbb{AFS}^{\Gamma}_{94}$ via sets of 94 concrete feature structures can be regarded as a complication of King's simpler abstraction function that maps entities in an interpretation directly to abstract feature structures. If the interpretation is an exhaustive Γ model, King's abstraction function is a surjection from the set of entities in the exhaustive Γ model to the abstract feature structures admitted by Γ. However, the detour that I have taken illuminates a number of terminological and conceptual issues in HPSG 94.

Firstly, two approaches to linguistic tokens are technically conceivable. The collection of possible tokens of a language is the intended model in the

[54] In order to prove this, King's definition of King feature structures must be replaced by 94 abstract feature structures, which causes some further minor adaptations in some mathematical details of King 1996.

realistic approach to the meaning of HPSG 94 grammars of King. Grammars directly specify the system of actual and non-actual tokens of a natural language. The natural language is one particular model in the class of exhaustive models of the grammar. Feature structures do not occur in this explanation of the meaning of HPSG 94 grammars at all. As we have seen, it is nevertheless possible to represent the possible tokens of a language as a collection of concrete feature structures. Any collection of 94 concrete feature structures determined by an exhaustive model of the grammar in question is appropriate to that task. In the set $\mathbb{CFS}^{\mathcal{L}}_{94}$ of 94 concrete feature structures that is determined by the natural language model, $\mathcal{L}$, in the class of exhaustive models of the grammar, every 94 concrete feature structure is a mathematical representation of a possible token of the language. The theoretical problem with this construction is that, under the perspective of a realist, mathematical representations of tokens are nonsensical, because a realist wants to specify tokens without any intervening modeling structure. From the representational perspective, specifying possible tokens as a collection of concrete feature structures is not reasonable either, because proponents of the representational approach aim at mathematical representations of the object types of the language as modeling structures and deny that it makes sense to see the meaning of a grammar in the collection of the possible tokens of the language. This position stands in the tradition of generative grammar, which sees the specification of linguistic competence as one of its major goals. Linguistic competence is understood as a knowledge system that is shared by the speakers of a language. The knowledge represented in the minds of the speakers of a language is assumed to be knowledge of linguistic object types rather than knowledge of individual linguistic utterance tokens.

Secondly, there are two approaches to types. In King's realistic view of HPSG 94, the object types of a natural language are the equivalence classes of indiscernible possible tokens of the natural language. They are an additional construct that is of no immediate interest, because the meaning of a grammar is given by its exhaustive models. In the eyes of Pollard and Sag (1994), this is completely different, because they see the meaning of a grammar in the set of object types admitted by the grammar. The object types of a natural language are represented by a set of potentially infinite, totally well-typed and sort resolved abstract feature structures. The set of 94 abstract feature structures admitted by a given grammar $\Gamma = \langle \Sigma, \theta \rangle$, $\mathbb{AFS}^{\Gamma}_{94}$, is the set of 94 abstract feature structures under Σ such that each (abstract) node of each one of them satisfies all descriptions in θ. $\mathbb{AFS}^{\Gamma}_{94}$ is thus determined directly by a feature structure admission function. Since the tokens of a language are the empirical basis of grammars, they are deemed as important in this approach as under the realistic perspective. However, in contrast to the realistic approach, tokens are not part of the linguistic formalism. The connection between the formalism and the data is made by the linguistic theory, which postulates a correspondence between certain observable aspects of tokens and certain properties of the structure of the object types that the grammars specifies.

We have seen two ways of relating the realistic and the representational interpretation of the object types of a natural language to each other. Firstly, I have summarized work by King, who defines abstraction functions that map each entity, u, in an exhaustive model, I, of a grammar to the abstract feature structure admitted by the grammar that represents the object type to which $\langle u, \mathsf{I}\rangle$ belongs. Secondly, I have sketched a way of first deriving the set of concrete feature structures determined by an exhaustive model of a grammar, and then abstracting the set of abstract feature structures admitted by the grammar from that set. Again, for each entity, u, in an exhaustive model, I, of a grammar, we obtain the abstract feature structure admitted by the grammar that represents the object type to which $\langle u, \mathsf{I}\rangle$ belongs. The second method is less direct, but it illuminates some aspects of the relationship between HPSG 87 and HPSG 94.

Comparing HPSG 94 with HPSG 87, we see a stark contrast in the role that feature structures play in the two formalisms. In HPSG 94, they are optional constructs that can be used as mathematical representations of tokens and types. But tokens and types can be defined without recourse to feature structures. From a realist's perspective, they are merely a convenient representation that does not play a role in explaining the meaning of an HPSG 94 grammar. Of course, that standpoint does not deny the mathematical or computational importance of feature structures in the broader context of an HPSG 94 formalism.[55] From a representational perspective, abstract feature structures are the only type of feature structures that are needed in HPSG 94, since the set of abstract feature structures admitted by an SRL grammar can be determined directly via an admission function.

In HPSG 87, feature structures are an integral part of the grammatical formalism. They function as indispensable representations of partial infor-

[55]For example, as we will see below, the existence proof of exhaustive models employs a canonical feature structure model.

mation. The fact that disjunctive feature structures form a Heyting algebra under subsumption provides the algebraic operations with which the grammatical principles and the way in which they are combined are formulated. The main source of the different status of feature structures in the two formalisms lies in a seemingly innocuous technical detail and the intuition about the task of grammars that it is meant to express: The feature structures of HPSG 87 are partial, because they are representations of partial information, and the unification of feature structures expresses the accumulation of partial information. The feature structures of HPSG 94 are totally well-typed and sort resolved, because, according to Pollard and Sag 1994, they are total representations of linguistic object types.

Coming back to the ongoing debate between the proponents of a realistic and of a representational interpretation of HPSG 94 grammars, it can be argued that the representational view of the relationship between the structures specified by a grammar and the empirical domain is currently more widely accepted among generative grammarians than the view proposed by King. It is not obvious, however, if the differences between the two, although philosophically deep, are significant for current linguistic practice in the HPSG framework. Since for each SRL grammar, both approaches make the same predictions about the empirical domain, only strong assumptions about the relationships between the possible tokens or the object types of a natural language and observable data, and arguments from philosophy of science can distinguish between the two positions.[56] An investigation of these issues is, however, well beyond the scope of my thesis, and I restrict myself to a very short summary of the arguments that have been put forth by King and Pollard.

King 1999 argues against the view that linguistic object types should be specified by HPSG 94 grammars and against explaining the meaning of grammars by means of feature structures. Regarding the notion of the object types of a natural language, which is fundamental in the theory of Pollard and Sag, King doubts that linguists share a commonly accepted intuition about what exactly a linguistic type is: "[It] is an amorphous, artificial notion that linguists can and do construe in a variety of ways to suit their particular purposes. Witness the different ways members of the same collec-

[56] For a very interesting discussion of the relationship between possible tokens and physical utterance events, and of the theoretical implications of different assumptions about this relationship, see Höhle 1999, pp. 74–77.

tion of speech events are divided into types as the concerns of a linguist are phonetic or phonological" (King, 1999, p. 313). Regarding the integration of object types into the HPSG 94 framework, King claims that object types are secondary mathematical entities in a theory of natural language, because they are not entities in the empirical domain of observable linguistic events. The object types of a grammar themselves are immune to empirical testing, and it is only their relation to tokens that makes them empirically accessible. Given this situation, King argues that the scientific principle of ontological parsimony suggests to avoid object types completely. The object types of a grammar would only become empirical objects if it could be shown that the mental representations of linguistic knowledge are indeed about object types, and if there were methods to observe and test these representations. From the realistic perspective of King, a similar argument of parsimony applies to the use of feature structures. Feature structures can be used as mathematical representations of tokens and types, but they are dispensable. In particular, nothing can be gained by representing language tokens as 94 concrete feature structures instead of assuming that the system of language tokens itself is the intended model of the grammar. In an early paper on feature logics, Gert Smolka already formulated this perception of the status of feature structures in grammar formalisms of the kind of HPSG 94. Contrasting the closed world approach, in which only feature structures are admitted in the interpretations, with an open world semantics, in which there is no commitment as to the nature of the entities in the interpretations of the expressions of the logical language, Smolka (1988, p. 38) observes: "Although the closed world approach is technically okay, it is unsatisfying philosophically [. . .]. The open world semantics presented in this paper renders the detour over feature structures superfluous: one can now view feature terms as directly denoting sets of linguistic objects and there is still no need for making precise what linguistic objects are."

Besides the philosophical argument, King cautions against using feature structures and feature structure admission as the basis of the content of a grammar for methodological reasons. Feature structure admission is defined with abstract feature structures. In the discussion of formalisms of HPSG 87 and HPSG 94, we have seen that many different definitions of abstract feature structures exist. When natural languages are modeled by feature structures, it is often not clear which linguistic intuitions correspond to particular restrictions on feature structures. Typical examples are the acyclicity condition or the finiteness condition on feature structures. But if the linguistic significance of that kind of technical decision is not transparent then a linguist who relies on a particular definition of feature structure admission cannot know which assumptions about language might be implicitly hidden in the formalism, and the linguist might or might not share the hidden underlying intuitions. In contrast to the purely technical definition of the meaning of a grammar by feature structure admission, the approach of King expresses three necessary conditions that say what it means for a grammar to be true of a natural language. A linguist who uses that formalism is thus informed about the assumptions about natural languages that are inherent in the formalism.

In accordance with the mainstream opinion about the meaning of grammars in generative linguistics, Pollard (personal communication) rejects the idea that a grammar specifies a collection of tokens. Most importantly, Pollard is deeply troubled by the necessity of postulating non-actual tokens, which is caused by the (intended) fact that a grammar of a natural language specifies infinitely many non-isomorphic phrases, whereas one can reasonably assume that only finitely many tokens of a language ever occur in our world. It follows immediately that almost all tokens specified by an SRL grammar are fictitious entities. To Pollard, the concept of non-actual tokens is contradictory and nonsensical, and a theory of language which must assume them is unacceptable. More technically, Pollard doubts that there are any reasons to assume that the collection of possible tokens of a natural language forms a set, as it is supposed by the models of SRL grammars. If one does not share this premise, the model theory of King implies an assumption that is not supported by intuitions about the concept of possible tokens. In contrast to King, Pollard agrees with the widespread assumption that a scientific theory of language should not describe individual linguistic tokens within the formal framework. Instead, mathematical representations of object types that abstract away from individual tokens are the right choice. Linguistic tokens, manifest in utterance events, are predicted by the theory via intervening modeling structures that stand in a conventional correspondence relation to utterance events.[57] For proponents of this architecture of linguistic theories, even the actual tokens among the possible tokens of a language are unfit to

[57]This view is strongly reminiscent of physical theory as it emerged at the turn of the century. It distinguishes an *experimental domain*, a *mathematical model*, and a *conventional interpretation*. See Abraham and Marsden 1967, pp. 1–4, for a brief overview. Note that this theory precedes the formulation by Tarski of model-theoretic semantics, which might introduce interesting new aspects that were not recognized at the time.

appear as modeling structures, because their fundamental nature cannot be precisely known. Given that King's account must concede the existence of non-actual tokens, it may even appear questionable if his account is more parsimonious than a representationalist view that eschews ontological assumptions and defines the meaning of a grammar as a certain mathematical entity.

In the next section, I present Pollard's definition of the strong generative capacity of a grammar. Like its intellectual ancestor, the account of Pollard and Sag 1994, Pollard's (1999) account is strictly representational, and similar arguments apply for and against Pollard 1999 as for and against the approach of Pollard and Sag 1994. However, there are some notable differences. First of all, Pollard 1999 avoids claims about the ontological status of object types and regards them as isomorphism classes of structures that include the structures that idealize tokens. Object types do not necessarily exist in the world. On the technical side, this slight change of emphasis finds its counterpart in new modeling structures. Although Pollard 1999 still refers to them as feature structures and abstract feature structures, they are no longer the automata theoretic concrete feature structures and abstract feature structures of the mathematical reconstruction of Pollard and Sag 1994 of the present section.

2.2.2.4 Strong Generative Capacity

The purpose of Pollard 1999 is to give a precise characterization of the meaning of HPSG 94 grammars that explains in which sense an HPSG 94 grammar is a generative grammar. In contrast to Pollard and Sag 1994, Pollard 1999 does not espouse the idea that an HPSG 94 grammar specifies the set of object types of a natural language; nor does it share the view that an HPSG 94 grammar specifies the possible tokens of a natural language. The goal of the paper is to formally define the strong generative capacity (SGC) of object-based HPSG. With the SGC of a grammar, Pollard intends to describe the set of structures that is "generated" by an HPSG 94 grammar in a similar sense in which a generative grammar generates a language. According to Pollard, the SGC of a grammar should capture two intuitions. Firstly, no two members of the SGC of a grammar are structurally isomorphic. Secondly, if a grammar is correct, exactly those utterance tokens that are structurally isomorphic to members of the SGC will be judged grammatical (Pollard, 1999, p. 282). Pollard's definition of the SGC of a grammar is related to representations of types and to representations of tokens. With types it shares the assumption that there is only one representative of a each collection of isomorphically configured linguistic entities. However, in contrast to the two characterizations of types in the previous section, where types have a different structure than tokens, it assumes that each token of the language is structurally isomorphic to a member of the SGC.

Pollard 1999 contains very precise descriptions of the mathematical constructs it envisages. In the present section, I spell out the technical details of the paper in the notation and terminology of the previous sections in order to fit its concepts seamlessly into the context of the other explications of HPSG 94. I proceed in three steps. Firstly, I will present Pollard's definition of a feature structure, which is a construct that we have already seen under a different name. Secondly, Pollard's abstract feature structures are introduced as a canonical form of Pollard's feature structures. Thirdly, I present Pollard's three ways of obtaining the SGC of a grammar as sets of abstract feature structures, and I discuss their relationship to the above formalizations of tokens and types as feature structures. Finally, I use the definitions of Pollard to formulate a definition of a canonical exhaustive model of a grammar that Pollard (1999, p. 295) indicates only briefly.

Pollard Feature Structures Pollard 1999 chooses to define feature structures in a non-standard fashion, exploiting the existence of a bijection between the set of 94 concrete feature structures determined by an interpretation, I, and the set of all pairs of an entity, u, of I and the subinterpretation of I in which all entities have u as their common ancestor. In order to signal the difference between Pollard's feature structures and the standard definitions of concrete feature structures, I will call the new feature structures *Pollard feature structures.* DEFINITION 33 introduces Pollard's terminology for what I have called the interpretation under an entity in an interpretation (see DEFINITION 23):

Definition 33 *For each SRL signature $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, for each $v \in \mathsf{U}$,*

$$\mathsf{I}_v = \langle \mathsf{U}_v, \mathsf{S}_v, \mathsf{A}_v \rangle \textit{ is the } \Sigma \textbf{\textit{ interpretation generated by }} v \textbf{\textit{ in }} \mathsf{I} \textit{ iff}$$

$$\mathsf{U}_v = \left\{ u \in \mathsf{U} \;\middle|\; \begin{array}{l} \textit{for some } \tau \in \mathcal{T}^{\Sigma} \\ T_{\mathsf{I}}(\tau)(v) \textit{ is defined, and } T_{\mathsf{I}}(\tau)(v) = u \end{array} \right\},$$

$\mathsf{S}_v = \mathsf{S} \cap (\mathsf{U}_v \times \mathcal{S})$,
$\mathsf{A}_v = \mathsf{A} \cap (\mathcal{A} \times (\mathsf{U}_v \times \mathsf{U}_v))$, *and*
$\mathsf{I}_v = \langle \mathsf{U}_v, \mathsf{S}_v, \mathsf{A}_v \rangle$.

A Pollard feature structure is simply a pair of an entity and an interpretation such that the interpretation is generated by the entity:

Definition 34 *For each SRL signature Σ, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, for each $v \in \mathsf{U}$,*

$\langle v, \mathsf{I} \rangle$ *is a* ***Pollard feature structure*** *iff*
I *is the Σ interpretation generated by v in I.*

A Pollard feature structure determined by v in I is what I have called a configuration of entities under v in I:

Definition 35 *For each SRL signature Σ, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, for each $v \in \mathsf{U}$,*

$\langle v, \mathsf{I}_v \rangle$ *is the* ***Pollard feature structure determined by*** v ***in*** I *iff*
I_v *is the Σ interpretation generated by v in I.*

In Section 2.2.2.3 we have seen that for every SRL signature Σ, for every Σ interpretation I, and for each $u \in \mathsf{I}$, exactly one 94 concrete feature structure with root node u is in the set of 94 concrete feature structures determined by I, and it corresponds to the Pollard feature structure determined by u in I in a natural way. It is because of this correspondence that Pollard 1999 chooses Pollard feature structures as its representation of concrete feature structures or tokens of linguistic entities in a natural language instead of a standard definition.

Pollard Abstract Feature Structures Pollard abstract feature structures stand in a similar relation to Pollard feature structures as standard abstract feature structures do to standard concrete feature structures. That means that the main idea remains to switch from nodes as concrete objects to an abstract representation of nodes as equivalence classes. The difference to the standard definition is that now the nodes that must be represented as equivalence classes are actually entities in an interpretation of an SRL signature. What is therefore needed is a canonical representation of entities in an interpretation. However, for a canonical representation of the entities in Pollard abstract feature structures, the only interpretations of interest are those whose entities are reachable by the interpretation of some term on a common ancestor, because the entities in a Pollard feature structure, $\langle v, \mathsf{I} \rangle$, are all components of v. For that reason, a variant of the familiar construction of an equivalence relation over paths can be applied: The entities in the interpretation of a feature structure generated by its root entity are identified with equivalence classes of terms. In preparation of this, DEFINITION 36 introduces an appropriate equivalence relation over terms for each SRL signature.

Definition 36 *For each SRL signature* $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, ν *is a* ***node abstraction in*** Σ *iff*

ν *is a pair* $\langle \beta, \varrho \rangle$,

$\beta \subseteq \mathcal{T}^{\Sigma}$,

$: \in \beta$,

for each $\tau \in \mathcal{T}^{\Sigma}$, *for each* $\varphi \in \mathcal{A}$, *if* $\tau\varphi \in \beta$ *then* $\tau \in \beta$,

ϱ *is an equivalence relation over* β,

for each $\tau_1 \in \mathcal{T}^{\Sigma}$, *for each* $\tau_2 \in \mathcal{T}^{\Sigma}$, *for each* $\varphi \in \mathcal{A}$,

if $\tau_1\varphi \in \beta$ *and* $\langle \tau_1, \tau_2 \rangle \in \varrho$ *then* $\langle \tau_1\varphi, \tau_2\varphi \rangle \in \varrho$.

I call ϱ the *re-entrancy relation in the node abstraction in* Σ. A node abstraction in Σ is a structure that can easily be used as a canonical representation of entities in an interpretation generated by an entity in the interpretation. The element that generates the interpretation will naturally be represented by the equivalence class $|:|_{\varrho}$. Therefore, ':' must always be in the carrier, β. Moreover, β is a prefix closed set of terms, because each prefix of a term with a nonempty denotation at an entity in an interpretation also has a nonempty denotation at that entity in the interpretation. The canonical representation of entities must therefore contain entities that are values of all prefixes of a term that is defined on the canonical root entity, $|:|_{\varrho}$, and they will be represented by the equivalence class of the node abstraction that contains that prefix. Finally, ϱ is a right invariant equivalence relation, because if two Σ terms τ_1 and τ_2 are defined at an entity u in an interpretation and yield the same value then for any Σ path π, $\tau_1\pi$ is defined on u exactly if $\tau_2\pi$ is defined on u, and if they are defined they yield the same value in the interpretation. The equivalence classes of the node abstraction must mirror this fact, i.e., $\tau_1\pi$ and $\tau_2\pi$ must be in the same equivalence class representing an entity if the interpretation of τ_1 and τ_2 at $|:|_\varrho$ yield the same value; right invariance guarantees that property.

A node abstraction in Σ is not related to the appropriateness function in Σ yet, and no species are assigned to the equivalence classes of terms in the node abstraction. The intended intuitive correlation between a node abstraction and appropriateness is achieved in the definition of Pollard abstract feature structures by tying the structure of the interpretation of Pollard abstract feature structures in a clever way to the node abstraction whose equivalence classes represent the entities in the interpretation in the abstract feature structure:

Definition 37 *For each SRL signature $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, for each $v \in \mathsf{U}$,*

$\langle v, \mathsf{I} \rangle$ *is a* ***Pollard abstract feature structure under*** Σ *iff*

$\langle v, \mathsf{I} \rangle$ *is a Pollard feature structure,*
for some node abstraction in Σ, $\langle \beta, \varrho \rangle$,
$\mathsf{U} = Quo(\langle \beta, \varrho \rangle)$,
$v = |:|_\varrho$,
for each $v_1 \in Quo(\langle \beta, \varrho \rangle)$, for each $v_2 \in Quo(\langle \beta, \varrho \rangle)$, for each $\varphi \in \mathcal{A}$,
$\mathsf{A}(\varphi)(v_1) = v_2$ *iff*
for some $\tau \in \beta$, $v_1 = |\tau|_\varrho$, and $v_2 = |\tau\varphi|_\varrho$.

Note that by construction, the collection of all Pollard abstract feature structures under a fixed SRL signature Σ is a set. I write $\mathbb{AFS}^{\Sigma}_{\mathrm{P}}$ for the set of Pollard abstract feature structures under Σ. If $\langle v, \mathsf{I} \rangle$ is a Pollard abstract feature structure, I will call v its *root entity*.

Since each Pollard abstract feature structure, $\langle v, \mathsf{I} \rangle$, under an SRL signature Σ is also a Pollard feature structure, its interpretation, I, is generated by its root entity, v. The entities in the interpretation I are the equivalence classes of a node abstraction in Σ, and the root entity, v, is the equivalence class containing the colon. The crucial trick is how the attribute interpretation function, A, of the Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$ of a Pollard abstract feature structure under Σ is related to the equivalence classes of terms that are its entities: An attribute φ is defined on an entity v_1, yielding v_2 as value, exactly if the equivalence class v_1 contains a term τ_1 that is a prefix of a term τ_2 in v_2, and $\tau_1\varphi = \tau_2$. It follows that if an entity v_n can be reached by interpreting the Σ term τ on the root entity then v_n is the equivalence class $|\tau|_\varrho$. Since I is a Σ interpretation, appropriateness follows by definition. More specifically, for each $\varphi \in \mathcal{A}$, for each $|\tau|_\varrho \in \mathsf{U}$, $\mathsf{A}(\varphi)(|\tau|_\varrho)$ is defined if and only if $\mathcal{F}\langle \mathsf{S}(|\tau|_\varrho), \varphi \rangle \neq \emptyset$, and if $\mathsf{A}(\varphi)(|\tau|_\varrho)$ is defined then $\mathsf{S}(|\tau\varphi|_\varrho) \in \mathcal{F}\langle \mathsf{S}(|\tau|_\varrho), \varphi \rangle$.

Since Pollard abstract feature structures have a canonical representation of their entities as equivalence classes of the terms whose interpretation on the root entity yields the respective entities, each one of them is a canonical representation of the Pollard feature structures that are isomorphic to it. PROPOSITION 4 confirms the existence of a unique isomorphic Pollard abstract feature structure for each Pollard feature structure.

Proposition 4 *For each SRL signature Σ, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$, for each $v \in \mathsf{U}$, if $\langle v, \mathsf{I} \rangle$ is a Pollard feature structure then there exists exactly one $\langle v', \mathsf{I}' \rangle \in \mathbb{AFS}^{\Sigma}_{\mathrm{P}}$, such that for each $\delta \in \mathcal{D}^{\Sigma}$,*

$$v \in D_{\mathsf{I}}(\delta) \text{ iff } v' \in D_{\mathsf{I}'}(\delta).$$

With previous results, the indiscernibility of the root entities of a Pollard feature structure and a Pollard abstract feature structure entails that the two feature structures mutually simulate each other. Since their interpretations are generated by their root entities, it follows that they are isomorphic.

Exploiting the canonical representations of entities, it is in fact not hard to define an abstraction function that maps each Pollard feature structure to the isomorphic Pollard abstract feature structure. This abstraction function is the counterpart of the abstraction functions from various kinds of concrete feature structures to their corresponding abstract feature structures in the previous sections. Since the three ways of obtaining the strong generative capacity of a grammar that Pollard 1999 presents do not involve Pollard feature structures that are not Pollard abstract feature structures, I omit it here.

Pollard abstract feature structures have the properties that express the intuitions of Pollard 1999 about the strong generative capacity of a grammar.

No two distinct Pollard abstract feature structures are isomorphic, but they can be isomorphic to language tokens, if the latter are construed as pairs of an entity, v, and the interpretation, I_v, generated by v in the natural language, I. The strong generative capacity of a grammar will be characterized as the set of Pollard abstract feature structures that the grammar generates, where the notion of generating a set of abstract feature structures corresponds to the notion of (abstract) feature structure admission implicit in Pollard and Sag 1994.

Strong Generative Capacity Pollard 1999 presents three equivalent definitions of the strong generative capacity of a grammar. All of them use the by now familiar relationships between the models of a grammar, the class of exhaustive models of a grammar, and the set of abstract feature structures that can be abstracted from them. Since Pollard uses canonical, non-standard concrete feature structures as abstract feature structures, there is no notion of the set of abstract feature structures admitted by a grammar that is defined independently of the models of the grammar. On the other hand, since an interpretation generated by the root entity is part of every Pollard abstract feature structure, the third definition of the strong generative capacity of a grammar will offer a new way of deriving the intended set of abstract feature structures. The first two definitions of the SGC of a grammar are similar to constructions that were discussed in Section 2.2.2.3.

For the first definition of the strong generative capacity of a grammar, Pollard defines an abstraction function, $\mathrm{Abst}_{\mathsf{I}}$. For each entity v in the interpretation I, it contains pairs of v and Pollard abstract feature structures whose root entities are indiscernible from v:

Definition 38 *For each SRL signature Σ, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$,*

> $\mathrm{Abst}_{\mathsf{I}}$ *is the binary relation on $\mathsf{U} \times \mathbb{AFS}^{\Sigma}_{\mathrm{P}}$ such that, for each $v \in \mathsf{U}$, for each $\langle v', \mathsf{I}' \rangle \in \mathbb{AFS}^{\Sigma}_{\mathrm{P}}$, $\langle v, \langle v', \mathsf{I}' \rangle \rangle \in \mathrm{Abst}_{\mathsf{I}}$ iff*
>
> > *for the Pollard feature structure $\langle v, \mathsf{I}_v \rangle$ determined by v in I, for each $\delta \in \mathcal{D}^{\Sigma}$, $v \in D_{\mathsf{I}_v}(\delta)$ iff $v' \in D_{\mathsf{I}'}(\delta)$.*

Since each entity, v, in an interpretation, I, determines exactly one Pollard feature structure, $\langle v, \mathsf{I}_v \rangle$, in the interpretation I, it follows with PROPOSITION 4 that each $\mathrm{Abst}_{\mathsf{I}}$ is functional:

Proposition 5 *For each SRL signature* Σ*, for each* Σ *interpretation* $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$*,* $\mathrm{Abst}_{\mathsf{I}}$ *is a total function from* U *to* $\mathbb{AFS}^{\Sigma}_{\mathrm{P}}$*.*

$\mathrm{Abst}_{\mathsf{I}}$ directly corresponds to King's abstraction function from the set of entities in an interpretation to abstract feature structures, and the same properties hold. Entities that are indiscernible in I abstract to the same Pollard abstract feature structure that is isomorphic to them. If v is an entity in an interpretation I, then I call $\mathrm{Abst}_{\mathsf{I}}(v)$ its *abstraction*. The strong generative capacity of a grammar is supposed to contain one isomorphic member to each utterance token in the natural language that is judged grammatical. King chooses exhaustive models for the explanation of the meaning of a grammar, because an exhaustive model contains indiscernible counterparts to all entities in all models of the grammar. Correspondingly, the SGC of a grammar must contain each Pollard abstract feature structure that can be abstracted from some entity in some model of the grammar. This leads to Pollard's first definition:

Definition 39 (SGC 1) *For each SRL signature* Σ*,* SGC *is the total function from the set of grammars to* $Pow(\mathbb{AFS}^{\Sigma}_{\mathrm{P}})$ *such that, for each* $\theta \subseteq \mathcal{D}^{\Sigma}$*,*

$$\mathsf{SGC}(\langle \Sigma, \theta \rangle) = \left\{ \alpha \in \mathbb{AFS}^{\Sigma}_{\mathrm{P}} \;\middle|\; \begin{array}{l} \textit{for some } \langle \Sigma, \theta \rangle \textit{ model } \mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle, \\ \textit{for some } v \in \mathsf{U}, \\ \alpha = \mathrm{Abst}_{\mathsf{I}}(v) \end{array} \right\}.$$

For each grammar Γ, Pollard calls the set $\mathsf{SGC}(\Gamma)$ the *strong generative capacity of* Γ.

Pollard's first definition of the SGC does not refer to exhaustive models, although looking at every possible model of the grammar is ultimately motivated by the concept of exhaustive models. However, exhaustive models provide an obvious alternative to obtaining the SGC of a grammar by taking the abstract feature structures that can be abstracted from the entities in all models. It is clearly sufficient to consider only one model if that model is big enough in the sense of containing SRL congruent copies of all possible configurations under entities in all models, and to take the set of abstract feature structures that can be abstracted from the entities in such a model. Choosing this way of defining the SGC presupposes a definition of the class of exhaustive models of a grammar. Pollard 1999 presents a definition that is different from the two previous characterizations of exhaustive models in

Definition 28 and Corollary 1, because it exploits the existence of a canonical Pollard feature structure for each entity v in an interpretation, namely the existence of $\mathrm{Abst}_{\mathsf{I}}(v)$:

Definition 40 *For each SRL signature Σ, for each $\theta \subseteq \mathcal{D}^{\Sigma}$,*

> $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$ *is an **exhaustive** $\langle \Sigma, \theta \rangle$ **model** iff*
>
> I *is a $\langle \Sigma, \theta \rangle$ model, and for each $\langle \Sigma, \theta \rangle$ model $\mathsf{I}' = \langle \mathsf{U}', \mathsf{S}', \mathsf{A}' \rangle$, for each $v' \in \mathsf{U}'$, there exists a $v \in \mathsf{U}$ such that* $\mathrm{Abst}_{\mathsf{I}}(v) = \mathrm{Abst}_{\mathsf{I}'}(v')$.

Pollard's definition of the exhaustive models of a grammar Γ is closely related to their definition in terms of simulation. In his definition, an exhaustive Γ model is required to contain an SRL congruent Pollard feature structure for all possible Pollard feature structures in all Γ models, because he demands that each abstraction of each entity in each model be also an abstraction of an entity in the exhaustive model. But that is equivalent to saying that a Γ model is an exhaustive Γ model if and only if it simulates all other Γ models.

The function that is defined next maps each Σ interpretation I to the smallest set of Pollard abstract feature structures under Σ that, for each Pollard feature structure determined by some v in I, contains an element that is indiscernible from $\langle v, \mathsf{I}_v \rangle$:

Definition 41 *For each SRL signature Σ,* ABST *is the large total function from the class of Σ interpretations to $Pow(\mathbb{AFS}^{\Sigma}_{\mathrm{P}})$ such that, for each Σ interpretation $\mathsf{I} = \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle$,*

$$\mathsf{ABST}(\mathsf{I}) = \left\{ \alpha \in \mathbb{AFS}^{\Sigma}_{\mathrm{P}} \,\middle|\, \begin{array}{l} \textit{for some } v \in \mathsf{U}, \\ \alpha = \mathrm{Abst}_{\mathsf{I}}(v) \end{array} \right\}.$$

For each Σ interpretation I, I call $\mathsf{ABST}(\mathsf{I})$ the *set of Pollard abstract feature structures determined by* I. The function ABST leads to the second definition of the strong generative capacity of a grammar:

Definition 42 (SGC 2) *For each SRL signature Σ,* SGC *is the total function from the set of grammars to $Pow(\mathbb{AFS}^{\Sigma}_{\mathrm{P}})$ such that, for each $\theta \subseteq \mathcal{D}^{\Sigma}$,*

$$\mathsf{SGC}(\langle \Sigma, \theta \rangle) = \left\{ \alpha \in \mathbb{AFS}^{\Sigma}_{\mathrm{P}} \,\middle|\, \begin{array}{l} \textit{for each (equivalently some)} \\ \textit{exhaustive } \langle \Sigma, \theta \rangle \textit{ model } \mathsf{I}, \\ \alpha \in \mathsf{ABST}(\mathsf{I}) \end{array} \right\}.$$

The third and final way of Pollard 1999 of defining the strong generative capacity of a grammar exploits the fact that a Pollard abstract feature structure, in contrast to the Moshier representation of abstract feature structures, contains an interpretation of a signature. In particular, it contains the interpretation with canonical representations of the elements of the universe that is generated by its root entity. Pollard says that a grammar Γ *generates* a Pollard abstract feature structure exactly if its interpretation is a Γ model:

Definition 43 *For each SRL signature Σ, for each $\theta \subseteq \mathcal{D}^{\Sigma}$, for each Σ interpretation I, for each $v \in \mathsf{U}$, $\langle \Sigma, \theta \rangle$* ***generates*** *$\langle v, \mathsf{I} \rangle$ iff*

$\langle v, \mathsf{I} \rangle \in \mathbb{AFS}^{\Sigma}_{\mathrm{P}}$, and

I is a $\langle \Sigma, \theta \rangle$ model.

This leads to the last characterization of the SGC of a grammar as the set of all Pollard abstract feature structures whose interpretation is a model of a grammar Γ:

Definition 44 (SGC 3) *For each SRL signature Σ, SGC is the total function from the set of grammars to $Pow(\mathbb{AFS}^{\Sigma}_{\mathrm{P}})$ such that, for each $\theta \subseteq \mathcal{D}^{\Sigma}$,*

$$\mathsf{SGC}(\langle \Sigma, \theta \rangle) = \left\{ \alpha \in \mathbb{AFS}^{\Sigma}_{\mathrm{P}} \;\middle|\; \langle \Sigma, \theta \rangle \textit{ generates } \alpha \right\}.$$

Just as standard definitions of abstract feature structure admission, **SGC 3** strictly limits its perspective to mathematical representations of linguistic entities. The relationship between the mathematical representations and the empirical domain of observable phenomena is left outside the formal framework. This architecture is fully in line with Pollard's view of the domain of scientific theories in general and of HPSG 94 grammars in particular: They do not talk directly about empirical phenomena, but about mathematical structures that are conventionally understood to correspond to (certain aspects of) empirical phenomena. With this interpretation of scientific theories, Pollard is in deliberate opposition to the realistic interpretation of King, who rejects the representational view of scientific theories of Pollard and holds that "the [empirical] domain is itself a mathematical structure that directly interprets the theory" (King, 1999, p. 310). For King, the natural language, $\mathcal{L}$, construed as a system of tokens, is the intended model of a grammar. $\mathcal{L}$ is in the class of exhaustive models of the grammar.

The philosophical incompatibility of the two positions means that the central concept of each position—the exhaustive models of a grammar and the strong generative capacity of a grammar, respectively—is conceptually meaningless to a proponent of the opposite philosophical persuasion. Ignoring the fundamental incompatibility at the conceptual level of the two positions, we can nevertheless show how they are mathematically tied to each other and how we can technically switch from one perspective to the other. The three definitions of the strong generative capacity of grammars are a good example of how different approaches to the same construct provide different insights into its properties. **SGC 1** considers every possible model of a grammar. According to the realistic perspective, the class of models of the grammar includes the model, $\mathcal{L}$, that is the natural language as well as every other conceivable model, and $\mathcal{L}$ simulates all of them. The relationship of **SGC 2** to the realistic perspective is given by the observation that the SGC of a grammar can be determined by taking the particular exhaustive model of the grammar that is the natural language and abstracting the intended set of Pollard abstract feature structures from it. Since **SGC 1** and **SGC 2** can thus be reinterpreted as referring to the model that is the natural language, the perspectives offered by them are more obviously linked to King's view of the meaning of HPSG 94 grammars than the perspective of **SGC 3**.[58] Since all three definitions are equivalent, it is of course also possible to relate **SGC 3** to a realistic view of the meaning of HPSG 94 grammars, just as it is possible to link feature structure admission to exhaustive models, but the relation is less direct than in the case of **SGC 1** and **SGC 2**. PROPOSITION 6 establishes a connection between the Pollard abstract feature structures generated by a grammar and the set of Pollard abstract feature structures determined by a model of the grammar.

Proposition 6 *For each SRL signature* Σ*, for each grammar* $\langle \Sigma, \theta \rangle$*, for each* $\langle v, \mathsf{I} \rangle \in \mathbb{AFS}^{\Sigma}_{\mathrm{P}}$*,*

> $\langle v, \mathsf{I} \rangle$ *is generated by* $\langle \Sigma, \theta \rangle$ *iff*
>
> *there exists a* Σ *model* I' *such that* $\langle v, \mathsf{I} \rangle \in \mathsf{ABST}(\mathsf{I}')$.

This result allows a proponent of the realistic interpretation of grammars to perceive the set of Pollard abstract feature structures generated by a

[58] The indicated reinterpretation is, of course, against the express intention of Pollard, just as for King it is completely superfluous to extract the strong generative capacity from a natural language.

grammar as the set of Pollard abstract feature structures abstracted from the natural language model.

Pollard 1999 does not make any claims about the ontological status of the strong generative capacity of an HPSG 94 grammar. The more modest aim of the paper is to give the intuitions behind the idea of the strong generative capacity of a grammar mathematical substance in object-based HPSG in order to relate the HPSG 94 framework to other formalisms of generative grammar and to their ideas about the structures that a grammar generates. In doing so, Pollard 1999 adds another interesting piece to the understanding of the role that different mathematical constructs play in determining the meaning of HPSG 94 grammars.

Regarding the terminology of Pollard 1999, it is important to be careful not to confuse Pollard feature structures with the standard, automata theoretic definition of concrete feature structures. Although it is easy to switch from an entity in an interpretation to the concrete feature structure that represents the entity, this change of perspective might have philosophical repercussions: A concrete feature structure is a complex mathematical object, but the ontological status of entities in an interpretation is open; they can be conceived as entities in the empirical domain. Pollard feature structures are entities in an interpretation, which leaves the possibility of a realistic perspective. The difference between Pollard abstract feature structures and standard abstract feature structures is also of technical importance. Standard abstract feature structures represent equivalence classes of tokens. They are not tokens themselves, and they are not isomorphic with tokens. They cannot be construed as entities in an interpretation, because they cannot model the situation in which two entities in the interpretation have congruent component structures but are still distinct tokens in the interpretation. Pollard abstract feature structures, however, are just a special kind of Pollard feature structures, namely Pollard feature structures whose entities are particular equivalence classes of terms. They are isomorphic with possible tokens. Therefore, they can well be understood as mathematical representations of idealizations of possible linguistic tokens, but there is only one mathematical representation for each collection of isomorphic tokens. These subtle differences between Pollard's feature structures and standard feature structures as well as the differences between the various different kinds of abstract and concrete feature structures that occur in the context of HPSG should caution linguists against an unspecific use of the term "feature structure."

Pollard abstract feature structures can serve yet another purpose. THEOREM 1 asserts the existence of an exhaustive model for each grammar. A natural way to prove this theorem is to construct a model for each grammar, and to show that it is exhaustive. King's (unpublished) proof uses feature structure admission and 94 abstract feature structures to construct the exhaustive models.[59] Pollard 1999, p. 295, briefly indicates that Pollard abstract feature structures and the set of Pollard abstract feature structures generated by a grammar can be used to the same end. In DEFINITION 45, I spell out this idea.

The entities in the interpretation are constructed as complex objects. They are pairs of a Pollard abstract feature structure generated by the grammar and an entity in the feature structure. Intuitively, these complex objects are Pollard abstract feature structures with distinguished nodes. The set of entities in the universe of the exhaustive models can be understood as all Pollard abstract feature structures that the grammar generates with one distinguished node, where each feature structure occurs once with each node as distinguished node. The sort assignment is obtained by assigning each entity the sort that the Pollard abstract feature structure in it assigns to the distinguished node. Similarly, the attribute interpretation mimics the attribute interpretations in the abstract feature structures. An attribute φ is defined on an entity u exactly if it is defined on the distinguished node v in the interpretation of the Pollard abstract feature structure in the entity u, and if φ is defined on an entity u then the value is the entity u' that contains the same Pollard abstract feature structure as u and the distinguished node that is the value of φ on v in the interpretation of the Pollard abstract feature structure in question. In effect, the interpretation that is obtained by this construction contains isomorphic copies of all interpretations in the Pollard abstract feature structures generated by the grammar.

Definition 45 *For each SRL signature $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, for each $\theta \subseteq \mathcal{D}^{\Sigma}$,*

$$\mathbf{U}^{\theta}_{\Sigma} = \left\{ \begin{array}{l|l} \langle v', \langle v, \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle \rangle \rangle & \langle \Sigma, \theta \rangle \textit{ generates } \langle v, \langle \mathsf{U}, \mathsf{S}, \mathsf{A} \rangle \rangle, \textit{ and} \\ \in Pow(\mathcal{T}^{\Sigma}) \times \mathbb{AFS}^{\Sigma}_{\mathrm{P}} & v' \in \mathsf{U} \end{array} \right\},$$

[59] We will see the details of King's construction below as part of the existence proof of the extended formalism.

$$\mathbf{S}^{\theta}_{\Sigma} = \left\{ \langle v', \sigma \rangle \in \mathbf{U}^{\theta}_{\Sigma} \times \mathcal{S} \;\middle|\; \begin{array}{l} \textit{for some } \langle v, \langle Quo(\langle \beta, \rho \rangle), \mathsf{S}, \mathsf{A} \rangle \rangle \textit{ generated} \\ \textit{by } \langle \Sigma, \theta \rangle, \\ \textit{for some } \tau \in \beta, \\ v' = \langle |\tau|_{\rho}, \langle v, \langle Quo(\langle \beta, \rho \rangle), \mathsf{S}, \mathsf{A} \rangle \rangle \rangle, \textit{ and} \\ \sigma = \mathsf{S}(|\tau|_{\rho}) \end{array} \right\},$$

$\mathbf{A}^{\theta}_{\Sigma}$ *is the total function from* $\mathcal{A}$ *to the power set of* $\mathbf{U}^{\theta}_{\Sigma} \times \mathbf{U}^{\theta}_{\Sigma}$ *such that, for each* $\varphi \in \mathcal{A}$,

$$\mathbf{A}^{\theta}_{\Sigma}(\varphi) = \left\{ \begin{array}{l} \langle v', v'' \rangle \\ \in \mathbf{U}^{\theta}_{\Sigma} \times \mathbf{U}^{\theta}_{\Sigma} \end{array} \;\middle|\; \begin{array}{l} \textit{for some } \langle v, \langle Quo(\langle \beta, \rho \rangle), \mathsf{S}, \mathsf{A} \rangle \rangle \textit{ generated} \\ \textit{by } \langle \Sigma, \theta \rangle, \\ \textit{for some } \tau \in \beta, \\ v' = \langle |\tau|_{\rho}, \langle v, \langle Quo(\langle \beta, \rho \rangle), \mathsf{S}, \mathsf{A} \rangle \rangle \rangle, \textit{ and} \\ v'' = \langle |\tau\varphi|_{\rho}, \langle v, \langle Quo(\langle \beta, \rho \rangle), \mathsf{S}, \mathsf{A} \rangle \rangle \rangle \end{array} \right\},$$

$$\mathbf{I}^{\theta}_{\Sigma} = \left\langle \mathbf{U}^{\theta}_{\Sigma}, \mathbf{S}^{\theta}_{\Sigma}, \mathbf{A}^{\theta}_{\Sigma} \right\rangle.$$

A standard proof shows that $\mathbf{I}^{\theta}_{\Sigma}$ is a Σ interpretation as PROPOSITION 7 claims:

Proposition 7 *For each SRL signature* $\Sigma = \langle \mathcal{S}, \mathcal{A}, \mathcal{F} \rangle$, *for each* $\theta \subseteq \mathcal{D}^{\Sigma}$,

$\mathbf{U}^{\theta}_{\Sigma}$ *is a set,*

$\mathbf{S}^{\theta}_{\Sigma}$ *is a total function from* $\mathbf{U}^{\theta}_{\Sigma}$ *to* $\mathcal{S}$,

$\mathbf{A}^{\theta}_{\Sigma}$ *is a total function from* $\mathcal{A}$ *to the set of partial functions from* $\mathbf{U}^{\theta}_{\Sigma}$ *to* $\mathbf{U}^{\theta}_{\Sigma}$, *and*

$\mathbf{I}^{\theta}_{\Sigma}$ *is a* Σ *interpretation.*

Moreover, it can be shown that for each SRL signature Σ, for each set of Σ descriptions θ, for each $\mathbf{I}^{\theta}_{\Sigma}$, for each $\langle \Sigma, \theta \rangle$ model I', and for each $v' \in \mathsf{I}'$, there exists a $v \in \mathbf{I}^{\theta}_{\Sigma}$ such that $\text{Abst}_{\mathbf{I}^{\theta}_{\Sigma}}(v) = \text{Abst}_{\mathsf{I}'}(v')$. In other words, $\mathbf{I}^{\theta}_{\Sigma}$ is not only a Σ interpretation, it is an exhaustive $\langle \Sigma, \theta \rangle$ model:

Proposition 8 *For each SRL signature* Σ, *for each* $\theta \subseteq \mathcal{D}^{\Sigma}$,

$\mathbf{I}^{\theta}_{\Sigma}$ *is an exhaustive* $\langle \Sigma, \theta \rangle$ *model.*

Intuitively, this is not hard to see: We know that the set of Pollard abstract feature structures generated by a grammar is the same as the set of Pollard abstract feature structures that we obtain by collecting all the abstractions of all entities in all models of the grammar, because the first is the strong generative capacity of the grammar according to **SGC 3**, and the second is the strong generative capacity of the grammar according to the equivalent definition **SGC 1**. But we have seen that $\mathbf{I}^{\theta}_{\Sigma}$ copies over all interpretations of all Pollard abstract feature structures in the entities in $\mathbf{U}^{\theta}_{\Sigma}$ that were obtained that way. It immediately follows that we get all possible abstractions of all entities in all possible models of the grammar back when we take the set of abstractions of all entities in $\mathbf{I}^{\theta}_{\Sigma}$. Therefore, $\mathbf{I}^{\theta}_{\Sigma}$ is an exhaustive model.

2.2.3 Problems: Missing Concepts[60]

In Section 2.2.2, I presented a mathematical formalism, SRL, and I have summarized how King 1999 and Pollard 1999 express grammars, the meaning of grammars, and the notions of linguistic tokens and linguistic types within SRL. The accounts of tokens and types presented in both papers differ deliberately from each other and from the original, informal account of Pollard and Sag 1994, which King 1996 formalized in SRL. We have seen that these differences are due to different views of what linguistic tokens and linguistic types are, and how they are mathematically reconstructed in the formalism. The main differences, however, concern the opinions about the ontological status of the mathematical structures that a grammar delineates, and the relationship between these structures and the natural languages that are the empirical domain of grammars. What King's account and Pollard's account hold in common is that they are expressed in the SRL formalism. The common underlying assumption is that SRL is an adequate formal foundation of HPSG 94, and that HPSG 94 grammars can be made precise and can be expressed as SRL grammars. In the words of King 1999, p. 335: "In summation, a HPSG grammar can be expressed in very large part, if not in entirety, as a SRL grammar, where a SRL **grammar** is a pair (Σ, θ) such that Σ is a [SRL] signature and $\theta \subseteq \mathcal{D}^{\Sigma}$."

In the present section I will challenge this view. While I accept that all parts of the SRL formalism are necessary in a formalization of HPSG 94, and while I think it is desirable to maintain the explanations of the meaning of

[60]This section is partially based on work previously published in Richter et al. 1999.