

Issues in Relating Syntax and Semantics

Daniel JURAFSKY

Berkeley Artificial Intelligence Research
Computer Science Division
University of California at Berkeley
Berkeley, CA 94720

1. Introduction

The design of a natural language understanding system is dependent on a rich and structured knowledge representation. Much recent work in computational and cognitive linguistics has focused on representations for syntactic and morphological knowledge. Similarly, the knowledge representation paradigm has proposed many conceptual tools for the representation of semantic knowledge. Combining the two in one system, a system for representing *linguistic* knowledge, is the essential next task for computational linguistics. This paper will explore some issues which arise in the foundational design of such a representational system. In using the fundamental building block of the *grammatical construction* (Fillmore 1987), I will reanalyze traditional problems like the passive construction and dative "movement", showing how a "construction" solution avoids common problems in more traditional grammatical theories. In addition, I will suggest how use of an inheritance hierarchy can replace other linguistic tools for abstraction like lexical rules, while others can be captured by appealing to shared semantic structures. I will also touch on the issues of lexicalization in theories of grammar - whether information is redundantly specified in the lexicon.

2. Representation of Syntactic and Semantic Knowledge

A complete and explicit description of a language, a grammar, may be represented in many ways. Rule-based descriptions, for example, describe knowledge of language as knowledge of a set of *rules*. In contrast, traditional descriptions in the pre-generative framework assume that knowledge of language can be expressed by a knowledge of *structures*. Fillmore (1987) proposes a return to the traditional notion of the grammatical construction. Construction grammar

aims at describing the grammar of a language directly in terms of a collection of *grammatical constructions* each of which represents a *pairing of a syntactic pattern with a meaning structure...*(p. 3)

Certainly this approach is standard for the representation of lexical knowledge. But extending this approach to larger constructions (the *passive* construction, the *subject-object* construction) means blurring the traditional boundaries between the lexicon and syntactic knowledge, a view also called for in Becker (1976), and Wilensky and Arens (1980). In the construction view, then, idioms, for example, are simply constructions with less variability than other, more productive constructions.

Given that a grammar is a large collection of these pairs, how are they to be represented? Here we turn to insights from the field of knowledge representation. The fundamental metaphor chosen as a foundation for a theory of grammar-meaning correspondences is that of the hierarchical stratified network. The use of inheritance hierarchies for semantic knowledge is well-established. Less common, however, is its use for syntactic knowledge (Although see Flickenger, Pollard and Wasow (1985) for inheritance in the lexicon). Fewer still are systems which have used a structured inheritance network as a representational

formalism for both syntactic *and* semantic knowledge. Notable exceptions are PSI-KLONE (Bobrow and Webber 1980), as well as Jacobs's (1985) ACE/KING system.

The economy and elegance derived from having a grammar defined by a collection of structures defined with a single formalism is a strong argument for this type of unified representation. But this approach, grammatical constructions represented in a unified knowledge-representation formalism, has a number of additional advantages:

- Use of a single representation scheme allows the correspondence between syntax and semantics to be represented in its own right.
- The two types of knowledge can be used together in natural language understanding or generation, thus aiding the development of truly integrated understanding or generation mechanisms.
- Using the same theoretical mechanisms to account for *generalizations* in both syntactic and semantic phenomena encourages the development of unified treatments of linguistic abstraction. Thus, for example, structured inheritance or other abstraction mechanisms can be used as a replacement for syntactic rules, lexical redundancy rules, and other earlier linguistic tools to capture generalizations. And finally, constructs that were originally motivated for the semantic domain, constructs such as prototypes and category theory can be applied to syntactic knowledge, following Lakoff (1987) and his arguments for cognitive grammar.

The explicit pairing of every syntactic form with a semantic one is extremely useful for such understanding techniques as "rule-by-rule" parsing.

It is interesting to note that the two main conceptual primitives of this approach, the doubly articulated construction and the semantic network, both arise directly from some of the earliest work in modern linguistics. Indeed, both are clearly expressed in Saussure's *Cours*. A construction is precisely a sign in the sense of Saussure. It consists of a syntactic and a semantic component, the correlates of his *signifiant* and *signifié*. As for the semantic network, the idea that a theory of meaning could best be expressed by describing the *relations* among concepts is simply Saussure's (1966/1916) claim that "...each linguistic term derives its value from its opposition to all the other terms." (p. 88)

2.1. Domains of Knowledge

I will refer to the two domains in the network which describe syntactic and semantic knowledge simply as the **Syntactic Domain** and the **Semantic Domain**. The semantic domain includes categorization of the lexical item's semantics with respect to the prototypical ontological primitives (as an **Action**, **Thing**, etc). If the construction is headed by a *valence-taking* type, such as most **Actions**, the choice of lexical entry will constrain the choice of *case roles* for each of the constituents. Each valence-bearing type in the lexicon has a *case frame* - a set of *case roles* such as *agent*, *patient*, or *instrument*, which it sub-categorizes (Fillmore 1968). Also associated with each entry is an image-scheme (Lakoff 1987) such as the **CONTAINER** Schema, or the **PART-WHOLE** Schema. In addition, a lexical entry is linked to a *frame*

*Thanks to Nigel Ward and Robert Wilensky. This research was sponsored in part by the Defense Advance Research Projects Agency (DOD), Arpa Order No. 4871, monitored by the Space and Naval Warfare Systems Command under Contract No. N00039-84-C-0089.

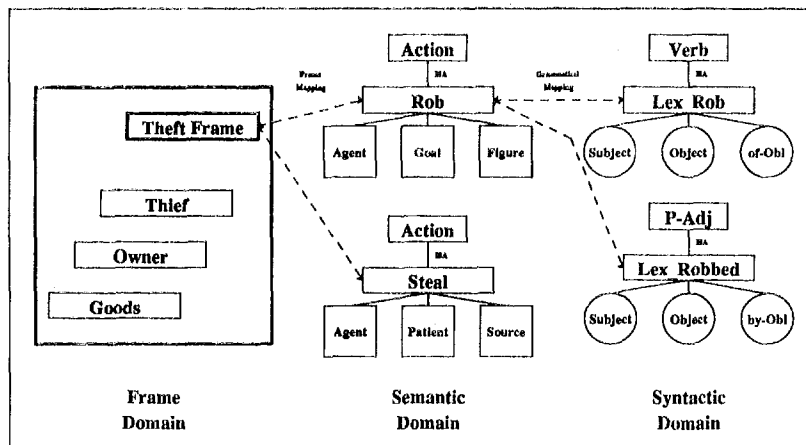


Figure 1 - Sketch of Syntax-Semantics Relation

(in the sense of Fillmore 1982). A frame describes the background constituents and concepts necessary for framing the lexical-semantic definition of a concept. Thus a lexical entry is linked to a frame which it evokes in the understanding process, and which helps in associating and organizing concept in the generation process.

The representation of the subcategorized constituent elements of syntactic patterns at the Syntactic domain is done by grammatical relations such as Subject, Object or Complement.

Grammatical Mapping links relate the Semantic and Syntactic domains, mapping the actual syntactic realization to each of the semantic constituents of the frame - specifying which constituent maps to the syntactic subject, object, and obliques. Following is a sketch of how some of these elements might be related. The grammatical mapping between the semantic entry for "rob" and its syntactic realization have been simplified for expository purposes - a fuller description will be suggested when the passive construction is discussed.

3. Accounting for Some Traditional Problems

Above I sketched an outline of a representation for semantic and grammatical knowledge. In this section I give some suggestions toward a more explicit characterization of the representation. I do this in the context of discussing some traditional linguistic problems which provide insight into and constraints on the design of a representational scheme.

3.1. Passive

The passive construction in English has been analyzed as a transformational rule, a lexical rule, a change in grammatical relations. In each of these views, an active sentence like (1) and a passive like (2)

- (1) Marlowe mixed that perfectly loathsome martini.
- (2) That perfectly loathsome martini was mixed by Marlowe.

are related by some sort of rule. Even in the lexical theories such as LFG, the passive and active lexical items are related by lexical rules which make reference to grammatical relations and other syntactic notions. Thus what these views have in common is an attempt to capture a generalization in language by appealing only to the syntactic component of the language.

In a construction-type theory, one can appeal to the semantic component to capture the requisite generalization. Thus the derived predicate adjective (ie "passivized" verb) is related to the verb it morphologically resembles, by sharing the same argument structure and ontological category at the Lexical-Semantic domain.

But how exactly does the meaning differ? And at what level do we express the productivity of the passive construction? For example, in a lexical theory like LFG, the passive is a redundancy rule that

applies to transitive verbs in the lexicon to produce new lexical entries with passive subcategorization frames. But as Rice (1987) points out, the criteria by which the passive can be used depend on much more information than can be stored in the lexicon. For example, the difference between (3) and (4) is not present in the lexical entry for *sleep*.

- (3) This bed has been slept in again by that flea-bitten dog.
- (4) *The living room was slept in by Mary.

In other words, the criteria for what makes an acceptable passive include more than just the subcategorization information of the verb. Rice gives many such examples, and argues that that the passive constructions makes use of a notion of transitivity which makes reference to the entire conceptualization that an utterance realizes.

A definition of the passive construction would have to include some characterization of the following constraints:

SEMANTIC FUNCTION:

- Focus on the most affected participant in the scene.
- Construe the verbal process as a property or modifier of this participant.

SYNTACTIC FORM:

- Realize this participant as the subject.
- Realize the verb (in the past participle form) as a predicate adjective.

Note that "most affected participant" must be defined with respect to the prototypical notion of transitivity. Additionally, the realization of the passive participle as a predicate adjective motivates its use with the copula (or with *got* - note that adjectives can also use this verb with a change-of-state reading).

This description of the passive construction has been described in such a way that it can be combined with other constructions in a model of language use. In other words, where redundancy rules in a lexical theory may describe an abstraction which has little to do with a performance model, here a construction is not present unless it is used. Redundancy-type abstraction can be captured by use of the abstraction hierarchy.

Thus in generation, for example, the speaker might choose a frame of participants, a verbal predicate which assigns thematic roles, and choose to focus on a participant which happens to be assigned a non-agentive role, or rather one which would prototypically be realized as a direct object. Then the passive construction might be chosen as one construction to include in the utterance. Such a model is used by McKeown (1985). She notes the use of the passive construction to focus on a semantic patient when answering question about missiles, producing "Missiles are carried by water-going vehicles and aircraft..." (p 78).

Similarly, in a parsing system, the presence of the passive con-

struction would indicate the focus on the subject, while the verbal root would enable reconstruction of the thematic roles assigned to the arguments.

3.2. Redundant Information in the Lexicon - Passive Lexical Entries

Although this representation of the passive as a construction which can be combined with other constructions in a model of use is general and elegant, should we still redundantly specify the passivized versions of verbs in the lexicon? This relates to the difficult question of when rules are simply redundancy rules, and when they are necessary for a theory. Two arguments for "compiling out" the lexicon, listing every entry redundantly, are given by Bresnan. The first is to avoid intricate feeding relations among rules. As I show later when I discuss the relation of the Passive and Dative constructions, this can be avoided by describing constructions with the understanding that they are to be embedded in a theory of use. The second is the now classic evidence against the *derivational theory of complexity* (Fodor, Bever, and Garrett 1974). The evidence that the understanding process is no longer for passive sentences than for active ones was strong evidence against a traditional transformational model. Of course, this can also be evidence against a derivational lexical model, if the passive lexical entry is derived from the active one by productive lexical rules. But this argument does not apply to the construction grammar model, for the passive sentence is not derived in any fashion from the active one, but from an underlying semantic conceptualization, in the same way as the active one.

However, there are certainly cases where even the most steadfast proponent of syntactic predictability from semantics must agree that the lexicon must include information that might be redundant. Indeed, as Lakoff (1987) points out "Our individual grammars may not be structured optimally... That is, a speaker may learn certain constructions by rote rather than fitting them into his grammar as a whole." If such lexicalizing of knowledge is necessary in a theory of grammar, we must have a way of representing it when it does occur. If, for example, the English lexicon does represent the passive predicate "mixed" (as opposed to simply including the passive construction and the semantic argument structure for *Mixing-Action*) how would it be represented? Roughly, using the inheritance hierarchy to instantiate a new node in the lexical-semantic domain of the network, and then having it multiply inherit from both *Mixing-Action* and *Passive*. Of course each of these concepts, *Mixing-Action* and *Passive* would have to be further fleshed out - *Passive* according to the construction specification sketched above, and *Mixing-Action* with some sort of image-schematic notions (Lakoff 1987 and Langacker 1987). But then the new concept *Passive-Mixing-Action* would inherit the conceptual structure of both of these.

For example, the ontological Event concerning Marlowe and that martini, and its two syntactic realizations, can be represented as in Figure 2. For clarity's sake I have represented the two different syntactic realizations on either side of the semantics.

Note that using structured inheritance networks allows the Subject-Agent mapping for *Bivalent Verbs* to be inherited from the more general Verb-Action mapping. Future directions for the representation system include the ability to assign *preferences* to different mappings, allowing as representation of the *agency hierarchy*, and other non-binary phenomena.

Talmy (1985) notes that just as the passive is a productive way of forming adjectives from semantic Actions, there are also many lexicalized adjectival forms of verbs, forms not derived from such a productive construction. Here there is no choice but to use the lexicalized version of the representation, as a general productive mapping does not exist. Figure 3 shows an example of how such a lexicalized representation might be characterized. In section 3.3 I will discuss the representation of similar such non-productive derived forms, non-productive

derived nominals.

3.3. Dative

The essential issue in the design of a grammatical representation is how the representation is to be used. In other words, how constructions are to be combined in generation, or how they are to be used in understanding. These issues have not tended to receive a lot of attention, but a few famous construction-combining puzzles exist. One such well-known example is the combination of the Passive with the Dative Shift. (or Dative Movement, or Dative Alternation). The dative shift, first noted in the early 60's, involves verbs such as *give* which seem to allow two different structuring of its complement arguments:

(5) The bartender gave a martini to Bond.

(6) The bartender gave Bond a martini.

A similar or related rule produces such sentences as (7)

(7) I baked my mother a cake.

In a lexical grammar, such as LFG, these rules are lexical redundancy rules which apply to some verbs like "give" and produce new lexical entries. In the new entries, the argument which fills the Goal thematic role is realized with the grammatical relation of direct object, while the Theme role takes the Second Object grammatical relation.

However, many recent semantic studies of these Dative and related examples show that there are strong semantic criteria to the application of the Dative construction. (Greene 1974, Oehle 1975, Baker 1979, Langacker 1986). These studies have shown that the felicitous use of the construction involves the Goal or Benefacted argument ending up with some sort of possession of or control over the Theme. Langacker (1986) notes:

(8) I sent a walrus to Antarctica.

(9) ?I sent Antarctica a walrus.

(10) I sent the zoo a walrus.

Here (9) is deviant because it is difficult to place a continent in the role of a possessor.

(11) I clear the floor for Bill.

(12) ?I cleared Bill the floor.

(13) I cleared Bill a place to sleep on the floor.

Similarly, (13) makes it much easier to conceive of the area of the floor coming under Bill's control than does (12).

The point of these examples is similar to the one made by Rice (1987) and noted above. The semantic constraints on when a Dative Direct Object can be used can not be expressed in the lexicon - the difference between (12) and (13) depends on a construal of the entire utterance conceptualization. Thus the Dative Direct Object construction is stated as a construction in the grammar whose combination with specific lexical items comes up only in *use*. A rough characterization of some aspects of this construction might be:

SEMANTIC FUNCTION:

- Focus on the affectedness of the Goal (or Beneficiary).

- Do so by emphasizing the final control or possession the Goal has over the Theme.

SYNTACTIC FORM:

- Realize the Goal as the Direct Object

- Realize the Theme as the Second Object.

Given the idea of this semantic constraint, and our representation of the passive construction, we are ready to explain a more complex problem, that of utterances with both Dative Alternation and Passive constructions, such as (14) or (15). Of course, in the transformation paradigm such examples were handled by *rule ordering* - if the Dative Movement rule moving "Bond" to object position took place before the Passive rule, a structure like (15) resulted, otherwise one like (14).

(14) Bond was given a martini by the bartender.

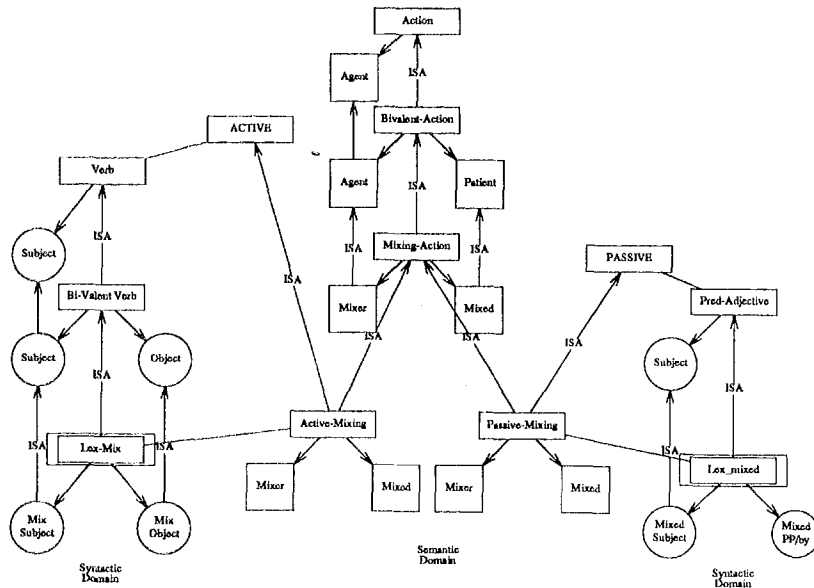


Figure 2 - Passive and Active Actions

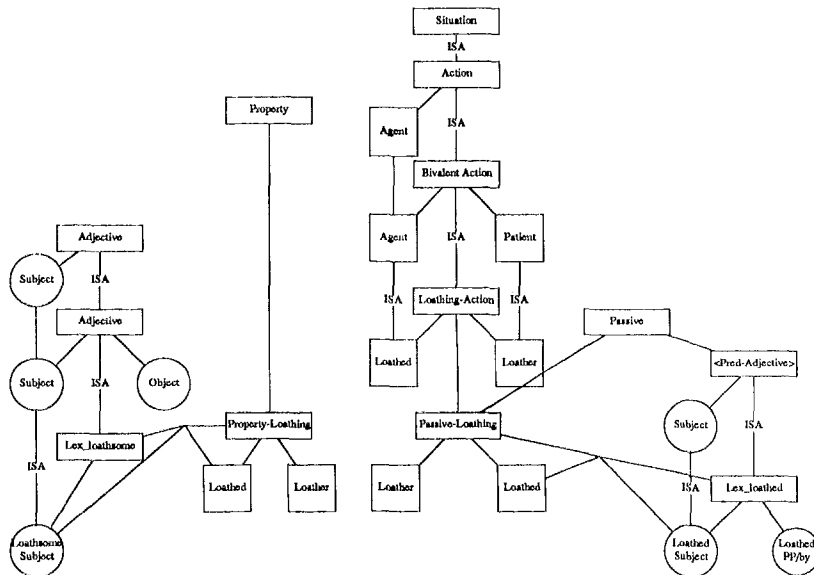


Figure 3 - Passive and Deverbal Adjectives

(15) A martini was given to Bond by the bartender.

However in a construction-based formalism, the interaction between the constructions can be predicted by examining the semantics of the constructions. In the Dative rule, the Goal participant must be focused on as having some sort of control over the Theme. The point is that what it means for a Goal participant to be affected in a Goal-transfer scenario is for the Goal to be transferred possession or control of the Theme.

For different types of scenarios (or Idealized Cognitive Model (Lakoff 1987)), then, transitivity means different things. Rice (1987) takes note of this when she proposes three cognitive domains with separate transitive event prototypes: the physical, mental, and social domains.

For example, for Locative-transfer scenarios, what it means to be affected by the action is somewhat different. As many writers have noted (among them Anderson 1971 and Foley and Van Valin 1984), in

the (b) sentences the locative arguments (*the truck* and *the wall*) can be understood as *completely* affected by the verbal action. In other words, the truck completely filled with hay, the wall completely covered with paint. The (a) sentences do not have this reading.

For example,

(16a) Bill loaded hay on the truck.

(16b) Bill loaded the truck with hay.

(17a) Harry sprayed paint on the wall.

(17b) Harry sprayed the wall with paint.

Thus "affected" here means something about completion. Of course, this is not surprising. Both Rice (1987) and Hopper and Thompson (1980) do list telic actions as more prototypically transitive than atelic ones. But the important point is that what it means to be affected by an action is different for different action-scenarios.

Given some such theory of affectedness in transfer-scenarios, we

can say that the Dative construction is used to mark the affectedness of the Goal in this way. But an affected object is exactly the sort of participant the the Passive is used for focusing on. Thus sentences like (14) and (15) can be characterized simply by whether or not they use certain constructions. Sentence (15) is not formed by syntactic feeding rules, but by the combination in *language use* of these two constructions.

3.4. Nominalizations

The correspondence between syntax and semantics is encoded by mapping relations which link groups of syntactic and semantic constituents in the network. To take a lexical example, the semantic primitive "dog" might be related to the syntactic entity which instantiates the noun "dog" by one of these relations. Likewise, more complex constructions are mapped between the syntax and semantics. So, for example, the syntactic construction [Adj Noun] might be related to a number of semantic constructions, specifying the various prototypical ways of combining the semantic information associated with adjectives and nouns. A complete theory of this correspondence would involve showing how every syntactic construction was the realization of some semantic object. But clearly the appropriate syntactic construction to begin our investigations is the lexical category. Note that a lexical category is a grammatical construction like any other, differing only in the very frequency of its use.

It has long been observed that a rough correspondence can be drawn between lexical categories and an ontological partitioning of the world. Everyone is familiar with the traditional grammarians' definition of a noun as a word representing a "person, place or thing". Naturally, the exceptions to this simplistic analysis are abundant. For example, nominalizations such as *destruction* seem to be more related to actions than to "things." But the many exceptions to this analysis caused the real semantic nature of lexical categories to receive less attention than it deserved. As Miller and Johnson-Laird (1976) note, "perhaps [traditional grammarians] did not really mean that everything labeled by a noun is a concrete object; perhaps they meant that when you use a noun to label something, you tend to conceptualize it as if it were a concrete object (p 527)". In the spirit of this Whorfian idea that lexical categories reflect a set of ontological categories, the **Lexical-Semantic Domain** includes a set of such ontological primitives as **Events, Actions, and Things**. Mapping relations in the network relate these primitives to syntactic constructions. In prototypical cases, like the noun "table", what is semantically a **Thing** is mapped to a constituent that is syntactically a noun (or more properly a noun phrase). Thus the lexical categories of language are the *prototypical realizations* of these underlying semantic regularities. These "prototypical realizations" are simply a version of what Lakoff (1987) called "central principles". Note that they are also similar to Chomsky's (1986) "Canonical Structural Realizations" of semantic concepts.

However, as each kind of ontological category can be mapped to different syntactic category, it is also possible to have cases of non-prototypical mapping. Thus what have traditionally been called productive nominalizations are non-prototypical mappings between **Actions** and different types of nouns. A nominalization like *John's destroying the city* is related semantically to an **Action**, but is mapped in form to a noun phrase. The **Action** "destroy" maps into the noun "destroying". Similarly, remember that the passive construction involved the reanalysis of an **Action** as a predicative adjective.

As for noun-phrase mappings, as Ross (1973) showed, there are a number of noun phrase constructions which bear different relations to the prototypical noun phrase. He discusses eight different types of noun phrases which he arranges along a *Nouniness Squish*. These NP's are *that-clauses, for NP to V, embedded questions, gerunds, possessive gerunds, action nominals, derived nominals, and Nouns*. An examination of the syntactic category **Noun-Phrase** is essential to any discussion of relations between **Actions** and **Nouns** (I use **Action** here only as an example - of course other ontological categories like **States** are

relevant as well). One of the great advantages to a unified representation for linguistic and conceptual knowledge is that such models as Lakoff's (1987) category theories for semantics are automatically applicable to syntactic knowledge.

But here I will only discuss a few of the lexical issues involved in nominalizations, and in particular the difference between what Ross called "action" and "derived" nominals, or the productive and non-productive nominalizations.

For example, the verb "destroy" has two lexical nominalizations - the productive or gerundive "destroying", and the derived or non-productive "destruction". The nominal use of the gerundive "destroying" is productive in the sense that if a new verb entered the language, say "to xerox", the native speaker would automatically be able to speak about "xeroxing".

Chomsky (1970) proposed that the relationship of these two types of nominals to their corresponding verb be captured in two different ways. The relationship between productive nominalizations and verbs was captured by means of a transformational rule. Non-productive nominalizations (like "destruction") are related to the verb by sharing similar phrase structures. To accomplish this, Chomsky proposed the X-Bar Convention, the use of generalized cross-categorical phrase structure rules.

In the network framework, the generalization between verbs and productive nominalizations is handled by structure sharing - they share descriptions of argument structure in the semantic domain, just as with the passive examples. So a network expressing the relation among "destroy" and "destroying" and their thematic structure would be similar to the passive example in Figure 2.

Indeed, the abstraction hierarchy allows a way to provide for the creativity of language without explicit rules. Imagine the invention of a new **Action**, *xerox*. Now the concept *xerox* is inserted under the more general concept **Action**, and inherits from it the mapping relation to the nominal form. A possible representation is outline in Figure 4.

Note that this method does not involve a "rule" to capture its generalizations. This may call into question the classic Chomskyan argument from creativity to the need for rule systems.

In this way we can handle non-productive nominalizations as well as productive ones. A characteristic feature of these non-productive nominalizations is that their semantics is rarely predictable from the semantics of the underlying verb - they tend to differ in idiosyncratic ways. From a diachronic element, we would say that the non-productive nominalizations have undergone semantic drift. The network representation of the synchronic results of semantic drift are simply the creation of a new node which inherits from **Thing** at the **Semantic level**. But note that although this new node may differ in various elements of its semantics, it will rarely drift far enough to change its thematic structure. Thus the similarity in the argument structure between non-productive nominalizations and verbs is not a syntactic fact, as predicted by Chomsky 1970, but a semantic one. This would predict that non-productive nominalizations which have drifted far enough to have changed their argument structure would also differ syntactically from the related verb. This is indeed found to be the case. So, for example although we are perfectly comfortable in saying (18) or (19), the meaning of "laughter" has drifted enough that it is uncomfortable to use it with an object in (20).

- (18) Who is stupid enough to laugh at Marlowe?
- (19) Laughing at Marlowe is a dangerous proposition.
- (20) *Laughter at Marlowe is rare.

A redundant lexicalized representation of the non-productive nominalizations might look much like the one for adjectives in Figure 3.

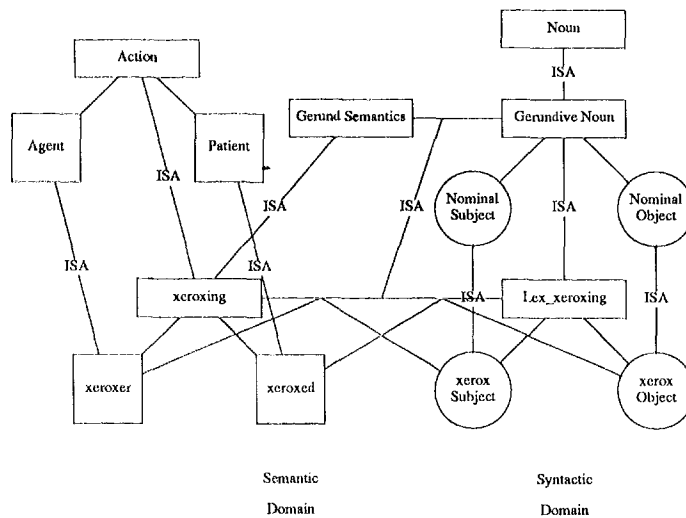


Figure 4 - Inheriting the Nominalization Relation

3.5. Problem and Future Directions

The derivation of verbs from nouns operates in a markedly different way from the derivation of nouns from verbs. As Clark and Clark (1979) note, a speaker in using a denominal verb denotes a situation where "the parent noun denotes one role in the situation, and the remaining surface arguments of the denominal verb denote other roles in the situation." A sketch of a solution would involve view-mappings which relate nouns to their derived verbs in the **Semantic** domain. In this domain, a heavily context-dependent process creates a new lexical entry through a metaphorical mapping from the underlying **Thing**. Thus this phenomenon may be much more process-oriented than the construction model currently allows for.

Another important issue is the ability to model prototype effects and other cognitive issues in category theory. The ISA links which I have sketched in the lexicalized models are only an impoverished model of categorization. A solution will probably involve some solution such as adding numerical weights to nodes and to links to express typicality weights.

4. Conclusion

This paper has suggested the barest skeleton of a representation system for linguistic knowledge. I have discussed the semantics of two important constructions, Passive and Dative, and have roughly them sketched out in such a way as to give us an insight into how they can be combined in a model of language use. In addition, I discussed how the information and constraints in these constructions might be redundantly specified in the lexicon. These problems have offered us a glimpse of what sort of constraints difficult linguistic problems will set for representational theories. In exploring these constraints we gain important insights into the requirements for a model of our use of natural language.

5. References

Andrews, Avery 1985. "The major functions of the noun phrase" in Shopen 1985
 Becker, Joseph 1976. "The Phrasal Lexicon". in *Proceedings of Interdisciplinary Workshop on Theoretical Issues in Natural Language Processing*. R. Schank and B. L. Nash-Webber (eds.). Cambridge, Mass. June 1975.
 Bobrow, D.G. and Allan Collins. 1975. *Representation and Understanding*. New York: Academic Press

Bobrow, R.J. and Bonnie Webber, "Knowledge Representation for Syntactic/Semantic Processing," AAAI 1980.
 Brachman, Ronald J. and James G. Schmolze. 1985. An Overview of the KL-ONE Knowledge Representation System *Cognitive Science* 9 pp 171-216
 Bresnan, Joan (1977) "The Passive in Lexical Theory", in *The Mental Representation of Grammatical Relations* Cambridge, Mass. MIT Press.
 Chomsky, Noam, 1970 "Remarks on Nominalizations." In R. Jacobs and P. Rosenbaum, eds., *Readings in English Transformational Grammar*. Ginn, Waltham, Mass.
 Chomsky, Noam. 1986. *Knowledge of Language*. New York, Praeger.
 Clark, Eve V. and Herbert J. (1979), "When Nouns Surface as Verbs," *Language*, Volume 55, 4.
 Cole, Peter & Jerrold Sadock, eds. 1977 *Grammatical Relations*. (Syntax and Semantics 8). New York: Academic Press
 Comrie, Bernard 1976. *Aspect* Cambridge, Cambridge Univ Press
 Fillmore, Charles, draft, *On Grammatical Constructions*. University of California, Berkeley 1987.
 Fillmore, Charles. 1982 "Frame Semantics." in *Linguistics in the Morning Calm*, Linguistic Society of Korea, Korea.
 Fillmore, Charles. 1977. The case for case reopened. in Cole & Sadock, eds. pp. 59-81.
 Fillmore, Charles. 1968 "The Case for Case". In *Universals in Linguistic Theory*. Edited by Emmon Bach and Robert T. Harms, 1-90. Chicago: Holt, Rinehart and Winston
 Fodor, J. A., T. G. Veber, and M. F. Garrett. 1974 *The Psychology of Language: An Introduction to Psycholinguistics and Generative Grammar*. New York: McGraw Hill.
 Foley, William A., and Robert D. Van Valin, Jr. 1985. "Information packaging in the clause" in Shopen 1985
 Foley, William A., and Robert D. Van Valin, Jr. 1984. *Functional Syntax and Universal Grammar* Chicago, University of Chicago Press
 Gazdar, G., E. Klein, G. K. Pullum and I. A. Sag. 1985. *Generalized Phrase Structure Grammar* Oxford, Basil Blackwell.
 Green, G. M. 1974. *Semantics and Syntactic Regularity*. Bloomington, Ind. Indiana University Press.
 Gruber, Jeffrey. 1965. *Studies in lexical relations*. MIT. Dissertation

- Jackendoff, Ray. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, Mass.: MIT Press.
- Jacobs, Paul S. 1985. A Knowledge-Based Approach to Language Generation. UC Berkeley Dissertation, also Report 86/254 Computer Science Div., University of California
- Kay, Martin 1979. "Functional Grammar" *Proc. 5th Ann. Meeting of the Berkeley Ling. Soc.* 142-158.
- Lamb, Sidney 1964. "Stratificational Linguistics as a Basis for Machine Translation" in Makkai and Lockwood 1973
- Langacker, Ronald. 1986. An Introduction to Cognitive Grammar *Cognitive Science* 10, 1
- Langacker, Ronald. 1987. *Foundations of Cognitive Grammar* Stanford University Press.
- Makkai, Adam, and David Lockwood 1973. *Readings in Stratificational Linguistics*. University of Alabama Press
- Marantz, Alec P. 1984. *On the Nature of Grammatical Relations* Cambridge, MIT Press
- Miller, George A. and Philip N. Johnson-Laird. 1976. *Language and Perception*. Cambridge, Mass. Harvard University Press
- Norvig, Peter. 1986. A Unified Theory of Inference for Text Understanding. UC Berkeley. Dissertation.
- Oehrle, R. T. 1976. The grammatical status the English dative alternation. Doctoral dissertation. MIT.
- Perlmutter, David ed. 1983. *Studies in Relational Grammar 1* Chicago, University of Chicago Press
- Perlmutter, David and Paul Postal. 1983. "Toward a Universal Characterization of Passivization" in Perlmutter 1983.
- Pollard, Carl and Ivan A. Sag. 1987. *Information-Based Syntax and Semantics: Volume 1: Fundamentals*, Stanford, CSLI.
- Rice, Sally Ann. 1987. Toward a Cognitive Model of Transitivity. Doctoral Dissertation. University of California at San Diego.
- Ross, John Robert. 1973. "Nouniness" in Osamu Fujimura, ed., *Three Dimensions of Linguistic Theory*, pp. 137-258. Tokyo: TEC Corporation
- Saussure, Ferdinand de 1915/1966. *Course in General Linguistics*, transl. Wade Baskin. McGraw Hill New York, 1966 orig 1915.
- Shopen, Timothy ed. 1985. *Language Typology and syntactic description, Vol. I*. Cambridge Cambridge University Press
- Talmy, Leonard. 1985. "Lexicalization Patterns: semantic structure in lexical forms" in Shopen (1985) Vol. III.
- Wilensky, R., and Arens, Y. 1980. PHRAN -- A Knowledge-based Approach to Natural Language Analysis. University of California at Berkeley, Electronics Research Laboratory Memorandum #UCB/ERL M80/34.
- Woods, William A. 1975. What's in a link? in Bobrow and Collins 1975. pp 35-82