Linking as constraints on word classes in a hierarchical lexicon *

Anthony R. Davis, University of California, Santa Barbara
Jean-Pierre Koenig, SUNY at Buffalo

June 16, 1999

Abstract

In this paper, we propose an account of linking patterns that does away with intermediary mechanisms such as thematic or actor/undergoer hierarchies. Instead, constraints on word classes, defined by both syntactic and semantic criteria, encode generalizations between semantic roles and syntactic arguments. We show that the generalizations a linking theory needs to capture can be modeled via the same mechanisms as other lexical generalizations, using conditions specified within the hierarchy of word classes. Each condition provides a partial specification of the mapping between semantic roles and syntactic arguments. We argue that this constraint-based, verb-class-based view of linking offers several empirical advantages: partial regularities and exceptions are easily accomodated, fine-grained semantic distinctions relevant to linking are countenanced, and cross-cutting similarities between semantic and syntactic verb classes are economically captured.

What are the general principles governing the mapping of semantic arguments onto syntactic functions? Theories of this correspondence (often referred to as linking theories or mapping theories) typically invoke intermediary mechanisms such as thematic or actor/undergoer hierarchies. Since the original work of Fillmore (1968) in case grammar, these hierarchies have featured prominantly in a wide range of models (see Foley and Van Valin (1984), Kiparsky (1987), Bresnan and Kanerva (1989), Jackendoff (1990), Grimshaw (1990) among others). By contrast, many recent theories rely on the notion of verb class and the semantic structure associated with sets of verbs (see Pinker (1989), Levin (1993), Goldberg (1995), Levin and Rappaport Hovav (1995)).

In this paper, we argue in favor of this second kind of theory and against the use of a thematic hierarchy in linking semantic arguments onto syntactic functions. We demonstrate that the generalizations a linking theory needs to capture can be modeled via the same mechanisms as other lexical generalizations and do not require intermediary mechanisms such as thematic or actor/undergoer hierarchies. In their place, we posit a set of

*The order of the authors' names is purely alphabetical; each contributed equally to this paper. We would like to thank Cleo Condoravdi, Gail Mauner, Sam Mchombo, Ivan Sag, Robert Van Valin and (at least) two anonymous referees for comments on the material included in this paper. All remaining errors are strictly our own.

hierarchically organized conditions, specified on word classes. Each condition provides a partial specification of the mapping between semantic roles and syntactic arguments. But our account differs from current verb-class-based theories in two crucial respects. First, our theory builds on the richer view of the lexicon which has emerged in recent years in several frameworks (see Bresnan (1982), Fillmore and Kay (forthcoming), Pollard and Sag (1987), Pollard and Sag (1994)). In particular, we take the view that word classes are organized in a subsumption hierarchy (Flickinger (1987), among others), with more specific classes inheriting information from more general ones. Secondly, we characterize the semantic generalizations on which linking depends in terms of an entailment-based theory of semantic roles (see Ladusaw and Dowty (1988), Dowty (1991), Wechsler (1991), Wechsler (1995)) rather than rely on the decomposition of verb meanings into a semantic metalanguage, as is typically the case in other verb-class-based accounts of linking. By adopting an entailment-based view of argument selection, we provide a more adequate semantic grounding for linking constraints than other verb-class-based approaches to linking. Three benefits stem from integrating a verb-class-based linking theory within a hierarchical organization of grammatical knowledge. First, we provide a clear theoretical foundation for the observations included in previous verb-class-based analyses; secondly, we easily model those aspects of subcategorization which do not follow from the semantics of the lexical items involved; finally, we can reduce the complexity of the argument selection process outlined in Dowty (1991).

In section 1, we address problems that arise from the use of the thematic hierarchy, arguing that verb and event classes must supplement it to handle some common linking patterns. Section 2 presents our theory of linking, based solely on verb classes and event types. The empirical advantages of this model are discussed in section 3, which includes a treatment of diathesis alternations, causative predicators,[1] and lexically determined exceptional cases.

1 Problems with thematic hierarchies.

Beginning with Fillmore (1968) and Jackendoff (1972), many linguists have assumed that semantic roles can be arrayed in a hierarchy to which various syntactic processes refer. Bresnan and Kanerva (1989), for example, propose the following hierarchy. Other researchers have proposed variants differing in minor respects.

(1) AGENT > BENEFICIARY > RECIPIENT/EXPERIENCER/GOAL > INSTRUMENT > THEME/PATIENT > LOCATIVE

The original, simple use of hierarchies similar to that in (1) was to account for the relative obliqueness of syntactic arguments on the basis of their ordering on the thematic hierarchy. In this section we discuss what we see as difficulties for a theory of linking based on a total ordering of thematic roles. Our empirical arguments rejoin the telling criticisms of Ladusaw and Dowty (1988), Dowty (1991), Wechsler (1991), Wechsler (1995)

[1]We use the term 'predicator' in this paper to label verbs or other relation-denoting words.

and others who question the validity of thematic roles in the first place. But our focus is different; we assume that the roles in (1) can be meaningfully defined and assigned to a verb's arguments, and show that a linking theory based on the thematic hierarchy *still* requires additional event-type specific or verb-class specific rules and constraints. We argue in the following section that such constraints on verb classes suffice to capture the behavior of lexical items, and therefore that the hierarchy proves superfluous.[2]

1.1 Empirical problems with the thematic hierarchy.

1.1.1 Causative verbs.

In many languages, causatives can be formed of verbs that already have an AGENT, as shown in (2) and (3). The result is a causative verb with two AGENT-like arguments. The thematic hierarchy as traditionally constituted tells us nothing about the ordering between these two, yet cross-linguistically, the pattern is clear: the causer is realized as the subject, and the "original" AGENT is realized as a direct object, as an oblique, or not at all.

(2) Vitsi naura-tt-i nais-i-a (Finnish)
joke laugh-CAUS-PAST woman-PL-PART
'The joke made the women laugh.'

(3) Hu na'-kanta si Pedro (Chamorro)
1sg CAUS-sing Pedro
'I made Pedro sing.' (Gibson and Raposo (1986))

By itself, the thematic hierarchy cannot account for why the causer is always the subject in accusative languages. Two related solutions have been proposed. The first is to assume a cyclic mapping whereby the semantic roles NPs bear to a CAUSE predicate override those they can bear to an embedded caused-event argument (Alsina and Joshi (1991), Alsina (1992)). This is depicted in (4), where a single participant (represented here by the variable y) plays a role both in the CAUSE predicate and in the embedded predicate:

(4) CAUSE $<$ x , y , EFFECT $<$ y , z $>>$

If we assume that 'the women' bears the role of causee/patient with respect to the cause predicate in (2), we account for its object status. Alternatively, we can posit a particular rule stating that causers outrank any other agent, as Grimshaw (1990) does with a parallel aspectual hierarchy, on which causers are the topmost element. Although these moves solve the problem posed by (2) and (3), they appeal implicitly to the notion of semantic class by making reference to the internal structure of the semantic representation in the case of causative verbs and *causative verbs only*. By itself, the thematic hierarchy is not enough to determine how arguments are realized syntactically; linking regularities must sometimes rely on a semantically-defined classification of verbs.

[2]Note that Van Valin (1993) presents a new interpretation of the nature of actor/undergoer selection which meets several of the challenges we level against the use of such hierarchies.

1.1.2 Alternate mappings.

Aside from causativized verbs, there exist several classes of exceptions to the predictions made by the hierarchy in (1) (modulo irrelevant variations in the ordering of roles). We present two examples in (5) and (6). In (5) the possessor (or experiencer) may be realized either as the subject of 'own' or 'miss', or as the prepositional object of the 'to' phrase subcategorized by 'belong' or the object of 'escape'.

(5) a. Mary owns many books. / This book belongs to Mary.

b. We missed the meaning of what he said. / The meaning of what he said escapes/eludes us.

(6) a. Oak trees plague/grace/dot the hillsides.

b. The hillsides boast/sport/feature oak trees.

In (6) are several verbs that lexicalize the mapping of location and theme in different ways. Those in (6a) map theme to subject and location to object, while those in (6b) map theme to object and location to subject. In all these cases, whichever ordering between thematic roles is predicted by the thematic hierarchy, at least one set of examples cannot be accounted for. The thematic hierarchy by itself cannot account for the full range of linking possibilities exhibited by the relevant classes of verbs.[3] More importantly, the need to add more and more thematic roles to address the various classes of examples we discuss in this section undermines the appeal of the thematic hierarchy in the first place: the existence of a universal, restricted set of roles to which linking principles can refer.

In addition to cases like those illustrated in (5) and (6), there are examples of fairly ordinary verbs for which assigning thematic roles proves quite difficult, and pragmatic factors come into play in the syntactic realization of their arguments. Again, we give a few examples, but they can easily be multiplied.

(7) a. His house borders Jamestown.

b. Zeta precedes eta in the Greek alphabet./Eta follows zeta in the Greek alphabet.

c. Brazil outplayed/outscored/outranked Italy.

The sentences in (7) provide examples of verbs where the two arguments are indistinguishable by normally invoked criteria for thematic roles and where the mapping of semantic arguments to syntactic complements is a matter of point of view, pragmatics, or lexical stipulation. The thematic hierarchy makes no predictions about their linking possibilities, In

[3]It is, of course, possible to assume that the thematic role borne by *the hillsides* in (6a) and (6b) differ, as one reviewer suggests. Perhaps *the hillsides* has an attribute-bearer thematic role in (6b). But lest we reduce thematic roles to syntactic diacritics, this difference in thematic role assignment must be grounded semantically or pragmatically, a difficult task for these particular examples, we surmise.

some cases, such as (7a), the mapping cannot be accounted for by the properties of the events and their participants. The two participants of 'border' are identical in terms of thematic roles or lexical entailments; animacy, affectedness or any other such property cannot tell them apart, a point made by Dowty (1991). Moreover, there are cases in which the mapping of these verbs is not merely exceptional or arbitrary. These include the verbs 'precede' and 'follow' as used in (7b), and verbs prefixed with 'out-', such as 'outplay', 'outscore', and 'outrank' in (7c).[4] In all such cases the relation is not symmetric, the two arguments seem semantically or pragmatically distinguishable, and the syntactic realization of the two provides a clue to the specific roles that each one plays. But the semantic/pragmatic difference cannot be characterized in terms of coarse thematic roles of the kind mentioned in (1). In such cases, the thematic hierarchy does not provide the appropriate descriptive tools needed for linking, weakening again the claim that appeal to a thematic hierarchy allows for a restricted linking theory.[5]

1.2 Ontological dependency of the thematic hierarchy.

Aside from the empirical shortcomings of the thematic hierarchy we have just reviewed, there are two reasons that a thematic hierarchy is theoretically questionable. First, why are there no attested predicators with certain combinations of roles, e.g., {BENEFACTIVE, GOAL}? Presumably, the answer centers on the meaninglessness of calling a participant a beneficiary in an event lacking another participant that performs some beneficial action. Likewise, a participant designated a goal is hard to imagine in the absence of another participant that moves toward the goal. Treating thematic roles as primitives that are unrelated to one another (except for their ranking in the thematic hierarchy) makes such patterns purely accidental; the excluded pattern could just as well be {AGENT, THEME}. What is lost is that, as we define the various thematic roles, we are implicitly defining types of events (to the extent that our definitions of thematic roles are semantically grounded).

[4] See Gawron (1983) for similar examples and brief discussion.

[5] The same argument we stress in this section applies to verbs with expletive or raised dependents, at least in monostratal theories. Consider the verbs in (8).

(8) a. John seems happy to me.

b. I believe John to be happy.

Since John in (8) plays no semantic role with respect to 'seem' or 'believe', the thematic hierarchy reveails nothing about its mapping to subject and object position respectively.

The same difficulty arises in the case of expletives. Consider example (9):

(9) It seems to me that Bill was right about that.

Assuming with most scholars that 'it' in (9) is an expletive NP, the question is: how can the thematic hierarchy account for its mapping onto subject position when there is at least one other, semantically non-empty argument (the so-called EXPERIENCER *to me*) situated in the hierarchy? Again, one could posit a new principle stating of this class of verbs that expletives are mapped onto subject position (and this is too simple in any case; see Postal and Pullum (1988)). Even if one could derive such a principle from other considerations (which we doubt), the thematic hierarchy plays no role in the mapping.

The thematic hierarchy is derivative from the notion of event class. Unattested theta-grids are unattested event-types.

Secondly, postulating a thematic hierarchy suggests the possibility of certain classes of rules which are either not attested or dubious cross-linguistically. To our knowledge, for example, no rule targets the lowest role of a verb's arguments, irrespective of its specific semantic role. To be sure, we do get rules that target patients or recipients, but in all cases we know of they target a specific semantic role, not the argument, whatever its semantic role, that occupies an abstract last position in the hierarchy. Thus, no single rule targets the content of experience of verbs of mental representation like 'believe', the patient of simple transitive verbs like 'eat', and the theme of triadic verbs like 'give'. The hierarchy thus provides us with more possibilities than are reliably attested cross-linguistically, suggesting that it does not sufficiently constrain possible theories of linking.[6]

1.3 Non-motivations for the thematic hierarchy.

Although the previous two subsections show that the thematic hierarchy is not needed for linking, one could suggest that it comes for free, since it is independently motivated. Under this view, the thematic hierarchy would not be motivated by linking phenomena, but could be used when convenient, since other data would support its existence. Kiparsky (1987), for instance, claims that the thematic hierarchy predicts the ease with which idioms can be lexicalized. But Bresnan (1982) and Nunberg *et al.* (1994) show that there are several counterexamples to this tendency and provide an alternative semantic explanation that does not appeal to the thematic hierarchy. It has also been noted for a long time (since at least Jackendoff (1972)) that the relationship between anaphors and antecedents is sometimes

[6] One candidate syntactic process which allegedly targets a verb's lowest argument on the hierarchy has been suggested to us by R. Van Valin. Consider the examples in (10):

(10) a. I gave a book to Joe to read.
b. I gave Joe a book to read.
c. I cooked a chicken to eat.

Assuming that Joe, as a recipient or goal, is the lowest argument of 'give' on the hierarchy of arguments susceptible to becoming an UNDERGOER (Van Valin 1993), the identification of the missing object of purpose clauses seems to be a process targeting a verb's lowest role on the hierarchy. But the purpose clause construction is severely restricted and does not uniformly target all lowest arguments on the undergoer hierarchy. Witness the ungrammatical sentences in (11):

(11) a. *The crowd entered the building to occupy.
b. *Sue mowed the grass to shorten.

Only two verb classes participate in the purpose clause construction, verbs of transfer of possession ('buy', 'receive') and verbs of creation ('build', 'cook'). The motivation for this grouping appears to be the following: both classes of verbs denote actions that result in a situation in which one participant has control over an entity (either by transferring ownership/control or by creating a new object or a new form of the same object). The purpose clause specifies what is to be done with this newly controlled object. The construction thus plainly targets particular classes of situations; characterizing it adequately requires reference to event classes, not to positions of arguments on the thematic hierarchy abstracted from these event classes.

governed not only by structural/grammatical function considerations but also by semantic considerations. Jackendoff's original suggestion is that the antecedent of an anaphor must outrank it on the thematic hierarchy. But again, alternative explanations exist for such data, be it the notion of center of perspective proposed in Engdahl (1990) or the notion of conceptual structure superiority of Jackendoff (1992), neither of which requires reference to a thematic hierarchy. Once again, the thematic hierarchy is not necessary to account for the generalizations mentioned in the literature.

1.4 Lexical Mapping Theory and the thematic hierarchy.

So far, we have discussed various difficulties that plague thematic hierarchies in general. In this section, we focus on a particular linking theory which employs thematic roles and hierarchies, LFG's Lexical Mapping Theory (henceforth LMT), presented in, among others, Alsina and Mchombo (1992), Alsina and Mchombo (1993), Bresnan and Kanerva (1989), Bresnan and Kanerva (1992), Bresnan and Moshi (1993). We focus on LMT because we believe it constitutes the best representative of the kind of linking theories against which we argue in this paper and presents some of the strongest arguments in favor of a linking theory based on the thematic hierarchy. Because of space considerations, we concentrate on the most compelling case they make in favor of the use of a thematic hierarchy, namely Bantu applicative and locative inversion constructions.

To understand the argumentation involved in this subsection, we must briefly summarize some relevant aspects and architectural choices of LMT. The basic goal of LMT is to formulate a restricted linking theory of wide cross-linguistic applicability. To achieve this goal, two hypotheses are proposed. First, linking generalizations only appeal to paired-down semantic information, namely thematic roles and the thematic hierarchy; secondly, linking generalizations require cross-classifying grammatical functions into various classes. Thus, LMT defines the grammatical functions subject (SUBJ), object (OBJ), restricted object (OBJ_θ), and oblique (OBL) through a two-way featural decomposition into the features ±r(estricted) and ±o(bjective). Table 1 depicts this classification.

	-r	**-r**
-o	SUBJ	OBL
+o	OBJ	OBJ_θ

Table 1: Classifiying Grammatical Functions

Once grammatical functions are decomposed, linking principles can apply to add information monotonically to partially underspecified lexical entries and determine their subcategorization properties. As described in Bresnan and Kanerva (1989), LMT effects the mapping from thematic roles to grammatical functions via the following ordered steps:

1. A universal intrinsic classification of arguments (henceforth IC) associates certain feature values with the thematic roles of a verb. For example, agents are assigned the IC

-o.

2. Morpholexical operations on a verb may add, suppress, or bind roles in an argument structure. For example, passivization is treated as suppressing the highest thematic role of a verb, while applicativization involves adding a role.

3. Default classifications are assigned if they do *not contradict features assigned by intrinsic or morpholexical classifications. The highest of a verb's roles is assigned -r (i.e., SUBJ or OBJ), then the next highest is assigned +o (OBJ or OBJ_θ), and the next highest after that is assigned +r (OBJ_θ or OBL).

The grammatical functions of each of a verb's roles are (partially) determined in this way, along with conditions that each verb has a subject and that each grammatical function (apart from OBL) is associated with only one thematic role.

Although this approach to linking has been applied to many languages, it provides a particularly elegant model of Bantu applicatives. We now turn to these data to assess how crucial the thematic hierarchy is to the account of the generalizations described in the work we cited at the beginning of this section. Work within LMT has suggested four generalizations about applicatives in Chicheŵa that arguably require reference to the thematic hierarchy:

1. Only applicatives introducing arguments below goal on the thematic hierarchy can be either primary or secondary objects, as demonstrated by the contrast between the beneficiary applicative in (12) and the instrument applicative in (13). (Data from Alsina and Mchombo (1992), examples (3)-(4).)

(12) a. Chitsîru chi-na-gúl-ír-a atsíkána mphâtso
7-fool 7s-PST-buy-AP-FV 2-girls 9-gift

b.*Chitsîru chi-na-gúl-ír-a mphâtso atsíkána
7-fool 7s-PST-buy-AP-FV 9-gift 2-girls
'The fool bought a gift for the girls'

(13) a. Anyani aku-phwányír-a mwala de-ngu.
2-baboons 2S-PR-break-AP-PV 3-stone 5-basket

b. Anyani aku-phwányír-a de-ngu mwala.
2-baboons 2S-PR-break-AP-PV 5-basket 3-stone
"The baboons are breaking the basket with a stone.'

2. Theme/patients must be realized as (primary or secondary) objects if an argument higher on the thematic hierarchy is realized as an object (primary or secondary). Thus, whereas the theme cannot be realized as a subject in (14b), because of the presence of the more highly ranked instrument, it can in (15b), since location is lower than theme on the hierarchy. (Data from Alsina and Mchombo (1992), examples (7)-(8).)

(14) a. Mwala u-ku-phwány-ír-idw-á de-ngu (ndí anyani)
3-stone 3S-PR-break-AP-PASS-FV 5-basket (by 2-baboons)

b.*De-ngu u-ku-phwány-ír-idw-á Mwala (ndí anyani)
5-basket 3S-PR-break-AP-PASS-FV 3-stone (by 2-baboons)
'The stone is being used (by the baboons) to break the basket'

(15) a. Pa-m-chenga pa-ku-lúk-ír-idw-á ú-konde (ndí á-sodzi)
16-3-sand 16S-PR-weave-AP-PASS-FV 14-net (by 2- fisherman)
'The sand is being woven on the net (by fishermen)'

b. Ú-konde pa-ku-lúk-ír-idw-á Pa-m-chenga (ndí á-sodzi)
14-net 16S-PR-weave-AP-PASS-FV 16-3-sand (by 2- fisherman)
'The net is being woven on the sand (by fishermen)'

3. The syntactic argument corresponding to beneficiary or goal applicatives can only be extracted when it is the highest expressed semantic role, as shown by the ungrammaticality of sentence (16b) in which the agent argument outranks the beneficiary argument. (See Baker (1988) and Alsina and Mchombo (1992), see also examples (23)-(24) and (31) in Alsina and Mchombo (1993).)

(16) a. Yi ndi mphátso iméné chítsîru chí- ná-gúl-ír-a atsík na
9-this be 9-gift 9-REL 7-fool 7S-PST-buy-AP-FV 2-girls
'This is the gift that the fool bought for the girls'

b.*Wa ndi atsíkána améné chítsûru chí-ná-gúl-ír-a mphâtso
2-these be 2-girls 2-REL 7-fool 7S-PST-buy-AP-FV 9-gifts
'These are the girls that the fool bought a gift for'

4. Applicatives must introduce a syntactic argument that corresponds to a semantic role lower than at least one other expressed role of the predicate. Examples (17) from Bresnan and Kanerva (their examples (22)) illustrate the generalization. Sentence (17a) is grammatical since applicative formation precedes passives and agents are higher than goals on the hierarchy. But sentences (17b)-(17c) are ungrammatical since the only remaining argument after passivization is the theme argument and goals are higher than themes on the thematic hierarchy.

(17) a. M-tsogoleri a-na-túmíz-ir-idw-á zí-pâtso (ndí ána)
1-leader 1-SU-PST-send-AP-PASS-FV 8-fruit (by 2-child)
'The leader was sent fruit (by the children)'

b.*M-tsogoleri a-na-túmíz-idw-ir-á zí-pâtso (ndí ána)
1-leader 1-SU-PST-send-PASS-AP-FV 8-fruit (by 2-child)

c.*Zí-pâtso a-na-túmíz-idw-ir-á m-tsogoleri (ndí ána)
8-fruit 1-SU-PST-send-PASS-AP-FV 1-leader (by 2-child)

In all four cases, the thematic hierarchy seems necessary to state the relevant generalizations in terms of a natural class. But they can be stated just as economically without it. On grounds of parsimony, then, there is no theoretical need to posit a thematic hierarchy, a conclusion in accord with our claim in the preceding subsections. We account for three of the four generalizations mentioned above with the constraints in (18), which refer to semantic roles (more precisely, semantic classes of verbs with arguments bearing these roles), but not to a thematic hierarchy. The fourth generalization does not in fact hold, as additional data in Alsina and Mchombo (1992) and data we present below reveal.

(18) a. Beneficiary and goal arguments of applicatives must be direct syntactic arguments (where direct denotes subjects and primary objects).

b. Instrument arguments of applicatives of passive verbs must be subjects.

c. Beneficiary and goal arguments can only be extracted if no agent is expressed.

Such statements do not make reference to the thematic hierarchy and do not require further complication or stipulations, assuming, as LMT scholars do, a cross-classification of GF's along two dimensions (see the [-r] column in table 1 above).[7]

As we noted above, the fourth generalization, an apparent startling confirmation of the role of the thematic hierarchy, does not in fact hold. It predicts that unaccusative verbs (those whose highest role is theme or patient) cannot receive a beneficiary applicative argument, since beneficiaries are higher on the thematic hierarchy than either patients or themes. Although Baker (1988) argues this is the case, Alsina and Mchombo (1992) convincingly show this prediction to be erroneous. Witness example (19) (A&M (38)).[8]

(19) a. Yêsu a-ná-wá-f-er-a (anthu)
Jesus 2S-PST-2-O-die-AP-FV 2-people
'Jesus died for them (the people)'

[7]The only potential advantage of using the hierarchy in stating the constraints lies in the fact that beneficiaries and goals constitute a natural class in a thematic hierarchy based account, i.e. a pair of adjacent roles. Note, though, that there are other ways of deriving this grouping of semantic roles. In particular, several scholars have grouped together goals and beneficiaries on the basis of metaphors relating causal/actor-undergoer structure to spatial structure (Croft (1991), Lakoff (1993), Goldberg (1995), Wechsler (1991), Wechsler (1995)). There thus seem to be good semantic reasons for assuming GOAL and BENEFICIARY arguments form a natural class independent of the thematic hierarchy.

[8]Garrett (1992) suggests that beneficiary is too broad a class and should be broken down into beneficiaries proper and substitutives ('in place of'). Only the latter may appear with unacccusatives. Bresnan and Kanerva could capture the facts in (19) by introducing a new thematic role at the appropriate point in the hierarchy. But the addition of such relation-specific roles erodes the original motivation for the thematic hierarchy (how many such additional roles will be needed?), whereas it suits a hierarchical verb-class-based model, where semantically defined verb classes may be fine-grained.

b. Anthu a-ná-f-ér-idw-a (ndí Yêsu)
2-people 2S-PST-die-AP-FV (for Jesus)
'The people were died for (by Jesus)'

The generalization is further undermined by examples like (20) (Sam Mchombo, p.c.):

(20) a.*	Mfumu	i-na-gw-ets-er-idw-er-a	miyala	nthiko
	chief	SM-PST-fall-CAUS-AP-PASS-AP-FV	stones	wooden.stick

b.*	Mfumu	i-na-gw-ets-er-idw-er-a	nthiko	miyala
	chief	SM-PST-fall-CAUS-AP-PASS-AP-FV	wooden.stick	stones

c.	Mfumu	i-na-gw-ets-er-idw-er-a	miyala	pa-tsindwi
	chief	SM-PST-fall-CAUS-AP-PASS-AP-FV	stones	on.the.roof

'The stones were made to fall on the roof for the chief.' (with 'the chief' as subject)

As the last example shows, Chicheŵa, in contrast to many other Bantu languages, allows more than one applicative per verb form (*pace* Baker (1988)). Generalization 4 makes a prediction concerning forms with more than one applicative. If the first applicative introduces a beneficiary, we would expect an instrument introduced by the second applicative to be possible, since beneficiaries outrank instruments on the hierarchy. Example (20) shows, however, that this prediction is not borne out. The correct descriptive generalization seems to be that only locative and reason arguments can be introduced by an applicative morpheme following a passive morpheme. Our second principle (which does not refer to the thematic hierarchy) directly accounts for this generalization: if instruments must be the subject of passive verbs, they cannot be introduced after the passive applies. Otherwise they would be mapped onto an object function. Furthermore, the language-specific nature of our second principle is confirmed by cross-linguistic data. Woolford (1993) presents data from the related Bantu language HiBena and the West Atlantic language Fula revealing the opposite pattern: in instrumental applicatives, the theme/patient invariably becomes the subject of passives, not the instrument.

We have shown that the range of intricate data that Bresnan and Kanerva (1989) and Bresnan and Kanerva (1992) attribute to the effects of the thematic hierarchy can be accounted for without making reference to it. Our discussion demonstrates that the Bantu data do not provide convincing evidence in favor ot the thematic hierarchy. Additionally, LMT faces some more general difficulties, which we will now briefly discuss. We mention these because they undermine the claim that a highly restrictive, monotonic linking theory, which makes minimal appeal to semantic information is feasible.

We see two major difficulties with the structure of LMT. First, it is crucial that the mapping proceeds in several cycles. The LMT thus relies on shadow levels of representation for which no independent linguistic motivation exists. To account for cross-linguistic variation regarding the number of 'object-like' arguments in Bantu languages, for example, Bresnan and Kanerva (1989), propose an Intrinsic Classification Parameter. Languages like Chicheŵa do not allow more than one [-r] IC, whereas no such restriction exists in Kichaga.

But there can be more than one [-r] argument per verb in Chicheŵa: all transitive verbs contain two [-r] arguments, one corresponding to the subject, one corresponding to the object. So, the IC level at which the IC parameter holds in Chicheŵa must be distinguished from the "surface" level at which it does not hold. The introduction of such multistratal mapping is not otherwise motivated, as far as we know.[9]

Secondly, the ordering of the mapping "cycles" seems problematic; in particular, an ordering paradox arises in the case of themes/patients. Themes/patients are intrinsically [-r], so one would predict them to be mapped onto either subjects or primary objects. Such is the case for ordinary active and passive sentences. But in the case of themes/patients occurring with locative arguments introduced by applicatives, they must be allowed to be classified as [+r]. Consider, for example, the following pair (Sam Mchombo, p.c.):

(21) a.	Asodzi	a-ku-pony-er-a	myala	pa-tsindwi
	fishermen	SM-PRES-throw-AP-FV	stones	on.the.roof
b.	Asodzi	a-ku-pony-er-a	pa-tsindwi	miyala
	fishermen	SM-PRES-throw-AP-FV	on.the.roof	stones
	'Fishermen are throwing stones on the roof'			

When the locative is realized as a secondary object, as in (21a), the patient can receive its intrinsic [-r] classification. But, in (21b), where the patient is the secondary object, it bears a [+r] classification. The only way for it to receive this classification is if the IC applies after morpholexical operations, since on the root "cycle" there is only one "object-like" argument, the patient, and nothing prevents its being assigned the IC [-r]. Thus, Bresnan and Kanerva must intend that morpholexical opreations apply prior to IC: once the applicative assigns [-r] to the locative applicative argument in (21b), the IC parameter forbidding more than one [-r] argument in asymmetric languages such as Chicheŵa, would ensure that the patient of (21b) would bear a [+r] classification. In this case, then, it seems that the morpholexical operation must apply before the IC applies, contrary to the order of application posited by the theory.[10]

We conclude that not only is the thematic hierarchy based account of applicatives in Bantu not necessary, it requires a linking theory that introduces unnecessary theoretical

[9] It is tempting to motivate the IC level by assuming that it corresponds to semantically defined argument types. But in fact, this is impossible. The IC can be overriden by specific lexical classes. Locatives, for example, are [-o] intrinsically, but when introduced by an applicative construction, they can receive the default classification [+o]. It is therefore doubtful that the IC level represents true semantic generalizations.

[10] Alsina and Mchombo's proposal avoids this difficulty by proposing that themes/patients can also be classified as [+o]. Two new issues arise from this proposal. First, it predicts that the patient argument of an active form of a $<$ AGENT, PATIENT $>$ verb can be realized either as a direct or secondary object. The introduction of a structural ambiguity without surface reflexes should be viewed with skepticism. Secondly, to account for the fact that instrument applicatives of passive verbs must be subjects, Alsina and Mchombo are forced to introduce a further default: themes/patients *must* be intrinsically classified as [+o] when the verb's argument structure includes a higher ranking non-agent argument. To our knowledge, this "default" is only needed to model the linking behavior of the passive forms of $<$ AGENT, INSTRUMENT, PATIENT $>$ verbs. The restricted application of this rule renders it equivalent to specifying a particular linking pattern for a specific verb or event class, and it thus provides little or no support for a thematic hierarchy.

costs. By contrast, the constraints we proposed in (18) to account for the Bantu facts, do not require any addition to the declarative, word-class-based approach to the mapping problem we present in the next section.[11] No ordering in the application of principles is needed, and consequently no ordering paradox can arise. The upshot of our discussion so far is three-fold:

1. The thematic hierarchy is insufficient to account for many linking regularities; it must be supplemented by many verb-class and event-type specific linking rules.

2. The thematic hierarchy is not constrained enough when it is separated from the event-types or word classes that motivate it.

3. Additional, non-linking related motivations generally advanced in favor of the hierarchy do not in fact require an appeal to it.

These three conclusions suggest that a thematic hierarchy does not furnish an optimal basis of a theory of linking. In the remainder of this paper, we present our alternative approach to linking, based exclusively on the notions of event type and verb class, which, as we have seen, are needed for an adequate theory of linking in any case.

2 An event-type and verb-class-based mapping theory.

2.1 The basics of HPSG.

Building upon data and intuitions of several linguists and parallel research in AI (see, e.g., Quillian (1968), Gawron (1983), Brachman and Schmolze (1985), Flickinger (1987), Wechsler (1991), Wechsler (1995), Koenig (1993), Goldberg (1995), and Davis (1996)), we base our theory of the mapping between semantic roles and syntactic complements on a classification of linguistic objects and the notion of semantically-defined word classes. In fact, our major contention in this paper is that such classifications of linguistic objects, and of the lexicon in particular, provide a suitable theoretical foundation for the notion of verb class that has figured prominently in several works on linking (see Green (1974), Pinker (1989), and Levin (1993) among others). Furthermore, we demonstrate that a hierarchical arrangement of linguistic information allows us to combine the benefits of both verb-class-based and entailment-based approaches to linking (see the collection of papers in Briscoe *et al.* (1993) and references therein for a discussion of work on hierarchical lexicons in general and Pustjovsky (1995) for a discussion of its role in lexical semantics aside from linking). A hierarchical structure to linguistic information provides a good model of the cross-cutting semantic similarities among sets of verbs that has been shown to underlie linking regularities. It also provides the appropriate formal structure to state what we claim is the appropriate

[11] For reasons of space, a detailed model of Bantu applicatives within the linking theory we propose in the next section must await another occasion. Our point in this section is merely to show that the Bantu facts do not provide firm evidence in favor of a thematic hierarchy.

correlation between lexical semantic categories and subcategorization properties of verbs. Finally, we show that such a hierarchy allows for a simpler entailment-based linking theory, compared to that of Dowty (1991). We will use Head-Driven Phrase Structure Grammar (henceforth HPSG) as the grammatical framework in which to present this model. (An introduction to HPSG can be found in Pollard and Sag (1987) and Pollard and Sag (1994).)[12]

The architecture of HPSG provides for a flexible encoding of word classes. All grammatical information in HPSG is encoded via a set of attribute-value pairs such as [NUM *sing*] or [PERS *1st*] (such sets are called Attribute Value Matrices, hereafter AVM), where a value may itself consist of a set of attribute-value pairs (for similar approaches, see Bresnan and Kaplan (1982) and Gazdar *et al.* (1985), among others). Figure 1, for instance, represents the semantic contribution of a pronoun such as 'she', namely that any use of 'she' introduces a discourse referent with a third, singular, feminine index.

$$\begin{bmatrix} \textit{nom-obj} & \\ \text{INDEX} & \begin{bmatrix} \textit{ref} & \\ \text{PERS} & \mathit{3} \\ \text{NUM} & \mathit{sg} \\ \text{GEND} & \mathit{fem} \end{bmatrix} \\ \text{RESTR} & \{\} \end{bmatrix}$$

Figure 1: The semantic contribution of 'she'

The italicized term at the top left of an AVM designates the class or type of linguistic object described by that AVM. It is fundamental to the HPSG approach that linguistic information is organized into a hierarchy of such classes of linguistic objects, which are referred to as types. Types include familiar classes such as words and phrases, part-of-speech (e.g., noun and verb) classes, valence-based classes (e.g. intransitive or transitive verbs), and traditional morphological classes (e.g. declension or conjugation classes). Capturing linguistic generalizations in such an approach consists in grouping the relevant linguistic objects under a single type. The properties characteristic of this class of objects (typically encoded through attribute value pairs) are associated with more general types and automatically passed down to specific instances of those types. In defining an individual lexical entry, any information that is inferred from the types to which it belongs need not be specified again for that entry. Moreover, linguistic objects typically belong to several cross-cutting classes at once; words can, for example, be members of the class of adjectives or verbs as well as members of the class of intransitive or transitive words. The hierarchy of word class types therefore serves to minimize redundancy in the lexicon. This property of linguistic classifications is referred to as multiple inheritance (see Carpenter (1992) among others) and is informally diagrammed in figure 2 (the AVM structures associated with each type is omitted for ease of presentation).

[12] Although the particular framework we use is HPSG, our theory can equally be couched within any unification-based framework using something like a type or inheritance hierarchy, such as Construction Grammar, see Fillmore and Kay (forthcoming).

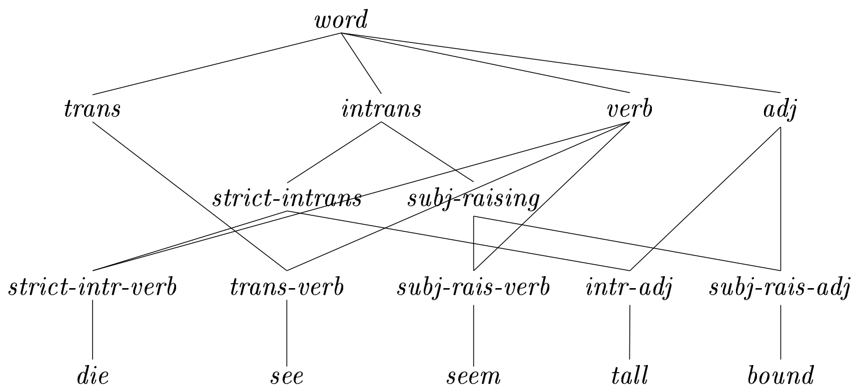

Figure 2: A simple hierarchy of words

To illustrate, the information associated with 'see' as used in example (22) derives from its membership in several classes: *verb* for part-of-speech information, *transitive* for subcategorization information, and so forth.

(22) Bill saw Harry.

Each verb class may be viewed as embodying generalizations about sets of lexical entries that share a certain amount of information, whether relating to subcategorization, semantics, morphology, or phonology. Thanks to multiple inheritance, a relatively small amount of information needs to be specified on the entry of a particular word such as 'see'; the great majority of a lexical item's properties are inherited from broader classes (types) in the hierarchical lexicon.

One important technical aspect of the type hierarchy concerns which attributes may appear within a feature structure of a given type. Each attribute is licensed by a single type in the hierarchy, which means that it may appear only in a structure of that type (thus avoiding any possible multiple inheritance conflict of the kind discussed by the various papers in Briscoe *et al.* (1993)). For instance, having an index is a property of the semantic content of nominals. We say that the type *nom-obj* licenses the INDEX attribute. Similarly, the type *index* licenses the PERS, NUM, GEND attributes. Crucially, all subtypes of a type t inherit all the attributes (properties) t licenses. Hence all subtypes of indices inherit the attributes PERS, NUM, GEND. Intuitively, being marked for person is a property of referential and expletive indices because it is a property of all indices (see Pollard and Sag (1994) for details).

The main contribution of this paper is to use such multidimensional classifications of linguistic objects to provide an empirically more satisfactory and more theoretically grounded account of linking. In particular, we claim that linking principles are constraints specified on certain lexical types. Each linking constraint provides a partial specification of the relationship between the semantics and the subcategorization of a class of words. Individual lexical entries must conform to all the constraints that hold of their supertypes. It is this multiple inheritance architecture of lexical knowledge, amply motivated elsewhere (see Flickinger (1987) and Pollard and Sag (1994), and references therein), not the thematic hierarchy, which provides an adequate foundation for linking theory. Our point is thus similar to the one made by Pinker (1989) and Levin (1993). But, we believe the hierarchical nature of our lexical representations provides an additional theoretical grounding as well as deriving some of the salient properties of linking regularities from a general hypothesis as to the architecture of lexical knowledge.

As in other lexically-oriented approaches to syntax, the phrase structure of sentences in HPSG derives from the subcategorization of lexical entries. The manner in which lexical information is projected into constituent structure is exemplified in figure 3. This figure covers sentences such as (22).[13]

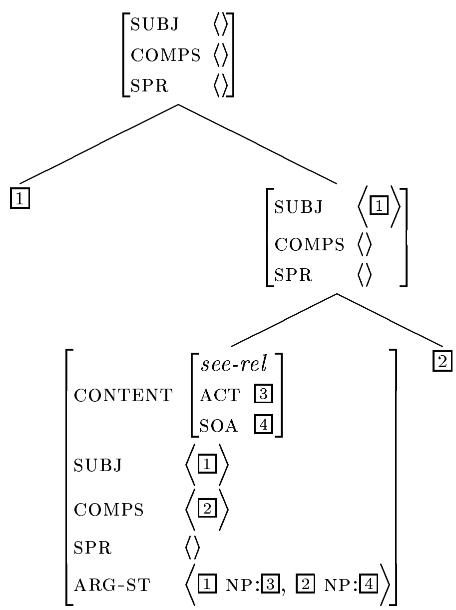

Figure 3: A simple example of the effects of the Valence Principle

Each lexical head specifies, in its ARG-ST (argument structure) attribute, a list of syntactic dependents (in figure 3, the head's ARG-ST contains two NPs). Each member of this ARG-ST list consists of the syntactic and semantic information of the head's syntactic dependents. Furthermore, the semantics of each of these dependents is appropriately identified

[13] HPSG-savvy readers will notice that the attributes shown in figure 3 are actually not attributes of *word* or *phrase*, but of the local portion of the *synsem* (syntax and semantics information) within each word or phrase. We omit these details for clarity. Note also that technically, there is no context-free backbone in HPSG. Internal constituency is represented via attributes on a par with other attributes, be they syntactic, semantic, or phonological. Since this detail is irrelevant to our discussion, we will use the better-known tree-structure representation for constituency.

with the value of a semantic argument of the head. The tags $\boxed{3}$-$\boxed{4}$ following colons in the figure represent this shared semantic information. (Roughly speaking, tags are variables, with identically numbered tags indicating identical pieces of structure; here identity between the semantics of the first and second NP on the ARG-ST list and the values of the ACT and UND attributes respectively.) Whereas the ARG-ST list comprises information about the possible dependents of the head, information about the actual subject, specifier, and complements for which a head subcategorizes is recorded on the SUBJ, SPR, and COMPS valence lists. In the typical case, these three lists taken together constitute the head's syntactic dependents, i.e. the ARG-ST list. This identity is indicated in figure 3 by the tags $\boxed{1}$ and $\boxed{2}$: the subject subcategorization requirement is identical to the first dependent on the ARG-ST list of the head verb while the first complement subcategorization requirement is identical to the second dependent on the ARG-ST list.

The projection of phrases from the valence information contained in lexical entries is accomplished by identifying the syntactic and semantic information of members of the valence lists of heads with the syntactic and semantic information of non-head daughters. The tag $\boxed{2}$ in figure 3, for example, identifies the syntax and semantics of the direct object with that of the first element of the COMPS list. Once realized, a valence requirement is cancelled from the relevant valence requirements of the projected phrase. As indicated in the figure, the COMPS list of the mother of the lexical head is that of the head-daughter node minus the subcategorized elements realized as complements. Hence the object requirement marked by tag $\boxed{2}$ is removed from the COMPS list: the VP node does not subcategorize for a direct object complement. This process is general and regulated by a general principle of HPSG, the valence principle stated below (see Pollard and Sag (1994) for more details):

(23) *Valence Principle* (adapted from Pollard and Sag (1994))
In a headed phrase, the list of valence requirements of the head-daughter (i.e. the value of the SUBJ, SPR, COMPS attributes) is the concatenation of the valence requirements of the mother and the list of SYNSEM values of the daughters in order of obliqueness.

2.2 A hierarchy of semantic relations.

Now that we have briefly introduced the relevant mechanics of HPSG, what is the place of a linking theory within such a framework? A theory of linking specifies the mapping between semantic entities (what corresponds to traditional semantic arguments in the typical case) and elements on the ARG-ST list. More precisely, it regulates the relationship between the values of a head's semantic arguments and the semantics of its syntactic dependents (i.e., of elements on the ARG-ST list). The theory we present in this paper uses two kinds of constraints on lexical information to that effect. First, each *linking constraint* partially specifies the mapping between semantic arguments and syntactic dependents for a given word class, as mentioned above. These linking constraints are inherited by all members of that class. Second, the structure of the word class hierarchy is subject to a principle ensuring that certain word classes are subtypes of others. The semantic relations within lexical entries denote types of situations, characterized by properties that one or more participants in such situations are entailed to bear. Accordingly, these semantic relations can themselves be arranged in a multiple-inheritance hierarchy, with more specific types of situations as subtypes of more general ones. The principle requires the word class hierarchy to mirror the semantic relations hierarchy. From this similarity in structure, it follows that words with a given semantics will be subject to certain linking constraints.

Since we contend that linking classes are defined using semantic types, we must first briefly describe the semantic relations we assume and the inheritance hierarchy they are structured in. Figure 4 illustrates a semantic relation type and the representation of participants in situations denoted by this type. In addition to this cross-classification by types, the semantic content of verbs indicates the semantic attributes of participant roles appropriate for each type.[14]

$$\begin{bmatrix} \textit{act-und-rel} & \\ \textsc{actor} & \textit{content} \\ \textsc{undergoer} & \textit{nom-obj} \end{bmatrix}$$

Figure 4: The *act-und-rel* class of semantic relations

In figure 4 the type *act-und-rel* is a relation involving an actor and an undergoer role, represented as attributes within feature structures of this type (see below for a discussion of these two roles). This relation covers the semantic content of verbs such as 'hit', 'build', and transitive 'break'; each of these verbs denotes a kind of situation that is a subtype of the kind of situation denoted by *act-und-rel*. Intuitively, all hitting, building, and agentive breaking events are instances of the more general class of events which contain both kinds of participants. As mentioned above, attributes are inherited by subtypes of the supertype that licenses them. So, the semantic content of 'hit', for instance, which is a subtype of *act-und-rel*, inherits both ACTOR and UNDERGOER attributes.

We can extend this classification of semantic relations in both directions. The relation *act-und-rel*, for instance, is in turn a subtype of two more general semantic relations, *act-rel* and *und-rel*, each of which licenses a single semantic attribute, ACTOR and UNDERGOER respectively. The relation *act-und-rel*, as a subtype of these two relations, inherits both, as shown in figure 5. Conversely, we can introduce subtypes (possibly licensing additional attributes) to represent the additional information carried by more specific *act-und-rel* semantic relations. The semantic relation denoted by transitive 'break', for example, is a particular kind of *act-und-rel*, a subtype of *cause-change-of-state-rel*. As a causal change-of-state relation, the denotatum of transitive 'break' contains an additional semantic role, the resulting state of the broken entity, which is denoted by the value of the attribute SOA. We represent this type of event with the type *cause-change-of-state-rel*, which inherits this

[14] We borrow the terms ACTOR and UNDERGOER from Foley and Van Valin (1984). Our usage of these terms, though inspired by their work, is different. See table 2 for our definition of these attributes. The types of the values of ACTOR and UNDERGOER, namely *content* and *nom-obj*, constitute the most general types the instantiations of these arguments can take.

additional attribute from the supertype that licenses it, *soa-rel*, as shown in figure 6(a).[15]

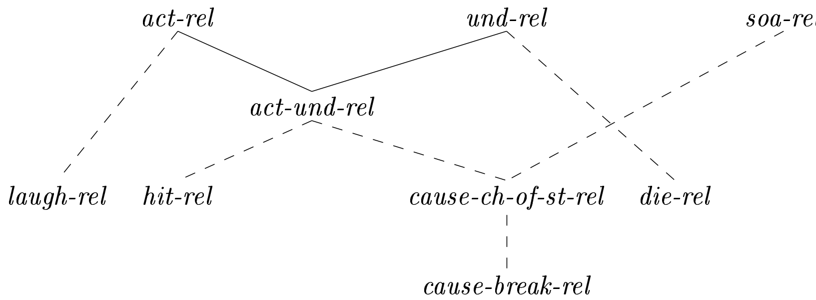

Figure 5: Classification of semantic relations using multiple inheritance

By factoring semantic attributes into different relation types we can straightforwardly represent cross-cutting semantic similarities. *Laugh-rel* denotes a subtype of *act-rel* that is not a subtype of *und-rel*. *Die-rel* denotes a subtype of *und-rel* that is not a subtype of *act-rel*. *Hit-rel* denotes a subtype of both *act-rel* and *und-rel* which is not a subtype of *soa-rel*. Finally, *cause-break-rel* denotes a subtype of both *act-und-rel* and *soa-rel*. It thus contains all three ACTOR, UNDERGOER, and SOA attributes, as shown in figure 6(b).

$$\begin{bmatrix} \textit{soa-rel} & \\ \text{SOA} & \textit{rel} \end{bmatrix}$$

(a) The *soa-rel*

$$\begin{bmatrix} \textit{cause-ch-of-st-rel} & \\ \text{ACTOR} & \textit{content} \\ \text{UNDERGOER} & \textit{nom-obj}\ \boxed{1} \\ \text{SOA} & \begin{bmatrix} \textit{rel} & \\ \text{ARG} & \boxed{1} \end{bmatrix} \end{bmatrix}$$

(b) The *cause-ch-of-st-rel*

Figure 6: The *soa-rel* and *cause-ch-of-st-rel* relations

The resulting semantic content of 'break' is represented in figure 7. Note that in figure 7, the value of the UNDERGOER attribute bears a tag identical to the value of the P-BEARER attribute.[16] The identity between the values of the two attributes represents the fact that a single participant fulfills two roles: the entity affected by the causal-event and the entity that is in a broken state subsequent to the event.

[15]SOA in the figure stands for 'state of affairs'; ARG stands for any semantic attribute. We use dashed lines to indicate that there are intermediate types, not represented in the diagram, between two types.

[16]We use P-BEARER as a cover term for whatever set of properties characterizes the semantic argument introduced by one-place stative predicates, such as 'red', 'round', 'broken', 'sparse', 'depressed'... We leave a more precise semantic characterization of the properties whose presence is entailed by one-place stative predicates to another occasion. The token-identity between the values of UNDERGOER and P-BEARER holds of a large class of semantic relations, namely those that denote causal change-of-state relations.

$$\begin{bmatrix} \textit{cause-break-rel} & \\ \text{ACTOR} & \textit{content} \\ \text{UNDERGOER} & \textit{nom-obj}\ \boxed{1} \\ \text{SOA} & \begin{bmatrix} \textit{broken-rel} & \\ \text{P-BEARER} & \boxed{1} \end{bmatrix} \end{bmatrix}$$

Figure 7: The lexical semantic representation of transitive 'break'

2.3 Proto-role entailments and proto-role attributes.

The hierarchy of semantic relations allows us to model semantic relatedness among words; the intuitive semantic similarities between the denotata of transitive 'break' and, say, 'freeze', is mirrored in the hierarchy by the existence of a fairly specific common supertype of their semantic contents, *cause-ch-of-st-rel*. We now need to provide some semantic grounding for postulating the particular semantic relation types and the semantic attributes within them. Each attribute appearing in a semantic relation should be there because of some entailment holding of the participant its value denotes. Conversely, event types with similar participant roles should be represented by semantic relations with similar sets of attributes. The participants in breaking and freezing events have entailments in common and therefore the semantic content of each verb is similar.

More technically, we associate each semantic role attribute with a set of *characteristic entailments*, at least one of which holds of any participant denoted by the value of that attribute. We further hypothesize that, for each of the characteristic entailments associated with a semantic role attribute, there is always a most general semantic relation requiring that entailment to hold of the participant denoted by its value; *prop-attitude-rel* (the semantic relation type common to verbs such as *believe*, *see*, ...), for example, specifies the ACTOR attribute on the basis of an entailment that holds of believers and seers: they have a notion of the referent of their object NP. Similarly, the UNDERGOER attribute in *cause-ch-of-st-rel* is present on the basis of an entailment that holds of frozen and broken entities: they undergo a change of state. Table 2 gives the attribute names, the semantic relation that licenses each attribute, and the (disjunctive) set of characteristic entailments defining each attribute. Inspired by Dowty (1991) and Wechsler (1995), it constitutes our current best hypothesis regarding the set of participant properties that affect the linking of semantic arguments onto syntactic functions.

We formally capture the model-theoretic relation between semantic attributes and their characteristic entailments by introducing the following constraint on semantic attributes:

(24) ATTRIBUTE-TO-ENTAILMENT CONDITION
If a semantic role attribute (ACTOR, UNDERGOER, and so forth) is present in a semantic relation r included in the lexical semantic structure of a predicator, then its value denotes a participant in the situation denoted by r that is entailed to bear one of the attribute's characteristic entailments (as they are listed in table 2).

Relation	*Licenses semantic role attribute(s)*	*Characteristic entailments*
act-rel	ACTOR	Causally affects or influences other participant(s) or event(s); Volitionally involved in event; Has a notion or perception of other participant(s) in event; Possesses an entity.
und-rel	UNDERGOER	Causally affected or influenced by another participant; Undergoes a change of state; Is an incremental theme; Possessed by entity.
soa-rel	SOA	Resulting state of affairs; Perceived or conceived of by another participant; A circumstance aspectually or temporally delimited by the relation
fig-grnd-rel	FIG	Entity located with respect to another participant; Moves with respect to another participant; Contains or constitutes another participant;
	GRND	Entity with respect to which another entity is located; Trajectory along which another participant moves; Is contained by or part of another participant.
property-rel	P-BEARER	Entity bearing property.

Table 2: Semantic roles and characteristic semantic entailments

It is important to notice that characteristic entailments are model-theoretic constraints which, by their very nature, are outside the metalanguage we use to describe the semantics encoded in lexical entries. They constitute constraints on appropriate models for the interpretation of the feature structures representing lexically encoded semantic information. As such, their satisfaction cannot be checked by looking at the metalanguage we use in our descriptions. It can only be checked by considering whether the model interpreting the metalanguage structures "fit" our intuitions regarding properties of event participants. As a reviewer correctly points out, this leaves some room for vagueness to creep in when assessing whether the referent of a given attribute's value bears the relevant entailment. But, ultimately, this vagueness attaches to any model-theoretic interpretation. Imagine defining a set of possible models for the interpretation of English. This class of models would consist (at least) of a set of relations, a set of individuals, and a set of constraints between relations which can be described by formulas of the form (using Dowty's (1989) ordered-argument style of representation for the description of the model):

(25) $\Box\ [\mathit{believe}(x,y) \Rightarrow \mathit{has\text{-}a\text{-}notion}(x,y)]$

The formula in (25) says that if a relation of belief holds between x and y, then x has a notion of y. The class of models compatible with our hypothesis regarding characteristic entailments can be described by such formulas. "Checking" whether a given hierarchy of semantic relations obeys the constraints described in table 2 would then be a matter of "translating" the feature structure representation of the lexical semantics of predicators into the relevant antecedents and insuring that no inconsistency arises. But, ultimately, whether the restrictions on appropriate models for the English (or any language's) lexicon embedded in table 2 is explanatory depends on whether we find entailments such as (25) "intuitively" true or not, and some vagueness is intrinsically present in these kinds of assessments.

A second point worth noting is that the Attribute-to-Entailment condition does not establish a one-to-one correlation between attributes and characteristic entailments. It merely says that if a certain attribute is present in a semantic relation, at least one of its characteristic entailments holds of the role played by the participant denoted by the attribute's value. The reverse implication does not hold; the condition does not require any attribute to be present in a semantic relation.[17] That the constraint runs only in one direction is consonant with the view expressed in Pinker (1989) or Van Valin and Lapolla (1997) that lexical semantics is not reduceable to language-independent conceptual structure. Finally, note that although it is difficult to provide a unifying characterization for each set of entailments, the ACTOR entailments relate to initiating an event and affecting other participants, while the UNDERGOER entailments typify "affected" participants. The entailments characteristic of the ACTOR attribute might then reduce to a general entailment roughly paraphrasable as: "has control over the unfolding of the situation" and similar reductions might be possible for other semantic attributes mentioned in the table. But the linking theory we present below does not require this. The semantic underpinnings of linking are the entailments themselves, borne by the participants (the denotata of the attributes' values). The names of the attributes are simply convenient reifications of these classes of entailments for the purposes of linking. Table 2, therefore, embodies two hypotheses and two hypotheses only.

1. Characteristic entailments constitute sets of properties of event participants on which linking constraints depend (see Davis (1996) for a more detailed discussion).

2. Semantic role attributes are needed to simplify the interface between semantic arguments and syntactic dependents.

We do not make the further claim that each semantic attribute is necessarily associated with a semantically coherent set of characteristic entailments.

The Attribute-to-Entailment Condition borrows the idea of defining semantic roles in terms of entailments shared across arguments of various predicates from Ladusaw and Dowty (1988) and Dowty (1991). In fact, the characteristic entailments associated with ACTOR and UNDERGOER closely resemble the proto-agent and proto-patient entailments of Dowty (1991). In contrast to Dowty (1991), however, semantic roles are defined absolutely, not

[17]However, the position of a semantic relation in the semantic relations hierarchy does require certain attributes to be present — namely, those that are present in its supertypes. For example, because *cause-break-rel* is a subtype of *ch-of-st-rel* (see figure 5), it is guaranteed to have the attributes ACT and UND.

through a numerical comparison of proto-agent and proto-patient entailments. Moreover, we reify proto-roles; that is, attributes such as ACTOR or UNDERGOER consitute a mediating level between semantic entailments and syntactic arguments. On both theoretical and empirical grounds, we reject Dowty's approach of directly comparing numbers of proto-agent and proto-patient entailments (as in (26) below) in favor of these mediating semantic attributes referenced in the attribute-to-entailment condition. We will now briefly discuss the motivations for this stance.

Metatheoretically, implementing a numerical comparison approach within declarative, monotonic approaches to grammars does not seem feasible (see Davis (1996) for details).[18] Empirically, we perceive two problems with Dowty's model. First, though the entailments he employs are semantically well grounded, his argument selection principle in (26) governs argument selection only "in predicates with grammatical subject and object".

(26) ARGUMENT SELECTION PRINCIPLE:
In predicates with grammatical subject and object, the argument for which the predicate entails the greatest number of Proto-Agent properties will be lexicalized as the subject of the predicate; the argument having the greatest number of Proto-Patient entailments will be lexicalized as the direct object.

Syntactically intransitive verbs are specifically excluded from the purview of this principle, as Dowty makes clear in a footnote (p. 576). Consequently, an important generalization cutting across transitivity is missed — namely, proto-agent roles are linked to subjects regardless of transitivity. To capture this generalization, we believe it is necessary to reify something like the notions of proto-agent and proto-patient, as we have done with the attributes in table 2, rather than treating them as epiphenomenal labels of convenience. Syntactically intransitive verbs denoting events and states with two participants, such as 'rely on', 'yearn for', 'suffice for', 'apply to', 'do without', 'deal with', and many more, illustrate the problem. If the argument selection principle does not apply to such verbs, we should expect for at least some verbs that the argument bearing more proto-patient properties can be realized as the subject and the argument bearing more proto-agent properties can be realized as an oblique, as shown in (27), on a par with the alternation in (28). (For the intended meanings of the examples in (27), exchange the two arguments in each sentence.)

(27) a.*Sleight of hand relies on/by/of/with the magician.

b.*A Porsche yearns for/by/of/with the president.

c.*The trip will suffice for/by/of/with 200 dollars.

[18] Monotonic, declarative approaches treat grammars as the order-independent application of various constraints. They present several advantages, as stressed by Bresnan and Kaplan (1982), and Sag (1991), among others.

The ungrammaticality of the examples in (27) suggests that transitivity per se is not the determining factor in subject selection. For the verbs in (27), like transitive verbs, the argument bearing the greater number of proto-agent properties is realized as the subject. We could broaden the scope of the argument selection principle in (26) so that it covers these cases, but some important modifications in Dowty's view of proto-roles will follow. If we try to extend the principle by saying that the argument bearing the greater number of proto-agent properties is realized as the subject (regardless of transitivity), then passive verbs and examples like (28b) (which are numerous) will violate it.

(28) a. Water filled the tank.

b. The tank filled with water.

An alternative is to partially decouple proto-role assignment from argument selection, in effect reifying the proto-roles. In other words, by introducing proto-role attributes into lexical representations as a mediating level between semantic entailments and syntactic arguments, we can avoid Dowty's difficulty in tying syntactic argument assignment directly to numbers of entailments. This is the option we have adopted in this paper, using the semantic attributes in Table 2.

Secondly, the numerical comparison procedure for argument selection that Dowty suggests encounters difficulties with some classes of verbs. In some cases, comparing the raw numbers of proto-agent and proto-patient entailments holding of two arguments is insufficient to capture a generalization about their syntactic realization. For example, Dowty acknowledges that "movement is apparently an agent property only when not caused by another participant in the event." That is, an entailment of causation between two arguments simply overrides considerations about motion in determining proto-role assignment. In (29) are Dowty's examples illustrating this point. Each argument of the verbs in (29) bears one of Dowty's proto-agent and one of his proto-patient properties. The wall causally affects the bullet; however, the latter is in relative motion while the former is stationary.[19] In case of a tie, one might expect lexicalization to be variable (as Dowty states), but with verbs of caused motion, it does not seem to be. There are no corresponding verbs which reverse the mapping of semantic arguments to syntactic ones, although verbs of motion that denote a non-causal relationship between a moving entity and a stationary one do realize the argument denoting the moving entity as their subjects, as Dowty says.

(29) a. John threw the ball.

b. The wall deflected the bullet. (cf. 'The bullet grazed the wall.')

Morphological causativization of verbs in many languages furnishes other verbs with two or three arguments that violate Dowty's argument selection principle. In many languages,

[19]Not literally stationary, of course. This is a somewhat vague entailment, having to do more with the frame of reference one adopts when talking of one object deflecting another than with the status of the objects in the world.

lexical causativization of intransitive verbs is more widespread than in English, and we find transitive constructions like the following, from Finnish (Arto Anttila, p.c.).

(30) Uutinen puhu-tt-i nais-i-a pitkään.
news-item talk-CAUS-PST woman-PL-PART long-ILL
'The news made the women talk for a long time.'

These verbs show a consistent pattern of mapping the causer to the subject, regardless of the number of other proto-agent properties holding of the causee, or of other arguments.[20]

The subject in this example, 'uutinen', bears fewer proto-agent entailments than the object 'naisia'. The latter is entailed to be sentient and volitional, but crucially the former causally affects it. As is the case with the English verbs of caused motion discussed previously, the causal structure of the predicator's semantics appears to take precedence over other entailments for purposes of linking. In Dowty's model, we would expect considerable variety in how causers and causees are lexicalized, depending on what other entailments hold of the causee. In reality, however, we observe that the causer is consistently realized as the subject.[21]

We conclude that semantic attributes in the semantic content of verbs are required to mediate the association between semantic entailments of argument positions and syntactic dependents, and that no numerical comparison between proto-entailments is necessary. A single entailment suffices to determine the semantic attribute whose value denotes a particular participant in the situation denoted by a verb (through the Attribute-to-Entailment Condition). And, as we demonstrate in the next section, the semantic attribute whose value denotes that participant suffices to determine its mapping to the corresponding syntactic dependent.

2.4 Linking constraints

Having briefly discussed the semantic classifications we assume in this paper, we now turn to linking proper. Linking, we claim, consists of a classification of words constraining the relations that can exist between (the type of) semantic content of verbs and (the type of) their subcategorization structure. Consider the transitive verb 'devour'. Its actor is mapped to the subject of its active form, while the undergoer is mapped to the object position.

[20] Singh (1992) criticizes Dowty's argument selection principles on similar grounds and offers an analysis of causation much like the one we propose, making use of Dowty's proto-role entailments. However, her replacement for Dowty's argument selection principle is a more complicated version of the numerical comparison procedure in (26), comprising several ordered subrules. Her argument selection procedures encounter some difficulties as well. For example, if no argument bears only proto-agent entailments, then "the argument with the greatest number of proto-agent entailments is realized as the subject." The example in (30), among others, violates this condition.

[21] More precisely, the causer is linked to the first member of the ARG-ST list, which, in syntactic absolutive languages, is realized as a subject, and in syntactically ergative languages is realized as an object. See Manning and Sag (1998) for a discussion of the need to distinguish between ARG-ST and valence lists (SUBJ, SPR, and COMPS lists) to model the distinction between absolutive and ergative languages and Croft (1991) and Van Valin and Lapolla (1997) for extensive cross-linguistic data supporting this point.

Crucially, the linking contraints that map from semantic attributes to the subject and object reference both the semantic attributes and the list of syntactic dependents or ARG-ST list. Not all verbs in a semantic class have the same subcategorization possibilities, however. 'Eat', for example, differs from 'devour' in that its undergoer need not be mapped to any syntactic argument. 'Eat' and 'devour' thus only share some of their linking properties. Our proposal is that a hierarchical organization of lexical knowledge constitutes the optimal architecture in which to specify these partial linking similarities among verbs, say the class of verbs with a semantics that is a subtype of *act-rel*. By itself, though, this claim is too weak. If different word classes could vary more or less arbitrarily in the mappings they specify betweeen proto-role attributes and subcategorization requirements, then linking would be free. In fact, the linking properties of a verb are highly constrained by the semantic relation it denotes; verbs denoting *act-rel* events map their actor onto subject position in the active form. To reflect such constraints, we restrict the kinds of word classes that exist in the lexicon via a universal condition on the relation between (linking-related) word classes and semantic relations. We have seen that both word classes and the semantic relations that form a part of their semantic content can be hierarchically arranged. The model we present in this section is based on a correspondance between these two hierarchies, which we call the Semantic Subtype Linking Condition.

(31) THE SEMANTIC SUBTYPE LINKING CONDITION (informal version)
If the semantics of one (linking-related) verb class is a subtype of the semantics of the second, then the first verb class is a subtype of the second verb class.

In other words, the hierarchy of semantic relations is mirrored in the hierarchy of (linking-related) word classes. The Semantic Subtype Linking Condition is thus a metagrammatical condition. While the Attribute-to-Entailment Condition constrains the lexical semantic representations we can posit for individual verb classes, the Semantic Subtype Linking Condition constrains the subtype/supertype relations that can exist between verb classes. Their combined effect is to limit the possible linking patterns for a verb class to those that are compatible with the linking patterns of all its superclasses.

In more formal terms, the relationship between the two hierarchies of semantic relation types and word class types can be expressed in a principle that imposes a homomorphism under subtypes from semantic relations to word classes.[22] This principle has the effect of ensuring that any lexical item's location in the lexical hierarchy is such that it inherits all linking constraints applicable to it, given its semantics.

(32) THE SEMANTIC SUBTYPE LINKING CONDITION (formal version)
If s is a type in the semantic relations hierarchy and there exists a type in the word class hierarchy with NUCLEUS value of type s, then there exists a type s-p in the word

[22]Technically, the homomorphism holds between the hierarchy of types of the value of the NUCLEUS attribute to subtypes of *word*. The NUCLEUS attribute in HPSG encodes the semantics of verbs excluding their quantificational structure. See Pollard and Sag (op.cit.) chapter 9 for more details. Since we are not concerned with quantification here, we disregard the CONTENT information outside of the NUCLEUS.

class hierarchy with NUCLEUS value of type *s* such that every type in the word class hierarchy with NUCLEUS a subtype of *s* is a subtype of *s-p*.

The effects of this homomorphism constraint are represented informally in the parallelism of the hierarchies of figure 8.

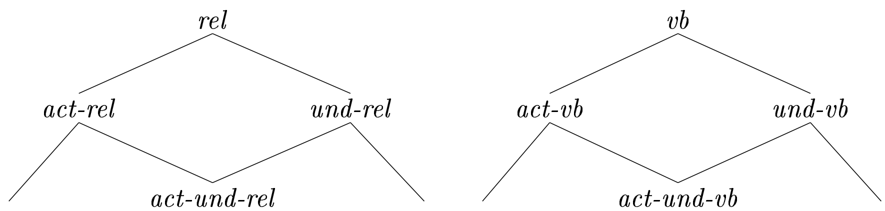

(a) Semantic relations   (b) Syntactic hierarchy

Figure 8: Homomorphism between semantic relation types and linking types

This condition expresses our main hypothesis regarding the semantic basis of linking, one similar to many previous proposals (see among others Goldberg (1995), Jackendoff (1990), Levin (1993), Pinker (1989), Van Valin and Lapolla (1997): if in a language some verbs with a semantic content containing an actor link actors in a particular way, any verb that denotes such an event *must* link its actor this way. This therefore rules out (33b), given the grammaticality of (33a), and assuming that the hypothetical verb *hypdrew* encodes a semantic relation which is a subtype of the relation licensing the linking of *drew* in (33a). As we will discuss explicitly in section 3.3, this condition is weaker than that assumed by scholars such as Pinker (1989) and allows for the appropriate amount of syntactic idiosyncrasies to affect linking. But, together with the Semantic Subtype Linking Condition which we introduce in the next subsection, it is strong enough to avoid the need for a thematic hierarchy, in contrast to what is proposed in Jackendoff (1990).

(33) a. John drew a picture.

b.*A picture hypdrew John.

The definitions of the types *act-vb* and *und-vb* mentioned in the hierarchy on the right are presented in figure 9, The *act-vb* class says that an actor argument is realized as the first NP on the ARG-ST list. The *und-vb* class says that an undergoer argument is realized as the last NP on the ARG-ST list (as indicated informally in the figure by the optional (XP, ...), where XP cannot be NP).[23]

[23] We do not address in this paper the issues of the categorial realization of semantic argument types. See Pesetzky's (1982) Canonical Realization Principles or Langacker (1987) for interesting suggestions in this regard.

$$\begin{bmatrix} \textit{act-vb} & \\ \text{CONTENT} & \begin{bmatrix} \textit{act-rel} & \\ \text{ACTOR} & \boxed{1} \end{bmatrix} \\ \text{ARG-ST} & \left\langle \text{NP:}\boxed{1}, \ldots \right\rangle \end{bmatrix}$$

(a) The type *act-vb*

$$\begin{bmatrix} \textit{und-vb} & \\ \text{CONTENT} & \begin{bmatrix} \textit{und-rel} & \\ \text{UND} & \boxed{1} \end{bmatrix} \\ \text{ARG-ST} & \left\langle \ldots, \text{NP:}\boxed{1}\ (,\text{XP} \ldots) \right\rangle \end{bmatrix}$$

(b) The type *und-vb*

Figure 9: The *act-vb* and *und-vb* linking classes

These two constraints ensure that actors are realized as subjects and undergoers as objects of active verbs in accusative languages. A few examples will illustrate this. Relevant verbs are in italics in the sentences in (34).

(34) a. Frank *ate* the soup.

b. The soup was *eaten*.

c. Helen *seems* to have forgotten her keys.

d. Helen *imagined* him to be more forgiving.

Let's consider 'ate' first, and more generally, the broad class of active forms of verbs denoting an *act-rel*. Members of this class are constrained to map the instigator of the event to the subject. Within the theory we propose, this mapping results from the independent application of two constraints (i.e. *act-rel* denoting verbs inherit from two different lexical classes). One class is the aforementioned *act-vb* class; the other is the *active-vb* class represented in figure 10(a), which requires the first element on the ARG-ST list to be identical with the lone element of the SUBJ list (see the tag $\boxed{1}$) and the remainder of the ARG-ST list to be identical with the COMPS list (see the tag $\boxed{2}$).

$$\begin{bmatrix} \textit{active-vb} & \\ \text{CAT} & \begin{bmatrix} \text{ARG-ST} & \left\langle \boxed{1} \right\rangle \oplus \boxed{2} \\ \text{VAL} & \begin{bmatrix} \text{SUBJ} & \left\langle \boxed{1} \right\rangle \\ \text{COMPS} & \boxed{2} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

(a) The active verb class

$$\begin{bmatrix} \textit{passive-vb} & \\ \text{CAT} & \begin{bmatrix} \text{VAL} & \begin{bmatrix} \text{SUBJ} & \left\langle \boxed{2} \right\rangle \\ \text{COMPS} & \boxed{3} \end{bmatrix} \\ \text{ARG-ST} & \left\langle \boxed{1} \right\rangle \oplus \boxed{2} \oplus \boxed{3} \end{bmatrix} \end{bmatrix}$$

(b) The passive verb class

Figure 10: The *active-vb* and *passive-vb* verb classes

Why *ate* in (34a) inherits these constraints is illustrated in figure 11.

The verb *ate* denotes an event that necessarily includes both an actor and an undergoer. The entailments holding of eaters and eaten things are associated only with these semantic

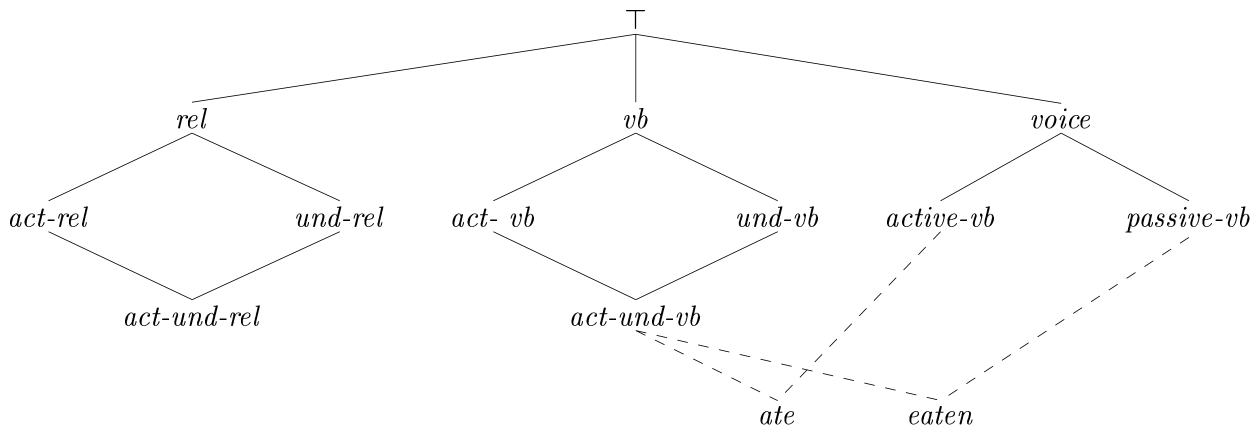

Figure 11: An illustration of multiple inheritance of lexical constraints

role attributes, so these are the only appropriate attributes for *eat-rel*, as per the Attribute-to-Entailment Condition. Hence, the lexical semantic relation of *ate* is a subtype of *act-und-rel*. Because of the Semantic Subtype Linking Condition, this means *ate* must obey the *act-vb* constraint (technically, it inherits the information encoded in the *act-vb* verb class), which holds of all verbs with semantics of supertype *act-rel*: the actor is mapped onto the first element of the ARG-ST list. Similarly, *ate* must obey, like all verbs with semantics of the sort *act-und–rel* the *und-vb* constraint, which holds of all verbs with semantics of supertype *und-rel*: the undergoer is mapped onto the last NP element of the ARG-ST list. Finally, because *ate* is morphologically a subtype of *active-vb* (see the diagram in figure 11), it inherits the general constraint on active verbs that the first member of their ARG-ST list be mapped onto the subject function and the rest of its ARG-ST list be mapped onto complements. The *act-vb* and *und-vb* constraints apply to all verbs that have an actor or undergoer argument; this includes not just transitive verbs, but also all intransitive actor verbs (such as *laugh*) and intransitive undergoer verbs (such as *die*). As mentioned earlier, the Semantic Subtype Condition not only permits identical linking of semantic arguments onto syntactic functions for all verbs which share the relevant semantic type (*act-rel* and *und-rel* in our example), it precludes alternative linking patterns for these verbs. There cannot be verbs that map the UNDERGOER to subject and the ACTOR to direct object. Such a mapping would require the existence of verbs whose lexical semantics is of type *act-und-rel*, but which are not members of the *act-und-vb* class, in violation of the Semantic Subtype Linking Condition.[24]

The passive counterpart to (34a) in (34b) illustrates the dissociation between logical subject and subject requirement typical of "non-canonical" verbs. The *passive-vb* class is

[24] Technically, this result follows from axioms of typed feature structure logic, namely that type hierarchies form a meet semi-lattice in which a unique greatest lower bound exist for any two types t_1 and t_2 (see Carpenter (1992) and Davis (1996) for details).

diagrammed in figure 10(b). As figure 10(b) shows, passive verbs differ from actives in how members of ARG-ST list align with the SUBJ and COMPS valence lists (see Manning and Sag (1998) and for other examples of misalignments between argument and valence structures). However, they do not differ in the linking between CONTENT and the ARG-ST list. The first element of the ARG-ST list is not included in any valence list of a passive verb. It is the second member of the ARG-ST list that is identical to the subject valence requirement, while the rest of the ARG-ST list comprises the COMPS list. Consequently, the dependent corresponding to the UNDERGOER is now assigned to the SUBJ list, hence it will be realized as the subject of the verb 'eaten', as shown in the entry for 'eaten' in figure 12.

$$\begin{bmatrix} \textit{passive-vb} & \\ \text{CAT} & \begin{bmatrix} \text{VAL} & \begin{bmatrix} \text{SUBJ} & \langle \boxed{2} \rangle \\ \text{COMPS} & \boxed{3} \end{bmatrix} \\ \text{ARG-ST} & \langle \boxed{1} \rangle \oplus \langle \boxed{2}\text{NP:}\boxed{4} \rangle \oplus \boxed{3} \end{bmatrix} \\ \text{CONTENT} & \begin{bmatrix} \textit{eat-rel} & \\ \text{ACT} & \boxed{1} \\ \text{UND} & \boxed{4} \end{bmatrix} \end{bmatrix}$$

Figure 12: The passive form 'eaten'

A few words justifying the inclusion of an ARG-ST element linked to the ACTOR in the definition of the passive verb class are in order. First, many phenomena, from binding in several languages (see Joshi (1989) for Marathi *aapaṇ*, for example), to control of purpose clauses and other adjunct clauses (see Ruwet (1972) for French *en* adjunct phrases) seem to target the logical subject of a verb, whatever the voice of the verb form is. By including this element we provide a unified representation of logical subject; it is the first element on the ARG-ST list. Secondly, psycholinguistic studies have recently shown that implicit as well as explicit agents (a subset of logical subjects) play an active role in sentence processing (see Mauner *et al.* (1995), Mauner *et al.* (1999)). Again, the constant presence of logical subjects as the first element of the ARG-ST list is a perspicuous representation of that fact.[25]

The linking properties of raising verbs such as 'seem' and 'imagine' in (34c) and (34d) follow from the set of principles we have seen together with the lexical analysis of raising common to HPSG and LFG. The argument structure of raising verbs like 'seem' and 'imagine' contains one element that is (token-)identical to the unexpressed subject of its VP complement (see Bresnan (1982), Pollard and Sag, op.cit.). This additional raised NP syntactic dependent does not correspond to a semantic argument of the raising verb or adjective and does not therefore affect the linking of the ACTOR argument to the first member of the ARG-ST list (see the constraints in figure 9).

[25]Readers familiar with HPSG have probably noticed that the presence of the logical subject on the ARG-ST list requires some changes in the definition of binding-theoretic notions such as O-command. We cannot pursue this matter here for lack of space. Suffice it to say that we assume a distinction between discourse-active and discourse-inert arguments suggested in Koenig (in press), and propose to model it by modifying the hierarchy of nominal objects.

We can already see in these examples two differences between the verb-class-based approach to mapping we advocate and a thematic hierarchy approach. First, the realization of syntactic dependents is achieved directly in the verb-class-based approach, by declarative constraints on verb classes, rather than through the intermediary of a thematic hierarchy. Secondly, our approach assumes that the mapping of semantic arguments onto syntactic dependents relies on the kind of event denoted by the verb. Rules that refer to a particular semantic relation type, whether very general or highly specific, are therefore easily stated and have the same status as other constraints on word classes. The causal verb rules mentioned in section 1.1 thus pose no embarrasment to the theory. The constraints of the *active-vb* class apply to active verbs whose semantic content is a subtype of *act-rel* (see figure 5). They must consequently hold of all active verbs denoting causal relations, since causal relations are a subtype of *act-rel.* No *ad hoc* mechanism is needed to ensure that the causer, rather than some other agent-like argument, is the first member of the ARG-ST list and therefore realized as the subject. The universal linking of causal verbs falls out from the Semantic Subtype Linking Condition and the architecture of the lexical hierarchy.

Our approach to linking also handles the well-known contrast between two classes of experiencer verbs in English, one with experiencer subjects ('fear', 'like', ...) and the other with experiencer objects ('frighten', 'please', etc.). As pointed out by Kiparsky (1987), Grimshaw (1990), and Croft (1991), only the second class is causal in its semantics. For 'frighten' and 'please', the stimulus, playing a causal role in the mental state, is realized as the subject. The semantics of these two types of verbs are contrasted in figure 13.

$$\begin{bmatrix} \textit{cause-notion-rel} & \\ \text{ACT} & \boxed{2} \\ \text{UND} & \boxed{1} \\ \text{SOA} & \begin{bmatrix} \textit{notion-rel} & \\ \text{ACT} & \boxed{1} \\ \text{SOA} & \boxed{2} \end{bmatrix} \end{bmatrix}$$

(a) Causal

$$\begin{bmatrix} \textit{notion-rel} & \\ \text{ACT} & \boxed{1} \\ \text{SOA} & \boxed{2} \end{bmatrix}$$

(b) Non-causal

Figure 13: Two kinds of psychological predicates

For the causative class, linking proceeds as for other causative verbs, with the stimulus realized as the subject and the experiencer as the direct object. For the 'fear'/'like' class the object of emotion is noncausal, and the top-level ACTOR corresponds to the participant who holds a mental representation, by virtue of the definition of semantic attributes summarized in table 2. The two mappings for verbs denoting emotional states or events do not therefore demand additional mechanisms or stipulated syntactic differences (as in Belletti and Rizzi (1988) or Legendre (1989)). They fall out from independently motivated differences in semantic structure and general linking classes.[26]

[26]As is well-known, there exists a third class of verbs denoting psychological states, the so- called 'piacere' class (see Belleti and Rizzi (1988)). We follow the lead of Filip (1996) here who argues that the dative expe-

The existence of pragmatically-based subject selection, exemplified in (6)-(7), also illustrates the benefits of the present theory of linking. Because argument selection may be stated in terms of different possible conceptualizations for the same basic scene, pragmatically-based linking regularities are not excluded from the purview of the theory. We can, for example, posit the verb class in figure 14, which maps lexically specified topics to the first NP on the ARG-ST list. This correctly models verbs like 'border', 'precede', and 'follow' for which no situational entailment relevant to linking exists.

$$\begin{bmatrix} \textit{topic-vb} & \\ \text{CAT} & \begin{bmatrix}\text{ARG-ST} \left\langle \text{NP:}\boxed{1} \right\rangle\end{bmatrix} \\ \text{CONTEXT} & \begin{bmatrix}\text{C-INDICES} \begin{bmatrix}\text{TOPIC } \boxed{1}\end{bmatrix}\end{bmatrix}\end{bmatrix}$$

Figure 14: The *topic-vb* linking class

As indicated in the figure, this verb class makes pragmatic information (the value of TOPIC) relevant to the selection of the subject. Verbs in this class will simply inherit this constraint.

Reviewing the proposals we have made in this section, we began with the hierarchical lexicon of HPSG. We have postulated alongside this word-class hierarchy a hierarchy of semantic relations, whose types serve as the values of NUCLEUS in predicators. The attributes within semantic relations replace the predicator-specific roles frequently used in HPSG (e.g. EATER or EATEN) with the more general semantic role attributes grounded in characteristic entailments. Our Attribute-to-Entailment Condition, licensing the presence of semantic role attributes in semantic relations, is new, as is the Semantic Subtype linking Condition, establishing a homomorphism between the semantic relations hierarchy and the word-class hierarchy. These conditions are the sole new mechanisms we have introduced. The linking constraints themselves have the same status as other constraints on types in HPSG.

In closing this section, we pose, but do not definitively answer, the question: how universal are the hierarchies of semantic relation types and linking types? We do not claim that they are cross-linguistically uniform— some languages lack ditransitives, for example— but they are plainly not free to vary without limit either (see Gerdts (1992) for some discussion of the cross-linguistic range of linking patterns). Whatever the answer to this question, the commonalities that exist can be viewed either as universals about the structure of the hierarchical lexicon and the presence of certain constraints within it, or as a consequence of general cognitive principles not requiring any explicit statement within grammatical theory (e.g., active participants are more salient to humans). We leave further discussion of this issue for another paper, and turn now to more specifics of our model.

---

riencers of the corresponding verbs in Czech are similar to beneficiaries rather than affected arguments. In our terminology, these verbs do not denote an actor-undergoer relation, but a source of affect and beneficiary relation. Their different linking properties then follow from their different semantics.

3 Empirical advantages of integrating linking and the hierarchical lexicon

Now that we have outlined how linking can be modeled with constraints on verb classes, we turn to other advantages of this view of linking, beyond providing linking constraints for specific event-types. The benefits of our hierarchical, word-class-based theory of linking fall in two groups. To a large extent, our linking theory shares the first group with other verb-class-based linking theories. The second group of benefits accrues exclusively to our proposal.

1. Because verb classes are often defined by the fine-grained, cross-classified semantic relations their members denote, many linking regularities necessarily outside the purview of a thematic-hierarchy-based theory can now be modeled.

2. Because linking regularities form an integral part of a hierarchical organization of lexical knowledge, subregularities, positive exceptions, and partial syntactic idiosyncrasies of an individual verb's linking properties can be accomodated alongside broader generalizations.

3.1 The fine-grained semantics of diatheses.

The kinds of diatheses (subcategorization alternations) that a verb displays are dependent on several fine-grained semantic criteria, a point forcefully made by Pinker (1989), Levin (1993), and Goldberg (1995), among others. Because linking, according to the view we have presented, relies on typed, fine-grained semantic representations, we can account for differences in the diatheses in which verbs can participate. Our proposal is no different from other verb-class-based accounts, in that respect, but for its integration of a model-theoretic grounding of the relevant semantics. (35) and (36) illustrate.

(35) a. Pat threw/tossed/kicked/rolled the groceries to me.

b. Pat threw/tossed/kicked/rolled me the groceries.

(36) a. Pat pulled/dragged/lifted/hauled the groceries to me.

b.*Pat pulled/dragged/lifted/hauled me the groceries.

The difference in grammaticality between (35b) and (36b) is attributed by Pinker, Levin, and Goldberg to a difference in event-type between the verbs of (35) and (36); they denote ballistic and entrained motion events, respectively. Proponents of a thematic hierarchy-based linking theory might counter that transitive and ditransitive 'throw' differ in their theta-grids. That of transitive throw contains AGENT, THEME and LOCATION (or GOAL), whereas its ditransitive counterpart has AGENT, RECIPIENT, and THEME on it. This often assumed difference in thematic roles would yield a different mapping of roles to syntactic dependents, thus accounting for the difference between (35a) and (35b). The difficulty for this counterproposal lies in preventing such an alternation from applying to the verbs in (36a), resulting in the ungrammatical (36b). Under a thematic hierarchy-based linking theory, the thematic roles assigned to the arguments of a verb like 'throw' in (35a) are presumably identical to those assigned to a verb like 'pull' in (36a)— namely, AGENT, THEME and LOCATION (or GOAL).[27] We would therefore expect 'pull' to have a ditransitive variant like 'throw'. We conclude that Pinker (1989), Levin (1993), and Goldberg (1995) are indeed correct, reference must be made to the fine-grained semantic classes in figure 15; the coarse-grained classes defined by theta-grids will not suffice.

How does an event-type based linking theory account for the dative alternation? First, the difference between so-called ballistic and entrained motion is represented by two different sorts in the semantic hierarchy, a subpart of which is represented in figure 15.

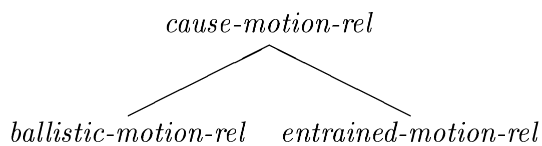

Figure 15: Two subtypes of the relation *cause-motion-rel*

To these two different sorts, different lexical entailments are attached which capture the semantic differences between the two events, as shown in (37).[28]

(37)	**ballistic-motion**	**entrained-motion**
	Force applied to moving object only at initial point of its trajectory.	(Force applied to moving object throughout its trajectory, or not at all)
	(Causer is not entailed to traverse the trajectory with the moving object)	Causer is entailed to traverse the trajectory with the moving object.

In other words, the semantics of a verb is of the sort *ballistic-motion-rel* only if the verb denotes an event for which the relevant entailments hold, and similarly for entrained-motion. Only those verbs with semantics of the first type may participate in the English dative alternation illustrated above.

[27] It is always possible to postulate some difference in thematic roles in such cases. The diathesis difference could then be attributed to the thematic role difference. The problem is that, as long as thematic roles are treated as unanalyzable symbols, such a move has the deleterious effect of obliterating any commonality between the corresponding arguments of, say, 'throw' and 'push'. If, on the other hand thematic roles such as ENTRAINED-THEME and BALLISTIC-THEME were introduced as subspecies of THEME, then thematic roles become essentially a way of covertly coding semantically-defined verb classes.

[28] The semantic distinctions in (37) could be modeled, at least in part, by adding information to *entrained-motion-rel.* The causer in *entrained-motion-rel*, as well as being the ACTOR, is also a FIGURE by virtue of its motion, and the GRND it (or some part of it) traverses is identical to that traversed by the UNDERGOER (which is also a FIGURE). This additional information is absent from *ballistic-motion-rel.*

Secondly, we assume, as Green (1974), Pinker (1989), and Goldberg (1995) have argued, that there is a semantic difference between the two alternants of a verb like ‘throw’, which the theta-grids implicitly do. In particular, the ditransitive alternant of verbs that display the dative alternation has the semantics of causing to possess, while the monotransitive alternant with ‘to’ has the semantics of caused motion.

$$\begin{bmatrix} \textit{str-trans-pp-vb} & \\ \text{CAT} & \begin{bmatrix}\text{ARG-ST} \left\langle \text{NP:}\boxed{1}, \ \text{NP:}\boxed{2} \ (\text{,PP:}\boxed{3}) \right\rangle\end{bmatrix} \\ \text{CONTENT} & \begin{bmatrix} \textit{ball-mot-rel} & \\ \text{ACT} & \boxed{1} \\ \text{UND} & \boxed{2} \\ \text{SOA} & \begin{bmatrix}\textit{mot-rel} & \\ \text{FIG} & \boxed{2} \\ \text{GRND} & \textit{path}\boxed{3}\end{bmatrix}\end{bmatrix}\end{bmatrix}$$

(a) The strict transitive verb class

$$\begin{bmatrix} \textit{ditrans-vb} & \\ \text{CAT} & \begin{bmatrix}\text{ARG-ST} \left\langle \text{NP:}\boxed{1}, \ \text{NP:}\boxed{2}, \ \text{NP:}\boxed{3} \right\rangle\end{bmatrix} \\ \text{CONTENT} & \begin{bmatrix} \textit{cause-und-rel} & \\ \text{ACT} & \boxed{1} \\ \text{UND} & \boxed{3} \\ \text{SOA} & \begin{bmatrix}\textit{poss-rel} & \\ \text{ACT} & \boxed{2} \\ \text{UND} & \boxed{3}\end{bmatrix}\end{bmatrix}\end{bmatrix}$$

(b) The ditransitive verb class

Figure 16: The strict transitive and ditransitive English verb classes for verbs of caused ballistic motion

For the prepositional alternants, we introduce the verb class on the left of figure 16. This is no different from other transitive verbs, except that the ARG-ST list optionally contains a PP complement linked to the GRND, a matter we will not pursue further in this paper (see Wechsler (1995) and Davis (1996) for discussion of PP complements along the lines we have in mind). For the ditransitive alternants, the constraints on linking are given in the AVM on the right in figure 16. In this AVM, the donor is denoted by the value of ACT in *cause-und-rel*, the theme by the value of UND, and the recipient by the value of ACT in the embedded *poss-rel*. The top-level ACT is the first element on the ARG-ST list (by the *actor-vb* class) and therefore the subject. The UND is linked to the last NP on the ARG-ST list (by the *und-vb* class). Finally, the recipient is linked to the second NP on the ARG-ST list; this is the only constraint that the type *ditrans-vb* needs to stipulate. To account for the range of verbs that participate in the alternation between the two verbs classes described in figure 16, we postulate, following Pinker (1989), that English speakers learn a rule permitting the same word to have either the semantics of *ball-mot-rel* (induced ballistic motion) or of *cause-poss-rel* (induced possession).[29] Given this rule and its targeted semantic classes, our fine-grained semantic relations hierarchy and the homomorphism condition requiring the hierarchy of (linking) word classes to parallel the hierarchy of semantic relations ensure that only verbs denoting these types of actions will exhibit this alternation. The two subcategorization alternants of these verbs arise because of the linking constraints operating on the corresponding two semantic alternants.

[29] Alternatively, speakers learn different *lexeme* classes, as in the architecture proposed in Koenig (1999a). The difference between a lexical rule approach or a Typed-Underspecified Hierarchial Lexicon approach is orthogonal to our point, and we adopt a lexical rule approach here for ease of exposition.

3.2 The mapping of lexical and morphological causatives.

We now turn to a benefit that arises from the hierarchical organization of lexical information upon which the Semantic Subtype Linking Condition depends. Our linking theory directly accounts for the often noted parallel between the linking of lexical and morphological causatives. As Kemmer and Verhagen (1994) note, morphological causatives are typically linked according to (at least) the following two rules:[30]

- Causatives of intransitives are linked like monomorphemic transitive verbs.
- In at least one of their diatheses, causatives of transitives are linked like monomorphemic three-place predicates, such as 'give'.

Examples (38) and (39) illustrate this second cross-linguistic tendency with data from Chamorro and Chicheŵa respectively (similar data for Dutch, French, and, several other languages are presented in Comrie (1981) and Kemmer and Verhagen (1994)):

(38) a. In nä'i si tata-n-mami nu i bäbui
1PL-EXCL give PN father-N-our OBL the pig
'We gave the pig to our father'

b. Ha na'-taitai häm i ma'estru ni esti na lebblu
3S CAUS-read 1PL-EXCL the teacher OBL this LNKER book
'The teacher made us read this book' (examples from Gibson (1980))

(39) a. A-na-wa-patsa ana
SM-PST-OM-give-FV children
'They gave them to the children' (Mchombo, p.c.)

b. Nungu i-na-phik-its-a kadzidi maungu
9-porcupine 9s-PS-cook-CAUS-FV 1a-owl 6-pumpkins
'The porcupine made the owl cook the pumpkins' (example from Alsina (1992))

Because mapping regularities are treated as constraints on correspondences between the semantic content and subcategorization properties of verb classes, we predict that verbs with identical semantic contents will display the same range of valence requirements, irrespective of whether they are morphologically derived or not. Now, the semantic content of causative verbs with causees, as we saw above, is an *act-und-rel*. In the case of intransitives, then, there is only one possible argument of the embedded state of affairs that can serve as this causee. This argument is thus realized the same way other undergoers are in the language. In the case of causatives of transitive verbs, on the other hand, the embedded state of affairs

[30] We will not consider here cases like Finnish, where the causee is encoded via an instrumental, since in such structures, the instrumental NP is typically an adjunct.

contains two arguments which can be identified with the causee, the embedded ACTOR and UNDERGOER respectively. Note that the same situation holds for verbs such as 'give', whose semantic content is also causal, as shown in figure 16. Whatever linking properties the counterpart to English 'give' exhibits in some language, our Semantic Subtype Linking Condition predicts that the morphological derived causatives must also exhibit it. If the language identifies the causee of its counterpart to English 'give' with the undergoer of the embedded state of affairs, we have a Chamorro-like language. If the language identifies the causee with the actor of the embedded state of affairs, we have a Chichewâ-like language. The Semantic Subtype linking Condition correctly predicts the parallelism between lexical and morphological causatives in each case.

3.3 Factorization of lexical properties.

The previous two benefits cash in the advantages of a grounded, fine-grained semantics for linking-related verb classes. But not all of a verb's subcategorization properties can be predicted from its semantic properties, *pace* what is sometimes suggested (Pinker (1989)). This is why we stated the Semantic Subtype Linking Condition in terms of a homomorphism rather than an isomorphism between the semantic relations and linking-related verb class hierarchies. What makes hierarchically-organized verb classes particularly useful in this respect is that they allow us to capture the semantic basis of linking without denying the existence of idiosyncrasies of subcategorization. To account for such cases, we rely on multiple inheritance from cross-cutting classes. Consider, for example, a control verb, such as 'force' or 'require' in (40).

(40) MP's forced/required the minister to resign.

The mapping of MP's to subject is determined by the membership of 'force' and 'require' in the class of causal verbs. Furthermore, as Farkas (1988) and Sag and Pollard (1991) have argued, the choice of controller in (40) is determined by the semantics as well— the influenced participant must be the controller of the complement VP. However, there is no requirement that the semantic roles be realized in precisely the fashion of 'force' and 'require'. As we see in (41a), the controller need not be a direct object, and in fact, there need not be a controlled VP at all, as shown in (41b).

(41) a. MP's prevailed on the minister to resign.

b. MP's insisted/required that the minister resign.

The verb 'prevail (on)' thus belongs to a different syntactically defined verb class than 'force', although they inherit the same semantic constraints; the subject is the causer in both cases and the controller of the VP complement of 'prevail (on)' is the influenced participant. Similarly, 'insist' and one alternant of 'require' do not inherit from the class of control verbs, though they do inherit some of the same semantics as 'force' and 'prevail on'. Through multiple inheritance of verb-class and event-type information, depicted in figure 17, we can efficiently capture the semantic similarities and sucategorization differences among these verbs.

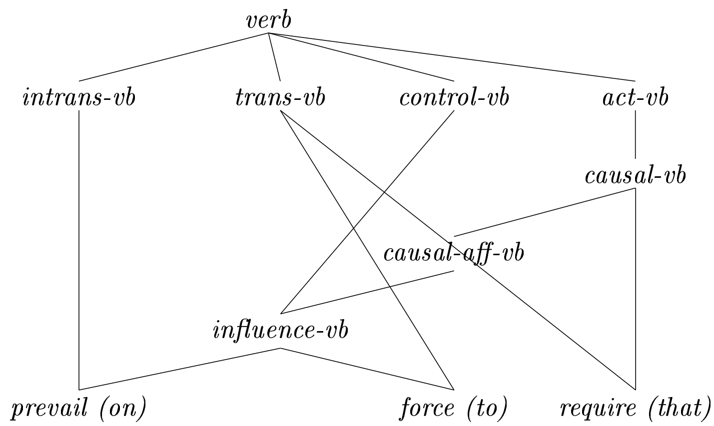

Figure 17: A multidimensional hierarchy of control verbs

3.4 An "optimal encoding" of lexical knowledge.

A related advantage of our theory is that, as in any theory relying on a multiple inheritance approach to lexical information, redundancies are removed from lexical specifications. Information shared by a set of entries is grouped in a single place in the classification. For example, the semantic information shared by all entries involving a causal event is represented only once in the semantic classification, at the node describing causal-events. All entries whose semantics are causal will simply inherit this information by virtue of their membership in this class.

Verbs denoting causal event types provide a good example of this benefit. Two kinds of causal relations, two-place and three-place, must be distinguished (see Comrie (1981), Alsina (1992), among others, for supporting evidence regarding this claim). Two-place causal relations involve the notion of causal force and effected event, whereas three-place causal relations involve an added causee argument corresponding to the participant in the event affected by the causal force. Examples of the two kinds of causal relations are given in (44)-(43) (and also, perhaps, in the controlled and non-controlled versions of 'require', in (40) and in (41b)):

(42) Dieu a fait pleuvoir
God have.PR make.PPT rain.INF
'God made (it) rain.'

(43) Marc prevented anything from happening.

(44) I made him fall.

Sentence (44) illustrates instances where a causal relation involves a causee on which the causal force impinges, whereas (42) or (43) denote a causal relation where no such causee is present. We can represent the similarities between the two conceptualizations of causation as shown in figure 18 (names of attribute values, such as *causer* or *causee* are used for mnemonic purposes only).

$$\begin{bmatrix} \textit{cause-rel} & \\ \textsc{act} & \textit{causer} \\ \textsc{soa} & \textit{effect} \end{bmatrix} \qquad \qquad \begin{bmatrix} \textit{act-und-rel} \\ \textsc{act} \\ \textsc{und} \end{bmatrix}$$

$$\begin{bmatrix} \textit{cause-und-rel} & \\ \textsc{act} & \textit{causer} \\ \textsc{und} & \textit{causee} \\ \textsc{soa} & \textit{effect} \end{bmatrix}$$

Figure 18: Two varieties of causal relations

Cause-rel represents the semantics common to both kinds of causation. The bottom type in figure 18 represents the notion of causation with causee where the participant attribute UNDERGOER is added (by virtue of its being a subtype of *und-rel*). The common entailments between the two kinds of causation are thus attached to a single node in the inheritance hierarchy, the general *cause-rel* type. Note further that the fact that the class *causal-und-vb* is a subtype of the class of *causal-vb* follows again from the universal Semantic Subtype Linking Condition, since members of the former denote a subset of the relations denoted by members of the second class. The architecture of our multiple inheritance network thus allows for an optimal statement of verb class and situation type regularities. Linking regularities shared by a set of verbs are only stated once. All verbs that denote more specific situation types will necessarily inherit those linking regularities.

3.5 Positive exceptions.

Finally, positive exceptions provide another example of the ability of our linking theory to capture the semantic basis of linking regularities without sacrificing the need to represent lexical idiosyncrasies. One example will be enough here. The pattern illustrated in (45) and discussed in Ruwet (1982) and Koenig (1993), Koenig (1994) applies productively to about 60 verbs of saying and mental representation. Although the pattern applies productively to all verbs denoting the relevant classes of events, it *must* apply to one sense of 'prêter' ('say' or lit. 'lend'), as shown in (45a). In other words, although all properties of 'prêter' qua Dative Predication verb are regular, 'prêter' is irregular in being an obligatory member of the class of Dative Predication verbs. In the verb-class-based account of linking presented here, we need only say that it is a member of the class of Dative Predication verbs. All of the relevant properties are then inherited from t$he general Dative Predication class. Because the linking theory we propose is verb-class based, the existence of positive exceptions, a hallmark of linking patterns, is thus easily accounted for.

(45) a. On croit/*prête qu' il a de l' enthousiasme
people believe.PRES/say.PRES that he have.PRES of the enthusiam

b. On lui croit/prête de l' enthousiasme
people to.him believe.PRES/say.PRES INDEF the enthusiasm
People say he is enthusiastic.

4 Conclusion.

We have pursued three goals in this paper, one reductional, one foundational, and one empirical. First, we argued for a theoretical reduction, the elimination of thematic hierarchies, which occupy center stage in many linking theories. We have shown that theories of linking based on the thematic hierarchy must supplement the hierarchy with various rules referencing specific situation types and/or verb classes. As we argued in the second part of this paper, we can make do simply with these additions and dispense with the hierarchy altogether. Secondly, we provided a theoretical foundation for those theories of linking that make use of verb classes. Our linking constraints are semantically grounded by means of the Attribute-to-Entailment Condition and sets of characteristic entailments. These mechanisms differ somewhat from previous verb-class and entailment-based theories of linking, but share with them a concern to exploit properties of event and situation types directly, rather than through a layer of thematic roles. Thirdly we illustrated the empirical advantages of integrating a verb-class-based linking theory within a larger, well-motivated architecture for lexical knowledge, the hierarchical lexicon. A hierarchical organization of linguistic knowledge (which includes both lexico-semantic representations and word classes) allows us to state what we believe to be the proper correlation between the semantic and subcategorization properties of verbs, the Semantic Subtype Linking Condition. Not only does this principle account for the absence of many doublets with similar semantics and reversed syntactic arguments, but it also predicts the parallelism between lexical and morphological causatives without breaching the Lexical Integrity Principle (see Bresnan and Mchombo (1995)). Furthermore, because the condition merely requires an homomorphism between the semantic relations and word class hierarchies, by contrast to stronger principles like the Universal Theta-Alignment Hypothesis of Baker (1988), enough flexibility remains to account for those aspects of subcategorization that do not follow from semantics. We can thus naturally model some well-known characteristics of linking and of lexical knowledge in general: fined-grained semantic distinctions relevant to linking, partial regularities and exceptions, and cross-cutting similarities between semantic and syntactic verb classes.

References

Alsina, Alex. 1992. On the argument structure of causatives. Linguistic Inquiry 23.517–556.

——, and Smita Joshi. 1991. Parameters of causative constructions. Papers from the 27th regional meeting of the chicago linguistics society. Chicago: University of Chicago.

——, and Sam Mchombo. 1992. The syntax of applicatives in chichewâ: problems for a theta theoretic asymmetry. Natural Language and Linguistic Theory 8.493–506.

——, and ——. 1993. Object asymmetries and the applicative construction. Theoretical aspects of bantu grammar, ed. by Sam Mchombo. Stanford: CSLI Publications.

Baker, Mark. 1988. Incorporation. Chicago: Chicago University Press.

Belletti, Adriana, and Luigi Rizzi. 1988. Psych-verbs and theta-theory. Natural Language and Linguistic Theory 6.291–352.

Brachman, Ronald, and James Schmolze. 1985. An overview of the kl-one knowledge representation system. Cognitive Science 9.171–216.

Bresnan, Joan. 1982. Control and complementation. The mental representation of grammatical relations, ed. by Joan Bresnan, 292–390. Cambridge, Mass.: MIT Press.

——, and Jonni Kanerva. 1989. Locative inversion in chichewâ: A case study of factorization in grammar. Linguistic Inquiry 20.1–50.

——, and ——. 1992. The thematic hierarchy and locative inversion in ug: A reply to paul schachter's comments. Syntax and semantics 26: Semantics and the lexicon, ed. by Tim Stowell and Eric Wehrli. New York: Academic Press.

——, and Ronald Kaplan. 1982. Introduction: Grammars as mental representations of language. The mental representation of grammatical relations, xvii–lii. Cambridge, Mass.: MIT Press.

——, and Sam Mchombo. 1995. The lexical integrity principle: Evidence from bantu. Natural Language and Linguistic Theory 13.181–254.

——, and L. Moshi. 1993. Object asymmetries in comparative bantu syntax. Theoretical aspects of bantu grammar, ed. by Sam Mchombo, 47–91. Stanford: CSLI Publications.

Briscoe, Ted, Ann Copestake, and Ann de Paiva. 1993. Inheritance, Defaults, and the Lexicon. Cambridge: Cambridge University Press.

Carpenter, Bob. 1992. The Logic of Typed Feature Structures. Cambridge: Cambridge University Press.

Comrie, Bernard. 1981. Language universals and linguistic typology. Chicago: Chicago University Press.

Croft, William. 1991. Syntactic Categories and Grammatical Relations. Chicago: Chicago University Press.

Davis, Anthony, 1996. Lexical Semantics and Linking in the Hierarchical Lexicon. Stanford University dissertation.

Dowty, David. 1989. On the semantic content of the notion of 'thematic role'. Properties, types, and meaning, ed. by Gennaro Chierchia, Barbara Partee, and Raymond Turner, volume 2, 69–129. Dordrecht: Kluwer.

——. 1991. Thematic proto-roles and argument selection. Language 67.547–619.

Engdahl, Elisabet. 1990. Argument roles and anaphora. Situation theory and its applications, ed. by Robin Cooper, Kuniaki Mukai, and John Perry, 379–393. Stanford: CSLI Publications.

Farkas, Donca. 1988. On obligatory control. Linguistics and Philosophy 11.27–58.

Filip, Hana. 1996. Psychological predicates and the syntax-semantics interface. Conceptual structure, discourse, and language, ed. by Adele Goldberg, 131–147. Stanford: CSLI Publications.

Fillmore, Charles. 1968. The case for case. Universals in linguistic theory, ed. by Emmon Bach and Robert Harms, 1–87. New York: Holt, Rinehart and Winston.

——, and Paul Kay. Forthcoming. Construction Grammar. Stanford: CSLI Publications.

Flickinger, Daniel, 1987. Lexical rules in the hierarchical lexicon. Stanford University dissertation.

Foley, William, and Robert Van Valin. 1984. Functional Syntax and Universal Grammar. Cambridge: Cambridge University Press.

Garrett, Andrew. 1992. Applicatives and preposition incorporation. Grammatical relations: A cross-theoretical perspective, ed. by K. Dziwirek, P. Farrell, and E. Mejías-Bikandi. Stanford: CSLI Publications.

Gawron, Jean Mark, 1983. Lexical Representations and the Semantics of Complementation. University of California at Berkeley dissertation.

Gazdar, Gerald, Ewan Klein, Geoffrey Pullum, and Ivan Sag. 1985. Generalized Phrase Structure Grammar. Cambridge, Mass.: Harvard University Press.

Gerdts, Donna B. 1992. Morphologically-mediated relational profiles. Proceedings of the 18^{th} annual meeting of the Berkeley Linguistic Society, ed. by Laura Buszard, Lionel Wee, and William Weigel, 322–337, Berkeley. Berkeley Linguistics Society.

GIBSON, JEANNE, 1980. Clause Union in Chamorro and in Universal Grammar. University of California San Diego dissertation.

GOLDBERG, ADELE. 1995. Constructions: A Construction Grammar Approach to Argument Structure. Chicago: Chicago University Press.

GREEN, GEORGIA. 1974. Semantics and Syntactic Regularity. Bloomington: Indiana University Press.

GRIMSHAW, JANE. 1990. Argument structure. Cambridge, Mass: MIT Press.

JACKENDOFF, RAY. 1972. Semantic Interpretation in Generative Grammar. Cambridge, Mass.: MIT Press.

——. 1990. Semantic Structures. Cambridge: MIT Press.

——. 1992. Mme tussaud meets the binding theory. Natural Language and Linguistic Theory 10.1–32.

JOSHI, SMITA. 1989. Logical subject in marathi grammar and the predicate argument structure. Proceedings of WCCFL 8, ed. by Jane Fee and Katherine Hunt, 207–219, Stanford. CSLI Publications.

KEMMER, SUZANNE, and ARIE VERHAGEN. 1994. The grammar of causatives and the conceptual structure of events. Cognitive Linguistics 5.115–156.

KIPARSKY, PAUL, 1987. Morphology and grammatical relations. Manuscript, Stanford University.

KOENIG, JEAN-PIERRE. 1993. Linking constructions vs. linking rules: Evidence from french. Proceedings of the 19^{th} annual meeting of the Berkeley Linguistics Society, ed. by Joshua Guenter, Barbara Kaiser, and Cheryl Zoll, 217–231, Berkeley. Berkeley Linguistics Society.

——, 1994. Lexical Underspecification and the Syntax/Semantics Interface. University of California at Berkeley dissertation.

——. 1999a. Lexical Relations. Stanford: CSLI Publications.

——. 1999b. On a tué le président! the nature of passives and ultra-indefinites. Cognition and function in language, ed. by Barbara Fox, Daniel Jurafsky, and Laura Michaelis, 256–272. Stanford: CSLI Publications.

LADUSAW, WILLIAM, and DAVID DOWTY. 1988. Toward a nongrammatical account of thematic roles. Syntax and semantics 21: Thematic relations, ed. by Wendy Wilkins, 62–74. New York: Academic Press.

Lakoff, George. 1993. The contemporary theory of metaphor. Metaphor and thought, ed. by A. Ortony, 2nd Edition. Cambridge: Cambridge University Press.

Langacker, Ronald. 1987. Foundations of Cognitive Grammar, vol.1. Stanford: Stanford University Press.

Legendre, Géraldine. 1989. Inversion with certain french experiencer verbs. Language 65.752–782.

Levin, Beth. 1993. English Verb Classes and Alternations. Chicago: Chicago University Press.

——, and Malka Rappaport Hovav. 1995. Unaccusativity At the Syntax-Lexical Semantics Interface. Cambridge, Mass.: MIT Press.

Manning, Christopher, and Ivan Sag. 1998. Dissociations between argument structure and grammatical relations. Lexical and constructional aspects of linguistic explanation, ed. by Gert Webelhuth, Jean-Pierre Koenig, and Andreas Kathol. Stanford: CSLI Publications.

Mauner, Gail, Alissa Melinger, and Jean-Pierre Koenig, 1999. The time course for encoding implicit agents: Evidence from cross-modal naming and eye-monitoring. Manuscript in preparation.

Mauner, Gail, Michael Tanenhaus, and Greg Carlson. 1995. Implicit arguments in sentence processing. Journal of Memory and Language 34.357–382.

Nunberg, Geoffrey, Ivan Sag, and Thomas Wasow. 1994. Idioms. Language 70.491–538.

Pesetzky, David, 1982. Paths and Categories. Cambridge, Mass.: MIT dissertation.

Pinker, Steven. 1989. Learnability and Cognition: the acquisition of argument structure. Cambridge, Mass.: MIT Press.

Pollard, Carl, and Ivan Sag. 1987. Information-based Syntax and Semantics, vol.1. Stanford: CSLI Publications.

——, and ——. 1994. Head-Driven Phrase-Structure Grammar. Chicago: Chicago University Press.

Postal, Paul, and Geoffrey Pullum. 1988. Expletive noun phrases in subcategorized positions. Linguistic Inquiry 19.635–670.

Pustjovsky, James. 1995. The Generative Lexicon. Cambridge, Mass.: MIT Press.

Quillian, M. Ross. 1968. Semantic memory. Semantic information processing, ed. by Marvin Minsky, 216–270. Cambridge, Mass.: MIT.

Ruwet, Nicolas. 1972. Théorie syntaxique et syntaxe du franais. Paris: Editions du Seuil.

——. 1982. Le datif épistémique en français et la condition d'opacité de chomsky. Grammaire des insultes et autres études, 172–204. Paris: Seuil.

Sag, Ivan. 1991. Linguistic theory and natural language processing. Natural language and speech, ed. by Ewan Klein and Frederik Veltman, 69–83. Berlin: Springer-Verlag.

——, and Carl Pollard. 1991. An integrated theory of complement control. Language 67.63–113.

Singh, Mona. 1992. An event-based analysis of causatives. Papers from the 28th regional meeting of the chicago linguistics society. Chicago: University of Chicago.

Van Valin, Robert. 1993. A synopsis of role and reference grammar. Advances in role and reference grammar, ed. by Robert Van Valin, 1–164. Amsterdam: John Benjamins.

——, and Randy Lapolla. 1997. Syntax: form, meaning, and function. Cambridge: Cambridge University Press.

Wechsler, Stephen, 1991. Argument Structure and Linking. Stanford University dissertation.

——. 1995. Preposition selection outside the lexicon. Proceedings of wccfl xiii, ed. by Raul Aranovich, William Byrne, Susanne Preuss, and Matha Senturia, 416–431. Stanford: CSLI Publications.

Woolford, Ellen. 1993. Symmetric and asymmetric passives. Natural Language and Linguistic Theory 11.679–728.