
Events and the Semantics of Opposition

JAMES PUSTEJOVSKY

13.1 Persistence and Change

13.1.1 Introduction

There has recently been a renewed interest in the explicit modeling of events in the semantics of natural language. This is more evident now than ever before, particularly with the interest in explaining the properties of syntactic linking in languages in terms of the event representations underlying sentential forms. The papers in this volume, for instance, are examples of this recent line of discussion. Most of this work assumes a logic of interpretation where events are associated with the tensed matrix verb of a sentence and sometimes with event-denoting nominal expressions, such as *war* and *arrival*. There has, however, been little serious discussion in the semantics literature of the logical consequences of adopting a stronger view of quantification over events in language, where the event structure representation makes explicit reference to object and property persistence for all the logical arguments in the sentence, and not merely the classical “theme” argument.¹

Typically, even with a binary event structure, the predicates associated with the individual events make reference to a unique change. For example, change-of-state predicates such as *break* and *die* affect a single argument position, as with *the glass* and *John*, respectively, in (1) below:

- (1) a. Mary broke the glass.
b. John died.

In the course of an activity or an event, however, the major predication

¹This remark also holds for the literature on object-event quantification and event plurality; see, for example, Link (1998), Krifka (1989), Barker (1999), and Schein (1993).

in a sentence may do more than change the properties associated with a single argument, or it may alternatively affect only part of the predicative content of the referring expression denoting the argument. Consider, for example, the sentences in (2).

- (2) a. Mary fixed every leaky faucet in the house.
 b. John mixed the powdered milk into the water.
 c. The child ate a cookie.

For (2a), the effect of the activity of fixing the faucets will render the description *leaky* applied to each faucet as contradictory; similarly, in (2b), the felicity of the description *powdered* applied to the milk is true only before the completion of the event of mixing and not after. Finally, for sentence (2c), the eating activity effectively makes quantification over the sortal term *cookie* impossible, if subeventual predicates are used to represent the overall event structure, as in Parsons (1990) and Pustejovsky (1991b).²

I will argue that a finer model of change is needed within semantic theory to handle such phenomena, than has conventionally been adopted. Furthermore, this model must incorporate the properties of *persistence* over the event as well as change. To characterize the persistence of objects and properties, the model developed here will make reference to the various modes of *predicate opposition* over objects and properties. We will study how opposition structure, lexically encoded in the semantics of verbs, adjectives, and nouns, is syntactically realized in the sentence.

It is usually assumed in most linguistic discussions of events that the expression should denote what effects follow from the event's occurrence: in other words, what changes have taken place as a result of this event happening. However, depending on how these changes are represented relative to the quantificational force of the sortal terms, this may in fact introduce contradictory information into the resulting logical form of the sentence. As we saw above in (2), some of the descriptive content of the referring expressions in a sentence does not persist as a result of the assertion of the changes invoked by the matrix predication. More examples illustrating this point are shown below in (3):

²Kowalski and Sergot (1986) and others have examined the problems associated with creating coherent stories when there are factive events that denote contradictory states-of-affairs, when not interpreted within a temporal model as in (i) below. In AI, these concerns are addressed in the context of solutions to the frame problem and limiting the application of reasoning mechanisms in restricted domains.

- (i) a. Mary was hired as lecturer on Tuesday.
 b. John left as CEO in November.
 c. Mary was fired as lecturer.
 d. Bill was promoted from programmer to director.

- (3) a. The father comforted the crying child.
 b. The woman on the boat jumped into the water.
 c. Mary rescued the drowning man.

In order to arrive at a representation that expresses both the change and the persistence of arguments in the sentence, I introduce the notion of a *gating function*. Viewed informally, these are linguistic expressions that operate on one or more arguments in a construction, acting to initiate or terminate a property of the object denoted by that argument. For example, the verb *die* in “A man died.” is a gating function for the animate individual denoted by the NP *a man*; that is, the predication of animacy of the man is de-activated by virtue of the proposition asserted by the sentence. In (3), the referring expressions “the crying child”, “the woman on the boat”, and “the drowning man” are *gated* by the very event they are participants in. What is important to notice about these examples is that the head terms (sortals) of the internal arguments in (3a,c) and the subject in (3b) are not themselves gated, but only the adjectival modifiers in (3a,c), and the stative description from the locative PP in (3b). For quantificational purposes, it is as though the descriptive content of these NPs has been split apart. I will refer to such cases as “contradictions of change”.

This illustrates only one way in which events have local *secondary effects* on other predicative structures in the sentence, as a result of that event’s occurrence. Studying the properties of gating functions helps us understand what secondary effects are relevant for computing the “maximally coherent event description” of a sentence. Furthermore, without knowing what local persistence and anti-persistence effects are at play in the computation of the meaning for an expression, it is impossible to create a coherent model of the event structure as it impacts all the arguments to a verb.

In the following sections, I reexamine the arguments for event individuation in the logical form for natural language, from the perspective of the issues presented above. After giving a brief treatment of opposition types, we then see how the semantics of opposition can be embedded into the event structure of a sentence. Next, I define the tool set we will need in order to address the contradictions of change mentioned above. This includes a discussion of the *Principle of Property Inertia* in natural language semantics and an analysis of the manner in which adjectives modify their heads as event predicates. I will then sketch out an algorithm that computes this representation from the semantics of the individual event expressions associated with a sentence; the result of this calculation will be called the *event persistence structure (EPS)* for that sentence. Finally, I demonstrate how the *EPS* representation can be computed for a variety of cases and extended to handle the treatment of stage-level nominals under changing conditions.

13.1.2 The Individuation of Events

Solid evidence for the individuation of events in natural language is intimately tied to the kinds of inferences that we expect to be able to make from the resulting logical expressions being constructed. The deeper that lexical semantics digs into the meanings of predicates and the more discourse semantics builds upward toward larger textual units, the closer linguistic semantics comes to facing the frame problem and the problems of circumscribed inference as studied in AI and logic (Hanks and McDermott, 1986, Harman, 1986, Morgenstern, 1995). What is the linguistic interface to commonsense reasoning over entities, properties, and relations? Is there a clean modular separation between language semantics and general reasoning? Most logicians and AI researchers would argue strongly against such a neat division and these questions still loom large over the study of meaning and language, and they are well beyond the scope of this paper. I will, however, offer some initial observations on how knowledge of change and persistence is computable from linguistic representations alone, as this pertains to events structures in language.

We begin our discussion by illustrating simple distinctions in event representations which facilitate inferences. Consider the two sentences below. For many reasoning situations, both sentences (4a) and (4b) could adequately be expressed as single atomic events, as shown in (5a) and (5b):

- (4) a. John kissed Mary.
 b. John painted a house.
- (5) a. $\exists e[kiss(e, j, m)]$
 b. $\exists x\exists e[paint(e, j, x) \wedge house(x)]$

For example, it may be the case that I require reference only to the assertable knowledge of kissings and house-paintings, without necessarily understanding their consequent states. The granularity of the representation—and in this case, the granularity of the event descriptions—is intimately linked to the requisite inferential demands and capabilities of the model that interprets the representation. Hence, updating the persistence of properties of arguments in a sentence could proceed in a number of ways. For example, the entailments from (4b) relative to the changed state of the house (i.e., being painted) can follow from meaning postulates associated with the matrix verb *paint*, as done in classical Carnapian Montague semantics (e.g., Dowty, 1979) and shown in (6a), or could follow directly through a richer event construction associated with the predicate *paint* (cf. (6b)), as accomplished in Generative Lexicon Theory (Pustejovsky, 1991a, 1995).

- (6) a. $\forall x, e, y \forall P \Box [(paint(e, y, x) \wedge P(x)) \rightarrow painted(x)]$
 b. $\exists x \exists e_1 \exists e_2 [paint_act(e_1, j, x) \wedge house(x) \wedge painted(e_2, x) \wedge e_1 < e_2]$

Both options are potentially adequate, depending on the specific reasoning tasks and associated model restrictions. Adopting the subeventual analysis for a sentence such as (6b) does have the advantage that it explicitly refers to the change inherent in the predicate, a state which may be useful or even necessary for subsequent reasoning.

Now consider the pair of sentences in (7) below.

- (7) a. John slept.
b. John ate a cookie.

For sentence (7a), the most direct event representation is simply the one given in (9):

$$(8) \exists e[\textit{sleep}(e, j)]$$

For (7b), things are a bit more complicated. Obvious semantic considerations tell us that there was something that John was eating, after which there is not that thing which John ate. Furthermore, John remains persistent throughout the event, relative to the assertion carried by the proposition. One question that immediately arises is that of how the relative persistence of the arguments is modelled, if at all, in the event structure for a sentence. Even a simple example such as (7b) above points out the tension in current hybrid treatments of event semantics. The representation in (9) is correct in a model that ignores the decompositional entailments inherent in the predicate *eat*.

$$(9) \exists x \exists e[\textit{eat}(e, j, x) \wedge \textit{cookie}(x)]$$

Once telicity and causation are associated with the internal structure of events, however, the semantics of change (and in this case, destruction) is built directly into the predication within the event semantics, making simple event and individual quantification as in (9) inadequate (cf. Parsons, 1990, Pustejovsky, 1991b). This is one reason why Dowty (1979) adopts a classical framework of predicate decomposition without events: he avoids this problem directly (see the discussion in the introduction to this volume). The difficulty, however, is that the specific entailments associated with a proposition must be built into the meaning postulates associated with lexical items and how they are compositionally deployed in a sentence.

The sentences in (4) and (7) illustrate how event individuation is associated with the matrix predicate of the sentence. The literature has discussed, of course, many other constructions that introduce individuated event interpretations into the resulting logical form for a sentence (cf. Alsina, 1999, Mohanan and Mohanan, 1999, Kratzer, 1995). Such constructions include resultatives (10), depictives (12), and other adjectival phrase adjuncts, as well as perception verb complementation (Higgin-

botham, 1985), and control constructions (ter Meulen, *this volume*).

- (10) a. John painted the white house blue.
 b. Mary cut her hair short.

The adjunct position instantiated by the adjectival phrase *blue* can be seen as denoting an individual terminus event, resulting from the completion of the individuated painting event; we might even think of (10a) as introducing a specific binding of the second event from (6b), as illustrated below:

$$(11) \exists x \exists e_1 \exists e_2 [paint_act(e_1, j, x) \wedge house(x) \wedge blue(e_2, x) \wedge e_1 < e_2]$$

Depictives, on the other hand, individuate states overlapping the event denoted by the matrix predication (cf. Rapoport, 1993).

- (12) a. Mary arrived in Boston drunk.
 b. John drank the whiskey undiluted.

The event structure for (12a), for example, is illustrated below:

$$(13) \exists e_1 \exists e_2 \exists e_3 [arrive_act(e_1, m, x) \wedge in(e_2, m, boston) \wedge drunk(e_3, m) \wedge e_1 < e_2 \wedge e_3 \circ \{e_1, e_2\}]$$

Further evidence for event individuation that is not associated with a matrix predicate alone can be illustrated with the classic paradigm of adjectival, inchoative, causative forms for stems such as *close*, as shown in (14)³:

- (14) a. Frances closed a window.
 b. A window closed.
 c. A window was closed.

For each of these sentences, what is common is that the resulting persistent state of the overall event denotes that a window is closed after a process of closing:

$$(15) \exists e_1 \exists e_2 \exists x [close_process(e_1, x) \wedge window(x) \wedge closed(e_2, x) \wedge e_1 < e_2]$$

The above strategy for event decomposition has proved useful in explaining the mapping from lexical semantic forms to predicate argument structure linking (cf. Grimshaw, 1990, Grimshaw and Vikner, 1993, Tenny, 1993, 1994, Ritter and Rosen, 1994). As a research programme, it has been integrally tied to work in aspect, but also to compositional and word formation processes. What is not represented consistently in the event structures above, however, is an explicit predicate opposition indicating the exact nature of the change of state, transformation, creation, destruction, and so

³In Pustejovsky and Busa (1995), it is claimed that two of the three forms below, (14a,b) have the same underlying event structure. The causative and inchoative forms share the same underspecified semantic form, and they are semantically and syntactically distinguished by virtue of an event focusing mechanism called *headedness*.

on, of the verbal argument. Often, an event is reified in the event structure largely for grammatical mapping purposes, while the logical consequences of such a reification are not worked through. If the concerns from the previous section are to be taken seriously, then we must move to a richer model of event structure, where change and persistence for all logical arguments in the sentence are modelled explicitly in order to arrive at a coherent event interpretation for a sentence.

By returning to the fundamental motivations for the existence of subeventual structure for natural language predicates, we hope to better understand how much of the verbal semantics is reflected in the event structure directly. I begin by articulating briefly the types of persistence and change that are expressed in natural language. We will see that there are two major parameters impacting the semantics of opposition as it is expressed in natural language:

- (16) a. What mode of opposition the predicate expresses;
 b. What aspect of the qualia structure the predicate operates over;

In the next section, we introduce the modes of semantic opposition and how they can be integrated in the event structure directly. This will provide us with another component needed for the computation of maximal coherence over event descriptions, as represented in the event persistence structure.

13.1.3 Modes of Opposition

In this section, we explore briefly how verbs express different types of change, and how this change is predicated of distinct and identifiable aspects of the entity undergoing the change. The basic framework of semantic analysis I will assume is Generative Lexicon Theory as outlined in Pustejovsky (1995, 1998), and it is the qualia structure which in part allows us to express these modes of change in an object.

In Generative Lexicon Theory (henceforth GL) it has been assumed that an essential component of semantic interpretation and composition is the manner in which predication is distributed over a complex event structure representation. The predicative force of a single relation or predicate (such as *build* or *die*) is distributed into distinct subpredicates, which are structurally positioned within an event tree annotated with temporal constraints. For example, rather than a neo-Davidsonian single event-place interpretation for the verb *build*, i.e., $\lambda y\lambda x\lambda e[build(e, x, y)]$, there are subevents which are associated with special subpredicates, each of which corresponds to some logical portion of the verb's meaning. Thus, for any predicate P , imagine that there are as many subpredicates, P_i as there are subevents that are distinguished in the event structure. This is illustrated below for the verb *build*.⁴

⁴This is equivalent to treating the predicate as a relation between events, e.g.,

$$(17) \lambda y \lambda z \lambda x \lambda e_2 \lambda e_1 [build_1(e_1, x, z) \wedge build_2(e_2, y) \wedge e_1 < e_2];$$

The number and nature of these subpredicates is inherently restricted by the *qualia structure*. The qualia are an interpretation of the Aristotelian “modes of explanation” for an entity or relation (Moravcsik, 1975), positioned within a type logic as defined in Pustejovsky (1995) with the following characteristics:

FORMAL: the basic category which distinguishes it within a larger domain;

CONSTITUTIVE: the relation between an object and its constituent parts;

TELIC: its purpose and function;

AGENTIVE: factors involved in its origin or “bringing it about”.

For relations, the qualia act in a capacity similar to thematic roles, where the individual qualia are possibly associated with entire event descriptions and not just individuals. For example, the qualia structure for the *constitutive causative* verb *build* in (17) divides into an initiating activity (the AGENTIVE) and a culminating stative terminus (the FORMAL). Hence, we can refer to *build*'s qualia structure as the pair [A,F], abstracting away the qualia values.⁵

The qualia structure is only one of three aspects of the lexical structure of a word which impacts the mapping of semantic information to syntax⁶:

- (18) a. ARGUMENT STRUCTURE: The specification of number and type of logical arguments.
 b. EVENT STRUCTURE: The identification of the event type of an expression and its subeventual structure.
 c. QUALIA STRUCTURE: A structural differentiation of the predicative force for a lexical item.

Simplifying slightly the formalism introduced in Pustejovsky (1995), the argument structure for the verb *build* can be represented as (19) below.

$$(19) \left[\begin{array}{l} \mathbf{build} \\ \text{ARGSTR} = \left[\begin{array}{l} \text{ARG}_1 = \mathbf{x:animate_individual} \\ \text{ARG}_2 = \mathbf{y:artifact} \\ \text{D-ARG}_1 = \mathbf{z:material} \end{array} \right] \end{array} \right]$$

$\lambda e_2 \lambda e_1 [build(e_1, e_2)]$. On this view, we would be able to refer to predicates by their event arity directly, should there be motivation to; e.g., specifying a predicate as an intransitive or transitive event description.

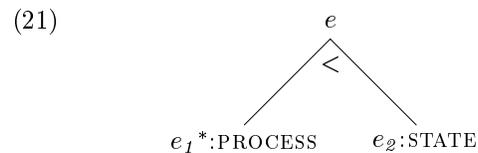
⁵See the discussion in Pustejovsky (1995) for details.

⁶Lexical Inheritance Structure is not relevant to our present discussion. This area of GL has been further elaborated in recent work, however. See Pustejovsky (2000b) and Asher and Pustejovsky (2000).

This is only partially correct, however. One additional constraint on the arguments is to establish the logical connection between the created artifact, ARG2, and the default argument of the material, D-ARG1. Namely, the created object is constituted of the material, a dependency captured directly in the qualia structure of the argument in the verbal semantic representation. Again, simplifying the structure for the present discussion, this can be represented as follows:

$$(20) \left[\begin{array}{l} \mathbf{build} \\ \text{ARGSTR} = \left[\begin{array}{l} \text{ARG}_1 = x:\mathbf{animate_individual} \\ \text{ARG}_2 = \left[\begin{array}{l} \mathbf{y:artifact} \\ \text{CONST} = z \end{array} \right] \\ \text{D-ARG}_1 = z:\mathbf{material} \end{array} \right] \end{array} \right]$$

The event structure for the verb *build* identifies it as a “left-headed” binary branching structure, with the initial event typed as a process and the terminus event typed as a state (the head event is marked with the diacritic *).⁷



The qualia are associated with parts of the overall event and are not uniquely associated with a single argument. Likewise, arguments may be associated with multiple subevents. In the present example, because *build* is a creation predicate, the final subevent introduces the sortal predication of a new object into the domain, illustrated in (22) below.⁸

⁷Within an event semantics defined not only by sorts but also by the internal configurational properties of the event, we need to represent the relation between an event and its proper subevents. Extending the constructions introduced in van Benthem (1983) and Kamp (1979), we interpret an “extended event structure” as a tuple, $\langle E, \preceq, <, \circ, \sqsubseteq, * \rangle$, where E is the set of events, \preceq is a partial order of *part-of*, $<$ is a strict partial order, \circ is overlap, \sqsubseteq is inclusion, and $*$ designates the “head” of an event. See the discussion in Pustejovsky, 1995 for details.

⁸I will assume that the mapping from qualia structure to syntax is constrained by the linking principles presented in Pustejovsky (1995). Briefly, these work as follows. The qualia of a lexical expression must be *saturated* by the syntax. That is, the semantic variables in the qualia structure must be fully interpreted in the resulting syntactic structure.

(i) *Qualia Saturation* :

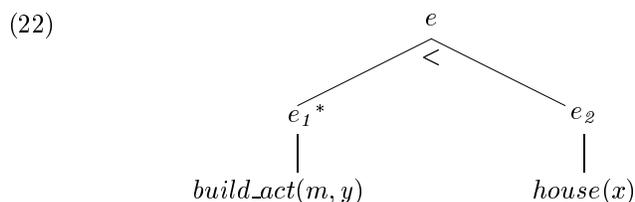
A qualia structure is saturated only if all arguments in the qualia are covered.

We define covering as follows:

(ii) *Covering*:

An argument x is covered only if:

- (a) x is linked to a position in s-structure; or
- (b) x is logically dependent on a covered argument y ; or



This abstracted event representation does not express the constitutive relation that exists between the material being acted upon and the resulting artifact created. More significantly, it fails to indicate the predicative structure of *build* as a *gating function*; i.e., the predicate opposition between a house not existing and then coming into existence and the way this is projected into a constrained event model for semantic interpretation. It is the elaboration of this notion of opposition to which we now turn.

In order to better understand the modes of opposition, let us look at some illustrative verb classes in English, each of which expresses a distinct mode of change. The partitioning of verbs into distinct lexical classes has provided us with a better understanding of what parameters of semantic representation help to determine syntactic form in language. The recent de facto standard of classification for English verbs is Levin's (1993) study of verb alternations. We will select four verb classes from this work to illustrate how they are analyzed in terms of different modes of opposition.

- (23) a. CREATION AND TRANSFORMATION: build, assemble, bake, cook, construct, design.
 b. DESTRUCTION: destroy, annihilate, decimate, demolish, ruin, wreck.
 c. CHANGE OF STATE: break, crack, crush, rip, tear, bend, fold, cook, bake, boil.
 d. CALIBRATABLE CHANGE OF STATE: climb, decline, decrease, fall, drop, increase, jump.

The names of these classes should not mislead us, since all of the predicates in the classes above involve some sort of change of state over an argument, whether that object is built, destroyed, broken, or decreased. Consider the sentences below, illustrating these verb classes and the various changes involved.

- (24) a. Mary assembled the table.
 b. Alice baked a cake.
- (25) a. The waves demolished the wall.
 b. The fire destroyed the building.

(c) *x* is skolemizable by virtue of its type.

- (26) a. Mary chipped the cup.
 b. John bent the photo.
- (27) a. The Dow climbed 2% in active trading.
 b. The temperature fell during the night.

In the sentences in (24), the change entails creation of an artifact where none existed before. Conversely, those events in (25) start off with reference to objects that are taken out of existence, as denoted by the sortal in the NP description. The change referred to in the sentences in (26) is of a specific aspect (property) of the object, and not of the object itself. Finally, the verbs shown in (27) refer to scales and the relative changes over these scales.

For the purpose of the present discussion, it will be useful to classify the predicates above according to the mode of opposition that an object undergoes. Put in terms of persistence, we will categorize predicates by the nature of their *gating* behavior.⁹

We will distinguish here between several distinct classes of predicate sorts, and the opposites that are constructable from them. The classic distinction between contradictories and contraries illustrates two modes of predicate opposition.

- (28) a. Bill is healthy
 a'. Bill is not healthy.
 b. Bill is sick.
 b'. Bill is not sick.
- (29) a. Jan is male.
 a'. Jan is not male.
 b. Jan is female.
 b'. Jan is not female.

Sentences (28a,b), involving *polar opposites* such as *healthy/sick*, are typically viewed as contraries, while (29a,a') are contradictories. While contradictories (28a,a') and (29a,a') usually follow from an interpretation of *not* as weak negation (cf. von Wright, 1963, Horn, 1989), the contradictories present in (29a,b) and (29a',b') cannot be the result of weak negation alone. As discussed in Pustejovsky (2000b), properties such as *male* and *female* are inherently contradictory when applied to its naturally predicated type, i.e., *animate-gendered* (or *gendered*). This being said, we will treat binary opposition as a two-element property semilattice:¹⁰

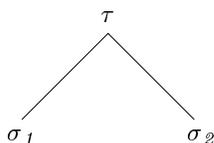
- (30) a. Binary Property:
 b. $\langle \sigma_1, \sigma_2, \tau, \sqcup, \sqsubseteq \rangle$ realizes a binary predicate P , where τ is a local top type for this sortal array, such that $\sigma_1, \sigma_2 \sqsubseteq \tau$, and

⁹Obviously, not all predication will refer to opposition; for example, *love* is a stative relation and *happy* is a stative property.

¹⁰Briefly, we define a property semilattice below:

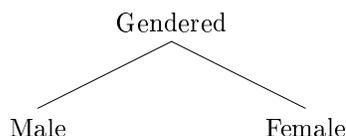
$\neg\exists\sigma[\sigma \sqsubseteq \tau \wedge [\sigma \neq \sigma_1 \vee \sigma \neq \sigma_2]]$. That is, σ_1 and σ_2 exhaustively partition τ .

c.



Other examples of adjective pairs with this behavior, besides *male/female* include *married/single* and *employed/unemployed*. These adjective pairs exhaustively partition the property that they are sorts of, as illustrated below.

(31)



BINARY OPPOSITION PREDICATE

For this type of binary predicate, P , there will be at least one opposition structure available in the language, viz. that arising from negation of the predicate, $\neg P$. If the language lexicalizes both forms in the opposition, then we of course have three unique opposition structures available as possible predications, $\langle P, \neg P \rangle$, $\langle P, Q \rangle$, $\langle \neg Q, Q \rangle$. For the binary adjective *dead* (and its antonym *alive*), $\lambda x \lambda e[\text{dead}(e, x)]$ is equivalent to $\lambda x \lambda e[\neg \text{alive}(e, x)]$, since there is no middle term.

For scalar properties such as *tall* and *short*, it has long been noted that they are measured relative to the same shared scale (cf. Hayes, 1979, Bierwisch and Lang, 1989). Relative to this scale, S , the polar adjectives are measured as positive and negative values or placements on this scale. Adopting Kennedy's (1999) recent discussion of degree adjectives, for degrees d_1 and d_2 , on the scale S , the following relation will hold for the antonymous adjectives ϕ_{pos} and ϕ_{neg} (*tall* and *short*, respectively):

$$(32) \quad d_1 \succ_{\phi_{pos}} d_2 \Leftrightarrow d_2 \succ_{\phi_{neg}} d_1$$

The poles on such as scale, however, are the *max* and *min* points and are uniquely predicable. The opposition structure for lexical pairs such as *dirty/clean* and *tall/short* has been referred to as polar opposites in the

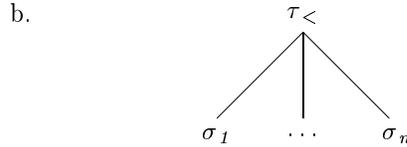
a. Property semilattice:

$\langle \Sigma, \tau, \sqsubseteq \rangle$ realizes a predicate P , where Σ is a sortal array of types, τ is a local top type for this sortal array, such that $\sigma_i \in \Sigma$ for $\sigma_i \sqsubseteq \tau$.

classical literature on predication (Lloyd, 1992), as well as more recently by Miller (1985,1989).

For our discussion, the nature of polar attributes can be defined in terms of a sortal array with distinguished elements (cf. Pustejovsky, 2000b).

- (33) a. $\langle \Sigma, \tau, \sqsubseteq, <, \sqsubset \rangle$ realizes a predicate P , where Σ is a sortal array of types, τ is a local top type for this sortal array, such that $\sigma_1, \dots, \sigma_n \in \Sigma$ for $\sigma_i \sqsubseteq \tau$, and $\sigma_i < \sigma_{i+1}$, and there are two poles, σ_1 , and σ_n , that are distinguished sorts.



In this paper, we will focus on the distinction between binary and scalar predicates and the oppositions they evoke. We will demonstrate how the opposition structure can be directly incorporated into the event structure of the predicate's semantics and what the effects of this are on interpretation.

Although *dirty/clean* are polar predicates over a scalar measure, a lexical item or phrase asserting one of these polarities of an argument is also construed as a binary predicate:¹¹

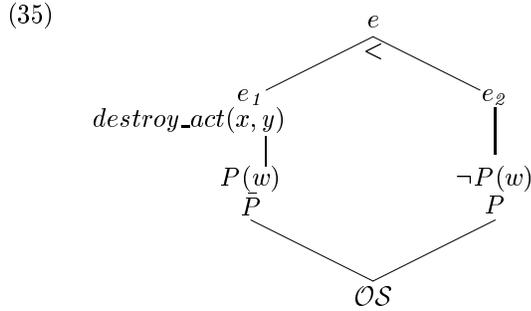
- (34) a. Mary cleaned an old car.
 b. Opposition Structure:
 $\exists x : \text{car} \exists e_1, e_2 [\neg \text{clean}(e_1, x) \wedge \text{clean}(e_2, x) \wedge e_1 < e_2]$

On the other hand, incremental theme verbs (cf. Tenny, 1994, Krifka, 1989) refer to the scale directly, and have no polar anchor to allow a binary predication over the changed object. Therefore, such verbs will denote a change of relation rather than a change of state for that argument. Given the above distinctions in predicate sorts, as characterized by modes of opposition, let us now examine what effects there are in how these modes map to syntax, as mediated by argument and event structures.

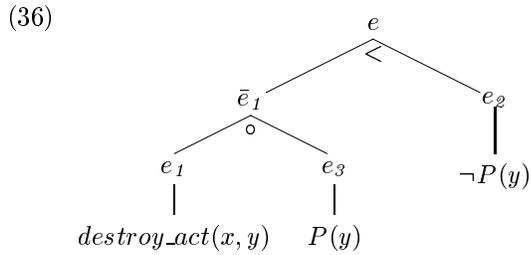
Let us begin with the simplest case, a *destruction*-predicate. As mentioned before, predicates such as *build* and *destroy* bring into and out of existence, respectively, the object denoted by the referring expression in that verb's internal argument position. Each of these is a gating functions for the type selected by the predicate for that argument. *Gating* will be defined as the introduction of termination or initiation conditions for the sort or properties of an argument; in this case, gating refers to the sort itself. For the opposition structures mentioned above, such as *dead/alive*,

¹¹Notice that the verb *clean* has a different assertoric force than the verb *wash*; if Mary washed a car, there is only an implicature that the car needed washing, i.e., that it was not clean. We return to this distinction in the next section.

the verb introducing the opposition will be a gating function for that predicate, $\langle P, \neg P \rangle$. For the verb *destroy*, the opposition structure introduces a termination condition on the Formal quale of the internal argument. This is illustrated below where the opposition structure (OS) for *destroy* is associated with its event structure.

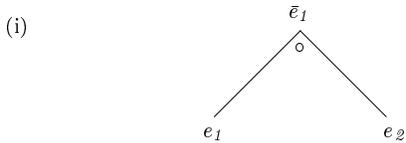


A more unified representation would be desirable, however, and opposition can be built directly into the event structure of the sentence by extending the calculus that defines the substructure of events in language.^{12,13}

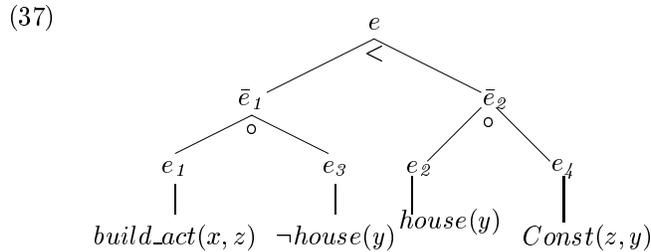


¹²The extensions to the event structure are fairly modest in fact, but we will not examine the consequences of these changes here. See Pustejovsky (forthcoming) for further elaborations.

¹³Headedness: For two events, e_1 and e_2 , where $e_1 \circ e_2$, if e_1 is the prominent element, and there is an event e_3 , which is the exhaustive sum of these two events, $\circ_\infty(\{e_1, e_2\}, e_3)$, then we will identify e_1 as the head of e_3 , and notate e_3 as a projection of e_1 , i.e., \bar{e}_1 . For convenience, we will sometimes represent this structure in tree notation as well, as shown below:



For a *creation*-verb such as *build*, we have the converse situation: a predication over the Formal quale of the internal argument is initiated by virtue of the opposition structure of the verb, i.e., $\langle \neg P, P \rangle$. This opposition can be incorporated directly into the event structure as shown in (37) below:



It should be noted that the constitution relation holding between the material and the object being created is expressed in the qualia structure of the arguments themselves, and not directly in the event structure. Therefore, $Const(z, y)$ actually holds of the entire spanning event and could be effectively factored out.

What we have done in the above discussion is to make explicit the semantic opposition of the object undergoing the change of state in predication. This is a minor modification to the event structure presented in Pustejovsky (1995), but allows us to address more substantially the concerns presented in Section One regarding the contradictions of change; that is, how can we model persistence as well as change in the event structure representation of a sentence, such that we provide the appropriate scope to the properties associated with an argument. We turn to this issue directly in the next section.

13.2 Event Persistence Structure

13.2.1 The Principle of Property Inertia

Most classical analyses of tense in logic and semantics have focused on the problems that sentential temporal operators create with displaced temporal reference, as in (38).

- (38) a. The President was born in 1946.
 b. Tom met his wife in 1988.
 c. All rich men were obnoxious children.

In each of these cases, the NP is interpreted in the same temporal frame as the predicate, but not by virtue of its definite description, but rather by its extension (cf. Kamp, 1979, Kamp and Reyle, 1993 for discussion). That

is, the person who is the current president was born in 1946 as an infant; when Tom met his wife, she wasn't yet his wife, and so on.

Much less studied in semantics are cases where the events that are denoted by the sentence act to either terminate or initiate the properties denoted by descriptions in the sentence. I will distinguish between the well-studied cases of temporally displaced reference, as in (38) above, and examples of *coherent event descriptions*, illustrated in (39) and (40) below:

- (39) a. John comforted the crying child.
 b. Cathie mended the torn dress.
- (40) a. The plumber fixed every leaky faucet.
 b. John cleaned every dirty dish.

As mentioned in Section One, the scope of the referring expressions in the object positions above must be split apart in order to not contradict the semantic opposition introduced by the matrix verb in each sentence. The temporal displacement involved in (39) and (40) is different from that encountered in (38) in several respects: (a) it is triggered by the semantics of the predicate governing it; and (b) the description is only *partially* displaced. In fact, if these were normal cases of temporal displacement, then we would expect the sentences in (41) to be acceptable.

- (41) a. !Mary cleaned the clean table.
 b. !John built the built house.
 c. !John drank the empty glass of milk.

The fact that they are not generally acceptable further supports the view here that contradiction of change is a distinct and more constrained phenomenon than temporal displacement. The goal of this paper is to explore how such interpretations are computed and how to constrain the application of such local displacement operations. In this section I outline a procedure for computing the maximally coherent event description over which the properties that are initiated and terminated by events denoted by a sentence can hold, continuously. When defined for the events within a natural language utterance, this will be represented in terms of the *event persistence structure (EPS)* for that sentence. The general observation relevant to our discussion of change is related to various attempts at solving aspects of the frame problem, and in particular, the formulation of circumscription and other devices. A simplifying assumption towards this goal is to assume some version of the principle of inertia (cf. McCarthy and Hayes, 1969, Reiter, 1991, Shoham, 1988). For our current concerns, I will state it as the principle of “property inertia”, and define it as follows:

(42) PRINCIPLE OF PROPERTY INERTIA:

1. No predicate (opposition structure) affects the sortal integrity of the type of an individual, as selected by the matrix predicate, unless explicitly asserted by a predication in the sentence.
2. No predicate affects the predicative integrity of a modification to an individual, unless explicitly asserted by a predication in the sentence.

In other words, the descriptions for objects will continue to hold throughout the lifetime of the event being described by the utterance, unless affected explicitly by the predicate itself or by virtue of computing the event persistence structure of a sentence. The event structure represents the basic inference of change, while the event persistence structure also represents the basic inference of persistence and secondary change. In addition to the specific type of predication involved, the scope of the affected aspect of the object is obviously crucial. This is often referred to as the problem of persistence. Briefly, there are two types of persistence which interest us here:

1. OBJECT PERSISTENCE: The integrity of the object as described (predicated) in the selection by a predication. An object is persistent relative to an eventuality (particular event type).
2. PROPERTY PERSISTENCE: The integrity of a property of an object as described (predicated) in the selection by a predication. A property is persistent relative to an eventuality.

Within the approach being outlined here, persistence and property inertia must hold over eventualities of some sort. Let us assume then, that all nouns and adjectives are treated as event descriptions, where even sorts such as *man*, *rock*, and *house* will be treated as relations between an individual and the state holding of that individual, as predicated by the sortal distinction itself;¹⁴

¹⁴The semantics of quantifiers must change to be a relation between two unsaturated event descriptions; namely, $\lambda\mathcal{P}\lambda\mathcal{F}[\mathbf{R}(\mathcal{P}, \mathcal{F})]$, where, for the specific quantifiers *some* and *every*, we have the following translations, where \mathcal{F} and \mathcal{P} are variables of type $\langle g, \langle e^\sigma, t \rangle \rangle$:

- (i) $[\textit{some}] = \lambda\mathcal{P}\lambda\mathcal{F}\lambda e\exists x[\mathcal{P}(e, x) \wedge \mathcal{F}(e, x)]$
- (ii) $[\textit{every}] = \lambda\mathcal{P}\lambda\mathcal{F}\lambda e\forall x[\mathcal{P}(e, x) \rightarrow \mathcal{F}(e, x)]$

See Pustejovsky (1995) for further discussion of event descriptions in a typed semantic derivation.

This move is not as radical as it may seem, since, in most cases, this interpretation of a noun or adjective is not exploited in composition. Rather, it is available as a resource to the logic. We could model this correspondence as a lexical type shifting rule (lexical rule), essentially adding an event variable to an expression where appropriate. For the current discussion, however, I assume the event variable is present for all expressions, and subsequently factored out by the computation of event persistence, as described

- (43) a. $\lambda x \lambda e [man(e, x)]$
 b. $\lambda x \lambda e [rock(e, x)]$
 c. $\lambda x \lambda e \exists e' \exists y [house(e, x) \wedge make(e', y, x) \wedge e' < e]$

The principle of property inertia states that such diverse nouns as *boy*, *rock*, and *house* are equally persistent without the interpretation of a context; furthermore, properties of these nouns are also equally interpretable relative to persistence, e.g., *big*, *efficient*, and *solid*.

In the next section, we examine how adjectives bind into the qualia structure of nouns, to select a narrow facet of the noun's meaning. As we will see, this has profound consequences on the adjective's subsequent persistence properties and how this figures into the computation of event persistence structure.

13.2.2 Adjectives as Events

We begin with a discussion of adjectives and the semantic classes they denote. In Pustejovsky (1995), I discussed the classic field-descriptive approach to adjective classes, as given in Dixon (1982), where a taxonomic classification is used to distinguish adjectives according to the general semantic fields associated with the term.

1. DIMENSION: big, little, large, small, long, short
2. PHYSICAL PROPERTY: hard, soft, heavy, light
3. COLOR: red, green, blue
4. HUMAN PROPENSITY: jealous, happy, kind, proud, cruel, gay
5. AGE: new, old, young
6. VALUE: good, bad, excellent, fine, delicious
7. SPEED: fast, quick, slow
8. DIFFICULTY: difficult, easy
9. SIMILARITY: alike, similar
10. QUALIFICATION: possible, probable, likely

Similarly, work from the computational literature such as WORDNET (Fellbaum, 1998) assumes that there are general, psychologically inspired categorizations for properties that are grammatically realized as adjectives.

While not discounting either of these approaches in spirit, the methodology taken here and in Generative Lexicon in general begins with a somewhat different set of discriminating features for analyzing adjectival categories. For example, let us assume, following Pustejovsky (1993) and Bouillon (1996), that evaluative adjectives such as *fast* and *good* are analyzed as event-denoting predicates. This analysis can be naturally extended to larger sets of adjectives, such as those listed above, by treating the qualia

below. Either solution is generally acceptable, as long as the phenomena of change and persistence can be adequately accounted for.

roles as temporally ordered relative to each other. Abstracting over the qualia in terms of their temporal properties gives the partial orderings below: $A < F$, $C \circ F$, and $F < T$. Now let us assume that any adjectival phrase, prepositional phrase, or relative clause modifying its head noun is bound to a specific qualia role of the head noun.¹⁵ Putting this principle together with the observations above regarding the temporal ordering of qualia values, we arrive at the thesis for qualia selection, stated below:

- (44) **QUALIA SELECTION THESIS :**
 Every Phrase, XP_i , occurring as a modifier to a nominal head, N , is associated with a specific qualia role, q_j , for that noun, according to the following constraints. If XP_i modifies:
- i. **FORMAL:** then the event for that phrase corresponds roughly to an overlap relation, 'o', with the head N ;
 - ii. **TELIC:** then the event for that phrase corresponds roughly to the '>' relation relative to the head N , but in fact is closer to a generic interpretation. o_g (see below);
 - iii. **AGENTIVE:** then the event for that phrase corresponds roughly to the '<' relation relative to N ;
 - iv. **CONST:** then the event for that phrase corresponds roughly to an overlap, 'o', relation with the head N .

The table below illustrates particular adjectives and the qualia they select for.

¹⁵This holds for most adjectives compositionally interpreted, but does not include noncompositional problems such as *an occasional sailor*, cf. Partee (1992).

(45)

ADJECTIVE	Qualia Selection
well-built	Agentive
unbaked	Agentive
red	Formal
stone	Constitutive
wooden	Constitutive
useful	Telic
carved	Agentive
effective	Telic
fast	Telic
heavy	Formal
dense	Const
large	Formal

TABLE OF QUALIA SELECTION PROPERTIES

As mentioned above, relative to predication and the ordering of the event descriptions within an entity intension, the qualia provide three relations: $<$, \circ , and $>$. Most adjectives appear to predicate of the formal role, and hence are overlapping event descriptions. For example, dimensional adjectives such as *small*, *long*, *wide*, and *tall* all refer to properties that hold of an entity while it persists as that entity. These are overlapping properties, and can be said to modify the formal qualia role. Nevertheless, some adjectives refer explicitly to AGENTIVE (46a), and others to TELIC (46c), or CONST (46d). Examples of each of these can be seen in the modifications in (46) below.¹⁶

- (46) a. a well-built (A_I) house ($[F, C, A_I, T]$)
 b. a two-story (F_I) house ($[F_I, C, A, T]$)

¹⁶Usually, the specified AGENTIVE of a type which has an AGENTIVE value (and this includes natural types as well) will not be allowed as a modifier to that entity. For example, consider the following interesting data:

- (i) an unwritten book / *a written book / a poorly written book
 (ii) an unbaked cake / *a baked cake / a half-baked cake

In fact, the ungrammaticality of such expressions is similar the argument shadowing effect seen with the verb *butter* as in **to butter toast with butter*. This is explored in Pustejovsky (2000a). It is not possible to express the shadow argument without further semantic content; in the case above, the AGENTIVE role which is expressed by the modification in the compositional structure *written book* is equivalent to the shadow from the Agentive in the NP. A similar relation holds between the TELIC role and the modification in (iii) and (iv) below:

- (iii) edible flowers / edible plants
 (iv) !edible food / !edible bread

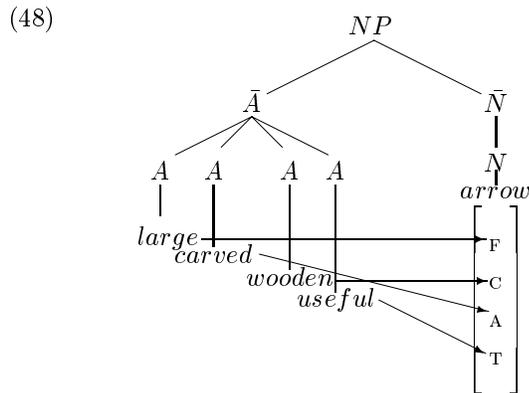
So, it appears that NPs have shadow argument behavior just as verbs do.

- c. a vacation (T_1) house ($[F, C, A, T_1]$)
 d. a brick (C_1) house ($[F, C_1, A, T]$)

All of these modifications might conceivably be present in the structure of a single NP, such as in (47) below.

- (47) a. a large carved wooden useful arrow
 b. a large (F_1) carved (A_2) wooden (C_3) useful (T_4)
 arrow ($[F_1, C_3, A_2, T_4]$)

The different types and bindings of the adjectival modification can perhaps be better understood if we examine the modification structure present in (47b).



Such examples illustrate the inherent richness of the qualia structure and how lexical items, denoting specific types, are typed to select individual qualia. The relevance of these data to the issues of change and persistence is this: the Qualia Selection Thesis will permit us to bind the behavior of potentially independent event descriptions to the persistence behavior of the head it modifies. This will greatly simplify our computation of the event persistence structure.

13.2.3 Computing Event Persistence Structure

In the previous sections we outlined the necessary assumptions for defining the procedure that computes the maximally coherent event description for a sentence, represented formally as something we will call the *event persistence structure* for a sentence. These assumptions include the following two principles:

- (49) a. The Principle of Property Inertia; objects and their properties tend to remain as they are unless explicitly affected;
 b. Qualia Selection Thesis; modifiers selectively bind to specific qualia of the head noun.

Given these preliminaries, we will now formulate our first approximation of how to compute the event persistence structure (EPS) for a sentence. Again, the purpose of the procedure is to construct the maximally coherent event description of the opposition structure for every predicate in a sentence. The strategy is to leverage the two principles of inertia and qualia selection thesis so as to factor out as many of the event-denoting predicates as possible from the sentence. The goal of the EPS is to represent not only what has changed by virtue of the matrix event description, but to also model secondary effects of the action, if they can be captured, as well as what has stayed the same.

The EPS is an annotated event structure, with event predicates showing the scope appropriate to their opposition structures, relative to the matrix event predication denoting change and persistence of the various arguments. Importantly, however, the resulting representation should be the minimally richest event-based model needed for arriving at a coherent event description of a sentence. Formally, the procedure is similar to the Skolemization procedure in theorem proving, where quantified expressions are simplified for deductive rule application.

In computing the EPS, we wish to factor out as much as possible from the semantic content of an expression pertaining to change and persistence, in order to return the simplest representation while still expressing the appropriate content. To this end, I will assume that any predicate, be it verbal, adjectival, or phrasal (PP), is assigned an independent event description, δ_i ; further, every sortal expression will be assigned an event description.¹⁷ For example, for a string *abcde*, regardless of composition and internal constituent structure, we assign each terminal an event description; $\{\delta_a, \delta_b, \delta_c, \delta_d, \delta_e\}$.¹⁸ The set of event descriptions will be referred to as Δ . We denote the event description assigned to the matrix predicate of the clause, *P*, as the *core event structure*. This is the representation which acts as the backbone in the construction of the event persistence structure; that is, all additional event predications in the clause are annotations to this core structure. The opposition structure that is carried by the core event structure is inviolable relative to other predicates that are subsequently

¹⁷It may be the case that clitics also introduce event descriptions or relations between events. In a recent paper, Castaño (2000) has argued that the Spanish clitic *se* is best analyzed as carrying its own event structure, effectively subordinating the event of the VP that it is in construction with into a higher event relation. If this is the case, then clitics would also be included in the set Δ .

¹⁸The constituent structure of the sentence is obviously relevant to the computation of the event persistence structure. We actually make use of it by virtue of the Qualia Selection Thesis, and the embedded temporal orderings this creates. In Pustejovsky (forthcoming), I explore the computation of EPS directly from the syntax.

added, and we will refer to these predicates as the ground terms.¹⁹

Given the event description set, Δ , and the construction of the core event structure, for each event-denoting predicate in the expression, we apply a single test, *gate*, defined as follows.

- (50) a. GATE: For an event description, $\delta \in \Delta$, in the domain of the matrix predicate P , δ is *gated* by P only if the property denoted by δ is either initiated or terminated by P .
- b. PERSIST: If δ is not gated, then it is said to *persist* relative to the matrix predicate P .

Next, we associate the gated event descriptions to the nodes identified with the gating function. Finally, all persisting predicates are factored out of the expression in the event structure. They will be said to take wide *persistence* scope (p-scope) over the event description. If a predicate does not take wide p-scope (such as all those that are gated), then it is narrow scope, and is associated only with the appropriate subevents.

Consider briefly the example of (51).

- (51) Mary cleaned the dirty table.

The predicate *dirty* is gated to \neg *dirty* because *clean* is the predicate ground and is not defeasible. All other predicates must be consistent with the ground term. Computing event persistence closure for each predicate is minimally to create opposition structures for each predicate in the expression and see if the new term in the opposition-structure is consistent with the ground.

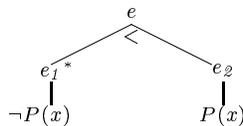
The *Principle of Inertia* has some important consequences for computing the EPS of a sentence. Most significantly, it states the following: the descriptions that are used for all objects appearing as arguments to a predicate are assumed to hold from the initial event of the predicate, unless otherwise “gated” by the predicate structure.

13.2.4 Examples of Event Persistence Structure

In this section I will briefly outline the manner in which arguments maintain their persistence, and catalogue how each argument behaves, relative to the matrix event denoted by the verb. There are essentially four situations that can arise for a referring expression:

¹⁹For example, in (i), the predicate P is the ground in e_2 , while $\neg P$ is the ground in e_1 .

(i)



- (52) a. The argument persists;
 b. The head of the argument does not persist;
 c. The head of the argument persists, but there are properties of the head introduced by predication that do not persist.
 d. The head of the argument persists, but there are properties of the head expressed in the referring expression that do not persist.

Examples of (52a) are given below.

- (53) a. Mary saw John.
 b. A man sat on a bench.

The predicate does not affect the persistence of either argument. Hence, the EPS is equivalent to a conventional event structure. The case of (52b) involves an argument that is gated by the predicate, as shown in (54) below.

- (54) a. Mary built a house.
 b. John became President of the club.
 c. Mary ate a cookie.

Situation (52c), on the other hand, arises when the verb introduces a gating of a property of the verb's argument, while the head stays persistent.

- (55) a. John closed the door.
 b. Mary cleaned the table.
 c. John painted the house.

Finally, situation (52d) arises when an argument expression is contradicted by the gating function, as illustrated in (56b) and (57b).

- (56) a. Mary cleaned the table.
 b. Mary cleaned the dirty table.
 (57) a. Mary fixed the tire.
 b. Mary fixed the flat tire.²⁰

Now let us step through the EPS algorithm with a specific example computation. Consider the interaction between the the predicate *clean* and its direct object head noun *table* in the sentence given in (58) below:

- (58) Mary cleaned the table.

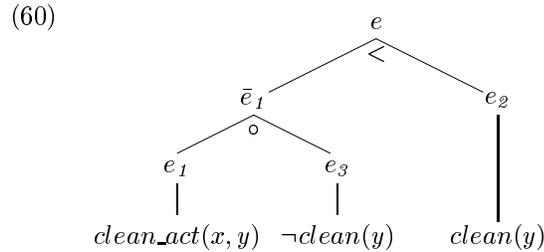
The set of event descriptions for the sentence in (58), Δ , is as follows:

- (59) $\Delta = \{ \text{mary}(e_1, x), \text{table}(e_2, y), \text{clean_act}(e_3, x, y), \neg \text{clean}(e_4, y), \text{clean}(e_5, y) \}$

²⁰It should be pointed out that we understand the event of fixing a tire on a bike to be an activity applied to *mending* the same object, while we typically are not so enabled with a car. The verb refers to a *replace* activity for the car rather than a mend.

Recall that we initially consider every predicative expression as a candidate event description. The EPS algorithm is designed to both (a) prune the events that are relevant to the interpretation of the sentence as well as to (b) provide the appropriate scope to the quantification associated with properties and sorts in the sentence.

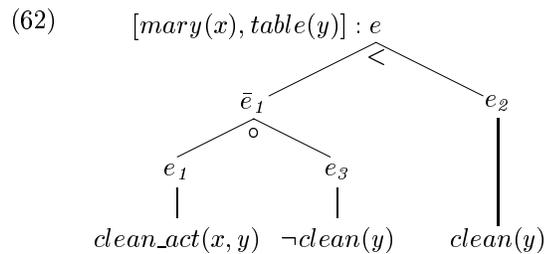
From Δ , we construct an event structure associated with the matrix predicate of the sentence, shown in (60):



Then we apply the operation *gate*:

- (61) a. *gate*(mary) fails;
b. *gate*(table) fails;

Hence, both these arguments persist, and the quantificational force for both the individual constant *Mary* and the NP *the table* are p-wide scope, since the predicate does not act to gate either. Following the general notational conventions in Discourse Representation Theory (Kamp and Reyle, 1993) for discourse referents, we will express wide scope of a term α_i relative to an event structure E as $[\alpha_i] : E$. For the current example, this gives the following:



The resulting EPS gives the correct scope to both arguments relative to the gating property of the predicate. In other words, neither expression is effected by the computation of the change of state denoted by *clean*.

Now consider the contradiction of change example shown below in (63).

(63) Mary cleaned the dirty table.

The set of initial event descriptions is the same as the previous sentence, with the addition of the predicate *dirty*:

$$(64) \quad \Delta = \{\text{mary}(e_1, x), \text{table}(e_2, y), \text{clean_act}(e_3, x, y), \neg \text{clean}(e_4, y), \\ \text{clean}(e_5, y), \text{dirty}(e_6, y)\}$$

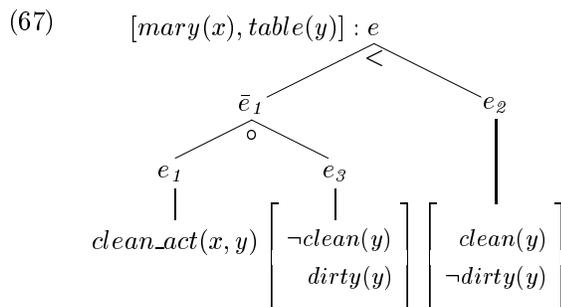
The core event structure is the same as that in (60) above. Again, we apply the operation *gate*:

- (65) a. *gate*(mary) fails;
 b. *gate*(table) fails;
 c. *gate*(dirty) succeeds;

Note that the predicate *clean* gates the predicate *dirty* in this example because the EPS must obey the logic of sortal predicates as discussed in the previous section. Recall that there are two opposition structures for an adjective like *dirty*:

- (66) a. $\langle \text{dirty}, \neg \text{dirty} \rangle$: Binary opposition
 b. $\langle \text{dirty}, \text{clean} \rangle$: Polar opposition

The core event structure introduces a terminating condition for the predicate *dirty*, hence gating it. Therefore, unlike *mary*(*x*) and *table*(*y*), the predication *dirty*(*y*) is not p-wide scope relative to the change denoted by the predicate. The resulting EPS is illustrated below in (67).



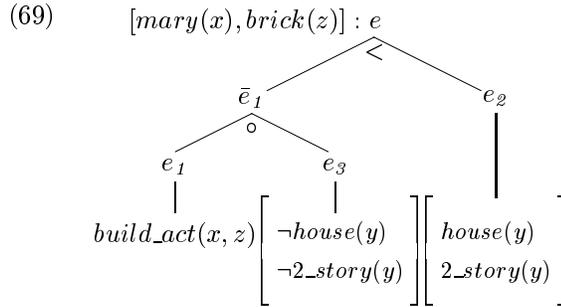
The predicative force of the NP has effectively been split apart according to the persistence properties of the predicates. This is a critical step in creating a logical form from which inference can be subsequently performed.

As another example of how the event persistence structure is computed, consider the sentence in (68a):

- (68) a. Mary built a two-story brick house.

As discussed in Section 2.2, adjectives such as *two-story* are FORMAL-qualia binders to the head noun. As a result, they will be subject to the gating conditions on the head noun introduced by any predication. In the case of a

creation-verb such as *build*, the predicate will gate not only the sortal *house* but anything bound to the FORMAL quale, such as the adjective *two-story*. The resulting event persistence structure for this sentence is interesting in another respect as well: the predicative use of the noun *brick* above binds to the CONST-quale role, and takes wide scope for persistence while the sortal description *house* does not;



All of the cases of contradictions of change discussed above can be treated in a similar fashion. As a final example, consider the split quantification cases encountered in earlier discussion, and shown below.

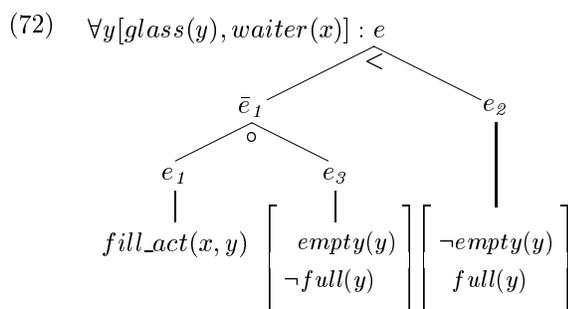
- (70) a. Mary fixed every leaky faucet in the house.
 b. The waiter filled every empty glass with wine.

Focusing here on just the gating portion of the algorithm, notice that the persistence of the predicates *leaky* and *empty* is gated by the core event structure of their respective governing verbs, *fix* and *fill*. The semantics of *fix* introduces an opposition structure that makes reference to the value of the TELIC of the internal argument.²¹ For the verb *fill*, the opposition structure predicates the FORMAL aspect of its argument, and the gating effects any modification of the noun *glass* that is not consistent with the predication. To illustrate this structurally, observe that there are two opposition structures for the adjective *empty*:

- (71) a. $\langle empty, \neg empty \rangle$: Binary opposition
 b. $\langle empty, full \rangle$: Polar opposition

The pair in (71b) satisfies the condition on gating, giving us the following event persistence structure:

²¹Elsewhere, I have argued that many verbs are dependent on aspects of the semantics of their arguments for full interpretation. Other examples of *functionally dependent* verbs are *break*, *open*, and *close* (Pustejovsky, 1995). Suffice it to say that the property *leaky*, relative to the qualia structure of *faucet* is gated by the core event structure of *fix*, which introduces the semantic opposition structure $\langle broken, \neg broken \rangle$ over the TELIC of the noun *faucet*.



What is important to observe here is the manner in which a description such as *every empty glass* undergoes a transformation, effectively pulling the gated event description of the adjective *empty* out of the expression, allowing wide-scope quantification of only the head sortal noun. In some respects, this is reminiscent of Kamp's *Now* operator (Kamp, 1979): that is, a reasonable paraphrase of (71b) is:

(73) The glasses that WERE empty are NOW filled by the waiter.

Importantly, the conditions under which such a transformation applies are completely determined by the predicate and how it acts to gate the quantified NP.

13.3 Persistence and Stage-Level Nominals

In this section we consider one final application of the event persistence algorithm discussed above. Namely, we will examine the behavior of a class of agentive nominals, called *stage-level nominals* (Pustejovsky, 1995), as studied in Enç (1981), Busa (1996), and Musan (1997). This class includes nouns such as *passenger*, *customer*, *audience*, and *pedestrian*, all of which refer to individuals that are presently engaged in an activity. Unlike the cases above, however, these nouns are either persistent or gated as a function of the computation of the EPS itself. Consider the following sentences:

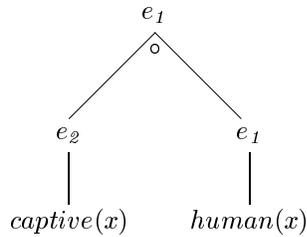
- (74) a. The prisoner escaped from the prison.
 b. The escapee has been put in police custody.
 c. The audience left the theatre.

Notice that, given what we have outlined above for the computation of event persistence, such examples are not unusual or difficult to model. They are similar in derivation to the contradictions of change involving the gating of adjectival descriptions. Assume that the qualia structure for stage-level nouns, such as *prisoner* and *audience* can be represented as shown below:

- (75) a.
$$\left[\begin{array}{l} \mathbf{prisoner} \\ \text{QS} = \left[\begin{array}{l} \text{F} = \text{human}(x) \\ \text{A} = \exists e[\text{captive}(e, x)] \end{array} \right] \end{array} \right]$$
- b.
$$\left[\begin{array}{l} \mathbf{audience} \\ \text{QS} = \left[\begin{array}{l} \text{F} = \text{human}(x) \\ \text{A} = \exists e, y[\text{attend}(e, x, y)] \end{array} \right] \end{array} \right]$$

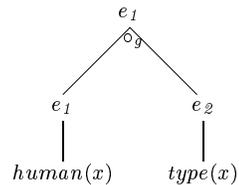
In terms of event structure, we have the following representations, respectively, where AGENTIVE-qualia modification in these nouns is taken as overlapping:²²

- (76) a.



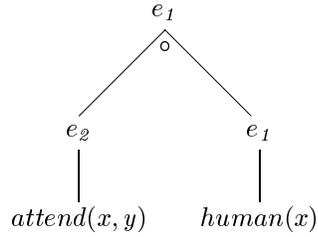
²²Within the representation of event persistence structure, what is expressed as a modal relationship in the TELIC for agentive nominals such as *linguist*, *violinist*, and *typist*, can be given in tree-form with the annotation of a temporal relation, \circ_g . Unlike \circ , which is transitive and symmetric, \circ_g is transitive, asymmetric, and introduces a modality. We can read $\circ_g(e_1, e_2)$ as individuating e_1 for some event description, which overlaps with the modal event description represented by e_2 .

- (i)



With Telic-modifying adjectives such as *good* and *fast*, it is clear how they take scope within a modal event description associated with the Telic of the noun (e.g., *a good typist*), acting as an event modifier. See Pustejovsky, 1995 for discussion.

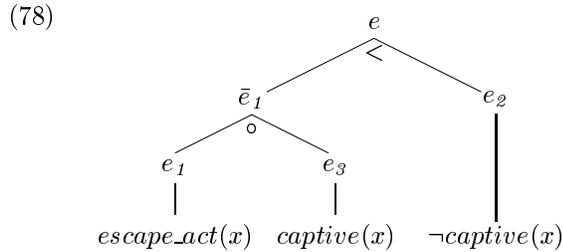
b.



To see how similar these nouns are to the cases studied above, consider the computation of the EPS for sentence (74a). The verb *escape* is an intransitive change-of-state predicate and introduces a binary opposition over its argument. Assume the set of event descriptions for this sentence is as follows:

$$(77) \Delta = \{\{\text{human}(e_1, x), \text{captive}(e_2, x)\}, \text{escape_act}(e_3, x), \neg\text{captive}(e_4, x), \text{captive}(e_5, x)\}$$

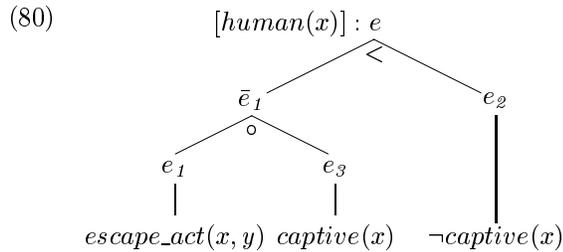
From Δ , we construct an event structure associated with the matrix predicate of the sentence, shown in (78):



Then we apply the operation *gate*:

- (79) a. *gate*(human) fails;
 b. *gate*(captive) succeeds;

The Formal content of *prisoner* persists globally, taking p-wide scope. The Agentive content, however, referring to the state of being captive, is gated by the predicate and takes p-narrow scope.



Hence, no contradictory interpretation results: certain predicates act to gate specific stage-level nominal interpretations in a systematic and predictable way. The above derivation illustrates why SLNs such as *prisoner* are able to participate in events that appear to contradict the very conditions that satisfy membership in that class of nouns. It is just those events that introduce the “boundary conditions” on the nominal expression itself, i.e., the gating functions for that noun that are acceptable as possible contradictions of predication.

As Enç (1981), Musan (1997) and Pustejovsky (1991a) point out, while these sentences are perfectly acceptable, use of the same nominal degrades if the predicate becomes semantically distant from the activity characterizing the individual noun. Furthermore, stage-level nouns cannot be used to introduce the property holding of the individual by virtue of the predication. This is illustrated in the pair below:

- (81) a. !The police arrested a prisoner in the bank.
 b. The police arrested the suspect / a man in the bank.

The NP *a prisoner* in sentence (81a) cannot be interpreted as “the individual who is now in custody for the first time”, without serious discourse context setting with the prisoner as the topic. This is because the opposition structure introduced by the verb *arrest* is in contradiction with the semantic content of *prisoner*. On the other hand, sortals such as *suspect* and *man* are unmarked relative to the semantic opposition of the verb, and are acceptable in this context.

There is another interesting consequence of the event persistence algorithm involving stage-level nominal interpretation. Namely, it does not require an event-based interpretation for nouns such as *pedestrian* unless the discourse context demands it. In other words, the lexical semantics of *pedestrian* can be seen as a “resource” for computation, but not necessarily exploited. Assume the event-related semantics of *pedestrian* is as shown in (82), where *s* is a type constant, referring to *street* and *sidewalk*:

$$(82) \left[\begin{array}{l} \mathbf{pedestrian} \\ \text{QS} = \left[\begin{array}{l} \text{F} = \text{human}(x) \\ \text{A} = \exists e[\text{walk}(e, x, s)] \end{array} \right] \end{array} \right]$$

Then notice how the noun is treated differently in the sentences illustrated in (83).

- (83) a. A pedestrian crossed the street. (*coherent* by EPS)
 b. A pedestrian went into Burger King. (predicate gates the persistence of the subject)
 c. A pedestrian is shopping in the store. (*incoherent* by EPS)

For (83a), the EPS algorithm allows simple individual-level quantification for the NP, giving an (abstracted) LF such as that in (84):

$$(84) \exists x \exists e_1 \exists e_2 [\text{cross_act}(e_1, x, s) \wedge \text{pedestrian}(x) \wedge \text{across}(e_2, x, s) \wedge e_1 < e_2]$$

For (83b), however, there is gating function present, and the event-based semantics of *pedestrian* must be exploited for the proper quantificational expression; that is, the FORMAL of the noun takes p-wide scope in the LF expression, but the characterization of the “walking” activity does not, since it is gated by the predicate. The event persistence structure for (83b) is given below:

$$(85) \begin{array}{c} [\text{human}(x)] : e \\ \swarrow \quad \searrow \\ \bar{e}_1 \quad e_2 \\ \swarrow \quad \searrow \\ e_1 \quad e_3 \\ \swarrow \quad \searrow \\ \text{go_act}(x) \quad \left[\begin{array}{l} \text{walk}(x, s) \\ \neg \text{in}(x, bk) \end{array} \right] \quad \left[\begin{array}{l} \neg \text{walk}(x, s) \\ \text{in}(x, bk) \end{array} \right] \end{array}$$

For (83c), there is no valid derivation, since the conditions on identifying the subject are not satisfied by the predicate. There is, however, the possibility of *coercing* a reading with a *NOW*-operator interpretation and the help of focusing information: “That pedestrian is *NOW* shopping in the store.”

Our final illustration of the interaction of gating functions and stage-level nouns involves the behavior of discourse anaphors. Consider the sentences shown in (86)-(88) below.

(86) The audience_{*i*} applauded to show its_{*i*}/their_{*i*} approval.

(87) The audience_{*i*} left the music hall.

- (88) a. *It_{*i*} then went home.
 b. They_{*i*} then went home.

c. $It_i/They_i$ had just heard Bernard Haitink's last performance.

Notice how the anaphor *it* in (88a) cannot refer to the antecedent NP *the audience*. This is because the persistence of this object has been gated; the object simply doesn't exist anymore. The FORMAL of the object, however, does still exist (as shown in (75b) above) and can be expressed as a plural anaphor *they*, viz., the individuals who formed the audience. In a way, the VP *leave the music hall* is acting as a grinding function over its subject, giving the plural component "parts", expressed as the value of the FORMAL qualia role.²³ Notice that either anaphor is acceptable in (88c), because the predicate is construed as referring to an event prior to the application of the gating function (cf. Asher and Lascarides, 1993).

13.4 Conclusion

In this paper, I have tried to motivate a richer notion of event structure for natural language semantics, based on data that prove difficult to model under current event theories. These data mostly involve contradictions of change, descriptions that, by virtue of the events they participate in, no longer hold. To solve these cases, I outlined an algorithm for computing the maximally coherent event description associated with a sentence. This resulted in a semantic representation called the event persistence structure, which, I argue, is a natural manifestation of the linguistically motivated entailments regarding change and persistence in a sentence, derived compositionally from sentential semantic interpretation. The result of the analysis is that the chain of states associated with an argument in discourse is initially projected from the lexical and compositional semantic properties of expressions in the sentence. This is a very different approach from that taken in Hobbs et al. (1993), for example, where abduction explores all possible derivations associated with the lexical items in a sentence. Probabilities may be assigned to rule applications in order to bias or weight particular derivations, but probabilities seem to have little to do with computing contradictions of change and other examples of event persistence; for implicatures it is perhaps appropriate, but entailments should not be expressed probabilistically.

Acknowledgments

I would like to thank José Castaño, Nicholas Asher, Jong Sup Jun, Bob Ingria, and Federica Busa for various discussions and comments on the material in this paper. All remaining inconsistencies are my own.

²³This is not entirely correct. José Castaño (p.c.) has observed that the grinding taking place in this example is more likely a shift from the FORMAL role to the CONST role of *audience*. This, however, does not change the basic observation regarding anaphora in these data.

Bibliography

- Alsina, A. 1999. On the Representation of Event Structure. In T. Mohanan and L. Wee (eds.), *Grammatical Semantics: Evidence for Structure in Meaning*, CSLI, Stanford.
- Asher, N. and Lascarides, A. (1993). Temporal Interpretation, Discourse Relations and Commonsense Entailment. *Linguistics and Philosophy* 16, 437–494.
- Asher, N. and Pustejovsky, J. (2000). The Metaphysics of Words. ms. Brandeis University and University of Texas.
- Bach, E. (1986) “The Algebra of Events,” *Linguistics and Philosophy* 9:5-16.
- Barker, C. 1999. Individuation and Quantification, *Linguistic Inquiry*, 30:4.
- van Benthem, J. (1983). *The Logic of Time*. Dordrecht: Reidel.
- Bierwisch, M. and E. Lang. 1989. *Dimensional Adjectives. Grammatical Structure and Conceptual Interpretation*, Springer Verlag, Berlin.
- Bouillon, P. (1997). *Polymorphie et sémantique lexicale : le cas des adjectifs*, Lille: Presses Universitaires du Spentrion.
- Bouillon, P. and Busa, F. (forthcoming). Where’s the polysemy? A study of adjective-noun constructions. *Proceedings of the Second Workshop on Lexical Semantics Systems*, Pisa, Italy.
- Busa, F. (1996). *Compositionality and the Semantics of Nominals*, Ph.D. Dissertation, Brandeis University.
- Busa, F., Calzolari, N., Lenci, A. and Pustejovsky, J. (1999). Building a Semantic Lexicon: Structuring and Generating Concepts. *Proceedings of IWCS-III*, Tilberg, The Netherlands.
- Castaño, José. (2000) The Spanish Clitic SE: Ambiguity, Underspecification and Co-Composition. ESSLLI-00 Workshop on: ”Paths and Telicity in Event Structure”. Hana Filip and Greg Carlson Ed.
- Chomsky, N. (1995). Language and Nature. *Mind* 104.
- Dixon, R.M.W. 1982. *Where Have All the Adjectives Gone? and Other Essays in Semantics and Syntax*, Mouton, Berlin.
- Dowty, D. R. (1979). *Word Meaning and Montague Grammar*, Dordrecht: Kluwer Academic Publishers
- Enç, M. 1981. *Tense without Scope*, Ph.D. Dissertation, University of Wisconsin, Madison.

- Enç, M. 1987. "Anchoring Conditions for Tense," *Linguistic Inquiry* 18:633-658.
- Fellbaum, C., (editor) (1998) *Wordnet : An Electronic Lexical Database*, MIT Press.
- Grimshaw, J. and S. Vikner 1993. "Obligatory Adjuncts and the Structure of Events," in E. Reuland and W. Abraham (eds.), *Knowledge and Language*, Vol. II, Kluwer Academic Publishers, Dordrecht, p. 143-155.
- Hay, J., C. Kennedy, and B. Levin (1999). "Scalar Structure underlies Telicity in "Degree Achievements"," in *Proceedings of SALT 1999*.
- Hanks, S., and D. McDermott. (1986). Default reasoning, nonmonotonic logic, and the frame problem. *Proceedings of the American Association for Artificial Intelligence* 328-333.
- Harman, G. (1986). *Change in View*. Cambridge, MA: MIT Press.
- Hobbs, J., M. Stickel, P. Martin, D. Edwards. 1993. "Interpretation as Abduction," *Artificial Intelligence* 63:69-142.
- Horn, L. R. 1989. *A Natural History of Negation*, University of Chicago Press, Chicago.
- Jackendoff, R. (1992). *Semantic Structures*, Cambridge, MA: MIT Press.
- Kamp, H. 1979. "Some Remarks on the Logic of Change: Part 1" in C. Rohrer, (ed.), *Time, Tense, and Quantifiers*, Tübingen, Niemeyer.
- Kamp, H. and U. Reyle. 1993. *From Discourse to Logic*, Kluwer Academic Publishers, Dordrecht.
- Kennedy, Christopher 1999. Polar Opposition and the Ontology of 'Degrees', forthcoming in *Linguistics and Philosophy*.
- Kowalski, R.A. and M.J. Sergot (1986) "A logic-based calculus of events," *New Generation Computing* 4:67-95.
- Kratzer, A. 1995. Stage-level and Individual-level Predicates. In G. Carlson and F. J. Pelletier (eds.), *The Generic Book*, 125-175, University of Chicago Press, Chicago.
- Krifka, M. 1989. "Nominal Reference, Temporal Constitution, and Quantification in Event Semantics," in R. Bartsch, J. van Ebthe, and P. van Emde Boas (eds.), *Semantics and Contextual Expressions*, Foris, Dordrecht, pg. 75-115.
- Krifka, M. 1992. "Thematic Relations as Links between Nominal Reference and Temporal Constitution," in I. Sag and A. Szabolcsi, (eds.), *Lexical Matters*, CSLI Lecture Notes, University of Chicago Press, Chicago.
- Levin, B., (1993) *Towards a Lexical Organization of English Verbs*, Chicago: University of Chicago Press.
- Levin, B. and Rappaport, M. (1995). *Unaccusatives: At the Syntax-Lexical Semantics Interface*, Cambridge: MIT Press.
- Link, G. (1998) *Algebraic Semantics in Language and Philosophy*, CSLI, Cambridge University Press.
- Lloyd, G. E. R. 1992. *Polarity and Analogy: Two Types of Argumentation in Early Greek Thought*, Hackett Publishing Company.
- McCarthy, J., and P. Hayes. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, Eds., *Machine Intelligence 4*. Edinburgh: Edinburgh University Press, pp. 463-502.

- Miller, G. 1985. "Dictionaries of the Mind" in *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics*, Chicago.
- Miller, G. 1989. "Contexts of Antonymous Adjectives," *Applied Psycholinguistics* 10:357-375.
- Mohanan, K. P. and T. Mohanan. 1999. On Representations in Grammatical Semantics. In T. Mohanan and L. Wee (eds.), *Grammatical Semantics: Evidence for Structure in Meaning*, CSLI, Stanford.
- Moravcsik, J. M. 1975. "Aitia as Generative Factor in Aristotle's Philosophy," *Dialogue* 14:622-36.
- Morgenstern, L. (1995). The problem with solutions to the frame problem. In K. Ford and Z. Pylyshyn, (Eds.), *The Robot's Dilemma Revisited : The Frame Problem in Artificial Intelligence (Theoretical Issues in Cognitive Science)*, Ablex Publishers, pp. 99-133.
- Musan, R. (1997). *On the Temporal Interpretation of Noun Phrases*, Garland Publishing, New York.
- Parsons, T. (1990) *Events in the Semantics of English*, MIT Press, Cambridge, MA.
- Partee, B. 1992. "Syntactic Categories and Semantic Type," in M. Rosner and R. Johnson (eds.), *Computational Linguistics and Formal Semantics*, Cambridge University Press
- Pustejovsky, J. (1988). "The Geometry of Events," in *Studies in Generative Approaches to Aspect*, C. Tenny, ed., Lexicon Project Working Papers 24, MIT, Cambridge, MA.
- Pustejovsky, J. 1991a. "The Generative Lexicon," *Computational Linguistics*, 17:409-441.
- Pustejovsky, J. 1991b. "The Syntax of Event Structure," *Cognition* 41:47-81.
- Pustejovsky, J. (1995). *The Generative Lexicon*, Cambridge, MA: MIT Press.
- Pustejovsky, J. (2000a) "Lexical Shadowing and Argument Closure," in Y. Ravin and C. Leacock (eds.) *Polysemy: Theoretical and Computational Approaches*, Oxford University Press, Oxford.
- Pustejovsky, J. (2000b) "Type Construction and the Logic of Concepts", in P. Bouillon and F. Busa (eds.) *The Syntax of Word Meaning*, Cambridge University Press, Cambridge.
- Pustejovsky, J. (forthcoming). *Language Meaning and The Logic of Concepts*, MIT Press.
- Pustejovsky, J. and F. Busa. 1995. "Unaccusativity and Event Composition," in P. M. Bertinetto, V. Binachi, J. Higginbotham, and M. Squartini (eds.), *Temporal Reference: Aspect and Actionality*, Rosenberg and Sellier, Turin.
- Rapoport, T. R. 1993. "Verbs in Depictives and Resultatives," in J. Pustejovsky (ed.), *Semantics and the Lexicon*, Kluwer Academic Publishers, Dordrecht.
- Reiter, R. (1991). The frame problem in the situation calculus: a simple solution (sometimes) and a completeness result for goal regression. In V. Lifschitz, Ed., *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*. Boston: Academic Press, pp. 359-380.

- Ritter, Elizabeth and Sara Thomas Rosen. 1994. The independence of external arguments. In Erin Duncan, Dokna Farcas, Philip Spaelti, eds. *The Proceedings of WCCFL XII*. CSLI Stanford. pp. 591-605.
- Schein, B. 1993. *Plurals and Events*. MIT Press, Cambridge.
- Shoham, Y. (1988). *Reasoning about Change*. Cambridge, MA: MIT Press.
- Schubert, Len (1999) Explanation Closure, Action Closure, and the Sandewall Test Suite for Reasoning about Change, in H. Levesque and F. Pirri (eds.) *Logical Foundations for Cognitive Agents*, Springer Verlag, Berlin.
- Tenny, Carol. 1994. *Aspectual Roles and the Syntax-Semantics Interface*. Dordrecht: Kluwer Academic Publishers.
- Tenny, Carol. 1995. Modularity in Thematic versus Aspectual Licensing. Paths and moved objects in motion verbs. *Canadian Journal of Linguistics* 40(2):201-234.
- Verkuyl, H. J. 1993. *A Theory of Aspectuality. The Interaction between Temporal and Atemporal Structure*. Cambridge University Press, Cambridge.
- von Wright, G. 1963. *Norm and Action: A Logical Inquiry*, Routledge and Kegan Paul, London.