VERBAL SEMANTICS VIA PETRI NETS

Miriam Butt and Alexandros Tantos
Universität Konstanz

Proceedings of the LFG04 Conference
University of Canterbury
Miriam Butt and Tracy Holloway King (Editors)

2004

CSLI Publications
http://csli-publications.stanford.edu/

Abstract

This paper introduces Petri Nets (Peterson 1981) into the vexed question of lexical semantic representation. We argue that Petri Nets can offer new insights into the organization of lexical meaning. Petri Nets can furthermore be translated directly into linear logic, thus ensuring compatibility with the glue logic semantic construction already defined within LFG (Dalrymple 1999). The potential usefulness of Petri Nets is demonstrated with respect to light verbs in Urdu.

1 Introduction

The study of complex predicates, both of the V-V and the N-V type, has been conducted quite intensely within Lexical-Functional Grammar (LFG) and related approaches (e.g., Alsina 1996, Alsina, Bresnan and Sells 1997, Butt 1995, Mohanan 1994). The major syntactic properties of complex predicates have been explored within, by now, well established analyses in terms of a(rgument)-structure fusion or merging. Complex predicates have been recognized to be complex in the sense that the two components of a complex predicate each contribute information about the a-structure properties of the construction. The main part of the predication is contributed by the noun or main/full verb. The second component of the complex predicate has generally been termed a *light verb* because its contributions tend to be semantically "lighter" than that of the main predicate. Some examples are shown in (1).

(1) a. ram=ne kɑhani=ko **yad** **k-i**
Ram.M.Sg=Erg story.F.Sg=Acc memory.F do-Perf.F.Sg
'Ram remembered the story.'

b. nadya=ne xɑt **lɪkʰ** **li-ya**
Nadya.F=Erg letter.M.Nom write take-Perf.M.Sg
'Nadya wrote a letter (completely).'

The light verb is generally finite and carries tense/aspect marking. Its contribution to the a-structure is so subtle that some researchers have been tempted to analyze it as being empty (e.g., Grimshaw and Mester 1988), though one can show that the light verb does indeed contribute at least one argument to the predication (usually the highest one) (e.g., Butt 1995, Ritter and Rosen 1993).

Since both the main verb/noun and the light verb contribute information to the a-structure, the predication is complex. Unlike in equi/raising or control constructions (see Bresnan 1982 for an overview), this complex predication corresponds to a simplex predication at f-structure. That is, although the a-structure is complex and potentially embedded, the f-structure is exactly that of a monoclausal predication (see Alsina 1996, Butt 1995, Mohanan 1994).

Although the syntactically relevant predicational aspects of constructions as in (1) have been investigated in some detail (see also Alsina 1996, Butt 1995, Kaplan and Wedekind 1993, Butt 1994, Butt, King and Maxwell 2003 on architectural issues within LFG), the precise semantic contribution of the light verb to the joint predication remains elusive. One classic difficulty with complex predication involving light verbs is that the meaning does not necessarily represent the sum of its parts. While light verbs are always form-identical to a main verb in the language, they do not predicate like the main verb version. This is especially true for V-V constructions as in (1b), where the light verb version of 'take' does not mean 'take', but rather seems to indicate some kind of completion. Additionally, it can be contrasted with the light verb 'give', as in (2), whereby 'give' indicates benefaction (for somebody else) and 'take' implies that the action was self-serving.

(2) a. nadya=ne gʰɑr **bɑna li-ya**
Nadya.F=Erg house.M.Nom make take-Perf.M.Sg
'Nadya built a house (completely, for herself).'

b. nadya=ne gʰɑr **bɑna di-ya**
Nadya.F=Erg house.M.Nom make give-Perf.M.Sg
'Nadya built a house (completely, for somebody else).'

However, this contrast between 'take' and 'give' does not always hold. Consider, for example, the sentences in (3). In neither of these examples can the light verb 'give' be interpreted as adding a benefactive meaning to the complex predication. Rather, as discussed in Butt and Geuder (2001), the main contribution of 'give' in (3b) is one of forcefulness, whereby in (3a), the use of 'give' implies responsibility for the loss on the part of the agent.

(3) a. kɪsi=ne baṭua **kʰo di-ya**
someone=Erg wallet.M.Sg.Nom lose give-Perf.M.Sg
'Somebody lost their wallet.' (from Hook 1974)

b. nadya=ne dʊʃmɑn=ko pani=mẽ **ḍub-a di-ya**
Nadya.F=Erg enemy=Acc water=in drown-Caus give-Perf.M.Sg
'Nadya drowned the enemy in the water (forcefully).'

Indeed, the range of usage documented for Urdu/Hindi[1] light verbs in V-V complex predicates (see Hook 1974 for a comprehensive discussion, for example) shows that the meanings contributed by the light verbs are manifold, varied, and contextually dependent. A popular analysis of light verbs has been that they represent a "semantically bleached" form of the main verb. However, it is not clear exactly what is being bleached into what and how the contextual dependence can be accounted for. Another popular analysis has been that the light verbs represent stages on the way to a grammaticalization of aspect (Hook 1991, 2001). While this would account for the completive/telic readings which are definitely part of the examples in (2) and (3), this analysis does not extend to the other subtle dimensions of meaning such as control, suddenness forcefulness, or benefaction.

This paper builds on the insights arrived at by Butt and Geuder (2001) and Butt and Ramchand (2003) as to the semantic composition of V-V complex predicates, but attempts to take things a step further towards a concrete formal understanding of the range and type of semantic predication that a light verb can have in interaction with a main verb. To this end, we introduce Petri Nets (Peterson 1981) into the vexed question of lexical semantic representation. In particular, we assume that the form-identical light verbs and main verbs must be derived from one and the same underlying, underspecified lexical entry. This underspecified lexical entry becomes specified as either a main verb or a light verb predication, depending on the surrounding syntactic context. The process of the semantic specification towards either a light or a main verb meaning is modeled in terms of a Petri Net, thus allowing one to concretely identify the semantic dimensions that are involved.

In what follows, we first lay out the reasons for assuming a semantically underspecified lexical representation for main verbs and their form-identical light verbs (section 2). We then briefly discuss Petri Nets and how they can be put to use in modeling linguistic processes (section 3). In section 4 we provide sample analyses involving the verb 'give' and try to account for the varied dimensions in meaning exemplified by (2) and (3). Finally, section 5 concludes the paper.

2 Light Verbs and Semantic Underspecification

This section briefly discusses why an analysis in terms of lexical underspecification should be adopted with respect to light verbs and their form-identical main verbs.

[1]The South Asian languages Urdu and Hindi are closely related. Differences are found mainly in the domain of the vocabulary. Both languages are among the 18 official languages of India and are spoken primarily in the north of India. Urdu is the national language of Pakistan.

2.1 Historical Interconnectedness

The primary reason for assuming underspecification has to do with the diachronic behavior of light verbs. Given the "semantic bleaching" or aspectual grammaticalization analyses alluded to above, one would expect that light verbs behave much like auxiliaries in terms of diachronic development. That is, given a verb like 'go', one expects a stage in which there is a 'go' which is used in a more temporal, rather than a concrete spatial sense. A current example of this in English is the "going to" future, as in *Peter is going to go home.* Over time, one expects that this 'go' will be reanalyzed as an auxiliary with future import and given more time, this auxiliary night then be further reanalyzed as tense inflection.

Indeed, this is exactly what can be found with the verb 'go' and the development of future morphology in Urdu. The table in (4) shows the paradigm for the Urdu/Hindi future.

(4) **Urdu Future Paradigm** (for mar- 'hit')

	Singular M/F	Plural M/F	Respect (ap) M/F	Familiar (tʊm) M/F
1st	mar-ũ-g-a/i	mar-ẽ-g-e/i		
2nd	mar-e-g-a/i		mar-ẽ-g-e/i	mar-o-g-e/i
3rd	mar-e-g-a/i	mar-ẽ-g-e/i		

The consensus in the literature is that the future *-g-* morpheme is derived from a Sanskrit participle of the verb *gā* 'go' (Kellog 1893:231, Beg 1988:191, McGregor 1968). The gender and number agreement morphology (*a/i/e*) exhibited by the future is regular synchronically in that exactly this agreement morphology is also found on the perfect, imperfect and progressive forms, all descended from participles. The appearance of this morphology follows unproblematically if the *-g-* is indeed associated with an old participle of 'go'. The person/number inflection of the future paradigm in (4) is identical to the inflections found in conjunction with the present tense paradigm of *ho* 'be'. There is some indication that these forms are indeed related, so that one can speculate that the modern Urdu future consists of a verb stem, some present tense inflections or a trapped present tense auxiliary, the remnants of the participle 'go', and the gender/number agreement inflections that belong with participles.

Up until a hundred years ago, the main verb+person/number (former present tense) morphology could be separated from the *g*+number/gender morphology (some speakers can still do this). This indicates that the change from periphrastic auxiliary to future inflection with respect to 'go' took place relatively recently and that the change from a main verb to a tense inflection took place quite rapidly. Now, if light verbs were to be derived from main verbs in a similar manner, one would expect to see similar patterns of historical change. However, a thorough scrutiny of the available diachronic data fails to turn up any such patterns (Butt and Lahiri 2004).[2] Instead, at every stage in the language where a light verb use can be identified (this is not always easy), the light verb is form-identical to a main verb in the language. This includes taking exactly the same inflections and participating in exactly the same paradigms as the main verb version. An example of a documented use of the light verb version of 'give' in Middle Indo-Aryan is shown in (5).

(5) a. . . . assamapadaṃ ānetvā aggiṃ **katvā** **adāsi**
hermitage.Acc lead.Gd fire.Acc.Sg make.Gd Aug.give.Impf.3.Sg
'. . . brought her to his hermitage and made a fire for her' (Pāli)
['having brought (her) to the hermitage, made a fire (for her)']
Jātaka Tales I.296.10, Sri Lanka (Hendriksen 1944:134)

[2] Hook 1991, 2001 documents an increase in the use of light verbs in South Asian languages, as well as a more definitive shift towards encoding aspectual differences, but nothing along the lines that has been established as typical of auxiliary formation.

b. daruni āharitvā aggiṁ **katvā** **dassati**
sticks bring.Gd fire.Acc.Sg make.Gd give.Fut.3.Sg
'Bringing wood he'll make a fire (benefactive use).' (Pāli)
(Trenckner 1879:77, cited by Hook 1993:97)

There is thus no documented development of the light verb *away* from the main verb version. Rather, the main and the light versions of a verb seem to be tied to one another in an intimate manner. This diachronic observation holds not just for Indo-Aryan languages like Urdu, but is one that has been documented across several language families (Germanic, Indo-Aryan, Dravidian; see Butt and Lahiri 2004 for a detailed discussion). Further evidence for the intimate connection between main and light verb versions is the observation that when a verb ceases to exist in a language, then both the main and the light verb usage disappear simultaneously (if both exist). For example, when the English *nimen* 'take' dropped out of the language, both main and light verb uses were taken over simultaneously by *taken* (Iglesias-Rábade 2001).

The available diachronic evidence thus shows that one never finds a light verb on its own: there is always a form-identical main verb in use as well. This situation stands in stark contrast to that of auxiliaries, which tend to develop away from the original main verb form until they are almost unrecognizable (e.g., the English preterite *-d* from *do* or the Urdu future *-g-* described above). This suggests a fundamental interconnectedness between the main and the light use of a verb. One could attempt to analyze the light verb as being derivative of the main verb, but then one would have to stipulate that the light verb must remain connected to the main verb in some way. Given what is known about historical change in general, this type of stipulation seems to be artificial and unexplanatory, to say the least.

On the other hand, if one assumed that both the main and the light uses of a verb are derived from one and the same underlying lexical entry, then the facts follow. If there is only one underlying lexical entry from which both are derived, then the main and light verb use should be form-identical. Furthermore, if the lexical entry is deleted from the grammar of the language, then both the main and the light verb use will cease to exist at the same time. Historical changes that apply to change the surface form of the verb (changes in morphology, form, etc.) will apply to both the light and the main verb uses, since there is just one underlying lexical entry, which these processes can access.

If one grants that an approach in terms of a single underlying lexical entry is on the right track, then the next question which arises is one of representation. Given that (at least) two uses must be derived from one lexical entry, one possible route to take would be to fully specify all the possibilities in the form of disjunctions. These disjunctive possibilities would then be simply associated with one and the same lemma (*de* 'give', for example). This type of "full listing" could lay no claims in terms of elegance of explanation or generalizability; however, if it did justice to the facts, one might be tempted to choose this approach.

The next section discusses a further set of observations that would seem to legislate against a "full listing" type of approach. Instead, a representation in terms of lexical underspecification emerges as potentially more feasible.

2.2 Defeasible Information

Recall that the use of a light verb like 'give' always entails that the action is completed. Furthermore, the light verb 'give' can potentially contribute the semantic dimensions of benefaction, control, or forcefulness to a given predication. Butt and Ramchand (2003) analyze the contribution of completion/telicity in terms of the internal lexical semantic structure of the predication. A given event is taken to have exactly three salient components: a cause/intiation, a process, and a result. These three semantic components tend to be grammatically encoded in languages via morphological or syntactic devices. Light verbs are seen as a syntactic device which interact with the event semantics of the main verb to produce a more complete event description. The more subtle semantic dimensions like benefaction or forcefulness are not dealt with under this approach, but are assumed to have a status that is akin to adverbial event modification, as proposed by Butt and Geuder (2001). That is, when the verb 'give' acts as a light verb, the joint event predication is one in which the agent had control over the event and the event is telic, happened in a forceful manner and had some benefit for a participant distinct from the agent of the event.

In something like (3), for example, repeated in (6), the lexical semantic contribution of the light verb to the predication results in a reading that the agent has control and that the event is telic and happened in a forceful manner. However, there is no sense of benefaction.

(6) nadya=ne dʊʃmɑn=ko pani=mẽ **ḍub-a** **di-ya**
Nadya.F=Erg enemy=Acc water=in drown-Caus give-Perf.M.Sg
'Nadya drowned the enemy in the water (forcefully).'

Now, if one were to take the disjunctive, full listing approach discussed above, then one would also have to anticipate all the situations in which the semantic dimensions of benefaction or forcefulness could or could not apply. Since this depends on contextual as well as lexical semantic factors, listing all the potential contexts which would license (or suppress) these additional event modifications is not a feasible solution.

Another question to consider with respect to a possible common underlying lexical representation for both main and light verb uses is whether one could derive event modificatory semantics such as benefaction or forcefulness from an abstract predicational force associated with the verb 'give'. That is, is there something about the predicational force of 'give' that would lead us to predict that it might be associated with benefaction and forcefulness, as opposed to a verb like 'take', for example (cf. the contrast in (2))? The answer that Butt and Geuder (2001) give to this is a "yes". Under their analysis, the meaning dimensions of the light verb 'give' are taken to be loosely based on the predicational force of an abstract action 'give'.

In conclusion, the diachronic and synchronic facts with respect to light verbs in Urdu lead us to the following realizations:

1. Both light and main verb uses must be derived from the same underlying lexical entry.
2. The event modificatory semantics contributed by the light verb are not random, but are to be associated with an abstract representation of the verb in question.
3. Not all of the meaning dimensions of a light verb must always apply to a given event modification.

Given these observations and realizations, what could such a single underlying, underspecified entry look like? Standard approaches to lexical representation in terms of lexical decomposition, inheritance hierarchies, cognitive semantics, or Davidsonian event semantics all seem to lack the repertoire necessary for a solution to this problem. Or to put it another way, there seems to be no room for the kind of phenomenon described here in the approaches to lexical semantics that have dominated the last few decades: it is very difficult to apply known techniques to the problem at hand, or even to come up with imaginative new ways within the existing approaches.

We do not think it is a coincidence that no detailed lexical semantics approaches exist that could capture the semantic contributions of light verbs, even though these semantic contributions have been observed for quite some time and have been well documented across several language families. At most, the semantics of light verbs is approached within an argument structure or grammaticalization approach (see discussion above), but this tends to be able to account only for a subset of the observed meaning dimensions.

Given that the existing approaches seem to be acting as limiters, rather than as enablers, we decided to explore alternative possibilities for the representation of lexical entries. With *Petri Nets*

(Peterson 1981), we found a model that seemed promising in terms of providing new insights with respect to the representation of lexical semantics. In particular, since the type of Petri Nets used here can be translated into the version of linear logic used for the glue semantic component of LFG (Dalrymple 1999), we anticipate that we would be able to translate the results we arrived at through an exploration of Petri Nets directly into consequences for the representation of the necessary semantic dimensions within glue semantics.

3 Petri Nets

In his dissertation, Carl Adam Petri (1962) introduced a kind of directed graph that opened up new ways of thinking about design issues in communication and computer systems. These so-called *Petri Nets* were accepted by researchers within such diverse fields as biology, astronomy, nuclear physics and sociology as a promising way of solving problematic modeling issues. The Petri Net community today is very active and organized (see, e.g. http://www.daimi.au.dk/˜petrinet/).

Petri Nets are used primarily to model the interactions between different components of systems, whereby each of the components is autonomous and functions independently from the others. This means that even while interacting with one another, each one of the components may execute different tasks at the same time. That is, Petri Nets deal with interacting concurrent components. Since the first presentation of Petri Net theory, a number of extensions for differing applications and purposes have been developed over the years. Even though many of the beneficial aspects of Petri Nets are irrelevant in terms of linguistic theorizing, we would argue that there are some extremely elegant properties of Petri Nets that can be used for a representation language for lexical semantics. In this section, we introduce the ordinary or black-white Petri Nets (based on the version first described by Petri 1962), whose architecture would seem to suffice for the linguistic modeling of lexical semantics.

A Petri Net is represented in the form of directed graphs and serves the purpose of maintaining simplicity and clarity in our analyses. Before looking at specific examples containing Petri Nets as means of representations (section 4), it is essential to present a brief overview of the formal properties of these graphs so that the importance and novelty of the current approach to lexical semantic representation becomes clear. One central characteristic of Petri Nets is that they constitute bipartite graphs. This means that the graphs are determined through the use of two distinct types of nodes: places and transitions. In addition to these two types of nodes and the arcs which assure the directivity of the graph, a fourth object is introduced in order to describe the dynamics of a Petri Net. This object is the token, expressed by a solid dot •, which resides inside the circles representing the places. In the framework we work in, namely, the ordinary Petri Nets, the tokens represent abstract information and are not distinguishable from one another. The property of Petri Nets of having two types of nodes defines the way the so-called ‘token game’ is conducted inside a Petri Net. The token game consists of a transfer of ‘tokens’ (represented as the black dots) through the graph, whereby the ‘event’, as the ‘complete’ execution (including the last possible transition) of a Petri Net can be intuitively called, can be described in different stages according to the positioning of the tokens. Below, the conception of a token game is explained, along with the procedural nature of Petri Nets.

A Petri Net structure, D, is a quadruple, D=(P,T,I,O), where $P=p_1, p_2,\ldots,p_n$ is a finite set of places, $n\geq0$. $T=t_1, t_2,\ldots,t_m$ is a finite set of transitions, $m\geq0$. The set of places and transitions are always disjoint, $P\cap T=\emptyset$; this represents formally that we are dealing with different syntactic types in the network. The mapping from the places to transitions and from transitions to places is undertaken by input ($I: T \rightarrow P\ \infty$) and output functions ($I: T \rightarrow P\ \infty$). These functions map transitions to bags of places. A graphical representation of Petri Nets as in Figures 1 and 2 illustrates their potential usefulness more clearly. A Petri Net graph consists of places as circles, transitions as bars, and arcs, which realize the input and output functions of the transitions and are always directed. Finally, the tokens are considered to be abstract entities and express the procedural nature of Petri Nets by carrying the relevant information and marking the network in the different stages of its execution.

It should be clear by now that a Petri Net is not just a diagram describing the relationships among the objects represented by the nodes. One could adopt other techniques to do this in a very convenient and simple way. An essential feature of Petri Nets is that they can be executed. Following the terminology adopted in Petri Net theory, the steps below express the algorithm for the execution of a Petri Net.

1. An initial marking is defined (by a marking μ we mean an assignment of tokens to the places of a Petri Net).
2. The set of eligible transitions is activated (eligible transitions are the ones whose input functions contain at least one token in their domains).
3. One of the eligible transitions fires and transfers the tokens to the places that belong to the range of its output functions.
4. Step 2 is returned to until all the eligible transitions have fired or else until the final marking state has been reached.

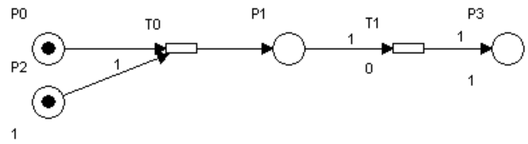

Figure 1: An Initial Marking

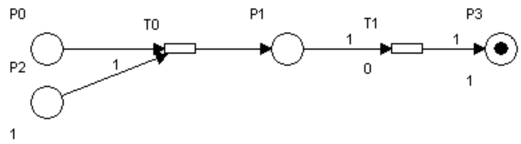

Figure 2: The Final Marking

Before we move on to an illustration of the more complex Petri Nets we explore for the representation of lexical semantics, some explanations of the basic notions just introduced are in order. The Figures in 1 and 2 represent a simple Petri Net in its initial and final marking stages. In what follows, we briefly go through how this Petri Net was executed.

The input function I(p,t) is defined as a mapping $P \times T \rightarrow \{0,1\}$ corresponding to the set of directed arcs from places to transitions and the output functions O(t,p) as a mapping $T \times P \rightarrow \{0,1\}$ corresponding to the set of directed arcs from transitions to places (0 and 1 represent the fact that the function is either realized or not). The structure of the Petri Net in Figure 1, i.e., the quadruple D that defines it, is represented as in (7).

(7) P: $\{p_0, p_1, p_2, p_3\}$
T: $\{t_0, t_1\}$
$I(p_0, t_0) = 1$, $I(p_2, t_0) = 1$, $I(p_1, t_0) = 0$, $I(p_3, t_1) = 0$
$I(p_0, t_1) = 0$, $I(p_2, t_1) = 0$, $I(p_1, t_1) = 1$, $I(p_3, t_1) = 0$
$O(t_0, p_0) = 0$, $O(t_1, p_0) = 0$
$O(t_0, p_2) = 0$, $O(t_1, p_2) = 0$
$O(t_0, p_1) = 1$, $O(t_1, p_1) = 0$
$O(t_0, p_3) = 0$, $O(t_1, p_3) = 1$

The formal representation of the structure of the Petri Net determines which arcs can be activated in the graph. This is based on the knowledge of which connections are allowed, given the structure in (7). A marked Petri Net MΠ is one where we additionally define a marking μ on it. This is a function that maps the members of the set of places p to the nonnegative integers n that represent the number of tokens. In other words, we assign a number of tokens to the places that carry the abstract information. The firing of a transition is achieved in two steps: first of all, the places that lead to the transition (these are all the places that result in a 1 in the input function of the transition) should be marked with at least one token. When this condition is met, the transition is enabled. The enabled transition then consumes the tokens of the 'input' places and, according to the 'power' of the arc, distributes tokens to each one of the appropriate 'output' places. Figure 1 illustrates the initial marking of the Petri Net, in which the places p_0 and p_2 include one token each. During the complete execution of the marked Petri Net, two firings are realized. The first firing is allowed by the enabled transition t_0. Both places p_0 and p_2 need to be marked with at least one token, otherwise the transition is not enabled. If only p_0 or p_2 had a token, then the necessary preconditions for the transition t_0 to be enabled would not have been met and the transition would not fire. Given that in Figure 1 the transition t_0 is properly enabled, it can fire. This results in position p_1 being marked with a token. The second possible firing in the net can now take place since the marking of p_1 enables the transition t_1. Figure 2 illustrates the final marking of the Petri Net, where there is no other possible enabling of a transition.

Now, how does this abstract modeling relate to lexical semantics? The procedural nature of Petri Nets and their usefulness for our purpose of semantic composition can be made clear by thinking of the token game as the consumption of available resources, whereby each new state of a Petri Net represents a step in the process of semantic composition. Sublexical and lexical semantic effects can be represented in terms of abstract tokens in certain initial positions. The interaction of these pieces of (sub)lexical semantic information with one another and with clausal morphosyntactic information as well as contextual information can then be modeled and reliably computed via a Petri Net. That is, the Petri Nets allow a detailed modeling of the subtleties and intricacies of semantic composition that are involved with respect to lexical semantics in general and light verbs in particular. The differing semantic interactions between the bits of semantic knowledge implied by a lexical item can be computed by just enabling one and the same Petri Net (which is meant to represent the lexical semantics of the verb) in different initial markings. These different initial markings correspond to differences in the morphosyntax or in the context.

This lexical semantic perspective on the procedural functioning of Petri Nets has potentially significant import for a treatment of lexical and clausal semantics because it works with fundamentally similar assumptions as *glue semantics*, one type of semantic analysis proposed within LFG. Within glue semantics, semantic composition is driven by the idea that linguistic resources are consumed until a final predication is reached in which all the arguments of the predicates are instantiated (Dalrymple 1999). Glue semantics relies on linear logic and, indeed, linear logic has also been suggested by Engberg and Winskel (1994) as a possible translation of Petri Nets.[3] Under their interpretation of Petri Nets, the multiplicative conjunction $\otimes$ takes as its arguments the places carrying at least one token. Multiplicative conjunction enables a transition in the following way: a given transition is only enabled through the 'conjunction' of the places connected to a transition. If the conjunction holds, the transition can fire. This firing or, better, 'consumption' of tokens is in turn realized by the combination of the multiplicative conjunction and the linear implication $\multimap$. The role of the linear implication is to activate the transitions' function of realizing tokens in the output places. The idea of consuming the appropriate places closed under the multiple conjunction $\otimes$ captures the expressibility of Petri Nets in a simple way.

The marked Petri Net illustrated in Figure 1 can thus be represented by the linear logic formula in (8). Given the interpretation of the multiplicative conjunction as defined above, this formula

[3] Engberg and Winskel (1994) even proposed a small ($\otimes$,$\multimap$) fragment of intuitionistic linear logic.

states that the places p_0 and p_2 enable the firing of a transition that provides us with p_1, which in its turn is consumed and results in p_3.

(8) $((p_0 \otimes p_2) \multimap p_1) \multimap p_3$

This connection of linear logic to Petri Nets concludes the discussion of the basics of Petri Nets. The next section moves on to an illustration of concrete examples with respect to the light verb *de* 'give'. In particular, we use places to represent bundles of semantic properties. Differences in the initial marking of the Petri Net capture the different ways the light verb can modify the event of the main predication and successfully express its subtle semantic contribution to the complex predicate.

4 Sample Analysis

Recall from section 2 that the data with respect to light verbs and their form-identical full verb counterparts points to a very intimate connection between the two. We propose to model this interconnectedness by assuming a single underspecified lexical entry from which both the main and the light verb readings can be derived. The question posed in section 2 was: how can this single underspecified entry be represented? The answer proposed in this section is that a representation in terms of Petri Nets can serve to clarify our ideas on the necessary representational components and their interactions with one another and with clausal and discoursal information.

Most of the familiar approaches to lexical semantics involve some kind of lexical decomposition (e.g., Dowty 1979, Jackendoff 1990, Hale and Keyser 2002). However, lexical decomposition is not at all helpful when trying to come to grips with contextually dependent meaning dimensions like benefaction or forcefulness. Cognitve Semantics (e.g., Newman 1996) appears to be more promising in this regard (see the discussion in Butt and Geuder 2001), but is ultimately not suitable for computational purposes. As a first approximation towards finding the right kind of underspecified lexical entry, we therefore decided to work with Dowty's (1991) Proto-Role entailments.

Dowty (1991) formulated a number of entailments which follow from the lexical semantics of a verb that could help with the identification of an argument as either a Proto-Agent, a Proto-Patient, or neither. As a reminder to the reader, Dowty's entailments are reproduced in (9).

(9) **Dowty's Proto-Role Entailments**

Proto-Agent

a. volitional involvement in the event or state
(Ex.: Kim in *Kim is ignoring Sandy.*)

b. sentience (and/or perception)
(Ex.: Kim in *Kim sees/fears Sandy.*)

c. causing an event or change of state in another participant
(Ex.: loneliness in *Loneliness causes unhappiness.*)

d. movement (relative to the position of another participant)
(Ex.: tumbleweed in *The tumbleweed passed the rock.*)

e. (exists independently of the event named by the verb)
(Ex.: Kim in *Kim needs a new car.*)

Proto-Patient

a. undergoes change of state
(Ex.: cake in *Kim baked a cake.*, error in *Kim erased the error.*)

b. incremental theme
(Ex.: apple in *Kim ate the apple.*)

c. causally affected by another participant
(Ex.: Sandy in *Kim kicked Sandy.*)

d. stationary relative to movement of another participant
(Ex.: rock in *The tumbleweed passed the rock.*)

e. (does not exist independently of the event, or not at all)
(Ex.: house in *Kim built a house.*)

Although some of Dowty's definitions/assumptions are problematic and more could be said on this matter, our interest for the moment is not in trying to go beyond Dowty's insights, but to investigate whether this way of looking at lexical semantics can help with modeling the semantic effects of light verbs.

Given the conclusion that a main verb and its corresponding light verb should be derived from the same underlying entry, one desideratum is that the underlying entry allow for the kind of flexibility in which either a full argument structure is instantiated (main verb reading), or where only some event modificatory meaning dimensions such as benefaction, control, or completion are instantiated (light verb reading). We think that this can be achieved by assuming that the lexical entry by itself only provides a bundling of properties which are typical for the kind of event that is described. That is, the lexical semantic representation primarily consists of an unordered collection of semantic properties that are akin to Dowty's Proto-Role entailments. This bundle of properties is only structured into the familiar lexically decomposed structures (e.g., an LCS) in interaction with syntactic properties that require a main verb predication.[4] When there is no call for a main verb predication, i.e., when the syntactic environment does not allow for one, the bundle of semantic properties is realized in terms of an event modificatory semantics. That is, the meaning dimensions are applied to modify the event semantics of the main verb in the clause.

The screen shots of working Petri Nets in Figures 3 to 6 illustrate this basic idea with respect to the light verb *de* 'give'.[5] As shown in Figures 3 and 5, the lexical entry for 'give' initially only consists of semantic properties such as "volitional involvement", "causation" (both derived from Dowty's Proto-Agent entailments), "change of state" (derived from Dowty's Proto-Patient entailments) and a general "change of location" (goal) component.

The fundamental property of Petri Nets that we make use of is that certain transitions can only be enabled if all of the places leading to that transition are 'armed'. That is, the armed places represent the necessary preconditions for a certain transition to fire. In terms of linear logic, these preconditions can be interpreted as the *resources* from which semantic conclusions are drawn. The idea in Figures 3–6 is that the underlying lexical semantics of the verb are invariant. The unordered bundle of lexical semantic properties is realized as tokens in the places p_0, p_1, p_{17} and p_2. Figures 3 and 5 differ in the additional kinds of resources. In Figure 3, for example, the initial markings indicate that the verb *de* 'give' is the sole verb in a clause with three NPs (p_{11}, p_{12} and p_{13}). In such a configuration, an execution of the Petri Net results in the final marking shown in Figure 4.

[4]Note that this idea shares many features with newer work by Marantz 1997 and Borer 2003 by which the syntactic type (and predicational power) of a lexical item is only determined once it has been inserted into syntax.

[5]The Petri Nets were modeled with the help of *PetrA* (http://computacion.cs.cinvestav.mx/˜ameneses/PetraPag/petra.html), one of the few Petri Net implementations that are compatible with a Macintosh operating system.

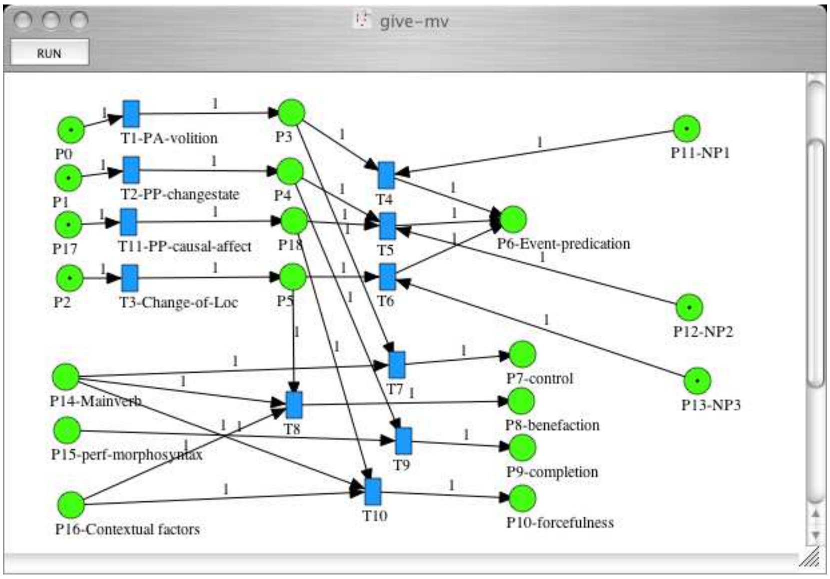

Figure 3: Underspecified entry for 'give', in main verb context (Initial)

In Figure 4, the transitions t_1, t_2, t_{11} and t_3 were all enabled and therefore all fired and distributed tokens to p_3, p_4, p_{18}, p_5. These places, in turn, in conjunction ($\otimes$) with the places p_{11}, p_{12} and p_{13} served to enable the transitions in t_4, t_5 and t_6, which then fired and led to a marking of the net in which p_6 is marked. The final marking of the Petri Net, where p_6 is the only place marked, is interpreted as the predication of the event semantics of a main verb with three arguments. These three arguments have the Proto-Role properties of volition (Proto-Agent), change-of-state (Proto-Patient), causal affectedness (Proto-Patient) and change of location (neither Proto-Agent or Proto-Patient, but a third argument).

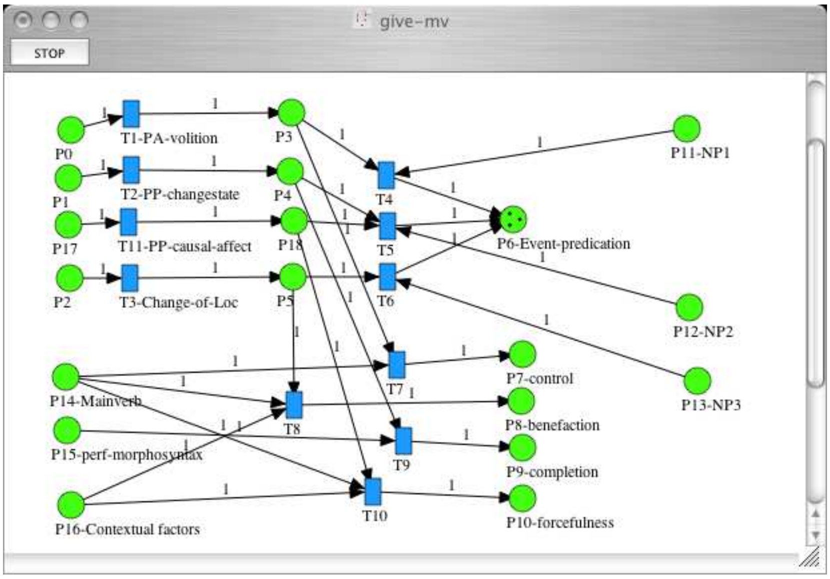

Figure 4: Interpretation for 'give', in main verb context (Final)

Contrast this with the initial markings that lead to a light verb interpretation. This is shown in Figure 5, where the places p_0, p_1, p_{17} and p_2 contain tokens representing the lexical semantics of *de* 'give', just as in Figure 3. The difference lies in contextual activation. The initial marking in Figure 5 indicates that there already is a main verb in the clause, it takes into account perfective morphology on the main verb (this is how the telic reading is actually licensed, see Butt and Ramchand 2003),[6] and it allows for contextual factors that would license the semantic dimensions of benefaction or forcefulness, for example.

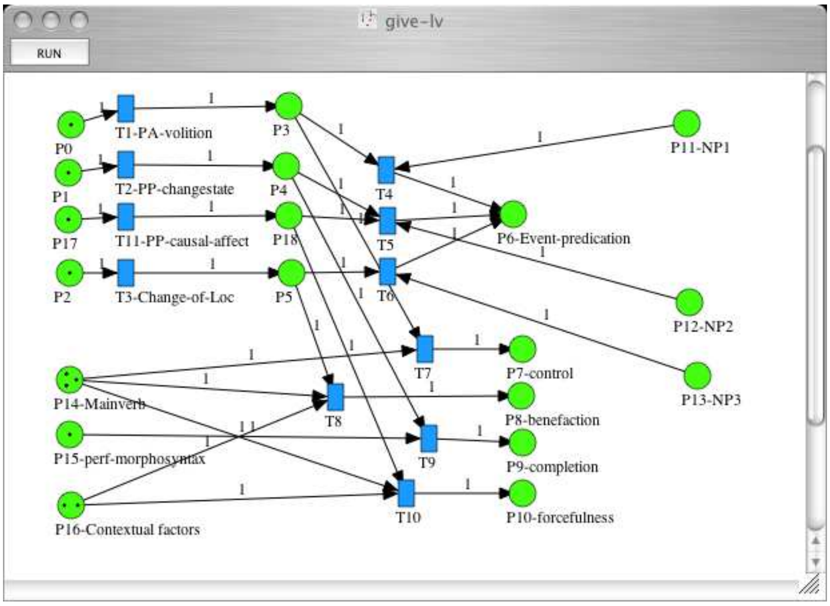

Figure 5: Underspecified entry for 'give', in light verb context (Initial)

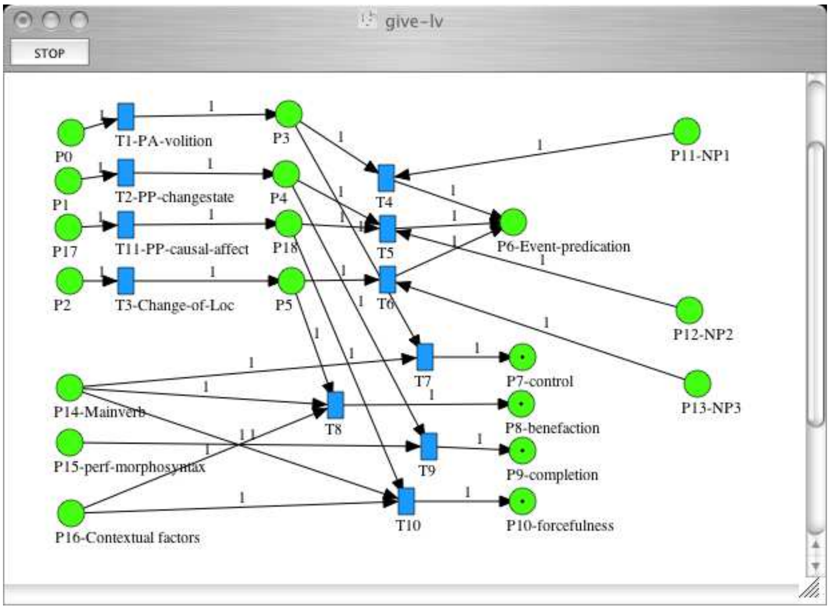

Figure 6: Interpretation of 'give', in light verb context (Final)

In Figure 5, the transitions t_1, t_2, t_{11} and t_3 were enabled and therefore fired. However, given that the positions p_{11}, p_{12} and p_{13} do not contain tokens in the initial marking in Figure 5, the transitions leading to a full event predication were not enabled and could not fire. Instead, the transitions t_7, t_8, t_9 and t_{10} were enabled and fired, resulting in the marking illustrated in Figure 6.

[6]Note that the "perfective" morphology this is taking into account was lost recently in Urdu/Hindi. Bengali, in contrast, still has an overt *-e* marking on the main verb. Discussing this issue in detail would take us too far afield, but see Butt and Ramchand 2003 for a detailed discussion.

Figure 6 thus shows a situation in which all of the meaning dimensions of *de* 'give' have been activated: control (p_7), benefaction (p_8), completion (p_9) and forcefulness (p_{10}). The predication of these meaning dimensions is not random, but must be seen as deriving or being licensed by the collection of initial lexical semantic properties. So, for example, the completion dimension that is the hallmark of transitive light verbs in Urdu can be derived from the change of state (Proto-Patient) property, control is licensed by volition (Proto-Agent), forcefulness by causal affectedness, and benefaction by the change of location.

Thus, the precise interpretation of the initial bundle of lexical semantic properties depends on the clausal and discoursal context. If 'give' were the sole verbal predicate in a clause, as in Figure 3, then the activated bundle of features could be realized as thematic roles if and only if the clausal (or discoursal in case of pro-drop) context licensed the arguments. In this case, the NP states contain tokens and the appropriate transitions (agent, patient, goal) are able to fire, leading to an event predication. On the other hand, if the 'give' is found in conjunction with another verbal predicate, it can be interpreted only as predicating a collection of semantic properties that must be applied to the main event predication of another verb. We assume with Butt and Geuder (2001) that these meaning components are analogous to adverbial modification.

A very nice feature of the Petri Net approach sketched here is that it fulfills exactly the requirements identified in section 2: 1) main and light verb readings should be derived from the same underlying lexical entry (indeed, *the same underlying lexical semantics*); 2) not all of the meaning dimensions supported by a light verb should necessarily always be enabled.

The first point has been illustrated with respect to the abstract situations represented by Figures 3–6. The second point can be illustrated with respect to the example shown in (10). This is a sentence uttered by a man in a context where the woman he has been promised to by his parents (arranged marriage) releases him from his obligation. As she explains that she realizes he does not love her (while she will always love him), he is moved almost to tears and he begs her to stop talking because otherwise he will have to cry.

(10) mɛ̃ ro d-ũ-g-a
I.Nom cry give-1.Sg-Fut-M.Sg
'I will cry.' (from the movie *Kabhi Khushi Kabhi Gham*)

In this situation, the light verb *de* 'give' carries no sense of benefaction or forcefulness. Rather, the use of 'give' overrides the default assumption that crying is more of an involuntary, rather than a controlled action (our hero is manfully controlling his crying, but might still decide to do it anyway). Again, this can be modelled via Petri Nets as shown in Figures 7–8. An execution of the net results in Figure 8. Here again the transitions t_1, t_2, t_{11} and t_3 were all enabled and therefore fired. However, of the transitions t_7, t_8, t_9 and t_{10}, only t_7 and t_9 were enabled and could fire. The tokens in positions p_{18} and p_5 are "stuck" and can be viewed as resources that were made available by the underlying lexical semantics of the verb, but which could not be consumed.[7]

Since these resources could not be consumed, no predication in terms of the semantic dimensions of benefaction and forcefulness is possible. Instead, the semantic import of the light verb in this situation is limited to contributing the information that the crying event was a controlled one and

[7] Note that this does not quite fit with the conception of glue semantics as articulated in Dalrymple 1999, whereby all resources must be consumed as part of arriving at a well formed semantic representation. Under our proposal, Petri Nets could be viewed as a procedural model of meaning whereby bits of semantic knowledge are consumed depending on the context. That is, we assume different stages in the procedural model. In an initial stage all the bits of semantic knowledge that are available are dealt with in some form or another and are "readied" for further possible semantic composition through the activation of p_3, p_4, p_{18} and p_5. In the next stage, only those bits of information which are suitable for consumption actually enter the semantic composition, resulting only in the marking of p_7 and p_9 in Figure 8. A more precise articulation of the relationship between the Petri Net model and glue semantics in this case, however, must be left to further work.

that it would have been completed, had our hero really commenced crying (recall that the utterance is in the future tense).

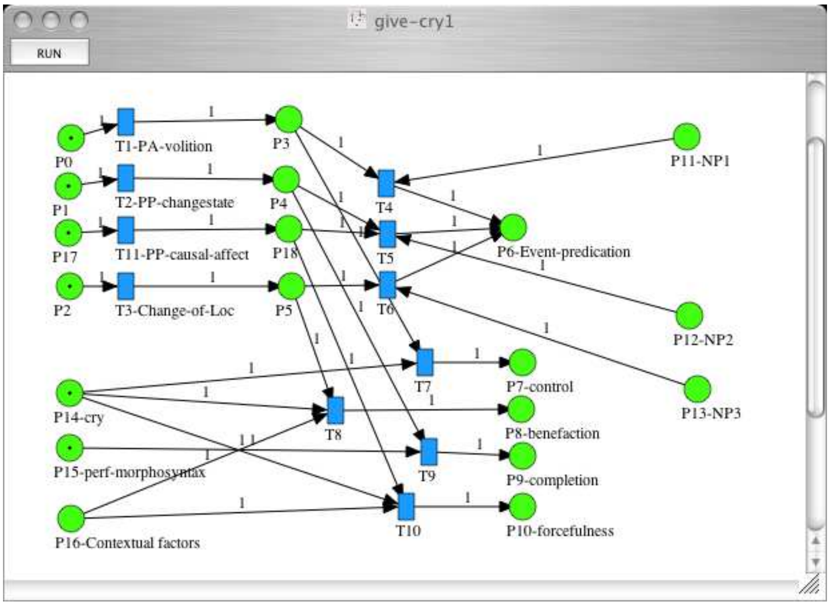

Figure 7: Underspecified entry for ‘give’, in light verb context with ‘cry’ (Initial)

In Figure 7, the initial marking for the underlying lexical semantics of ‘give’ is again the same. And, again, it is the contextually supplied initial information that differs. So, in its initial state, the Petri Net representing the analysis of (10) is marked with the information that there is a main verb ‘cry’. Since ‘cry’ only has one argument which could perform a volitional action, it is associated with transition t_7.

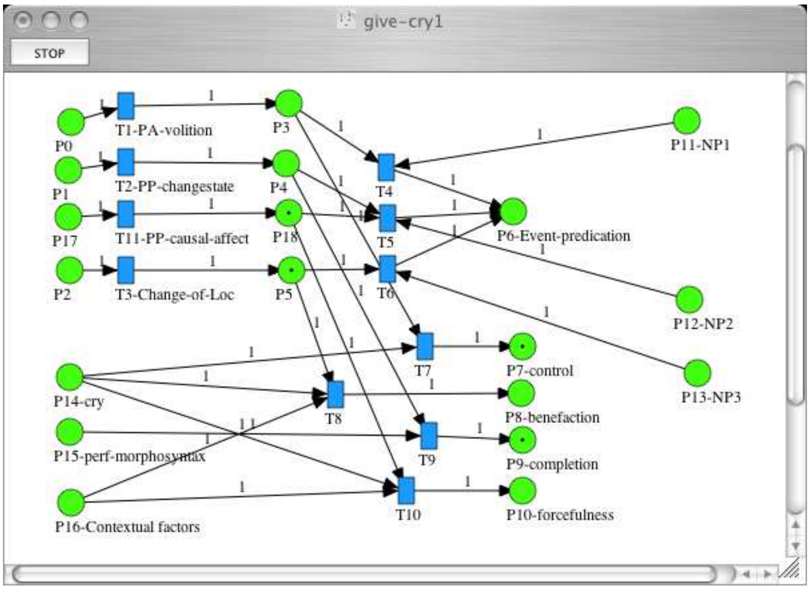

Figure 8: Underspecified entry for ‘give’, in light verb context with ‘cry’ (Final)

This concludes the presentation of sample lexical semantic analyses in terms of Petri Nets. There are several other problems to tackle, such as noun-verb constructions (e.g., *give a wash*) and why not all verbs allow light verb meanings. However, our initial explorations have shown that the Petri Net model can deal with some difficult problems with respect to the semantics of complex predicates quite nicely. Petri Nets would thus seem to provide a promising model of lexical semantic representation, particularly in interaction with context related factors.

5 Conclusions

In conclusion, Petri Nets not only provide an intuitive model for the encoding of lexically underspecified information, they also allow for the necessary flexibility in semantic interpretation. As was shown above, the meaning dimensions of benefaction and forcefulness with respect to the light verb use of *de* 'give' in Urdu are highly context dependent. The relevant transition will only be activated if a corresponding token has been found in the discourse context (encoded as *contextual factor* in the Petri Nets).[8] When the "right" context for benefaction and/or forcefulness is not available, then these meaning dimensions cannot be computed. That is, they are resources made available by the lexical semantics of 'give' that are not always necessarily consumed as part of the computation. The resource sensitivity inherent to Petri Nets is highly reminiscent of glue semantics within LFG (Dalrymple 1999). And indeed, the version of Petri Nets we have explored in this paper has been shown to be equivalent to the version of linear logic assumed within glue semantics. As such, the insights gained from modeling lexical semantics within Petri Nets should be directly translatable into LFG's glue semantics. This remains to be done in future work.

References

Alsina, Alex. 1996. *The Role of Argument Structure in Grammar*. Stanford, California: CSLI Publications.

Alsina, Alex, Joan Bresnan, and Peter Sells (ed.). 1997. *Complex Predicates*. Stanford, California: CSLI Publications.

Beg, Mirza Khalil A. 1988. *Urdu Grammar: History and Structure*. New Delhi: Bahri Publications.

Borer, Hagit. 2003. *Structuring Sense: An Exo-Skeletal Trilogy*. Oxford: Oxford University Press. In Press.

Bresnan, Joan. 1982. Control and Complementation. In *The Mental Representation of Grammatical Relations*, ed. Joan Bresnan. 282–390. Cambridge, Massachusetts: The MIT Press.

Butt, Miriam. 1994. Machine Translation and Complex Predicates. In *Proceedings of KONVENS 94*, 62–71.

Butt, Miriam. 1995. *The Structure of Complex Predicates in Urdu*. Stanford, California: CSLI Publications.

Butt, Miriam. 1998. Constraining Argument Merger Through Aspect. In *Complex Predicates in Nonderivational Syntax*, ed. A. Kathol E. Hinrichs and T. Nakazawa. 73–113. New York, New York: Academic Press.

Butt, Miriam, and Wilhelm Geuder. 2001. On the (Semi)Lexical Status of Light Verbs. In *Semi-lexical Categories: On the content of function words and the function of content words*, ed. Norbert Corver and Henk van Riemsdijk. 323–370. Berlin: Mouton de Gruyter.

[8]Note that in languages like Korean, benefaction is always entailed by the light verb 'give', regardless of context. In these languages, no contextual enabling would be required in order for the benefaction meaning component to flow into the semantic interpretation of the clause.

Butt, Miriam, Tracy Holloway King, and John Maxwell III. 2003. Complex Predicates via Restriction. In *On-line Proceedings of the LFG03 Conference*. CSLI Publications. University at Albany, State University of New York.

Butt, Miriam, and Aditi Lahiri. 2004. Verbal Passepartouts. Unpublished Ms.

Butt, Miriam, and Gillian Ramchand. 2003. Complex Aspectual Structure in Hindi/Urdu. In *The Syntax of Aspect*, ed. Nomi Ertishik-Shir and Tova Rappaport. Oxford: Oxford University Press. To appear.

Dalrymple, Mary (ed.). 1999. *Semantics and Syntax in Lexical Functional Grammar*. Cambridge, Massachussetts: The MIT Press.

Dowty, David. 1979. *Word Meaning and Montague Grammar*. Dordrecht: Reidel.

Dowty, David. 1991. Thematic proto-roles and argument selection. *Language* 67(3):547–619.

Engberg, Uffe, and Glynn Winskel. 1990. Petri Nets as Models of Linear Logic. In *CAAP '90, 15th Colloquium on Trees in Algebra and Programming*. Copenhagen, Denmark.

Grimshaw, Jane, and Armin Mester. 1988. Light Verbs and Θ-Marking. *Linguistic Inquiry* 19(2):205–232.

Hale, Kenneth, and Samuel Jay Keyser. 2002. *Prolegomenon to a Theory of Argument Structure*. Cambridge, Massachusetts: The MIT Press.

Hendriksen, Hans. 1944. *Syntax of the Infinite Verb Forms of Pali*. Copenhagen: Einar Munksgaard.

Hook, Peter Edwin. 1974. *The Compound Verb in Hindi*. Center for South and Southeast Asian Studies: The University of Michigan.

Hook, Peter Edwin. 1991. The Emergence of Perfective Aspect in Indo-Aryan Languages. In *Approaches to Grammaticalization*, ed. Elizabeth Traugott and B. Heine. 59–89. Amsterdam: Benjamins.

Hook, Peter Edwin. 1993. Aspectogenesis and the Compound Verb in Indo-Aryan. In *Complex Predicates in South Asian Languages*, ed. Manindra Verma. 97–113. Delhi: Manohar Publishers & Distributors.

Hook, Peter Edwin. 2001. Where do Compound Verbs Come From? (And Where are They Going?). In *The Yearbook of South Asian Languages and Linguistics*, ed. Peri Bhaskararao and K.V. Subbarao. 101–130. New Delhi: Sage Publications.

Iglesias-Rábade, Luis. 2001. Composite Predicates in Middle English with the verbs *nimen* and *taken*. *Studia Neophilologica* 73:143–163.

Jackendoff, Ray. 1990. *Semantic Structures*. Cambridge, Massachusetts: The MIT Press.

Kaplan, Ronald, and Jürgen Wedekind. 1993. Restriction and Correspondence-based Translation. In *Proceedings of the Sixth European Conference of the Association for Computational Linguistics*, 193–202.

Kellogg, S. H. 1893. *Grammar of the Hindi Language*. Delhi: Munshiram Manoharlal Publishers Pvt. Ltd. Second Edition, reprinted 1990.

Marantz, Alec. 1997. No escape from syntax: Don't try morphological analysis in the privacy of your own lexicon. In *UPenn Working Papers in Linguistics*, ed. L. Siegel and A. Dimitriadis. 201–225. University of Pennsylvania.

McGregor, R.S. 1968. *The Language of Indrajit of Orchā*. Cambridge: Cambridge University Press.

Mohanan, Tara. 1994. *Argument Structure in Hindi*. Stanford, California: CSLI Publications.

Newman, John. 1996. *Give: A Cognitive Linguistic Study*. Berlin: Mouton de Gruyter.

Peterson, John L. 1981. *Petri Net Theory and the Modeling of Systems*. London: Prentice Hall.

Petri, Carl Adam. 1962. *Kommunikation mit Automaten*. Bonn: Institut für Instrumentelle Mathematik.

Ritter, Elizabeth, and Sara Thomas Rosen. 1993. Deriving Causation. *Natural Language and Linguistic Theory* 11:519–555.

Trenckner, V. 1879. *Pali Miscellany*. London: Williams & Norgate.