Contextually-Dependent Lexical Semantics

Cornelia Maria Verspoor

PhD
The University of Edinburgh
1997

Abstract

This thesis is an investigation of phenomena at the interface between syntax, semantics, and pragmatics, with the aim of arguing for a view of semantic interpretation as lexically-driven yet contextually dependent. I examine regular, generative processes which operate over the lexicon to induce verbal sense shifts, and discuss the interaction of these processes with the linguistic or discourse context. I concentrate on phenomena where only an interaction between all three linguistic knowledge sources can explain the constraints on verb use: conventionalised lexical semantic knowledge constrains productive syntactic processes, while pragmatic reasoning is both constrained by and constrains the potential interpretations given to certain verbs. The phenomena which are closely examined are the behaviour of PP sentential modifiers (specifically dative and directional PPs) with respect to the lexical semantic representation of the verb phrases they modify, resultative constructions, and logical metonymy.

The analysis is couched in terms of a lexical semantic representation drawing on Davis (1995), Jackendoff (1983, 1990), and Pustejovsky (1991, 1995) which aims to capture "linguistically relevant" components of meaning. The representation is shown to have utility for modeling of the interaction between the syntactic form of an utterance and its meaning. I introduce a formalisation of the representation within the framework of Head Driven Phrase Structure Grammar (Pollard and Sag 1994), and rely on the model of discourse coherence proposed by Lascarides and Asher (1992), Discourse in Commonsense Entailment. I furthermore discuss the implications of the contextual dependency of semantic interpretation for lexicon design and computational processing in Natural Language Understanding systems.

Declaration

I declare that this thesis has been composed by myself and that the research reported herein is my own. This thesis complies with all the regulations for the degree of PhD at the University of Edinburgh, and falls below the requisite word limit specified.

Cornelia Maria Verspoor

August 1997

Acknowledgements

I have been lucky to have had opportunities to work with many stimulating people during the course of researching and writing this work. This has surely helped to shape my thesis in ways I would never have been able to achieve on my own.

First and foremost I must thank my primary supervisor Alex Lascarides for her detailed and probing comments, delivered with amazing efficiency, and her constant encouragements. Claire Grover, my second supervisor, also deserves thanks for working through my representations and examples with a fine-toothed comb. Frank Schilder, Bernie Jones, and Alice Drewery have admirably tolerated my rantings about lexical semantics, and have even managed to solve a few sticking points for me. I would also like to thank Paul Buitelaar, Marc Verhagen, Jesse Tseng, Janet Hitzeman, Massimo Poesio, Gosse Bouma, and Ann Copestake for giving me feedback on various parts of this work.

In the early stages of this research, I spent some time at the Research Institute for Speech and Langauge (Onderzoeksinstituut voor Taal en Spraak, OTS) in Utrecht, the Netherlands, and I would like to thank Henk Verkuyl for discussing some underdeveloped ideas with me, and the OTS itself for making it possible. I also spent a few months at Brandeis University in Boston in 1996; this was an invaluable experience and I am grateful to Ray Jackendoff, Joan Maling, and James Pustejovsky for inviting me, making time to meet with me, and including me in everything. My colleagues in Boston, Vittorio Di Tomaso and Bhuvana Narasimhan, spent many tiring hours with me thinking about strange patterns in data, and were the inspiration for my work on manner of motion verbs. I thank the Doctor Catharine van Tussenbroek Foundation in the Netherlands and the Centre for Cognitive Science for making that trip possible.

Many thanks to the Engineering and Physical Sciences Research Council, the World Bank, and my parents who have provided financial support to me during my time in Edinburgh. My parents, Adriaan and Antoinette, although very far away in India, always managed to keep me motivated. I am sure that I would never have achieved so much in life without their open-minded encouragement to explore.

Finally, I would like to thank my husband Vicente Uceda, for his soothing effect on me, and the love, care, support, and understanding which have carried me through some crazy times.

Contents

List of Figures

List of Tables

one Introduction

> *What emerges very clearly from the recent work on the interface between lexical and non-lexical semantic information is that polysemy is not a single, monolithic phenomenon. Rather, it is the result of both compositional operations in the semantics [...] and of contextual effects, such as the structure of rhetorical relations in discourse and pragmatic constraints on co-reference. What is necessary is for research to tackle the difficult question of how other components in the natural language interpretation process interact with the lexicon to disambiguate and fully determine the semantics of the words in context. This work, emerging from very different traditions, illustrates how lexical semantics can be made sensitive to sentence level compositional processes as well as discourse level inference mechanisms, reacting to the diverse and multiple causes of lexical ambiguity.*
>
> (Pustejovsky 1995a:236)

Traditionally, representations of meaning have been used for the purpose of modeling inferences which can be made on the basis of the meaning of a phrase or sentence. Logical representations of meaning are rigorously defined and conclusions drawn from them can be proven within the formal framework established for the logic. As a linguistic representation, however, purely logical representations (i.e. representations devoid of syntactic information) are inadequate. Traditional approaches to the problem of how the meaning of a sentence can be built up from the meanings of the components of the sentence (e.g. Montague Grammar) are flawed in their dependency on strict rules of combination and their ignorance of the distinction between syntax and meaning. So, for example, a verb such as *drink* is semantically bivalent — it has two semantic arguments, the drinker and what is drunk — yet can be used in both the transitive (*John drinks milk*) or intransitive (*John drinks*) syntactic frames. A semantic grammar which assumes a strict semantics-syntax correspondence cannot account both for the two different syntactic forms in which this verb can appear and for the constancy of its semantic valency across these forms (the intransitive form still suggests that something is drunk). Furthermore, the inference mechanisms associated with classical logical representations are hampered by the fact that inference is not always a matter of drawing logically

valid conclusions; it might depend on a particular person's model of the world, the context in which that inference takes place, or on his knowledge. The interpretation of the intransitive *John drinks*, for example, can depend on the discourse or situational context in which it is uttered, or might be interpreted as *John drinks alcohol* without any information to the contrary (by default).

On the other hand, a theory of syntax which says nothing about the meaning associated with the syntactic structures entirely misses the point of language — what are language and its varied structures for if not to convey meaning? Furthermore, such a theory will contribute nothing to the problem of explaining how language is naturally used, and cannot account for the distinction between ungrammatical sentences (e.g. **John ice cream likes not*), infelicitous sentences (e.g. the famous *#Curious green ideas sleep furiously*), and fully grammatical and felicitous ones (e.g. *The baby sleeps quietly*).

To address such issues, recent work in linguistics and computational linguistics has attempted to construct representations which accommodate both syntactic and semantic information about a word. The representations of the lexicalist (word-driven) grammar framework Head-Driven Phrase Structure Grammar (HPSG) (Pollard and Sag 1994) provide an example of such a representation. These representations reflect an important step in the development of linguistic theory towards investigations of phenomena which lie at the intersection of syntax and semantics. The next step, as suggested in the quote from Pustejovsky above and as will be argued in this thesis, is to consider how word-level representations can be used to capture word usage and how these structures might interact with pragmatic reasoning.

In this thesis I will advocate a view of linguistic research which challenges modular investigations of language — i.e. research concentrating on some aspect of language use (syntax, semantics, discourse structure) in isolation from the others. I will show, through the example of several linguistic phenomena, that a model of language use must take the interactivity of distinct (linguistic and non-linguistic) knowledge sources into account in order to achieve a full explanation of the range of usage. Syntactic structure does not exist in isolation of meaning, while meaning is influenced through context. I will argue that a distinction does exist between syntax, semantics, and pragmatics, and that attempts to explain language understanding in a non-modular framework that conflates linguistic and non-linguistic knowledge (e.g. Hobbs *et al.* 1993) will not be able to account for conventionalised language use, but that the structure and functions of these modules must be considered in relation to each of the others. We will see in Chapter 5, for example, that the representation assumed for a verb depends on both the syntactic structure in which it is used, the representational structure of its complement, and its behaviour in the discourse context. Thus only through knowledge of a word's use *in context* — both the linguistic context (syntactic frame) and the discourse context — can an appropriate representational structure for it be established.

I am certainly not the first to make arguments for interactivity in linguistic process. A body of research exists which focuses on various aspects of linguistic interactivity: the theory of Conceptual Structures (Jackendoff 1983, 1990) investigates regularities in the relationship between meaning and syntax; Construction Grammar (e.g. Fillmore 1988; Goldberg 1995) argues that not all meaning is built up compositionally but that certain syntactic forms are

associated with particular interpretations and that these interpretations merge in consistent ways with the meanings of the constituents; the Generative Lexicon (Pustejovsky 1991, 1995a) is a theory based on the idea that the interpretations of words are generated through interactions between rich word-level representations; Persistent Default Unification (Lascarides *et al.* 1996) and related work (e.g. Lascarides and Copestake 1995; Lascarides, Copestake, and Briscoe 1996) looks at the interaction of default interpretations of words with world knowledge. This thesis therefore is a reflection of a more general shift in focus which has recently occurred in linguistics. The work presented here in fact builds on the research cited above, integrating ideas from each theory in order to create a research framework which allows investigation of the structure of word meaning required to explain language use in context.

I will examine phenomena where only an interaction of semantic, syntactic, and pragmatic information can fully explain the range of felicity of the phenomena, and I will argue along the way that structured information about word meaning plays a key role in this interaction. The examination of each phenomenon will result in proposals focusing on different aspects of word-level representation, but each proposal fits into a single framework, outlined in Chapter 2. The proposals are data-driven; reflecting as far as possible the range of usage of the constructions examined. To that end, the thesis contains many example sentences and even some corpus analysis (Chapter 5).

1.1 Definitions and Assumptions

This thesis will employ a certain amount of linguistic terminology, which I will clarify before continuing. Here I define several key terms:

- *Lexicon:* a collection of representations for words used by a linguistic processor as a source of word-specific information. These representations might contain information on the morphology, phonology, syntactic argument structure, and semantics (e.g. valency, thematic roles) of the words.
- *Semantics:* the context-independent meaning of something (a word, phrase, sentence); its truth conditions.
- *Lexical Semantics:* the semantic structure associated with a word, as represented in the lexicon.
- *Pragmatics:* the study of the relations between language and context that are basic to an account of language understanding; including the interaction between specifically linguistic knowledge and more general world knowledge.
- *Discourse:* conversational interactions which are governed by principles such as coherence and cooperation.

Implicit in these definitions are various assumptions. Most important is the assumption that there is a distinction between linguistic and world knowledge, as other assumptions follow on from this. This point will be discussed in detail in Chapter 2; the fundamental

motivation for this distinction derives from evidence of idiosyncratic uses and interpretations of words which could not follow from general reasoning on world knowledge and a desire to isolate language-specific knowledge from general knowledge which does not vary among speakers of different languages. The implication of this point for the lexicon and lexical semantics is that words in the lexicon do not necessarily correspond directly to concepts, that the structure of the lexicon (i.e. relationships between lexical items) will not necessarily reflect the structure of the world knowledge associated with the denotations of the words in the lexicon, and that the information in lexical entries is restricted to information which has direct relevance for truth-conditional meaning and syntax. The last point means that there is justification for postulating that a certain piece of knowledge is lexicalised *only* when a particular linguistic phenomenon or a particular interpretation of something cannot be explained purely in terms of world knowledge. The full spectrum of knowledge associated with the concept referred to by a word should not be present in the lexicon; only conventionalised linguistic information should appear there. This assumption will have implications in particular for my analysis in Chapter 5.

An additional assumption which follows from the first is that pragmatic reasoning works with the *result* of linguistic processing — that is, it has access to the logical form which results from syntactic and semantic combination but not to information encoded in the lexicon or syntactic constraints. This means that the interaction between linguistic knowledge and world knowledge occurs through manipulation of a logical form, and that any linguistic information which is relevant for this interaction must somehow be represented in this logical form. Any syntactic or lexical semantic constraints on the interpretation of a sentence must be taken into consideration during the construction of a logical form. This assumption will inform the analyses utilised in Chapters 3 and 4, in which particular syntactic combinations of elements in a sentence will affect the interpretation of that sentence.

Much of the dicussion of discourse constraints and pragmatics in this thesis comes from an intuitive, data-driven perspective. What formal elements I introduce derive largely from work by Asher (1993b, 1993a), Asher and Lascarides (1995a, 1995b), and Lascarides and Copestake (1995), but formal models of discourse coherence are still in the early stages of development and do not yet fully reflect the complexity of discourse-level reasoning. I have tried to harness the intuitions which underpin these models in my discussions, in order to argue for consideration of the influence of discourse context on interpretation. My goal has not been to identify weaknesses in these models or to develop an alternative model; I aim solely to suggest how a model of discourse coherence might interact with the results of syntactic and lexical semantic processing to explain difficult or apparently idiosyncratic data.

1.2 Sense variation

The core chapters of this thesis (Chapters 3-5) examine particular linguistic phenomena which involve interactions of lexical semantic structure. Each phenomenon examined involves subtle shifts in the meaning conveyed by the main verb in a sentence, due to the influence of complements or adjuncts. In each case, meaning is inferred which does not appear to be explicitly

present in the sentence. The main question which the chapters address for each distinct phenomenon is, where does this extended meaning come from? In each case, the solution will involve lexical specification of semantic structure. Furthermore, that lexical semantic structure will be shown to interact with constraints deriving from discourse coherence and pragmatic reasoning.

This thesis essentially investigates sense variation of verbs and how that sense variation can be modeled through use of the lexicon and the information formalised therein. Words are often *polysemous* — they often have several different but related meanings — and it has been recognised that the variations in meaning which a particular word shows often can be generalised to entire classes of words (see e.g. Levin 1993 for a thorough overview of verbal syntactic alternations and their corresponding semantic alternations). Recent lexical research has focused on characterising those classes and modeling the alternations in terms of lexical semantic structure (Pinker 1989, Jackendoff 1990, Markantonatou and Sadler 1995, Sanfilippo 1995, Pustejovsky 1995a, *inter alia*) but little of the insights from this work has influenced lexicon design for computational systems or constraint-based grammars.

There are various ways in which sense variation has been captured:

- *Enumeration:* Different senses of the same word can be listed as distinct lexical entries. This could also include listing of all the different subcategorisation frames in which a word can appear, with each frame associated with its own specific semantics, which might ignore the independent contribution of the subcategorised complements. This approach might be combined with constraints specified in the lexical entries about the kind of arguments a word expects, which help to resolve the sense of a word intended on a particular use.

- *Lexical rules:* A rule can be added to the lexical component which captures a regular meaning shift of a certain class of words. The rule has an *input form* which specifies the kind of word which undergoes the shift, and an *output form* which is the shifted version of the word. The input and output forms are distinct lexical entries for the same word, with different semantic and possibly syntactic information. Lexical rules capture generalisations and can be utilised productively for words which match the required input structure, to produce unestablished or novel word senses.

- *Type coercion:* Pustejovsky (1991, 1995a) suggests that some sense variation can be captured using operators which shift the semantic type of an entity, turning for example an object-referring noun into an event-referring noun by a well-defined generative operation. This operation might be triggered by lexical constraints (e.g. a conflict in the actual type of a verbal argument and the type that the verb expects it to be). This approach differs from the lexical rule approach in that the operators are generally available (not specific to particular word classes) and will only be applied if necessary to resolve a constraint violation.

The enumeration approach fails to record any regularities in sense variation which exist and is a weak lexical model in that it is not productive. The lexical rule approach does allow the

capture of these regularities, but as it is restricted to manipulation solely of lexical entries for individual words it cannot be applied in cases in which sense variation stems from properties of non-subcategorised elements in a sentence (e.g. adjuncts). The type coercive approach suffers from a lack of constraints — it is not always desirable to attempt to shift the type of a word when there is a conflict in types, as the word may simply be infelicitous in that context. These approaches are, however, complementary rather than competing, since each handles some aspect of the creativity of language use.

The approach taken in the thesis will build on the lexical rule and type coercive approaches, in order to achieve maximum generality in the representation. I will utilise lexical rules which can operate at both the word and phrasal levels, by replacing the standard HPSG mechanisms for adjunct addition with lexical rules which add adjuncts to the SUBCAT lists (see Section 3.5). I will identify specific phenomena in which type coercion does seem to occur, and will discuss how the coercion can be represented in the lexicon, and what constrains the coercion. In addition, the use of a hierarchical typed lexicon allows generalisations over related groups of words to be stated at one node in the lexicon and passed down to lower nodes through inheritance.

Furthermore, the implementation of a lexical semantic representation which builds in elements which have been identified in the previous research to be important for modeling alternations and fits directly into the grammatical framework of HPSG (subject to the modification mentioned above) opens up the possibility of capturing generalisations about at least some sense variations in computational systems. As I will argue in Chapter 6, this is important for development of robust systems which can adequately handle the interpretation of natural language, given the flexibility of language use. The proposals I make are for productive processes generating extended senses for verbs, which are constrained by syntactic, lexical semantic, or discourse factors. These constraints must be in place in order to prevent *overgeneration* of meanings, that is, generation of interpretations for a sentence which are infelicitous or unattested in the language.

1.3 Chapter summaries

Chapter 2 discusses issues of lexical semantic representation, arguing for a certain amount of decomposition in lexical representation in order to efficiently capture generalisations about the behaviour of semantically and/or syntactically related words. In this chapter I argue for a distinction between linguistic knowledge and world knowledge, and suggest that the ontology underlying lexical representation should reflect word-specific linguistic knowledge, including information about conventional usages and default interpretations.

I propose a lexical representation for which the level of decomposition stems from insights by Jackendoff (1983, 1990) and Pustejovsky (1991, 1995a), and for which the formal details derive from proposals by Davis (1995). The representation integrates with the framework of HPSG, overriding the standard representation of semantic relations in that theory with one which captures relationships between word meanings in a multiple-inheritance hierarchy.

Chapter 3 considers prepositional phrases (PPs) and variations in their contribution to the

meaning of the sentences they appear in. This chapter discusses the status of various PPs, focusing on dative PPs such as those in (1.1) in contrast to adjunct PPs such as those in (1.2), and investigates how to account for the different ways a particular PP can interact with the meaning expressed in the modified main clause. An example of this variation is the contrast in the interpretation of the PP *for Mary* in (1.1b) and (1.2c). In the former example the PP specifies someone who benefits from an object (*a cake*), and in the latter it specifies the beneficiary of the full event specified in the main clause (the *jogging a mile* event).

(1.1) a. John gave a book to Mary.

b. John baked a cake for Mary.

(1.2) a. John jogged in the park yesterday.

b. John jogged for twenty minutes twice a day.

c. John jogged a mile for Mary.

I review proposals by Kasper (1993) and van Noord and Bouma (1994) for adjunct incorporation in HPSG, and present an analysis of the contribution of the PPs which depends on the lexical semantic structure of both the modified verbs and the modifying PPs. The analysis is justified in terms of the syntactic and semantic properties of the PP data. The proposals here will also provide the basis for accounting for the sense variation of manner of motion verbs in the following chapter (compare (1.2c) with (1.3)).

(1.3) John jogged a mile to the store.

Chapter 4 investigates manner of motion verbs which can acquire a directed motion interpretation, shown in (1.4), comparing and contrasting them with constructions known as *resultatives*, as exemplified by the sentences in (1.5).

(1.4) a. The horse jumped.

b. The horse jumped over the fence.

(1.5) a. John hammered the metal flat.

b. The river froze solid.

c. John sneezed the tissue off the table.

d. John laughed himself silly.

These constructions are associated with a resultative interpretation which is not inherent in the meaning of the main verb (so (1.5a), for example, means *John hammered the metal and as a result the metal became flat*), and often violate normal subcategorisation requirements of the main verb (contrast (1.5c) and (1.5d) with **John sneezed the tissue* and **John laughed himself*).

The two constructions exemplified in (1.4)-(1.5) have in the past been modeled with a single analysis (e.g. Levin and Rappaport Hovav 1995). I will argue on the basis of semantic intuitions and contextualised data that these constructions are different, and that their behaviour must be explained in distinct ways. I will show how the manner of motion

verbs acquire their extended sense due to semantic properties encoded in their lexical representations, and how the resultative construction can be explained in terms of a combination of construction-specific semantics, conventionalisation, and contextual constraints.

Chapter 5 considers the phenomenon of *logical metonymy*, constructions in which a word (verbs in the following examples) requiring an eventive complement appears with an object complement. In these cases the event which fills the relevant semantic argument position of the word is not explicit in the sentence and must be inferred. For the sentences (1.6)-(1.7), for example, the (a) sentences can be interpreted as having a meaning identical to that explicitly expressed in the (b) sentences.

(1.6) a. John began the book.

b. John began reading / to read the book.

(1.7) a. John enjoyed the book.

b. John enjoyed reading the book.

The chapter discusses how that interpretation process might proceed, and argues that this phenomenon is governed by lexically specified usage conventions. The content of individual lexical entries is shown to depend on linguistic knowledge, including discourse-level knowledge, particular to the words they represent.

Chapter 6 discusses the implications of the contextual dependency of semantic interpretation for lexicon design and computational processing in Natural Language Processing (NLP) systems. It focuses on the problem of polysemy, and the issues of lexical representation and lexical acquisition for NLP systems which stem from the sense variation which words can exhibit. I will discuss automatic lexicon acquisition techniques from machine-readable dictionaries and corpora, and will argue that this acquisition must necessarily be guided by a theoretical linguistic framework.

Chapter 7 brings the thesis to a close, highlighting the main conclusions of the thesis and suggesting a few areas of future investigation. The broad result which can be drawn from the analyses in this thesis is that syntactic and semantic phenomena cannot be considered in isolation of one another, or in isolation of the context in which they occur. An account of linguistic phenomena which aims to accurately model the full range of usage — from ungrammatical and infelicitous constructions to grammatical and fully felicitous constructions — must consider the interaction of the distinct knowledge sources contributing to the determination of the meaning of a sentence in context.

two Lexical Representation

2.1 Introduction

The line between lexical semantics and world knowledge has been notoriously difficult to draw. As observed by Copestake (1992:1-2), however,

> [I]t is methodologically important to distinguish between linguistic and non-linguistic representation, even though the two have to be interrelated so that linguistic utterances can be interpreted as having some connection with the real world... we wish to provide a testable constrained theory, and a formal representation language, and to avoid problems which arise in knowledge representation which do not have a linguistic dimension.

The view of lexical semantics advocated by Copestake is that lexical representation must include information about a word which is necessary to account for its grammatical use, but which is not entirely predictable from real world knowledge of the entity denoted by the word. Furthermore, Copestake argues that a certain amount of real world information must be lexically encoded in order to describe interactions of lexical semantic variations with syntax (for example, Copestake extensively discusses the count / mass distinction which has a syntactic reflex and which in some cases depends on lexical specification[1] and in others depends on regular lexical sense extension properties[2]).

In this thesis I will adopt a rich lexical semantic structure, involving a certain amount of lexical decomposition and hierarchical relations between lexical semantic structures, in order to explain a variety of linguistic phenomena. My point of departure in choosing an appropriate representation is the assumption that the inferences which can be drawn about the use

[1] Copestake's examples of this are the fact that a pile of small feathers can be referred to as *feathers* or *down* and that in Italian *spaghetti* is count while in English it is mass.

[2] For example, the lexical process of *grinding* underlies the use of *rabbit* on an animal use in *The rabbit ran across the field* (count) and on a meat use in *We had rabbit for dinner* (mass).

of a particular word in a particular context must play a role in determining lexical semantic structure. That is, there are regularities in the interaction between syntactic structure and word meaning which stem from critical semantic features of the word. Without a lexical semantic representation of such features it is impossible to capture those regularities. This point has been extensively argued by Davis (1995), who develops a theory of word order based entirely on semantic relations encoded in the lexical entries of verbs (i.e. a linking theory).

A purely pragmatic account of word meaning cannot explain the syntactic reflexes of semantic shifts, unless the pragmatic reasoning is assumed to have access to purely syntactic information. This is the line taken by Hobbs *et al* (1993). However, on such a view the pragmatic module cannot be language independent since it will be influenced by and would depend on syntactic properties which vary between languages. A pragmatic component which interacts with language dependent syntactic information is undesirable for several reasons.

Firstly, it forces the linguistic system to be modeled nonmodularly since syntactic, semantic and pragmatic information must be considered simultaneously in the construction of an interpretation of a sentence or discourse. This obscures the distinction between general conceptual knowledge and specifically linguistic knowledge, critical for explaining certain aspects of language use. Consider the sentences in (2.1).

(2.1) a. John met Sue on / *in Tuesday.

b. John met Sue in / *on the morning.

c. John met Sue in / *on March.

d. John met Sue (*in / *on) last night.

e. John met Sue at / *in / *on ten o'clock.

From a world knowledge perspective, phrases like *March*, *Tuesday* and *the morning* pick out particular portions of time, of varying lengths. Periods during the day are subintervals of days of the week, which are subintervals of the weeks of the month, which are in turn subintervals of the months, etc. These interval relations can be conceived of hierarchically, with the days being subtypes of the week, the weeks of the month, and so on. In this knowledge, however, there is nothing to explain the fact evidenced by the data in (2.1) that the intervals of time corresponding to days are picked out by a different preposition from other such intervals (*on* vs. *in*). It cannot follow simply from differences between the relationship of the described event *John met Sue* to the specific time interval in the prepositional phrase (PP); in each case that event is interpreted to take place during a sub-interval of the interval picked out by the PP. The choice of preposition is instead purely a matter of linguistic convention (in English). Days and months are typed differently at a linguistic level but not in the world knowledge ontology. The distinction between language-dependent usage facts and facts about the world must be made explicit in the language model.

Secondly, a nonmodular approach leads to overgeneration of word senses within a productive linguistic system. Imagine that speakers of English know that it is possible to refer to the meat of an animal by using the name of the animal, as shown in (2.2a). This can be considered a productive generalisation as it can be applied in novel instances (2.2b).

(2.2) a. We ate chicken / lamb / turkey for dinner last night.

b. We ate kangaroo / aardvark for dinner last night.

c. #We ate cow for dinner last night.

d. We ate beef for dinner last night.

(2.2c-d) exemplify, however, a process known as *lexical blocking* (Copestake and Briscoe 1995): the application of a productive rule can be blocked by the presence of a lexeme with the same meaning (but different form) as the lexeme which would be generated by the rule. In this case, the use of *cow* for the meat of a cow is blocked by the independent existence of the word *beef* in the lexicon. Lexical blocking must be considered a linguistic phenomenon which occurs in isolation of world knowledge in that it depends purely on lexical information and not on concepts associated with the denotation of a word. A purely pragmatic account of word meaning could therefore not account for such processes which restrict sense generation.

The necessity of representing lexical semantic information is clear: simply treating sentential constituents as syntactic entities cannot result in an analysis which is fine-grained enough to account for the range of grammaticality judgements about sentences. It is necessary to take into account the semantic nature of entities combining into larger constituents in order to adequately capture grammatical sentences. These observations have influenced the definition of constraint-based grammatical theories: HPSG (Pollard and Sag 1994), for example, explicitly allows for the encoding of both semantic and contextual information in the structures representing words and phrases. This information interacts with syntactic specificatons in the construction of grammatical sentences.

An example of a grammaticality difference which cannot be captured through syntax alone is shown in (2.3)-(2.4). Why is it possible for one spatial prepositional phrase to modify *close* and not a different one, while *push* does not show this restriction? There is clearly no surface difference in the structures of (2.4b) and (2.4c),[3] so the difference must be due to something semantic.

(2.3) a. John pushed the filing cabinet.

b. John pushed the filing cabinet at the office.

c. John pushed the filing cabinet into the wall.

(2.4) a. John closed the filing cabinet.

b. John closed the filing cabinet at the office.

c. *John closed the filing cabinet into the wall.

As I will argue below, Davis' (1995) representation cannot be used to explain all linguistic phenomena which depend on lexical semantic features, because it concentrates entirely on the entailments relevant to the lexical semantics-syntax interface. I will therefore extend his representations with further semantic decompositions derived from the work of Jackendoff

[3]Note that there is structural ambiguity in both (2.3b) and (2.4b). The prepositional phrase *at the office* can be either a NP or a VP modifier. We assume the VP attachment reading for the purposes of this discussion.

(e.g. 1983, 1990). This work has provided many insights with respect to the identification of components of lexical semantic structure which influence syntax and the encoding of generalisations about inferences to be drawn from uses of a wide range of verbs and prepositions.

In this chapter, I will outline the approach to the representation of lexical semantic structure, *conceptual semantics*, which Jackendoff has proposed and the constraint-based version of lexical semantic structure put forth by Davis. I will argue that the framework and analytical methodology Jackendoff develops will give clear structure to the endeavour of capturing linguistically-relevant semantics, and will show how Davis builds on this framework and brings the representations in line with the constraint-based grammar framework of Head-driven Phrase Structure Grammar (HPSG). I assume that conceptual structures are representations of those aspects of the meaning of words and phrases which are relevant to syntax and over which particular, regular inferences can be made, as it is in this sense that these structures can be utilised within a computational framework. I therefore avoid a complete representation of knowledge about entities denoted by a word in that word's representation in the lexicon. The attempt to capture all knowledge associated with a word at the lexical level of representation would place too much pressure on the lexicon, as follows from the quote from Copestake above. The representation I settle on will explicitly reflect a linguistic, and not a world knowledge, ontology.

2.2 Jackendoff and syntactically-relevant semantics

Ray Jackendoff (e.g. 1983, 1990) develops a representation for *conceptual structure*, or entities which he argues reflect human knowledge of the world. His methodology for identifying the constituents of such structure includes analysing alternations in the linguistic context in which particular words (mainly verbs and prepositions) can be used, to identify generalisations over relations between alternate uses of lexical items. These generalisations can be captured in regular lexical structure. Components of the lexical representation are assigned a consistent semantics and can be combined in constrained ways.

Given this methodology, it is clear that Jackendoff focuses primarily on syntactically-relevant lexical semantics, although he argues at length in favor of viewing conceptual structures as entities which capture *concepts* or *thoughts*. His analyses, then, can be used to address the problem of capturing regularities in the lexicon which have a direct influence on syntactic structure. Below, I will introduce some of these analyses and argue that current approaches in computational lexical semantics could benefit by incorporating them.

2.2.1 Components of Jackendoff's Conceptual Semantics

Semantic Decomposition

It is widely accepted at this stage of investigation into lexical semantics that some form of lexical decomposition is necessary to capture generalisations about the relationship between syntactic form and intended meaning (Jackendoff 1983, 1990; Dowty 1979; Pinker 1989; Pustejovsky 1991, 1995a; Davis 1995; *inter alia*). It is equally well accepted that it is impossible to

decompose meaning into necessary and sufficient conditions for identification of the entities words correspond to, for it will always be possible to discover some additional element of meaning which is needed to distinguish between two arbitrarily closely related words (Pulman 1983; Medin 1989; Keil 1989).

The position which Jackendoff advocates with respect to semantic decomposition is that word meanings must have internal structure, due to the creativity of language use (i.e. the ability of speakers of a language to understand and create an indefinitely large number of sentences with which they have no prior experience) and the regularities which accompany that use, and that part of this structure is the specification of necessary conditions for the application of a word. He then goes on to attempt to identify these necessary conditions. The decomposition of word meaning into smaller semantic elements allows specification of a generative, compositional system which constrains the way such elements can be related and thereby constrains the ways in which sentences can be constructed (to prevent semantically anomalous sentences), while not attempting to predict *a priori* what structures will actually be created and used. It is clearly desirable from a computational perspective to develop a system with such generative properties (this point is strongly argued by Pustejovsky 1991, 1995a), and thus an analysis of the framework which Jackendoff proposes is warranted. In particular, since his framework largely derives from the analysis of linguistic alternations, it is likely to provide precisely the grain of decomposition necessary for developing the linguistic representation Copestake so persuasively advocates in the quote at the start of this chapter. This follows from the fact that linguistic alternations, such as the locative alternation exemplified by (2.5), have subtle semantic distinctions (Levin 1993, Levin and Rappapport Hovav 1995) which determine the syntactic position of the arguments of the verb, and can be used to predict when a verb will not participate in an alternation (Davis 1995, Markantonatou and Sadler 1995). These semantic distinctions therefore directly interact with syntactic form and must be considered a part of lexical knowledge rather than world knowledge about the verb forms.

(2.5) a. Bill loaded books onto the truck.

b. Bill loaded the truck with books.

A Semantic Ontology

Jackendoff (1983, 1990) proposes an ontology which provides the building blocks of conceptual representation. Each representational constituent corresponds to one of the ontological categories: Thing, Event, State, Action, Place, Path, Property, Manner, or Amount. These categories are motivated linguistically in several ways.

1. Each of the ontological categories permits the formation of a wh-question (2.6) (from Jackendoff 1983:53), except State. States are characterised by their contrast with Events, as shown in (2.7).

(2.6) a. What did you buy? Thing

b. What happened next? Event

c. What did you do? Action

d. Where is my coat? Place
e. Where did they go? Path
f. What was she like? Property
g. How did you cook the eggs? Manner
h. How long was the fish? Amount

(2.7) a. What happened was that { Bill flew around the pole. / the rock fell off the table. } Events

b. ?What happened was that { Max was in Africa. / the statue stood in the park. } States

2. Each major syntactic constituent of a sentence corresponds to a conceptual constituent. Thus in *John ran toward the house*, *John* and *the house* are Things, the PP *toward the house* is a Path, and the entire sentence is an Event (example from Jackendoff 1990:22).

3. Several of the ontological categories support quantification, as in (2.8) (from Jackendoff 1990:23).

(2.8) a. Every dinosaur had a brain. Things
b. Everything you can do, I can do better. Actions
c. Anyplace you can go, I can go too. Places
d. Bill can shuffle cards in every way that Mary can. Manner

These categories overlap to a certain degree with the types already assumed in standard HPSG, although there is not an exact equivalence as Jackendoff's categories concentrate on more fine-grained semantic distinctions than are incorporated in HPSG. They can, however, be integrated into HPSG quite straightforwardly, as we will see in Section 2.3.1.

Jackendoff's conceptual structures are built up using constituents in the ontology. Each constituent can be decomposed into function-argument structure. Functions impose conceptual constraints on the nature of and relations between function argument(s), which are themselves required to correspond to conceptual constituents. So, for example, a Place may be characterised as *Place-function(Thing)* and a Path might be characterised as *Path-function(Place)*. Jackendoff seeks to identify functions which explain grammatical patterns of combination, using the ontology as a reference point. The existence of such an ontology provides a starting point for the identification of generalisations about how words can combine to form phrasal constituents and how phrasal constituents can combine to form larger constituents. Due to the creative capacity of language, it is clear that we do not simply learn lexically specific ways of putting words together — there are regular, productive patterns of combination which apply to classes of words and phrases. By allowing for the mapping of words and phrases to general classes, the ontology supports the identification of such patterns of combination.

Why is such an ontology useful? It divides the linguistic domain up into types of entities to which each sentential constituent may correspond, thereby enabling the definition of broad constraints over the mapping between syntactic constituents and semantic entities. It further allows the definition of functions ranging over ontological entities, which indicate the ways

in which entities can combine semantically. These functions are not fixed in advance by the ontology, but rather can be determined via linguistic analysis, in order to capture the range of relations expressible in language. The output of such functions is constrained to correspond to an entity in the ontology, thereby ensuring that there are at least very broad restrictions on what functions are possible. Thus the ontology reflects the types of entities which are expressible in language and guides the identification of coherent relations among these entities.

The definition of the ontology which Jackendoff uses stems in part from psychological claims about the projection of entities in the real world to a mental representation of those entities. Although the cognitive processes of categorisation and so forth are possibly not of direct concern to the computational lexical semanticist, the linguistic motivations for a semantic ontology as introduced here and the framework which such an ontology provides are reasons for adopting one. There are in addition semantic reasons for the adoption of an ontology, stemming from "pragmatic anaphora" (Hankamer and Sag 1976), in which anaphors like *this* and *that* can refer to entities of specific ontological categories compatible with the linguistic context,[4] and psycholinguistic and developmental studies (see Pinker 1989 for an overview). Abstraction over cognitive structures is necessary to account for children's ability to learn language. Furthermore, Talmy (1983, 1988) has shown that most verb meanings cross-linguistically are built around recurring elements of meaning and their combinations. These studies strongly suggest that the adoption of an ontology is necessary to capture generalisations about language use.

The ontology as Jackendoff has proposed it, however, does not assume any explicit subdivision of the ontological categories, although he implicitly assumes a richer structure. These kinds of relations can easily be captured in a hierarchical ontology which has the conceptual categories introduced by Jackendoff at its top. This ontology would reflect categorisations of entities, events, etc. and relations between them. It therefore captures (a part of) world knowledge. Such an ontology is basic to constraint-based theories of grammar (e.g. HPSG) and have been widely used in computational approaches to natural language processing (see Copestake 1993 for an example). I will assume that an inheritance-based hierarchy is a critical

[4]When one utters sentences such as those in (2.9) (from Jackendoff 1983:49), there are constraints as to what each of the italicised elements can refer to.

(2.9)
a. Your coat is *here*. [pointing]
b. He went *thataway*. [pointing]
c. Can you *do that* / *do this*? [pointing / demonstrating]
d. You shuffle cards *this way*. [demonstrating]
e. The fish that got away was *this* long. [demonstrating]

These constraints are imposed by the lexical semantic contexts in which the elements appear. The hearer of such sentences can only interpret the anaphors in a way compatible with the type of entity required by the context. This "entity type" must correspond to a particular ontological category. This is evidence for use of the ontological cateogories as building blocks for conceptual structure – without such general categories there would be no way to constrain the entities that could fulfill function argument positions in structures capturing the meaning of particular expressions.

representational component.

However, the lexical ontology which I will assume will be a *linguistic* ontology, following HPSG, as opposed to a *world knowledge* ontology, as emphasised in the introduction to this chapter. This means that information which is specifically relevant for the relationship between word forms and meanings will be represented here. For *drink*, for example, we must encode the information that the intransitive form of this verb has a specific default interpretation. That is, without more specific information (2.10a) means (2.10b) (Lascarides and Copestake 1997).

(2.10) a. John drinks all the time.

b. John drinks alcohol all the time.

This default interpretation cannot be explained solely on the basis of world knowledge since probabilistically *drinking alcohol* is not prima facie more likely than the drinking of other kinds of liquids. I will take advantage of the generalisations that can be made using an ontology, incorporating the notion of defaults which are associated with individual (groups of) lexical items.

Thematic role relations

Much of Jackendoff's analysis is motivated by thematic role relations, i.e. the actor vs. patient relation and motion and location roles. He argues that actor / patient relations are orthogonal to other thematic relations, and formalises this distinction explicitly by introducing an Action tier and a Thematic tier into the representation of Events. The Action tier captures the actor / patient relations, while the Thematic tier captures the motion and location relations. The question for a computational linguist is whether it is necessary, or, more weakly, simply useful, to represent such relations explicitly in lexical semantic structure.

There has been much research, particularly within the Government and Binding grammatical framework, focused on how thematic relations can be used to explain syntactic phenomena such as binding and control (e.g. Gruber (1965), Fillmore (1968), Jackendoff (1983, 1990), Grimshaw (1990)). There have been other attempts made at explaining these phenomena (e.g. within HPSG), but the results of research in this area are far from conclusive. Thematic relations may prove to be the best way of capturing the phenomena.[5] In that case, it will clearly be in the interest of anyone interested in the syntax / semantics interface to include thematic roles in lexical semantic representation.

An explicit example in which the representation of, at least, Actor / Patient relations seems to be critical can be found in the sentences in (2.11) (from Jackendoff 1990:130).

(2.11) a. What Bill did to the books was load them on the truck.

b. ?What Bill did to the truck was load the books onto it.

c. *What Bill did to the books was load the truck with them.

[5]See discussion in Pollard and Sag (1994), pp. 275-277.

d. What Bill did to the truck was load it with books.

In each case, the expressed event is essentially the same: the books go onto the truck. What changes, as reflected through the syntactic form utilised, is which entity is viewed as the Patient — which entity is most directly "affected" by Bill's action. (2.11a-b) identify *the books* as the Patient, as that element (or an anaphoric reference to it) is in direct object position in the embedded sentence, while (2.11c-d) identify *the truck* as the Patient. The reason for the reduced acceptability of (2.11b-c), then, is that there is a conflict between the entity (X) explicitly identified as being acted upon by Bill, *Bill did something to* X, and the entity suggested to be the Patient by the syntax of the embedded sentence. This example suggests that at least the semantic notion of Patient is relevant to the mapping between syntax and semantics, and that this thematic relation must be represented in the semantics to account for the distinction between the pairs of parallel sentences above. Parallel examples can be found to support the necessity of representing other thematic relations.

An additional active area of research with respect to the syntax / semantics interface is that of linking theory, a theory attempting to explain regularities in the mapping between semantic and syntactic arguments. So, for example, it is in the domain of a linking theory to explain why there is no verb *quain* in English such that (2.12b) is synomous with (2.12a) (Davis 1995).

(2.12) a. The child trained / called / petted / fed / kicked the dog.

b. The dog quained the child.

Jackendoff (1990) and Davis (1995) *inter alia* argue that generalisations about linking depend on a lexical semantic representation of thematic roles.[6] Elements of verb meaning, such as causation, can influence the syntactic realisation of a verb's arguments, and it is therefore necessary to embed the notion of thematic roles into a richly structured semantic representation. Linking rules can then use this structural information to determine how semantic arguments surface syntactically.

A principled explanation of the syntax / semantics interface, then, seems to depend on the notion of thematic relations within lexical semantic representations. Adopting linking theory within a computational framework would free the lexicographer from having to specify explicitly in the lexicon how arguments in a verb's subcategorisation frame correspond to its semantic arguments. In fact, it is possible that subcategorisation may be almost entirely semantically determined[7], thereby largely eliminating the need for ad hoc specification of subcategorisation frames for individual verbs.[8] The predictive capacity of linking theory would therefore clearly reduce the amount of word-specific syntactic information which would have to appear in the computational lexicon. A linking theory is therefore desirable from both a theoretical and a practical perspective.

[6] Other researchers, e.g. Dowty (1991) and Wechsler (1991), deny that an explicit representation is necessary, instead arguing that judgements made on the basis of certain lexical entailments are sufficient for explaining linking.

[7] Idiomatic expressions may in some cases violate general principles, and there may be some lexical exceptions which behave unpredictably.

[8] An example of such an approach as applied to the causative alternation can be found in Johnston (1995). Furthermore, Davis (1995) introduces thorough inheritance-based proposals for constraining the relationship between a verb's semantic and syntactic arguments.

In sum, the representation of thematic role relations within the lexicon seems to be warranted by its necessity within explanatory frameworks for various syntactic phenomenon, and ultimately for a clearly defined, principled syntax / semantics interface.

Functions relating entities

As discussed above, Jackendoff's work aims toward identification of function-argument structures which capture semantic relations between entities. In particular, he wishes to distinguish at the lexical semantic level between related verbs or uses of verbs which differ in syntactic realisation or in the types of modifiers which can appear with the verbs in order to account for the differences in their usages.

If we examine verbs conveying spatial information, we find a consistent distinction between verbs conveying motion of something along a path (e.g. (2.13)) and verbs specifying the location of something (e.g. (2.14)). The prepositional phrases in (2.13) contrast with those in (2.14) in that the former are Paths while the latter are Places. Additionally, while the verbs of motion can appear with Place PPs (e.g. (2.15)),[9] the verbs of location are wholly incompatible with Path PPs (e.g. (2.16)).

(2.13) a. The fly flew around the room.
b. The balloon floated out the window.
c. John ran through the tunnel.

(2.14) a. The book is in(side) the room.
b. The statue stands on the floor.
c. The picture hangs in(side) the tunnel.

(2.15) a. The fly flew in(side) the room.
b. The balloon floated outside the window.
c. John ran in(side) the tunnel.

(2.16) a. *The book is around the room.
b. *The statue stands onto the floor.
c. *The picture hangs through the tunnel.

These verbs must be differentiated lexically in order to account for the distinction in the types of PP arguments each class takes as arguments, as there is no way to distinguish them syntactically.[10] Jackendoff proposes (2.17a) as the underlying representation for the motion verbs and (2.17b) for the verbs of location. The representations constrain the type of PP arguments each verb type can have, and specify a relation between the two verbal arguments.

[9] Note that on the interpretation I am interested in for the sentences in (2.15), however, the PP behaves as an adjunct rather than a verbal argument. The location is perceived as the location where the event takes place rather than the location of the entity referenced by the subject NP.

[10] Unless we assume extremely fine-grained syntactic categories which would miss an important generalisation about the relation between this data and its semantic basis.

(2.17) a. $[_{\text{Event}}\ \text{GO}([\text{THING}], [\text{PATH}])]$

b. $[_{\text{State}}\ \text{BE}([\text{THING}], [\text{PLACE}])]$

This example illustrates how identification of distinctions in function-argument structure can aid in the syntax / semantics mapping and the modeling of appropriate interaction between verbs and their arguments. A full presentation of Jackendoff's proposed function-argument structures is not necessary in this context; I simply wish to convey that the decomposition of lexical meanings in terms of such structures can be used to account for the potential range of grammatical and semantically felicitous sentences.

2.2.2 Jackendoff and inferences

Jackendoff is also interested in capturing inference patterns associated with groups of verbs in a systematic way. For example, he associates different rules of inference with each conceptual function. This gives the functions "meaning", capturing generalisations of inference patterns associated with words related in some way. An example from Jackendoff (1990, p. 39) illustrates how decomposition of certain aspects of meaning into regular structural relations can capture such generalisations.

(2.18) a. i. X killed $Y \rightarrow Y$ died

ii. X lifted $Y \rightarrow Y$ rose

iii. X gave Z to $Y \rightarrow Y$ received Z

iv. X persuaded Y that $P \rightarrow Y$ came to believe that P

b. i. X kill Y : X CAUSE [Y die]

ii. X lift Y : X CAUSE [Y rise]

iii. X give Z to Y : X CAUSE [Y receive Z]

iv. X persuade Y that P : X CAUSE [Y come to believe that P]

c. X CAUSE [E to occur] $\rightarrow E$ occurs

In (2.18a), each verbal item is associated with a lexically specific inference rule. Each of these rules, however, is closely related to the others. Recasting the verbs in terms of a general function CAUSE, as done in (2.18b), allows one inference rule (2.18c) to be stated over the conceptual structure associated with each of these verbs.

Two questions which we must ask ourselves at this junction are, 1) Why do we want to represent such inferences in terms of lexical structure? and 2) Is it motivated to tailor our lexical structures in order to capture certain inferential distinctions between lexical items (i.e. should we postulate particular function-argument structures in order to capture semantic variations which may not have a syntactic effect)?

The answer to the first question is quite clear with reference to the example just discussed: specification of inferences which hold for elements of lexical structure which are common to many words allows an inference shared by these words to be captured in a general way. A lexical semantic relation (e.g. function-argument structure) with independent motivation from

syntax which has a consistent inference associated with it thus serves an explanatory role in the specification of the semantic relationship within groups of words.

The second question is a bit more tricky. It is clear that there are often inferences which hold for groups of words semantically related in some way, while the way in which they are related cannot be shown to be a factor affecting syntax. Consider the distinction between the sentences in (2.19), for example. The verbs *drag* and *throw* are related in that they both express the motion of something somewhere caused by someone's (*Bill's*) action. However, they differ in the temporal relations they express and in the entailments they have — in (2.19a), the causing action and the effected motion are temporarily coextensive and *Bill* ends up at the endpoint of the path, while in (2.19b), the cause precedes and initiates the effect, and Bill does not move along the path.

(2.19) a. Bill dragged the ball into the field (?and ran to the field to get it back).

b. Bill threw the ball into the field (and ran to the field to get it back).

Would we wish to argue that the semantic element on which they differ (entraining vs. launching causation, in Jackendoff's terminology) should be captured in the lexical semantics of the words, since the inference affects what can subsequently be stated? If the answer is no, then we seem to be excluding the capture of a potential generalisation, for then there is no way to reflect the fact that words such as *drag*, *tow*, *pull*, etc. are related in some way and therefore no way to specify that they share an inference. Furthermore, the semantic difference between the verbs in (2.19a) and (2.19b) constrains what can felicitously follow the sentences containing these verbs in the discourse. The generalisation could only be captured if a rich non-linguistic knowledge representation were assumed which would express the semantic relationship between the words in question and allow the specification of inferences associated with such semantically related words. We would have to assume, for instance, a hierarchical knowledge representation in which inferences were inherited from supertypes. This hierarchy would have to be very carefully constructed in order to capture all the possible semantic relationships which can be shared among groups of words.

If the answer is yes, we are in danger of introducing non-linguistic information into the lexicon and then the amount of information which must potentially be included in the lexicon could not be constrained and the types of inference procedures we would have to have in the lexicon would have to be extremely powerful.[11] The problem is complicated by the fact that there is no clear test — not even ungrammaticality — which can distinguish between "linguistic" and "non-linguistic" information. Consider the following example from Copestake (1992):

[11]In fact, in the case of the example (2.19), there does seem to be syntactic evidence for including the launching/entraining causation distinction in the lexical semantics: the recipient dative construction is only possible with launching causation and not entraining causation, as shown in (2.20).

(2.20) a. Throw Bill the ball.

b. *Dribble Bill the ball.

c. *Drag Bill the ball.

(2.21) a. The rabbit bit John.

b. *?The newborn rabbit bit John.

c. *The rabbit with no teeth bit John.

Here we have a grammaticality difference which depends on arbitrary bits of world knowledge, e.g. that something must have teeth in order to bite, and that newborn rabbits have no teeth but rabbits in general do. We clearly would not wish to attempt to capture this semantic anomaly through lexical restrictions, or to claim that this knowledge is linguistic. There is no obvious test which can be used to distinguish between infelicities that depend on world knowledge and infelicities that can best be captured through lexical semantic structure and the constraints upon this structure (cf. the discussion of linguistic vs. world knowledge in Section 2.1).

The issues raised here really boil down to whether we should view semantic relationships between lexical items as lexical knowledge or world knowledge. The decision taken with respect to this question must be made very carefully and may have to be made on a case-by-case basis. Given that the main point of developing a lexical semantic framework is to capture generalisations about the way that language is used, it is worthwhile to reflect semantic relationships between words which have *general* application — i.e. for which the inferences hold across many words and for which the semantic element involved in the relationship can be shown to influence the interpretation and use of several word classes. It seems particularly motivated to attempt to capture semantic elements lexically when the relationships they express hold across semantic domains.

In this section, I will review the ways in which Jackendoff's theory allows inferences to be captured. Although there may be other ways of capturing the inferences we would like to capture about sentences, it should become clear from the examples to be introduced below that Jackendoff's approach provides a clearly-defined framework within which inference patterns can be captured in regular ways. His work goes quite a long way towards identifying which relationships between words could be considered *general* and therefore which should be implemented lexically.

Thematic role relations

Thematic relations capture many of the basic inferences which we can draw about participants in an event. On Jackendoff's account, thematic relations are not associated with syntactic positions but rather correspond to particular configurations in conceptual structure. Thematic relations are therefore essentially particular semantic relations which may hold among elements of verb meaning. They do not have independent status within the theoretical framework, but there are generalisations over certain semantic relations which can be made.

For example, Jackendoff (1983, pp. 206-207) shows that the understood subject of a complement infinitive can depend on the thematic relations among elements in the sentence. In example (2.22), the relations between both the main verb and its arguments and the infinitive and the noun governing the infinitive, heavily influence the inferences drawn. Thus

in (2.22a,d) the person understood to be leaving is Bill while in (2.22b,c) that person is John despite the syntactic parallelism between (2.22a,c) and (2.22b,d).

(2.22) a. John gave Bill orders to leave.

b. John got from Bill orders to leave.

c. John gave Bill a promise to leave.

d. John got from Bill a promise to leave.

Furthermore, thematic relations allow the capture of different inferences from syntactically parallel sentences (from Jackendoff (1990:54):

(2.23) a. Harry buttered the bread.

b. Joe pocketed the money.

These convey opposite notions — in (2.23a), the butter goes onto the entity picked out by the direct object, the bread, while in (2.23b), the entity picked out by the direct object, the money, goes into the pocket. These differing inferences must be derived from the meaning of the verbs and the roles which the (syntactic) direct objects play semantically. Since according to Jackendoff thematic relations are not associated with syntactic positions but rather with structural positions, thematic relations in his framework provide a mechanism for accounting for the differences between (2.23a&b). Specifically, different thematic relations may correspond to the same syntactic position. The two direct objects in these sentences are associated with different structural positions in the representation of the meaning of these related verbs (both verbs essentially conveying that someone caused something to go somewhere, but differing on what goes where), thereby accounting for the differing interpretations we get of the sentences. Which thematic relation an entity fulfills depends on the underlying lexical semantics of the verb for which the entity is an argument. Representation of the relations between the verb and its arguments is therefore critical to accurate interpretation.

Functions relating entities

As discussed in the introduction to this section, each of Jackendoff's function-argument structures is associated with a particular set of inferences. The functions are therefore used to capture similarities and differences in meaning conveyed by syntactically similar sentences. Although I will not endeavour to provide a complete overview of the function distinctions Jackendoff introduces in order to capture general inference patterns, the examples should give a flavour of the kind of analysis he undertakes.

The semantic distinction between the two different readings of the sentence in (2.24a), from Jackendoff (1983:166) is captured by associating the preposition *under* with two different Path functions.

(2.24) a. The mouse went under the table.

b. The mouse arrived under the table.

c. The mouse passed under the table.

d. The mouse was under the table.

The first reading, that the mouse ran under the table and stayed there, is expressed by the Path function TO, while the second reading, that the mouse passed under the table, is expressed by the Path function VIA. Only the first reading of the prepositional phrase is available in (2.24b), while only the second reading is available in (2.24c). Both of these uses of *under* must be distinguished from the use in (2.24d), in which the prepositional phrase conveys a Place rather than a Path. The determination of which interpretation(s) of the preposition is (are) possible in each sentence depends on the meaning of the verb — crucially captured in its lexical semantic representation, as shown in (2.25). *Under* itself has three different interpretations, shown in (2.26).

(2.25) a. $go = [_{\text{Event}}\ \text{GO}([\text{THING}], [\text{PATH}])]$

b. $arrive = [_{\text{Event}}\ \text{GO}([\text{THING}], [_{\text{Path}}\ \text{TO}([\text{PLACE}])])]$

c. $pass = [_{\text{Event}}\ \text{GO}([\text{THING}], [_{\text{Path}}\ \text{VIA}([\text{PLACE}])])]$

d. $be = [_{\text{State}}\ \text{BE}([\text{THING}], [\text{PLACE}])]$

(2.26) a. $under = [_{\text{Path}}\ \text{TO}([_{\text{Place}}\ \text{UNDER}([\text{THING}])])]$

b. $under = [_{\text{Path}}\ \text{VIA}([_{\text{Place}}\ \text{UNDER}([\text{THING}])])]$

c. $under = [_{\text{Place}}\ \text{UNDER}([\text{THING}])]$

The Path argument of *arrive* is restricted to being a goal, i.e. a TO Path and can therefore only combine with the (2.26a) interpretation of *under*. The Path argument of *pass* is restricted to being a route, i.e. a VIA Path, and can therefore only combine with the (2.26b) interpretation of *under*. *Go* is simply underspecified for what kind of Path argument it expects, thus allowing both interpretations of *under* to fill that argument position and accounting for the ambiguity of (2.24a). The PP in (2.24d) is restricted to being a Place since the construction conveys that the PP is predicated of the subject; as it is not possible for *the mouse* to be the argument of a Path function but it is possible for it to be the argument of a Place function, due to ontological restrictions on what may serve as that argument, the PP must be interpreted as a Place, under the (2.26c) interpretation of *under*.

The discussion in Section 2.2.1 highlighted a difference between verbs of motion and verbs of location and introduced the representation Jackendoff proposes as underlying the distinction. The functions in the representation do not simply constrain the kind of PP arguments the different verb types can appear with, they also suggest something about the relationship between the PP argument and the subject of these sentences. From the function-argument structure $[\text{GO}([\text{THING}], [\text{PATH}])]$, for example, we infer that the THING travels along the PATH. Similarly, from $[\text{BE}([\text{THING}], [\text{PLACE}])]$ we infer that the THING is in the location PLACE. Thus the functions each convey a particular meaning.

Additional support for the utility of the function-argument structures and their associated inference patterns comes from the fact that many of the functions which Jackendoff identifies can be extended from the spatial domain to other domains. Thus the GO and BE functions can be utilised within, for example, the domain of possession. It makes sense to utilise the same

functions because the verbs represented with them can be used in the same syntactic frames as the spatial verbs and because the inferences associated with their use are simple extensions of the spatial inferences. Consider the sentences in (2.27).

(2.27) a. The doll belongs to Beth.

b. Beth received the doll.

(2.28) a. $[\text{BE}_{\text{Poss}}\,(\textit{the doll}, [\text{AT}(\textit{Beth})])]$

b. $[\text{GO}_{\text{Poss}}\,(\textit{the doll}, [\text{TO}(\textit{Beth})])]$

Sentence (2.27a) can be represented by (2.28a) and (2.27b) by (2.28b). The inference in the former case is that the doll is possessed by Beth (a possessional extension of being located at Beth), and in the latter that the doll transfers possession to become Beth's (the possessional analogue of travelling along a path which ends at Beth). These and other extensions of the basic spatial functions suggest the utility of the functions identified by Jackendoff and reflect regularities in use of language components across domains (and, hopefully, across languages). More strongly, the theory of the extensibility of spatial functions to other domains results in a claim that many semantic fields have essentially the same structure,[12] and that the spatial domain defines the terms in which many kinds of discourse must be framed.

Selectional restrictions

Although it is clearly impossible to attempt to rule out all semantically anomalous combinations of lexical items at the lexical level (not to mention undesirable due the fact that arbitrary bits of world knowledge can affect the felicity of a sentence — see Copestake (1992)), the definition of function-argument structures in terms of ontological categories does rule out certain combinations. For example, the sentence in (2.29) would be ruled out due to the fact that the GO function underlying the semantics of *walk* requires its first argument to be a THING, not a PROPERTY.[13]

(2.29) *Happiness walked down the street.

Similarly, some selectional restrictions are captured by explicitly specifying function values within a verb's lexical structure. For example, the verb *pass* may only appear with a PP complement specifying a route (e.g. *John passed by the office*). Goals, directions, etc. are incompatible with the meaning of the verb. Therefore, the lexical semantics of *pass* is specified as in (2.30). The semantics of its PP complement must be compatible with the Path type specified therein.

(2.30) $[\text{GO}([\text{THING}], [_{\text{Path}}\ \text{VIA}([\text{PLACE}])])]$

[12] This of course only holds at a certain general level of detail, since for example the inferences from a function in the domain of possession would be different than in the spatial domain.

[13] Note that we are only interested in the treatment of literal sentences. The word *happiness* could possibly be a metonymic reference to a THING — Jackendoff's theory correctly predicts that under such a type-shifting interpretation the sentence in (2.29) would be grammatical, as in that case there would be no violation of the constraints on the argument of the GO function.

Jackendoff also attempts to capture more fine-grained selectional restrictions, such as the fact that the argument of *drink* must be a liquid. He accomplishes this by specifying the value of the semantic argument at the appropriate argument position within the lexical semantics of the verb, as LIQUID. He further assumes an operation of *fusion* which merges the value of the verb semantic argument with the semantic value of the argument only if the two are compatible.

A similar approach is used to capture some lexically-specific inferences, such as the inference associated with the sentence (2.23a) that it is butter (and not anything else) which goes onto the bread. In this case, the first argument of the GO function is specified to be BUTTER in the lexical semantics of the verb *butter* (e.g. [GO(BUTTER, [$_{\text{Path}}$ TO([$_{\text{Place}}$ ON([THING])])])]).

2.2.3 Conclusions

Jackendoff's analysis of lexical semantic structure could be extremely useful to the computational linguist interested in creating a highly constrained lexical framework which aims at capturing as many generalities as possible about linguistic and semantic relationships among words. Jackendoff does not attempt to characterise the distinction between many verbs within a class, e.g. verbs of locomotion — run, swim, fly, etc. — within his proposed lexical structure, unless the distinction is syntactically relevant or seems to be associated with a general pattern of inference. This means that the information to be lexically represented does not attempt to model the full range of meaning of words, but rather is restricted to capturing generalisations.

The adoption of a semantic ontology with respect to which all lexical semantic structure is developed, a notion of thematic roles which allows certain syntactic phenomena to be elegantly modeled, and the identification of function-argument structures which reflect distinctions in both syntactic use and semantic inference among words together form a framework for lexical semantics which is constrained, extendable, and implementable. Jackendoff's approach seems to be precisely the kind of approach to lexical semantics which is necessary for maintaining a distinction between lexical and world knowledge, and for adequately modeling the ways in which words can be used.

Jackendoff's work concentrates on capturing generalisations about word use. It must also, however, be emphasised that the lexicon is a repository of word-specific information which should include idiosyncrasies of the use of individual lexical items. A balance must be struck between being stipulative and being explanatory in lexical representations; generalisations must be reflected in order to predict word use and interpretation and to address the generativity of the lexicon, while apparent "accidents" — behaviour or interpretation which does not seem to follow from general (semantic) principles — must also be encoded to account for the full range of use and interpretation of individual words. Lexicon theory therefore differs from syntactic and pragmatic theories, which aim to be fully explanatory. Not everything that a person knows about words can be motivated; there are idiosyncrasies which cannot be explained. We will see this clearly in Chapter 5 with respect to the phenomenon of logical metonymy.

2.3 Conceptual Structure and HPSG

2.3.1 Lexical Semantics in HPSG

In this section I will provide a brief overview of the standard semantic representation of nominal and verbal semantics in HPSG[14] and will discuss its application to the purpose of modeling of verbal sense extensions.

The semantics of a word is primarily specified in the SYNSEM|LOCAL|CONTENT value of the sign which captures the lexical entry of the word. The SYNSEM|LOCAL|CONTEXT value also may contain information relevant to its interpretation, although this feature is used to specify background conditions on the utterance context rather than truth-conditional meaning. In what follows I will concentrate on the CONTENT values.

Nouns in HPSG

The CONTENT of a noun is represented in HPSG by a feature structure of type *nom-obj*, which is structured as in (2.31).

(2.31)
$$\textit{nom-obj}\left[\begin{array}{ll}\text{INDEX} & \textit{index}\left[\begin{array}{l}\text{PERSON } \textit{person}\\ \text{NUMBER } \textit{number}\\ \text{GENDER } \textit{gender}\end{array}\right]\\ \text{RESTR} & \textit{set (psoa)}\end{array}\right]$$

The INDEX value is a reference marker, to which semantic roles are assigned. The RESTRiction attribute specifies any semantic restrictions on the referent of the noun. These restrictions are in the form of parametrised states of affairs (*psoa*s) which convey relations between referential indices and/or other *psoa*s. Different relations are realised as subsorts of the type *qfpsoa* (quantifier-free *psoa*). These subsorts are defined for particular numbers and types of arguments. Following Situation Semantics (Barwise and Perry (1983)), which provided the inspiration for the HPSG notion of *psoa*s, *qfpsoa* relations are basic, corresponding to individuations of relations in the world made by cognitive agents. I will discuss the relations utilised in HPSG in further detail below.

Verbs in HPSG

The CONTENT of a verb is represented in HPSG by a feature structure of type *psoa*, which is structured as in (2.32).

(2.32)
$$\textit{psoa}\left[\begin{array}{ll}\text{QUANTS} & \textit{list (quantifiers)}\\ \text{NUCLEUS} & \textit{qfpsoa}\end{array}\right]$$

As indicated, the value of the NUCLEUS feature is a *qfpsoa*. This type is divided into subsorts corresponding to particular semantic relations, each defined for appropriate semantic roles

[14] HPSG is *Head-Driven Phrase Structure Grammar*, as defined by Pollard and Sag (1987, 1994).

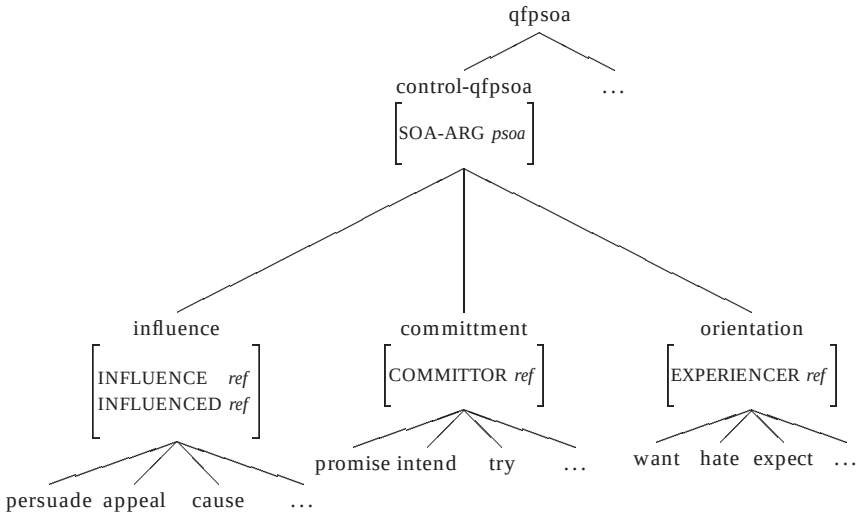

Figure 2.1: A part of the HPSG sort hierarchy for semantic relations

which pick out either referential indices or *psoa*s. The only subsort of *qfpsoa* which Pollard and Sag (1994) investigate in any detail is that of verbs which take controlled complements, for which controller assignment is argued to be based (following Jackendoff 1972) on semantic roles. The subsorts are defined as indicated in Figure 2.1. Semantic arguments are associated to syntactic arguments of the word introducing a relation via structure sharing between the indices of the semantic arguments of the relation and of elements on the word's subcategorisation list (SYNSEM|LOCAL|CATEGORY|SUBCAT).

An example of the CONTENT|NUCLEUS value of the sentence *They persuaded me to leave* can be seen in (2.33). This representation includes reference to a *qfpsoa*, *leave*, which is not specified in the hierarchy in Figure 2.1, but which can be viewed as a one-place relation which specifies that a LEAVER leaves. The structure sharing between the INFLUENCED value and the LEAVER value is achieved by the HPSG control theory.

(2.33)
$$\textit{persuade}\left[\begin{array}{l}\text{INFLUENCE} \quad \textit{ref (they)} \\ \text{INFLUENCED } \boxed{1} \textit{ref (me)} \\ \text{SOA-ARG} \quad \textit{leave}\left[\text{LEAVER } \boxed{1}\right]\end{array}\right]$$

Utility of the representation

The approach to lexical semantics within HPSG as outlined above has several positive properties.

1. Semantic relations are constrained according to their definition in the subsumption hierarchy. This definition specifies how many (semantic) arguments the relation has and associates each argument with a particular semantic role.

2. Semantic relation types are organised via an inheritance hierarchy which allows similar

types to be grouped together. Generalisations over these groups can therefore be directly represented in the hierarchy.

3. The semantic roles filled by the syntactic arguments of a word can be specified directly in the lexical entry of the word, via structure-sharing.

However, there are also some shortcomings to the system as presented in Pollard and Sag (1994), mainly due to the fact that a complete description of lexical semantic representation was by no means a goal of that work. The shortcomings identified here relate directly to the problem of efficiently capturing sense extensions which is the main focus of this thesis.

1. Alternations in the surface order of syntactic arguments would have to be captured by distinct lexical entries. These lexical entries would differ only in the mapping between the syntactic arguments and the semantic arguments, thereby missing subtle semantic distinctions between the alternate forms.

   The alternate forms could be generated via lexical rules, but this would either require one form to be the base form (see arguments against this in Markantonatou and Sadler 1995), or would require one lexical rule to generate each of the alternate forms from a verbal stem in the lexicon. Whichever approach is chosen, these lexical rules would not be constrained enough given the existing grain of semantic representation in HPSG to block the application of the rules to non-alternating verbs.[15]

2. Each sense of a word must correspond to a distinct semantic relation type in the sort hierarchy, and must be defined by a lexical entry specific to that sense. This fails to adequately capture relationships among the senses. Polysemous words, for which the different senses of the verb share a certain amount of meaning, are treated equivalently to truly ambiguous words, for which one word corresponds to multiple entirely distinct meanings (e.g. *bank* as in *edge of a river* and *financial institution*). This problem could be addressed by introducing polymorphic types (along the lines of e.g. Sanfilippo 1995).

3. No allowance is made for the specification of selectional restrictions, or for reference to the *kinds* of entities which instantiate the semantic arguments of the relations. This is a problem for, for example, a proper treatment of logical metonymy, to be addressed in Chapter 5, which requires information about the referent of a complement noun in order to appropriately specify the *psoa* expressed by a sentence.

These problems could be alleviated through the incorporation of a certain amount of semantic decomposition, along the lines of what Jackendoff (1983, 1990) has proposed for verb semantics and what Pustejovsky (1991, 1995a) has proposed for nominal semantics.

[15] Of course, some mechanism could be invented for this purpose, such as grouping non-alternating verbs together in the subsumption hierarchy or adding a binary feature specifying whether the verb alternates or not, but these options are not motivated from a theoretical standpoint given the semantic basis of these alternations (Pinker 1989).

2.3.2 Incorporation of Jackendoff's theory into HPSG

It would be quite straightforward to incorporate Jackendoff's theory of conceptual structures into a typed feature-based grammatical framework like HPSG. Some notions within Jackendoff's theory transfer almost directly to HPSG, and many of the notational and operational details of the theory are actually captured more consistently within HPSG. We discuss how below.

The core of Jackendoff's theory consists of the semantic ontology and the definition of function-argument structures in terms of elements in this ontology. This corresponds directly to the notion of type constraints on the structure of signs in HPSG. The ontology would be defined within the sort hierarchy, anchored at the *content* sort. Each function-argument structure can be defined as a feature structure of a particular subtype of *qfpsoa* whose arguments are restricted to be of a particular type. Words functioning as lexical heads would be given a semantic specification in their lexical entries corresponding to one of the ontological types; in particular, verb meanings would be decomposed into a function-argument structure for either an Event or State.

Jackendoff has argued that a single entity may correspond to multiple semantic roles within a relation. This can be easily modeled in HPSG with structure-sharing between the values of each semantic role which the entity fills and the index of the entity. Any type constraints (i.e. both of ontological category and more specific selectional restrictions) imposed on a subcategorised argument can be enforced through required type compatibility during unification. No special operation needs to be defined to handle this, assuming that the sort hierarchy has been adequately defined. The use of HPSG to handle this mechanism has an additional benefit: HPSG makes explicit the existence of a hierarchy which Jackendoff implicitly assumes. A semantic hierarchy is fundamental to HPSG and therefore the notion of "compatibility" between entities on which Jackendoff depends can be handled directly.

To generalise over the semantic contribution certain adjuncts make, Jackendoff introduces *adjunct rules* which specify how the semantics of an adjunct is to be combined with the semantics of the verb phrase it modifies. These rules specify both the syntactic and semantic structure of the verb phrase which a particular adjunct can modify. The actual combination of the semantics of the VP and the adjunct is accomplished via an *adjunct fusion* rule. This rule is essentially analogous to the function the HEAD-ADJUNCT schema performs in HPSG, while the details of particular adjunct rules can be captured in the signs for each adjunct, in the SYNSEM|LOCAL|CATEGORY|HEAD|MOD and the CONTENT features. Thus the mechanisms required for Jackendoff's theory already exist within HPSG.

2.4 Davis' multiple-inheritance lexical semantics

Anthony Davis introduces a linking theory in his thesis (Davis 1995) which utilises to a limited extent insights made by Jackendoff (1983, 1990) about lexical decomposition. He formalises this decomposition in terms of feature structures, capturing verb meanings in a multiple-inheritance lexical semantics compatible with HPSG. I will take Davis' representation as a

starting point for the development of the representation to be used in subsequent chapters of this thesis. I will augment it with further decomposition, in order to address the phenomena relevant to this thesis. In this section I will introduce his representation and in the following I will discuss my extensions to it. I will not go into the details of Davis' proposals for the mapping between semantics and syntax, for although the issues raised by that mapping drive his choice of the level of lexical semantic decomposition in his representations, there is nothing in his proposals which precludes further detail in the representation.

Davis begins his discussion of lexical semantic representation with a critique of the 'traditional' conception of thematic roles as introduced by Gruber (1965) and Fillmore (1968). He argues, drawing on Dowty (1991), that it is difficult to identify a (relatively small) set of thematic roles which can be used to classify each argument of every predicate, and furthermore that argument mapping cannot easily be explained in terms of such a set. He points out that the traditionally assumed set of thematic roles runs into difficulty in accounting for the primacy of causation relations: the causer of a causative verb is realised as the subject of a sentence despite the presence of other agentive arguments. Dowty's (1991) solution to this problem is to eliminate an explicit representation of thematic roles, instead explaining the mapping in terms of a numerical comparison of the lexical entailments associated with the arguments of a predicate which leads to identification of proto-agents which map to subject position and proto-patients which map to object position. Davis argues that this explanation is too dependent on surface transitivity and identifies weaknesses with the numerical comparison approach. He goes on to show how explicit representation of certain lexical entailments can better model linking.

The lexical entailments which Davis represents are derived from suggestions by Dowty (1991), Wechsler (1991), and Jackendoff (1983, 1990). He defines a multiple-inheritance hierarchy of lexical semantic relations, each of which specifies the proto-role properties which hold of the relation's arguments. The relations characterise those properties that are relevant for linking and argument selection. The sort hierarchy of lexical semantic relations and the list of proto-roles are presented in Figures 2.2 and 2.3, respectively. In the sort hierarchy, I specify the features which are defined for each type. The values on some features are left unspecified — this indicates that Davis has not specified a typal restriction on the values. For others, a specific type is specified (e.g. the value of GRND for a *mot-rel* must be a *path*); this is indicated by a non-parenthetical italicised type. Values given in parentheses are intended to provide only an intuitive indication of the entailment associated with that feature in a particular relation; they do not correspond to specific types in the hierarchy. So for the *cause-und-rel* relation, the value of the ACT feature should be interpreted as a "causer", while the value of the UND feature should be interpreted as "caused".

Every verb in the lexicon will correspond to at least one sort in the sort hierarchy. This sort aims to capture basic entailments associated with the verb, minimally identifying the participant roles, designated in terms of the proto-roles in Figure 2.3, of the kind of event that is picked out by the verb. There is in addition a constraint, the *entailment-to-attribute-condition* (Davis 1995, ch. 3:27), which requires that participants in an event "of which a particular proto-role entailment holds (in virtue of playing a particular participant-role in an event) must be

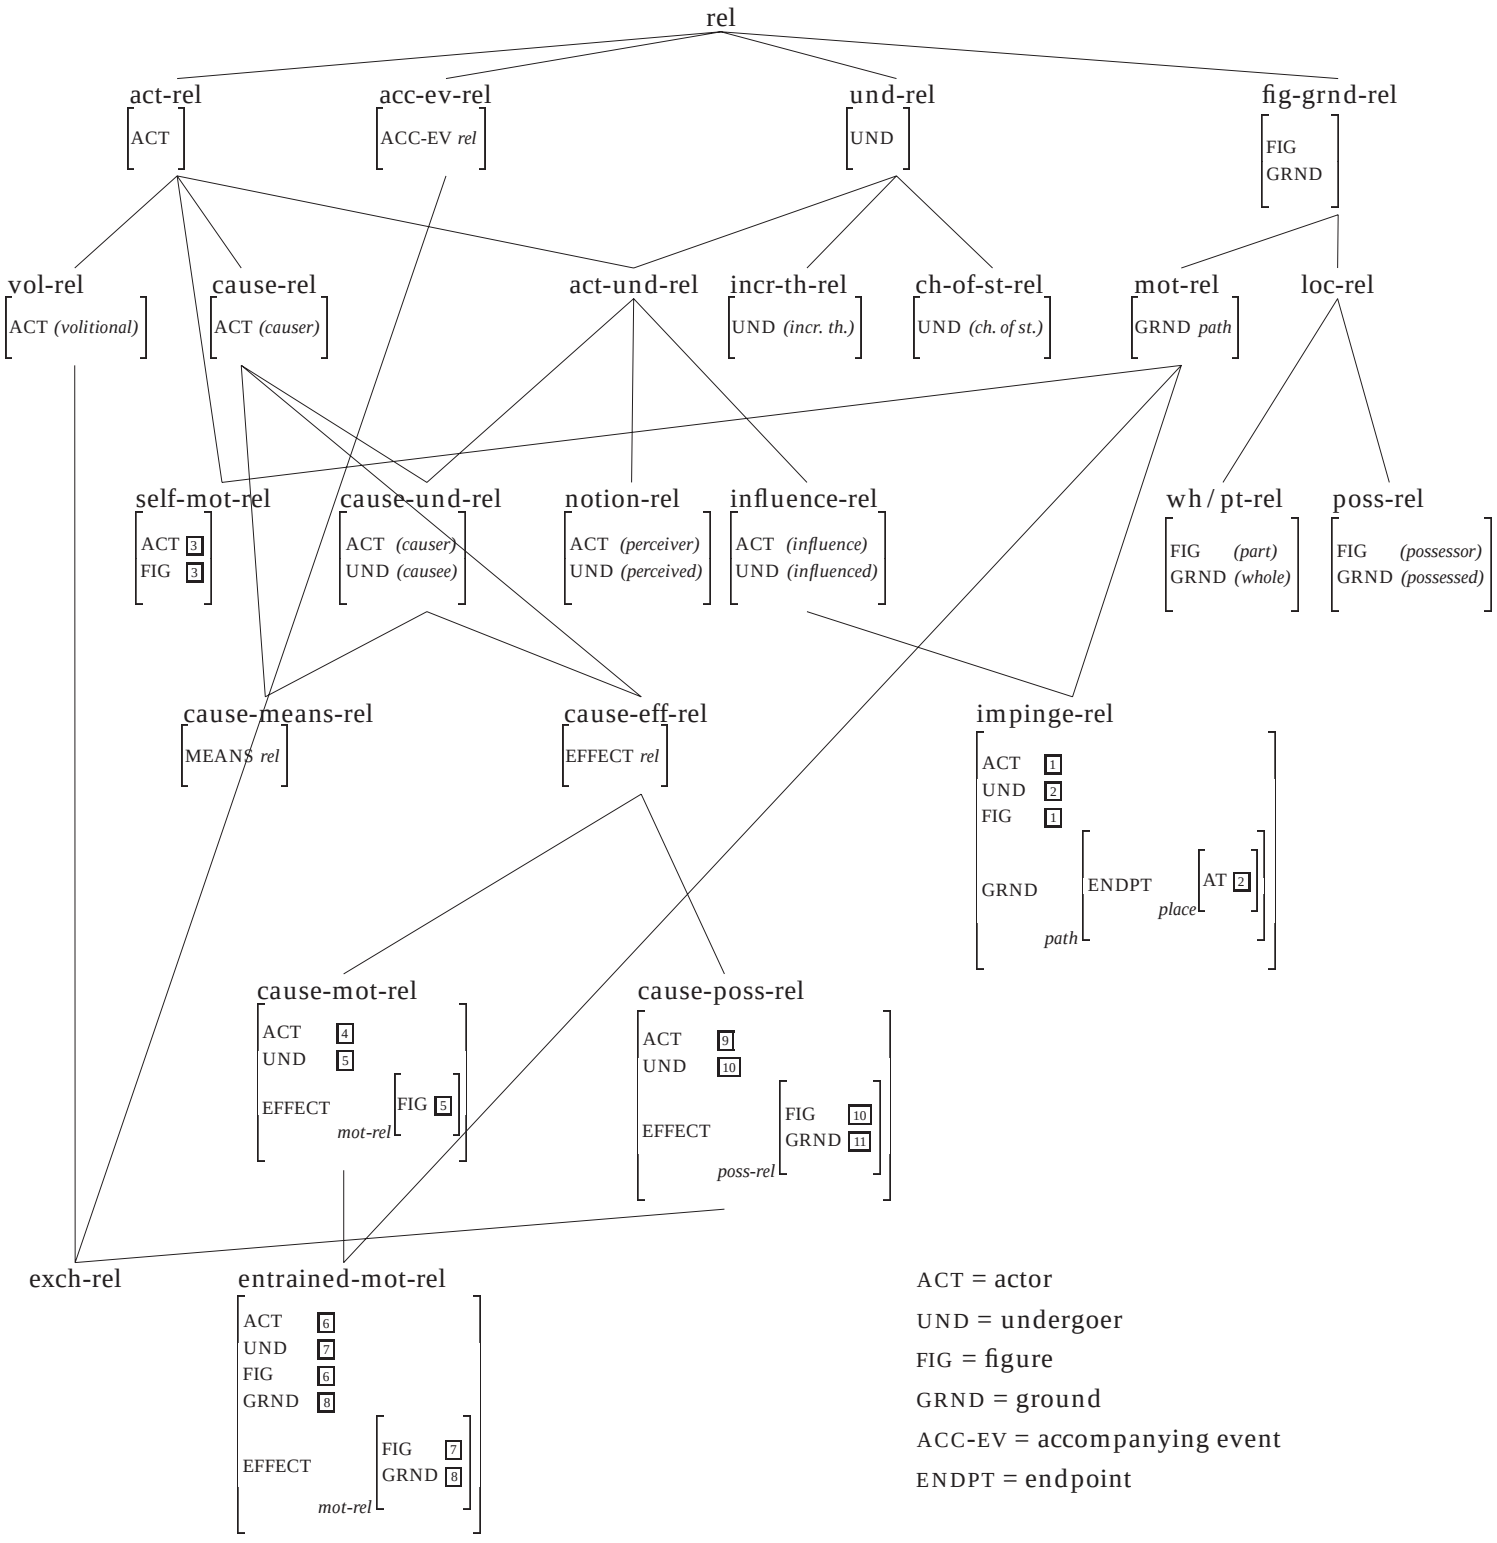

Figure 2.2: Sort hierarchy of proto-roles in Davis (1995)

Proto-Role	Associated Entailments
ACT	Causally affects or influences other participant(s) or event(s) Volitionally involved in event Has a notion or perception of other participant(s) in event
UND	Causally affected or influenced by another participant Undergoes change of state Is an incremental theme
FIG	Moves with respect to another participant Contains or constitutes another participant Possesses another participant
GRND	Path traversed by another participant Is contained by or part of another participant Is possessed by another participant
EFFECT	Is an event or state caused by another event
MEANS	Is an event that intervenes within another event and enables it
ACC-EV	Is an event that necessarily accompanies another event

Figure 2.3: Summary of proto-roles in Davis (1995)

denoted by the value of a particular proto-role (ACT, UND, or some other attribute) in lexical semantic structure". This condition ensures that the values of features in lexical semantic structure will correspond to the appropriate participants in an event.

Participant roles in different events are related to one another by virtue of the inheritance relationships between the sorts. These relationships provide the basis for Davis' generalisations about the mapping between semantic and syntactic roles.

2.4.1 Davis' relation types

The relation types in Figure 2.2 require further explanation. They are based, as should be clear by now, on the entailments that hold of the participants in each particular relation. Davis draws mostly on Dowty (1991) and Wechsler (1991) for the basic entailments which he identifies.

I begin with the relations defined in terms of ACT and UND. The proto-role ACT corresponds generally to an actor in an event, and UND to an undergoer. These are specified to more particular entailments in subsorts of *act-rel* and *und-rel*. In *vol-rel* the ACT participant is required to be volitional with respect to the corresponding event, in *incr-th-rel* the UND participant is required to be an incremental theme, and in *ch-of-st-rel* the UND participant is entailed to undergo a change of state.

A relation in which one entity causally affects another is represented by a *cause-rel* relation, for which the causer corresponds to the ACT proto-role and the affected entity to the UND participant role. This is a subtype of the *act-und* relation because it requires (at least) two participants and associates particular entailments with those participants. The subsort *cause-eff-rel* applies to causative verbs which express that some event is caused by the causing event.

For these verbs, the entity affected by the causing event is also the actor of the caused event. The representation Davis assumes of this follows Pinker (1989) who in turn follows Jackendoff (1990): an event characterised by some relation *rel* is embedded under an EFFECT attribute. Davis gives the example (1995, ch. 3:38,40) in (2.34) of the relation *cause-jump-rel*, a subsort of *cause-eff-rel* for which the embedded event is a *jump-rel* event.

(2.34) a. Chris jumped the horse over the gate.

b.
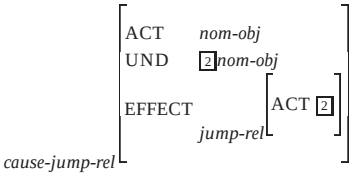

Similarly, *cause-means-rel* is meant to characterise relations which entail a particular event which is interpreted as the means by which the expressed causation is achieved.

The *notion-rel* derives from a suggestion by Wechsler (1991) that there are verbs which require a sentient participant which necessarily has a notion of (perceives) the other participant, but does not require the perceived object to be sentient with respect to the event. Consider an example from Wechsler (1991, ch. 2, ex. 81), also cited by Davis (1995, ch. 3:36), shown in (2.35). Although normally when two people marry each person has a notion of the other, it is not entailed by the semantics of the verb but rather seems to be constrained by world knowledge.

(2.35) a. The duke married the two year old princess.

b. #The two year old princess married the duke.

This relation subsumes the semantics of mental state, perception, and volitional action verbs, such as *remember*, *like*, *see*, *hear*, *chase*, and *murder*.

The *influence-rel* stems from the discussion of control phenomena in Pollard and Sag (1994), as introduced in Section 2.3, and is meant to reflect the semantics of verbs like *urge* and *persuade*, for which one participant influences another in some way. The entailments associated with these roles are left rather vague by Davis, because there do not seem to be many clear necessary entailments associated with the INFLUENCE and INFLUENCED roles. This is suggested by (2.36) (Davis 1995, ch. 3:47), which shows that the urger / persuader need not be sentient or have a notion of the urged / persuaded individual.

(2.36) The letter / police persuaded / urged Sandy to leave.

The relations rooted at *fig-grnd-rel* are intended to correspond to the thematic tier proposed by Jackendoff (1983, 1990) and introduced in Section 2.2.1 (page 16). That is, these relations capture the motion and location of participants in an event. Davis thus encodes Jackendoff's two tiers through the distinction between *act-rel*s / *und-rel*s and *fig-grnd-rel*s. The two tiers are brought together, to capture both actor / patient relations and motion / location relations in a single lexical structure, in types which inherit from both kinds of relations. This merging is motivated by Davis in terms of linking: for transitive verbs such as *pass* and *cross* motion is

associated with the argument in subject position while for causative verbs such as *throw* and *deflect* motion is associated with the noun in object position (Davis 1995, ch. 3:43), suggesting that both the entailments of motion and causation must be taken into consideration for accurate mapping of arguments to syntactic position.

The *fig-grnd-rel* captures all figure / ground relations. Its subtype *mot-rel*, motion relation, entails relative motion of two participants — the object in motion is denoted by the value of the feature FIG, the stationary object by the value of GRND which must be a *path*. Davis adopts Jackendoff's representation of paths and places. The semantics of *enter*, for example, is represented by Davis as in (2.37). The *enter-rel* is a subtype of *mot-rel*, and indicates that the entity denoted by the value of FIG enters the place denoted by the value of IN.

(2.37)

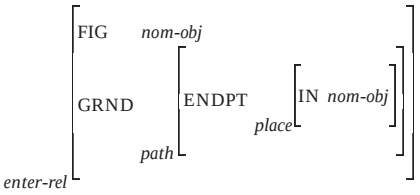

The type *mot-rel* can therefore be seen to correspond directly to Jackendoff's GO function. The type *self-mot-rel*, which inherits from both *mot-rel* and *act-rel* and constrains the value of FIG to be structure-shared with the value of ACT, captures the semantics of verbs which express that one of the participants in the event acts agentively and moves.

Verbs of impingement, e.g. *hit*, *poke*, *tap*, express contact of their subject with their object as a result of motion. So *impinge-rel* inherits from both *mot-rel* and *influence-rel* to capture both the contact and motion aspects of this meaning. The feature structure associated with the sort *impinge-rel* in Figure 2.2 expresses these two aspects and requires the actor to the participant that moves, and the influenced entity to be the point of contact.

Verbs of launching vs. entrained motion, such as those discussed in Section 2.2.2 with respect to example (2.19), are captured by Davis in terms of the sorts *cause-mot-rel* and *entrained-mot-rel*, respectively. Both specify that an effect of motion is caused, and the latter relation additionally indicates that the causer accompanies the causee along its path.

The relation *loc-rel* is for location relations, for which Davis argues (1995, ch. 3:57-61) there do not seem to be strong entailments for the locatum and location. In fact, Davis suggests that individual verbs idiosyncratically specify which of its arguments corresponds to FIG (locatum) and which to GRND (location). In some cases this assignment must even be determined pragmatically.

Possession (e.g. as expressed by *have*, *own*, *acquire*) can be considered a kind of location relation — the possessed entity is located at the possessor. Similarly, the whole / part relations expressed by verbs of inclusion (*contain*, *include*, etc.) also correspond to location — the part is located at the whole. Thus the relations *poss-rel* and *wh/pt-rel* are subsorts of *loc-rel*. The entailments associated with FIG and GRND are in these cases, however, stronger than for simple location relations.

The last sort in Figure 2.2 which needs to be discussed is *acc-ev-rel*. This sort introduces a feature ACC-EV, which corresponds to 'accompanying event'. This feature is instantiated

in a relation which entails that two specific subevents occur simultaneously. An example of such a relation is the *exch-rel*, for verbs of exchange such as *buy* and *sell*. In the case of these specific verbs, one subevent involves the transfer of goods and one subevent involves the transfer of money. This is represented by Davis (1995, ch. 3:63) for the case of *sell* as in (2.38a). Here the annotations specify what participants in the selling event the various semantic roles correspond to. The *cause-poss-rel* is a subsort of *caus-eff-rel*, as indicated in the sort hierarchy.

(2.38) a.

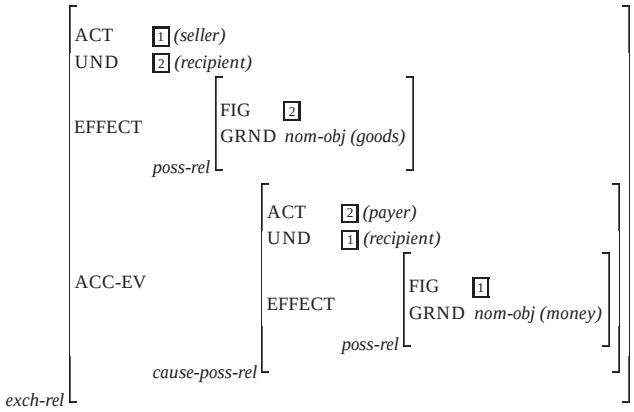

b.

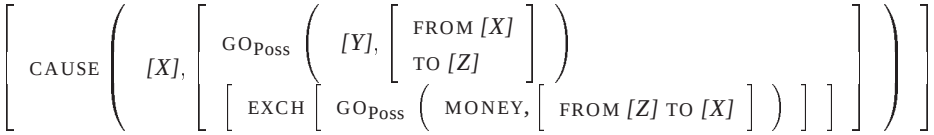

Note the correspondence between Davis' representation and Jackendoff's (1990:191) representation of the same verb in (2.38b). Jackendoff's CAUSE function is captured by the fact that *exch-rel* is a subsort of *caus-poss-rel*. The GO function is captured by the subevent embedded under Davis' EFFECT attribute, and Jackendoff's EXCH (a modifying function in his system meaning "in exchange for") introducing a transfer of money event corresponds to the embedding of a *cause-poss-rel* event under Davis' ACC-EV attribute.

2.4.2 Comparison with Jackendoff (1983, 1990)

The granularity of lexical semantic decomposition in Davis' representations generally is based on what is required for linking semantic arguments to syntactic positions. The representation corresponds to a great extent almost directly to Jackendoff's Conceptual Structure representation, since Jackendoff also looks to variations in syntactic form as an indicator of what needs to be represented in the underlying semantics. What Jackendoff represents in terms of a limited set of functions which relate conceptual constituents, Davis formalises as semantic relation sorts.

Functions vs. Relations

The functions which Jackendoff (1983, 1990) assumes are summarised in Figure 2.4. A direct comparison of these functions with the semantic relations proposed by Davis can be found in Figure 2.5. It is obvious from this comparison that Jackendoff's representation has wider

Function	Associated meaning and argument structure
GO	an Event-function which denotes motion along a path; two arguments: the Thing in motion and the Path it traverses
AFF	a formal elaboration of an Event which specifies that an actor "affects" a patient; two arguments: the actor and the patient.
STAY	an Event-function which denotes stasis over a period of time; two arguments: the Thing standing still and its location.
CAUSE	an Event-function specifying cause and effect relations; two arguments: a Thing which is the agent or an Event which is the cause, and an Event which is the effect.
INCH	an Event-function specifying the "inchoative", a change taking place which has a final state; one argument: the State in which the event terminates.
MOVE	an Event-function which specifies that an object is moving / moves; one argument: the Thing which is moving.
BE	a State-function for specifying the location of objects.
ORIENT	a State-function for specifying the orientation of objects.
EXT	a State-function for the spatial extension of linear objects along a path.
CONF	a State-function that expresses that a verb describes the internal spatial configuration of an object; one argument: the Thing that is in the spatial configuration.
EXCH	a modifying function specifying an event which is "in exchange for" the entity being modified.
REACT	a formal elaboration of an Event which specifies that one Thing reacts to another; two arguments: X, Y; X "reacts to" Y.
place functions	at, on, in, under, ... functions expressing location.
path functions	to, from, toward, away-from, via; functions expressing direction.

Figure 2.4: Summary of functions in Jackendoff (1990)

Jackendoff function	Davis relation(s)
GO	spatial: *mot-rel* possessional: *cause-poss-rel*
AFF	*act-rel*, *und-rel*, *act-und-rel*
STAY	subsort of *und-rel*, but no exact correlate
CAUSE	*cause-eff-rel*
INCH	no exact correlate, similar to *ch-of-st-rel*
MOVE	*mot-rel* with an underspecified GRND value (cf. GO)
BE	*loc-rel*
ORIENT	no correlate
EXT	no correlate
CONF	no correlate
EXCH	*cause-poss-rel* event embedded under a ACC-EV attribute
REACT	*notion-rel* or *influence-rel*

Figure 2.5: Correspondences between Jackendoff functions and Davis relations

coverage than that discussed by Davis. The functions which have no correlates in Davis' representation have particularly to do with the internal structure of objects — ORIENT is for verbs like *point* in *The sign points toward New York* (Jackendoff 1990:44) for which an object is oriented along a path, EXT is for verbs like *extend* and stative *go* as in *The road goes from New York to San Francisco* which describes the extension of an object along a path, and CONF is for verbs like *stand*, and *sit* (*Sally stood for hours on end*, Jackendoff 1990:91) which express the internal spatial configuration of their argument.

The verbs which are represented using these functions impose very specific entailments on their arguments, ones which do not correlate well with the proto-roles and associated entailments Davis proposes, as shown in (2.3). Those proto-roles are appropriate only for events rather than states, with the exception of the FIG and GRND features in certain relations (e.g. *loc-rel*). However, the entailments of FIG / GRND as presented by Davis do not quite reflect the same entailments as the functions ORIENT, EXTEND and CONF. In the case of ORIENT, the ORIENTED entity is not located at specific location or contained within a path but has an ORIENTATION along that path. Similarly for EXTEND, the EXTENDING entity covers the full extent of a path rather than being contained within the path. The function CONF seems to convey a kind of *fig-grnd-rel* in which the single argument fills both the FIG and GRND roles. I propose to introduce three subsorts of *fig-grnd-rel* to capture these specific entailments.

There are certain other discrepancies of coverage between Jackendoff's functions and Davis' relations. The function STAY, for example, does not have a clear correlate in Davis' sort hierarchy. However, Jackendoff (1990:44) suggests that this function might be unnecessary, since instances of this function can correspond to standard relations with the additional requirement that the relation holds over a certain period of time. So the sentence *Bill stayed in the kitchen* can simply be represented by a subsort of *loc-rel* which adds the entailment of stasis. This function will therefore not be added to Davis' hierarchy as an independent relation.

The function INCH can easily be accommodated in Davis' framework. As indicated in

Figure 2.5, this function is similar to *ch-of-st-rel* in that a particular state results from the event. However, the subject of a sentence instantiating that relation is entailed to be an UNDergoer. The inchoative interpretation of sentences like those in (2.39), which expresses an event of a particular state occurring (i.e. (2.39a) could be paraphrased as *Bill came to be standing on the table*), seems to be more accurately represented by a subsort of *act-rel* in that the subjects of the sentences act, effecting the change of state.

(2.39) a. Bill stood on the table.

b. Snow covered the hills.

In the event interpretation of these sentences, *Bill* and *snow* not only participate in the resulting state, but they seem to be causally involved with achieving the result. Contrast this with the examples in (2.40) which only have a stative interpretation: the event interpretation is impossible because the subjects of these sentences cannot be construed as ACTors.[16]

(2.40) a. The vase stood on the table.

b. The blanket covered the table.

I therefore will introduce a subsort of *act-rel*, *act-inch-rel*, corresponding to Jackendoff's INCH function which is defined for an ACTor and a stative EFFECT. The semantics of sentence (2.39a) can be represented as shown in (2.41).

(2.41)

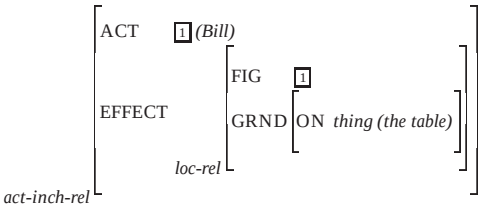

Jackendoff also extends the use of his functions to semantic fields other than the spatial domain, where some concept is abstracted yet conceptually related to spatial motion. He provides the following examples (Jackendoff 1990:25), arguing that the sentences in (2.42), in the spatial semantic field, are similar to those in (2.43), in the possessional semantic field.

(2.42) a. The bird went to the tree.

b. The bird is in the tree.

c. Harry kept the bird in the cage.

(2.43) a. The inheritance went to Philip.

b. The money is Philip's.

c. Susan kept the money.

[16] Note that the interpretation of *snow* as an actor in (2.39b) depends on world knowledge that snow falls and that through this falling snow can have certain effects. No such world knowledge of associated events is available for *vase* or *blanket*.

There is a basic sense of the verbs which seems to be conveyed in each of the parallel sentences: the (a) sentences convey a change (of location in (2.42) and of possession in (2.43)), expressing a GO function, the (b) sentences express the result state of the (a) sentences, a BE function, and the (c) sentences "denote the causation of a state that endures over a period of time" (Jackendoff 1990:26), treated as a STAY function in the original Jackendoff analysis. These abstract concepts are made precise in the different semantic fields, and associated with field-specific inferences, but certain generalisations follow from the abstract concepts.

To a certain extent Davis acknowledges these relationships, for example by the fact that his *poss-rel* is a subsort of *loc-rel*: so possession is a more specific kind of location. However, there is no relation in Davis' hierarchy at the level of generality of Jackendoff's GO function. As is clear from Figure 2.5, the spatial and possessional variants of Jackendoff's GO function correspond to entirely distinct relations in Davis' hierarchy. This treatment causes Davis to miss certain generalisations. He proposes, for example, a representation for the semantics of the verbs *give* and *send* as a *cause-poss-rel*, as in (2.44) (based on Green 1974 and Pinker 1989).

(2.44)

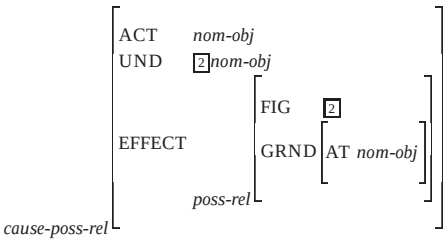

Jackendoff, in contrast, treats these verbs as conveying a caused GO_{Poss} function. The Davis / Green / Pinker semantics of *give* / *send* could be paraphrased as "cause to possess" whereas Jackendoff's semantics of these verbs is "cause to receive". In some sense these are almost equivalent, in that it can in most cases be inferred from a GO_{Poss} function that the person at which the "go-ing" entity ends up possesses that entity. But the Davis / Green / Pinker semantics has stronger entailments and therefore runs into difficulties.

Firstly, it is only an *intended* effect of a *send* event that the entity to which something was sent possess it. This is not an entailment of the verb. Consider (2.45).

(2.45) John sent a letter to Mary, but it got lost in the post and she never received it.

In this case to say that Mary possesses the letter is patently false, although the sending event is completed. It is more accurate to say that the effect of sending is that the letter started along an abstract path (ending at Mary) than that the effect is that Mary possesses the letter.

Secondly, things do not have to be sent to entities which are capable of possessing things. Consider (2.46).

(2.46) a. John sent a package to New York.

b. John gave $100 to the cancer research fund.

For (2.46a), it can be inferred that the package will be acquired by someone in New York but this is not expressed by the lexical semantics: there is no participant in the sentence which necessarily comes to possess the package. Similarly for (2.46b), in what sense can a fund

possess a contribution? Such cases are more appropriately analyzed by Jackendoff's abstract GO function, leaving further reasoning for pragmatics.

Thirdly, the semantics in (2.44) completely ignores the contribution of the *to* prepositional phrase. It contributes a path which is not evident in the representation in (2.44). I will argue in Chapter 3 that the semantics of prepositional phrases should in many cases not be ignored, even for verbs which obligatorily subcategorise for them. This has also been suggested by Gawron (1986). Doing so misses generalisations which can be made about the PPs, and fails to capture a relationship between the verbs which subcategorise for these PPs.

For these reasons, I will introduce a relation *go-rel* under *fig-grnd-rel* with subsorts *mot-rel* and *ch-poss-rel* (change-possession-relation). The former subsort will be defined as in Figure 2.2 and will correspond to the spatial GO function. The latter will be defined as in (2.47) and will correspond to the possessional GO function. Other subsorts could be added for other semantic fields, as necessary. The relation *poss-rel* will remain in the hierarchy for verbs (such as *have*) which really do convey this relation, but will not be used for change-of-possession relations.

(2.47)
$$\begin{bmatrix} \textit{ch-poss-rel} \\ \text{FIG} \quad \textit{nom-obj} \\ \text{GRND} \begin{bmatrix} \text{ENDPT} \begin{bmatrix} \textit{at-place} \\ \text{LOCATION} \quad \textit{nom-obj} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

States and Events

A critical difference between Davis' representation and Jackendoff's representation is that Davis does not make an explicit distinction between events and states. In fact, as is clear from examination of which of Jackendoff's functions have no good correlates in Davis' hierarchy, states are generally neglected in his representation. This event / state distinction is left implicit in the entailments associated with proto-roles in the semantic relations. However, it seems critical to explicitly represent this difference in the ontology due to the linguistic influence of this distinction.

In Section 2.2.1, I made reference to Jackendoff's arguments in favor of the state / event distinction, in particular the constrast between how states and events can be referred to in a discourse. There are in addition further semantic and syntactic distinctions between them, and they have different effects on the temporal relations which can be established in a narrative discourse (Kamp and Reyle 1993). For instance, when a state is expressed with the present simple tense in English, the sentence is interpreted as a referring to a situation in which the state holds, while events expressed in this tense receive a habitual interpretation or have a 'newspaper headline' quality (Jackendoff 1983). Consider the data in (2.48).

(2.48) a. John loves Mary. (*state*)

b. John walks. (*activity*, habitual)

c. John reaches the summit. (*achievement*, habitual)

d. John walks to the station. (*accomplishment*, habitual)

States also differ from events in that they cannot normally occur in the progressive form in English, as show in (2.49).

(2.49) a. *John is loving Mary.

b. John is walking.

c. John is reaching the summit.

d. John is walking to the station.

Partee (1984) and Hinrichs (1986) provide proposals for updating time in narrative discourse. Roughly put, events move narrative time forward and states do not. This is shown in (2.50), where the state of the music being very loud in the pub is interpreted (by default) as being true before, during, and after the entering event.

(2.50) Peter entered the pub. The music was very loud.

There are, however, instances in which a state is interpreted as moving time forward, and / or as having an explicit starting point. Consider the discourse in (2.51), for example.

(2.51) Mary switched off the light. The room was pitch black.

These cases would seem to require the reinterpretation of a state as a bounded event, but Schilder (1997) has argued that instead the interpretation is the result of discourse processing which forces the introduction of a boundary for the interval characterised by the state. The state does not shift aspectual class, its associated interval merely acquires a starting point. This analysis follows from the fact that the conceptualisation of the state and its aspectual properties do not change when a discourse adds a boundary. So for my purposes these cases do not provide a counter-example to the need for a lexical distinction between states and events.

These examples show clearly that the state / event distinction has implications for linguistic processing at the syntactic, semantic, and discourse levels. This evidence points to the need for explicit representation of the distinction. I will therefore incorporate this distinction into the semantic relations hierarchy.

Extended relations

Jackendoff also makes use of subordinating functions which can be freely added to conceptual structures to capture additional entailments or to incorporate information contributed by, for example, an adjunctive modifier. We have already seen the use of one such function, EXCH, in (2.38b) above. These functions take as argument an eventuality X, which is subordinated to an eventuality Y. Those discussed by Jackendoff are listed below, with the corresponding features introducing subevents as used by Davis listed to the right.

- BY X is the means to accomplish Y. (MEANS)
- FROM X is the reason why Y occurs. (CAUSE)
- FOR X is the intended result of the action in Y. (EFFECT)

- WITH The event / state X accompanies the event / state Y. (ACC-EV)
- EXCH Y occurs in exchange for X. (EXCH)

These subordinating functions broaden the applicability of Jackendoff's representations. His conceptual structures are not only used to capture the 'core' lexical semantics of a verb, but can be extended to include the contributions of sentential modifiers to the meaning of a sentence as a whole. Davis' representations, independent of other mechanisms, only reflect the semantics of the main verb of a sentence plus the contribution of its arguments. Thus the subevents associated with the features above can only be subevents entailed by the semantics of a represented verb. This is natural since Davis' goal is to explain the linking of the semantic arguments to syntactic positions, and not primarily to provide a representation from which semantic entailments can be derived. Some mechanism for extending Davis' relations with semantic information not directly relevant to linking and capturing subevents introduced by phrases not participating in the core semantics of the matrix verb in a sentence — i.e. for accounting for adjuncts which augment the semantic relation expressed by a sentence — must be implemented.

Polysemy

For the purposes of this thesis, Davis' emphasis on linking relations means that his representational architecture is not powerful enough. An issue which fails to be addressed adequately is that of *complementary polysemy* (Pustejovsky 1995a), i.e. cases in which a single verb has multiple senses which are related to one another in some predictable way.

There are essentially two cases of polysemy relevant to this point. The first is that a verb can have a sense which derives from the interaction of its base sense with the semantics of a complement. This is the case for *bake*, for example, which can be interpreted as a change-of-state verb or as a creation verb with particular complements, as shown by (2.52) (Pustejovsky 1995a:122). This process is called *co-composition*.

(2.52) a. John baked the potato. (*change-of-state*)

b. John baked a cake. (*creation*)

The creation sense arises from the fact that a cake comes into existence by baking, and so there is "co-specification" between the verb and its complement. The difference between these two senses cannot be captured in Davis' hierarchy as the semantic roles played by the various semantic arguments are identical in each case, at least as far as linking is concerned. There is no syntactic difference displayed by these two forms. There is however a semantic difference which is not reflected. This could easily be remedied by introducing another subsort of *und-rel*, *creation-rel*, to correspond to the entailment of creation of an entity.[17] Specific composition rules

[17] In fact, the difference between *ch-of-st-rel* and *creation-rel* has implications for determination of aspect. Consider the contrast between (2.53a) and (2.53b):

(2.53) a. John baked a potato for 90 minutes / ?in 90 minutes.

are then needed to control the co-composition between the verb and complement to induce the shift from *ch-of-st-rel* to *creation-rel*.

The second kind of complementary polysemy occurs when the primary semantic relationship between the semantic arguments is different in the different senses of a verb. In most cases of this kind of polysemy, the verb also appears in a distinct syntactic pattern for each sense. An example of this kind of polysemy is the causative alternation. Consider the sentences in (2.54).

(2.54) a. The ball rolled down the hill.

b. John rolled the ball down the hill.

In (2.54a) we have an instance of a *move-displacement-rel*,[18] while in (2.54b) we have an instance of a *cause-move-rel*. So each sense of *roll* receives an entirely distinct relation under Davis' architecture. Davis reflects correspondences between two senses with statements about the relationship between two lexical semantic structures; for the alternation in (2.54) this would be something like the statement in (2.55). Note that Davis does not assume any directionality in such relationships.

(2.55)

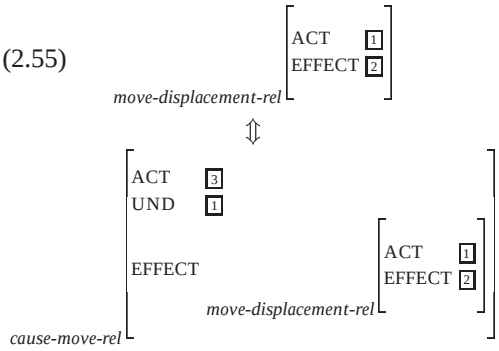

Davis states that it is not necessary to view such statements as lexical rules (ch. 3, p. 76), suggesting that verbs which exhibit alternations are systematically polysemous. He argues (ch. 7, pp. 3-6) that alternating verbs can be characterised as "two senses sharing a lexeme", for which the knowledge of relationships such as (2.55) means that the presence of a lexical entry capturing either of its senses implies the existence of a lexical entry capturing the other sense. This perspective is directly in line with how I will suggest sense shifts should be viewed, but Davis does not explicitly address how he intends for this perspective to be formalised. In each of the following chapters of this thesis I will address this issue of the representation of polysemous verbs. The solutions will involve further use of the power of the multiple inheritance hierarchy and underspecification in lexical entries, following Bredenkamp *et al* (1996) and Markantonatou and Sadler (1995).

---

b. John baked a cake ?for 90 minutes / in 90 minutes.

The sentence *John baked a cake* has a much stronger accomplishment sense than *John baked a potato.*

[18] A *move-displacement-rel* is a relation expressing that the ACTor moves along some path. See Section 4.7.2.

2.4.3 Necessary Extensions

Some advantages of Davis' formalisation over Jackendoff's are listed below.

1. Relation arguments are associated with specific, consistent entailments.
2. It takes advantage of inheritance to capture relationships among semantic relations and to effectively group relations together in various ways (where one relation can belong to several 'groups' via multiple inheritance).
3. The formalism can be directly inserted into the grammar formalism of HPSG, leaving in place a framework for parsing sentences and building up a representation of their semantics.

For these reasons I will build on Davis' formalisation, augmenting it with certain representational elements. Several additions have been motivated in the previous section on the basis of Jackendoff's proposals. I will suggest several more here which do not derive from Jackendoff's work but from the issues to be examined in this thesis. My goal in proposing these extensions is to extend the functionality of the representation language to allow for modeling of interactions between lexical semantic information in compositional semantics, as well as the interactions with syntactic realisation.

Nominals

The interpretation of a verb may depend on certain properties of its arguments. Pustejovsky (1995a) has argued, for example, that a process of *co-composition* can occur between a verb's semantics and the semantics of a complement noun phrase which shifts the sense of the verb. One of his examples of this process is exemplified by the *bake* sentences introduced above in (2.52). The solution he proposes requires that the complement noun be associated with a lexical semantic structure representing core aspects of the noun's meaning. This semantic structure interacts with the semantic structure of the verb in order to make precise the intended meaning of the verb, relative to its complement. I will outline in greater detail in Section 2.6 the semantic structure Pustejovsky (1991, 1995a) proposes for nouns. I will adopt his representation and will show in Chapter 5 how its use can further my goal of capturing the generative processes which underly sense extensions.

Preposition semantics

As will become clear in Section 2.5.1 and Chapter 3, the semantics of a modifying prepositional phrase can interact in various ways with the semantics of what it modifies. I will therefore introduce a hierarchy of preposition types which will be used to type prepositions according to how they interact with the phrases they modify. This typing will be used in the compositional semantics to control the interaction of modifier with modfied.

A situation index

Davidson (1967) observes that the logical form of action sentences should include some kind of index, or reference marker, in order to characterise an action as a singular term. This index can then be referred to or elaborated in subsequent sentences. He provides the discourse in (2.56) to illustrate his point (Davidson 1980:105).

(2.56) Strange goings on! Jones did it slowly, deliberately, in the bathroom, with a knife, at midnight. What he did was butter a piece of toast.

This he paraphrases into an informal logical form of 'There is an action *x* such that Jones did *x* slowly and Jones did *x* deliberately and Jones did *x* in the bathroom…'. He proposes that actions are things and that, analogously to nominals which are associated with an index, the term which must be substituted for *x* is a variable associated with an event predicate. So in the case of (2.56) we might represent the action and its modifying information as in (2.57). Davidson argues that this representation allows us to maintain appropriate entailment relations among actions.

(2.57) $\exists x \textit{butter}(\textit{Jones}, \textit{a piece of toast}, x) \wedge \textit{in}(\textit{the bathroom}, x) \wedge \textit{deliberately}(x) \ldots$

Kamp and Reyle (1993) extend this idea from events to states, arguing that particular states can also be referred to in discourse. They distinguish between discourse referents (indices) for events and states explicitly by using es to refer to events and ss to refer to states (e, e', e'', …and s, s', s'', …). Schilder (1997) argues against this, preferring an analysis in which there is only one type of discourse referent, referring to *situations*. A situation can be either a state or an event. This allows Schilder to define a logic of temporal relations over situations in general, rather than having a separate treatment of states and events.

I will assume that there is a sort *situation* in the sort hierarchy which subsumes *event* and *state*, and which is defined for an index SIT-IND.

Davis, in contrast, opts for having structure-sharing of CONTENT values instead of making use of a referential index. He adopts this approach not only for events but also for nominal objects. His motivation for this stems largely from linking issues (including some data on reflexives and the fact that only one syntactic argument should realise a given semantic role), but also because he desires a uniform statement of the subcategorisations of verbs which allow either NP or VP/CP complements. In the original HPSG analysis of VP complements, there is full sharing of the CONTENT values of the VP complement and the SOA-ARG of the main verb's CONTENT, and Davis simply extends this mechanism to nouns for uniformity. The main criticism of this approach stems from the logical interpretation of the resulting feature structures: having full CONTENT values rather than indices serving as arguments to predicates results in a higher-order logic, so that rather than representing (2.56) as (2.57), we would have the interpretation (2.58).

(2.58) $\textit{in}(\textit{the bathroom}, \textit{butter}(\textit{Jones}, \textit{a piece of toast})) \wedge \textit{deliberately}(\textit{butter}(\textit{Jones}, \textit{a piece of toast}))$ …

The logical complexity of the feature structure is therefore dramatically increased, making reasoning computationally less efficient. The entailment relations of the logical form in (2.58)

are also incorrect as there is nothing which forces the *buttering* event in the first conjunct to be the same as the one referred to in the second conjunct. I therefore prefer using indices for both nominals and events. All the feature structures previously given for Davis' relations will be modified to use indices rather than full CONTENT values as the values of the attributes (i.e. the values will be of type *index* rather than *nom-obj*). This change also allows for the fact that some nouns may refer to events and can therefore be associated with an event index. The problems which my choice creates for linking will need to be re-examined in future research, in the light of the logical complexity which Davis' choice leads to.

Defaults

Regularities often have exceptions. It is therefore important to have mechanisms both for expressing regularities and their exceptions. Inheritance gives a way of straightforwardly expressing regularities: subtypes of a more general type each have something, some regularity, in common. These subtypes may however differ from their supertype in some way that does not eliminate the usefulness of representing the regularity, and for these cases it is important to incorporate the notion of *defaults*. Some very well-established examples of the need for such defaults in the lexicon come from morphology: consider a rule for past participle formation which states that *-ed* is added to the root of the verb. This rule can be applied productively to generate past participle forms of many verbs, including novel verbs. However, there are exceptions. The past participle form of *hold*, for example, is not *holded* but *held*. Other examples can be found in Gazdar (1987), and some examples of the application of defaults specifically in lexical semantics are discussed in Copestake (1992). In Section 2.6 I will introduce a representation for nominals which incorporates defaults.

Summary

To sum up, the extensions to Davis' architecture which are needed are as follows:

- A further elaborated semantic relation hierarchy, including subsorts for *orient-rel*, *extend-rel* and *conf-rel*.
- An explicit ontological distinction between states and events.
- Mechanisms for addressing polysemous verbs.
- Mechanisms for incorporating the semantic contribution of modifiers.
- The introduction of a situation index.
- A further developed representation for nominals (THINGs).
- A description of the semantics of prepositions.
- Incorporation of defaults.

These will be incorporated as described in the following two sections.

2.5 Verb semantics

In this section I will present the representation of verb semantics which I will assume in the remainder of this thesis. It is based on Davis' (1995) architecture for lexical semantics, but is expanded to accommodate the extensions proposed in the previous section and to provide the basis for interaction of this lexical semantics with the semantics of sentence modifiers. I will outline precisely how the representation presented integrates into the standard HPSG lexical entries for verbs (see Section 2.3.1), and give examples of that integration. In the following discussion, I will assume basic familiarity with HPSG as described in Pollard and Sag (1994).

2.5.1 Internal and External verb semantics

I propose to make a distinction in the semantic representation of a sentence between *internal* and *external* semantics.

The internal semantics of a verb reflects the meaning expressed by the verb itself. This includes specification of the verb's semantic arguments and all of the relations involving these arguments: the roles they play, and any events / subevents which are entailed by the verb. This structure corresponds at the word level to a semantic relation in the sort hierarchy rooted at *situation*.

The external semantics reflects meaning particular to a particular situation expressed by the verb on a particular use. Examples of elements of external semantics include location, time, and thematic information (contributed to a situation by thematic adjuncts). The latter might include purpose clauses and temporal adverbial phrases like frequentatives and duratives. These elements derive from explicit components of sentences, and therefore can be considered part of the truth-conditional meaning of these sentences, in contrast to information in the BACKGROUND conditions of the sentences which reflect presuppositions or conventional implicatures not directly contributed through the meaning of phrases in the sentence. This information cannot therefore simply be pushed into the BACKGROUND.

This distinction allows us to identify and isolate the kinds of effects of modifying phrases can have on the interpretation of a sentence. Consider the sentence in (2.59).

(2.59) John saw Mary in the park.

This sentence can be interpreted in two ways. The preferred reading is that "John seeing Mary" event took place in the park, which entails that both John and Mary were in the park. Another available reading is that Mary was in the park, and John saw this, without indicating anything specific about John's location. This reading may actually be the preferred reading under certain circumstances, given world knowledge. This is the case for (2.60). We wouldn't want to say that the PP in this example is part of the complement noun phrase, due to the oddness of (2.61).

(2.60) John saw Mary in the ladies' room.

(2.61) *Mary in the ladies' room was seen by John.

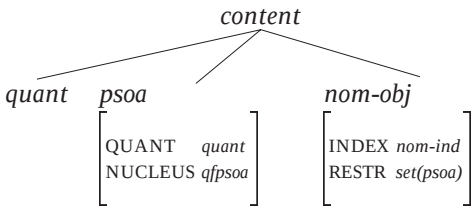

Figure 2.6: Sort hierarchy of *content* objects

There is, however, no available reading for (2.59) which places John in the park and Mary somewhere outside of the park. So a locative phrase can be taken as restricting the location of an event as a whole, or of particular participants in the event (in this case, the direct object). Traditional treatments of locative adverbials (e.g. Davidson 1967) would only account for the located-event reading, as they are only associated with event variables. Introducing a distinction between internal semantics and external (situational) semantics opens up the possibility of an analysis where the adverbial is interpreted as either an external (event-oriented) modifier, or an internal (participant-oriented) modifier (see also Johnston 1994 for discussion of the event- vs. participant- orientation of purpose clauses). The distinct readings can be predicted through an ambiguity in the compositional semantics rather than through complex inferencing: the locative can modify either the event index, or pick out individuals from the internal semantics.

Another example for which this internal / external modification ambiguity holds is found in (2.62).

(2.62) The mouse ran under the table.

This sentence has three available readings: either the "mouse running" event is located underneath the table, or the mouse ran along a path with its endpoint under the table, or the mouse ran along a trajectory which includes a point under the table. The first reading is an event-oriented reading of the locative, while the other two readings are both participant-oriented readings (in Davis' terms, the mouse is FIG and the path is GRND in a *mot-rel*). The two participant-oriented readings stem from ambiguity of the preposition *under*.

I will return to these issues in Chapter 3, where I will explore more fully the status of modifying prepositional phrases. For the current purposes of defining a representation for verb semantics, however, the examples given here should suffice to justify the internal / external distinction I introduce. In Chapter 3, we will see how this distinction can lead to a solution for the modification ambiguities presented here.

2.5.2 Formalisation of the representation within HPSG

Following standard HPSG, I assume that the CONTENT value of a verbal lexical entry is of type *psoa* and that a *psoa* is defined for a NUCLEUS feature of type *qfpsoa*, as summarised in Figure 2.6. I depart from the standard definition of *qfpsoa*, however, instead assuming a hierarchy as indicated in Figure 2.7. The sort *qfpsoa* is divided into four main subsorts: *sit-desc*

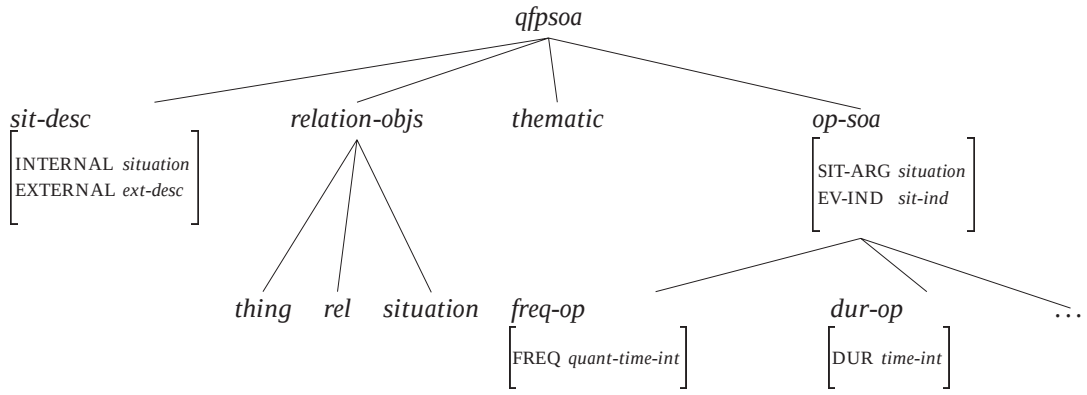

Figure 2.7: Sort hierarchy of *qfpsoa* objects

(situation description) which will be used to reflect the semantic representation of a situation, *relation-objs* which is the type of all basic semantic relations including those for both nouns (*thing*) and verbs (*situation*), *thematic* which will capture thematic elements of either internal or external semantics, and *op-soa* which will capture the semantics of different kinds of operator adverbials, such as frequentatives (*freq-op*) and duratives (*dur-op*).

The type *sit-desc* is defined for attributes INTERNAL, with value of type *situation*, and EXTERNAL, with value of type *ext-desc*. I assume a lexical constraint on the feature structure associated with a verb which requires the SYNSEM:CONTENT:NUCLEUS field of a verb to be of type *sit-desc*. In this structure the internal semantics of a situation is held distinct from its external semantics. This must be represented in the lexical entry for verbs, and not just at the sentence level, due to the mechanisms of the Semantics Principle of HPSG. This principle is shown in (2.63) (Pollard and Sag 1994:322).[19]

(2.63) In a headed phrase,
(Case 1 – e.g. for a constituent headed by a verb)
if the semantic head's CONTENT value is of sort *psoa*, then its NUCLEUS is token-identical to the NUCLEUS of the mother;
(Case 2 – e.g. for a constituent headed by a noun or preposition)
otherwise, the CONTENT of the semantic head is token-identical to the CONTENT of the mother.

Through this principle, the semantics of a sentence is projected from the semantics of the main verb of the sentence, since a sentence is a constituent headed by a verb. This means that if the internal/external situation semantics distinction is made within the CONTENT:NUCLEUS field of a verb at the lexical level, it will be passed up to the sentence level via the Semantics Principle. As a verb combines with its arguments, their CONTENTs will be incorporated into the internal semantics characterising the situation expressed in the sentence. Adjuncts will be

[19] I will ignore issues of quantification in this thesis. Therefore the clause of the Semantics Principle handling quantification is not included in (2.63).

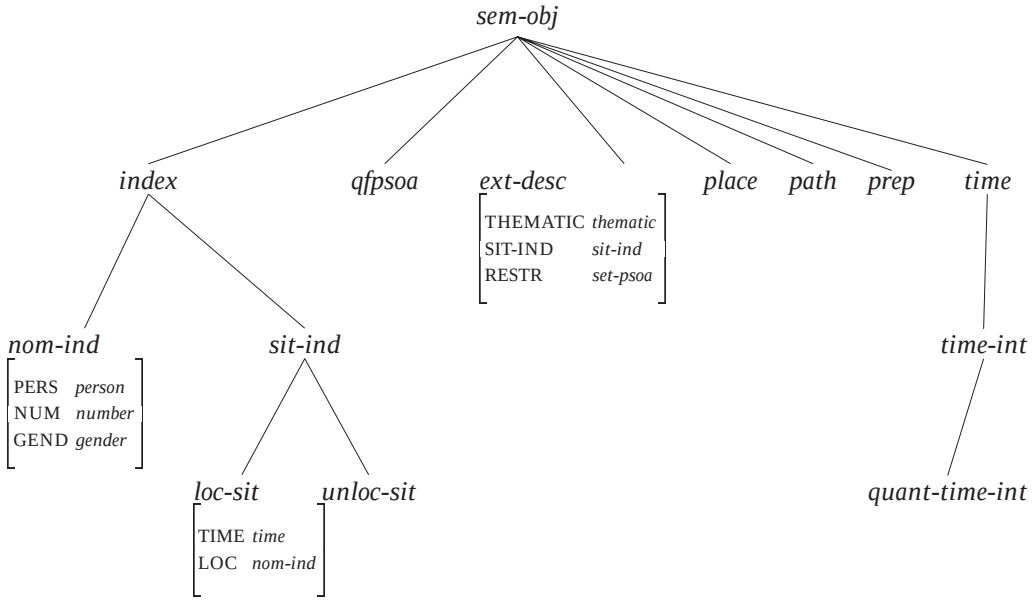

Figure 2.8: Sort hierarchy for Semantic Objects

treated in such a way as to account for the possibility of their incorporation into either the internal semantics or the external semantics of the situation (see Chapter 3).

The type *qfpsoa* is rooted at *sem-objs*, along with other semantic types, as indicated in Figure 2.8. Included in this (sub-)hierarchy is the definition of the type *ext-desc*. This type is defined for an attribute, THEMATIC, which will capture thematic elements of a situation. The hierarchy of *thematic* objects is partially specified in Figure 2.9. The attributes for which a thematic object can be defined correspond to Jackendoff's subordinating relations, and reflect different kinds of situational information which can be added to the description of a situation via clausal modifiers.

The use of an inheritance network for characterising this thematic information is inspired by the subsumption hierarchy defined in Wechsler (1991) for sorting psoas according to valency, and has two basic motivations: (a) it allows for explicit specification of restrictions on combinations of thematic PPs via missing links in the *thematic* type hierarchy (a theoretical possibility, although I haven't explicitly investigated it) and (b) the type of a THEMATIC attribute will precisely reflect which thematic entities have already been added to a sentence. The latter occurs as a result of moving down the type hierarchy (to more specific subsorts) as thematic elements are added. Multiple modifying phrases of the same thematic type do not appear to be possible, and this approach allows redundancy constraints on thematic entities to be implemented. Sentences such as (2.64), containing two *for-to* phrases, would be ruled out.

(2.64) *Peter sang to Mary to please her to make her happy.

I will assume, although it does not appear in the partial *thematic* hierarchy represented in Figure 2.9, that each type will also have a "grounded" subtype which cannot be extended. For example, the type *for-to* will have a subtype *for-to_only* which itself does not have any

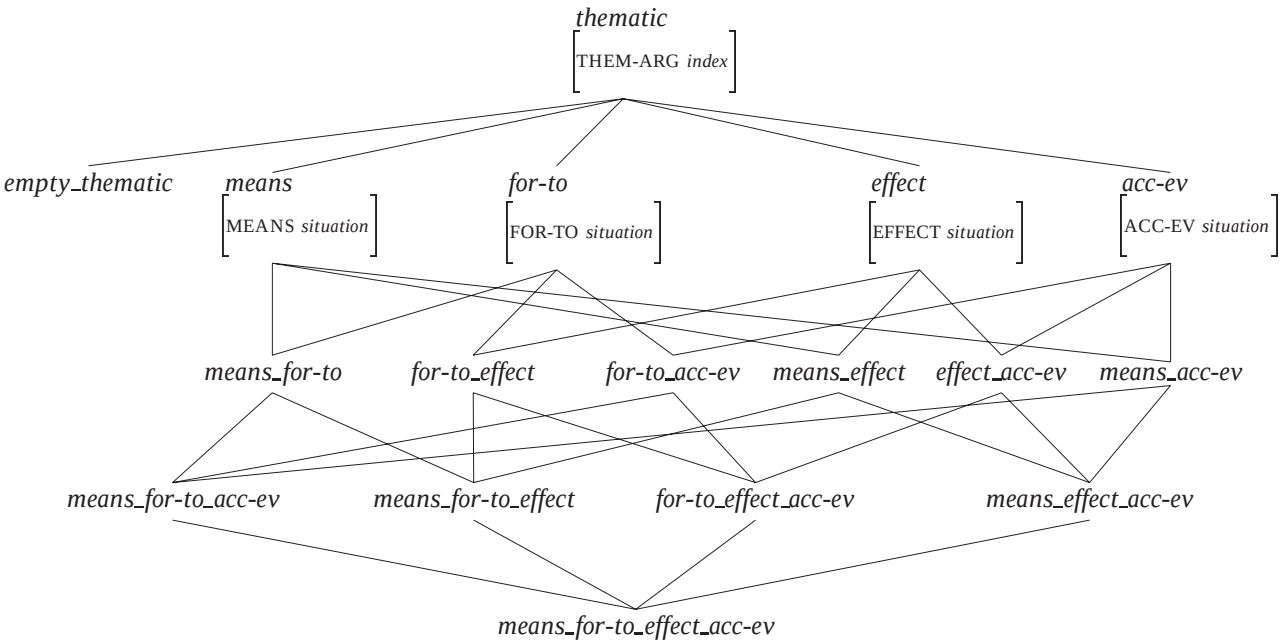

Figure 2.9: Sort hierarchy for thematic relations

subtypes. It will inherit only the attributes of its supertype, in this case only the attribute FOR-TO. Similarly, the type *for-to_effect* will have a grounded subtype *for-to_effect_only* defined for attributes FOR-TO and EFFECT but with no subtypes. These grounded subtypes will be used in the verbal relation hierarchy to constrain the thematic elements which can be associated with particular relations (see Section 3.3.3).

The type *thematic* has an attribute THEM-ARG (thematic argument) which will reflect which index from the internal or external semantics plays a role in the thematic situations. This will be used in Chapter 3 in accommodating prepositions which can modify either the internal or the external semantics.

The type *ext-desc* is also defined for two additional attributes: SIT-IND and RESTR. The former attribute captures the situation index which is associated with the situation expressed by the sentence as a whole. This will be anchored to the main eventuality expressed by the verbal relation, through the constraint on the type *sit-desc* specified in (2.65).

(2.65)
$$\begin{bmatrix} \text{INTERNAL:EV-IND } \boxed{1} \\ \text{EXTERNAL:SIT-IND } \boxed{1} \end{bmatrix}$$

The RESTR is a restriction feature, analogous to the restriction feature for nominal objects (type *nom-obj*), which will reflect any semantic restrictions on the situation. These will largely come from restrictive adjuncts such as temporal and locative adjuncts and will restrict some aspect of the SIT-IND.

The *sem-objs* hierarchy also contains the declaration of the *index* type. This type is divided into two subsorts. The first, *nom-ind*, reflects the standard HPSG nominal index declaration. The second, *sit-ind* is the type of the indices which will be associated with situations. Building

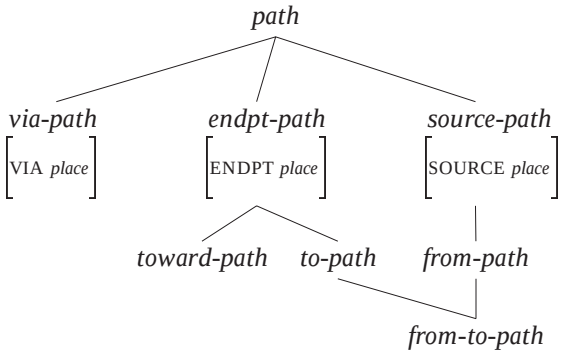

Figure 2.10: Sorts associated with *path*s

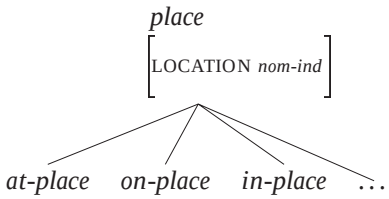

Figure 2.11: Sorts associated with *place*s

on ideas in Kasper (1993), which treats some situations as located in space and time and others as unlocated, *sit-ind* is divided into two subtypes: *loc-sit*, defined for attributes LOCATION and TIME,[20] and *unloc-desc*, not defined for either attribute.

The types *path* and *place* correspond directly to the ontological categories PATH and PLACE proposed by Jackendoff (1983, 1990). They are integrated into a typed inheritance framework as shown in Figures 2.10 and 2.11. Different kinds of paths (subsorts of type *path*) are defined for features indicating which part of the path is in focus: the source (*source-path*), the endpoint (*endpt-path*), or a location at some other point along the path (*via-path*). Each feature takes an object of type *place* as an argument. The *place* hierarchy is divided into subsorts corresponding to different areas relative to a LOCATION of *thing* which are relevant: the place of interest could be *at* the location, *on* the location, etc. This hierarchy clearly does not exhaust the possible relations that can hold between a location and a located entity, but simply serves as an approximation for the purposes of this thesis. For a fuller investigation of spatial prepositions, see e.g. Sablayrolles (1995).

A hierarchy of preposition types, rooted at *prep* is shown in Figure 2.12. The motivation for these types will be outlined in full detail in Chapter 3, but it stems from the observations made in Section 2.5.1 that modifiers (including modifying prepositional phrases) can be either external or internal. Prepositions which can be used for external modification will be of type *adjunct*. Prepositions which can be used for internal modification will be of type

[20] I have not explored the representation of temporal information and will leave the precise definition of the *time* type unspecified.

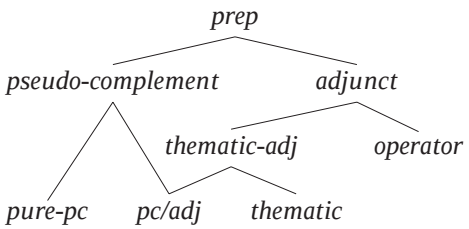

Figure 2.12: Prepositions defined in the sort hierarchy

pseudo-complement (this terminology will be explained in Chapter 3). Operator adverbials (e.g. frequentatives, duratives) will be headed by prepositions of type *operator* and will be constrained to have a CONTENT:NUCLEUS value of type *op-soa*.

In Section 2.4.3 I outlined extensions to Davis' (1995) relation hierarchy which are necessary for achieving the goals of this thesis. A portion of the resulting hierarchy can be found in Figure 2.13. I have maintained the top-level relation sort assumed by Davis, *rel*, for compatibility with his linking theory and because not all relation sorts correspond directly to a particular kind of situation (either a state or event). I have, however, added another top-level relation, *situation* which has subsorts *state* and *event*. The hierarchy rooted at *situation* cross-cuts the *rel* hierarchy, essentially categorising particular relations in the *rel* hierarchy as either *states* or *events*. This categorisation was motivated in Section 2.4.2.

The hierarchy also reflects the introduction of various sorts discussed above, derived from Jackendoff (1983, 1990): *conf-rel*, *extend-rel*, *orient-rel*, *go-rel* (and its subsort *ch-poss-rel*) and typeact-inch-rel. It also shows the addition of *creation-rel* as distinct from *ch-of-st-rel*, given the differing presuppositions of a change-of-state verb as compared with a creation verb, mentioned above in the discussion of complementary polysemy.

In proposing this sort hierarchy, I adopt the basic architecture as assumed by Davis, with the changes to the relation argument type discussed in Section 2.4.3 (*A situation index*). The features which are defined for the various sorts rooted at *rel* in Davis' hierarchy will also be defined for those sorts here. I assume that *act-rel* is defined for a feature ACT of type *index*, *und-rel* for a feature UND of type *index*, and so on. The thematic elements (ACC-EV, EFFECT, MEANS, etc.) are treated differently than in Davis' framework in that they are bundled together under a single attribute THEMATIC which can be associated with an *event* but not a *state*. The value of this feature is of type *thematic*, as described previously, shown in Figure 2.9. Internal verb semantics can thus incorporate the same kind of thematic relations as external verb semantics.

This sort hierarchy would be elaborated to represent further verbal relations, including relations such as *notion-rel* and *influence-rel* as depicted in Figure 2.2 but not explicitly shown here. Each additional relation would inherit from (at least one of) the subsorts represented in Figure 2.13 and would have specific entailments associated with its arguments. In the remainder of this thesis I will introduce as necessary subsorts of *situation* and *rel* to reflect the semantics of particular verbs, and I will assume that they are inserted into the hierarchy in the appropriate place.

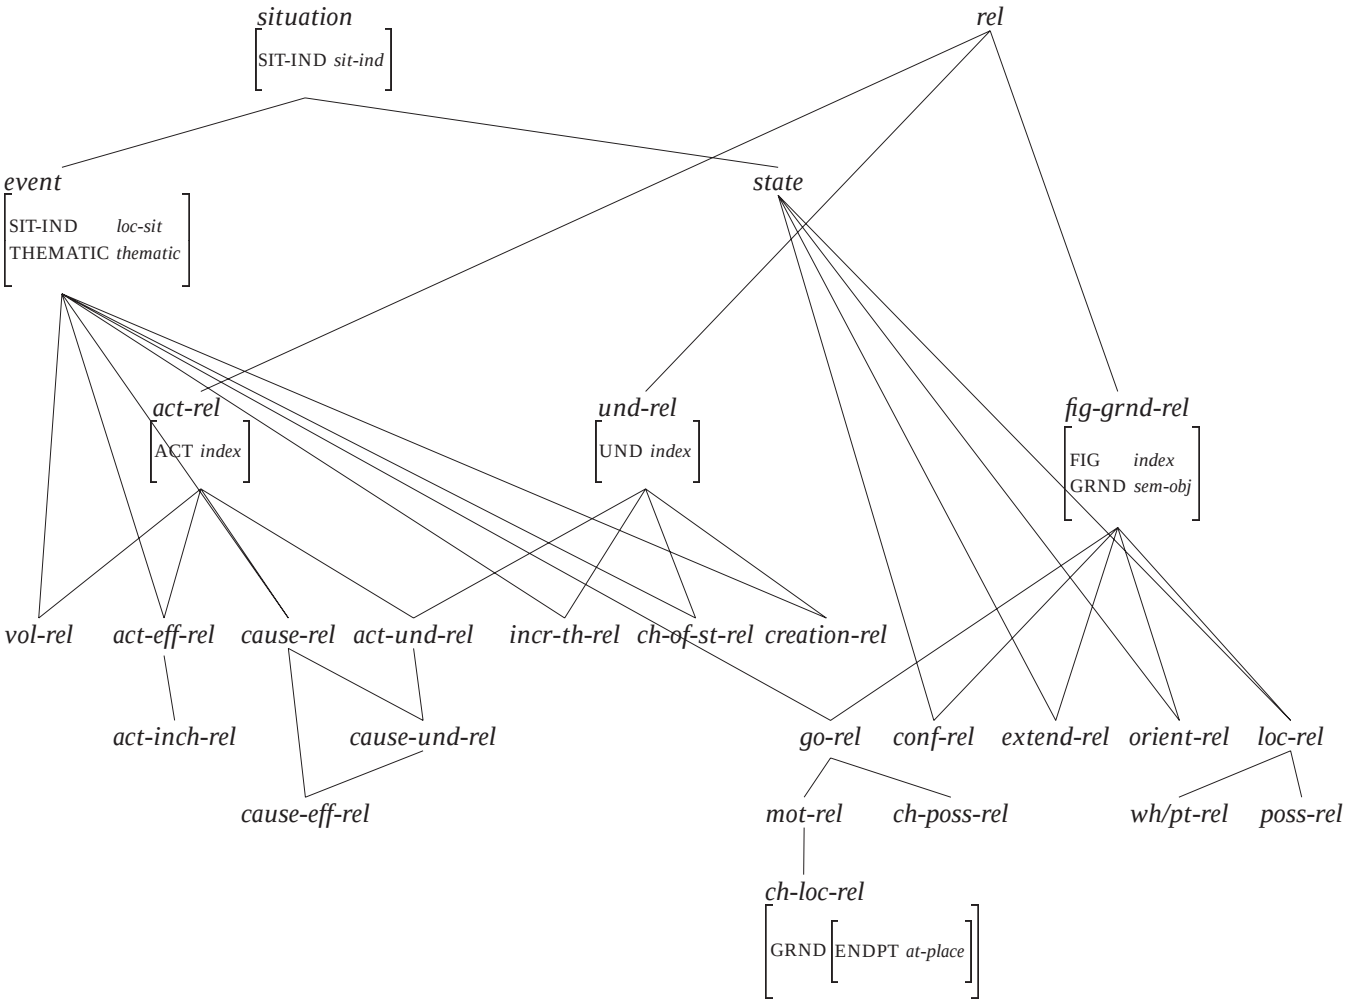

Figure 2.13: Proposed *situation* and *rel* sort hierarchy

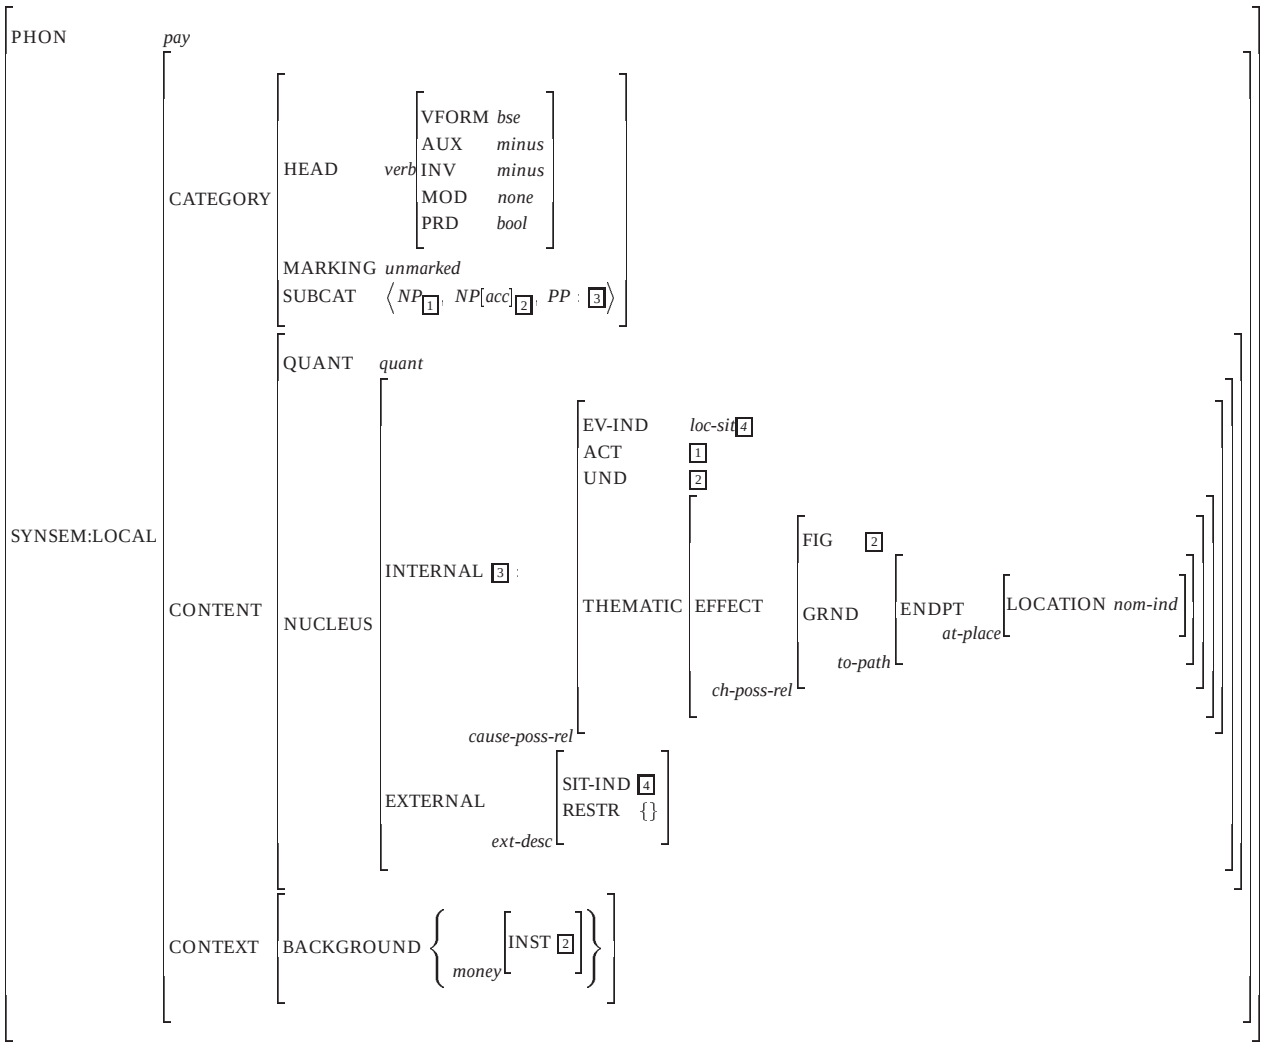

Figure 2.14: The lexical entry for the verb *pay* on the sense expressed in *John paid $100 to Mary*

2.5.3 An example of the representation

A sample HPSG lexical entry incorporating the representation as described above for the verb *pay* in the sense of *John paid $100 to Mary* can be found in Figure 2.14. The semantics expressed in this entry can be paraphrased as *John acts on $100 with the effect that $100 goes to Mary's possession*.

I will briefly review what the various fields of this representation correspond to. The phonological (PHON) feature has as its value the written word to which the entry corresponds, for lack of a more precise phonological transcription. The SYNSEM feature contains the syntactic and semantic information associated with the word being represented. Only the LOCAL information will be relevant in this thesis, specifically the CATEGORY and CONT (content) information. CATEGORY includes the HEAD features of the verb, all of the features

defining the form of the verb and how it can be used (AUX specifies whether the verb is an auxilliary verb, INV specifies whether the verb can appear in inverted form, MOD contains verb modifier information, PRD specifies whether the verb is predicative). The CATEGORY also contains a MARKING feature which indicates whether the verb is being used within a complementised clause (see Pollard and Sag 1994, pp. 45-47), and the SUBCAT feature which has a list of synsem objects, corresponding to the SYNSEM values of the signs with which the verb must combine to become "saturated". Here I use the notation "NP$_{\boxed{i}}$" as a shorthand notation for the feature structure representing a noun phrase, with an index value of $\boxed{i}$, as shown in (2.66) for both object- and event-referring NPs.

(2.66)

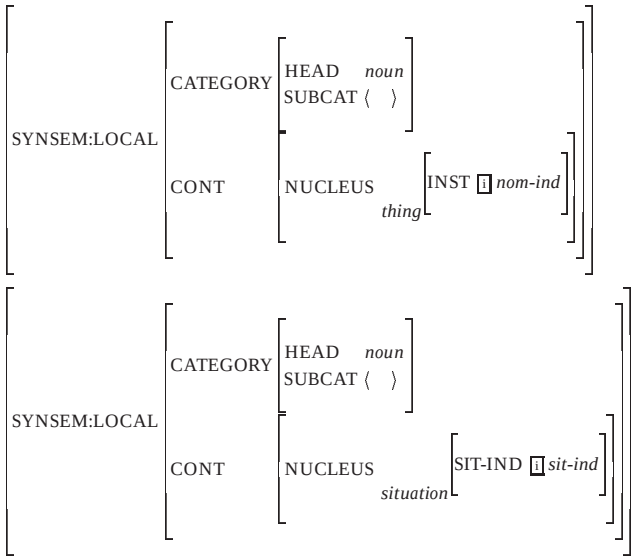

The CONT (content) field has two features: the NUCLEUS, containing the core of the semantic information, and QUANTS, used in the HPSG treatment of quantification (see Pollard and Sag 1994, ch. 8). The value of the nucleus field in this work differs dramatically from what appears in Pollard and Sag's original HPSG work. The semantic representation described in this section is integrated into HPSG in this field.

The value of the NUCLEUS feature reflects the semantic structure of the verb. It is of type *sit-desc* and as such is defined for an INTERNAL feature of type *situation* and an EXTERNAL feature of type *ext-desc*. For *pay*, the INTERNAL semantics is an *event* of type *cause-poss-rel*, a subsort of *cause-eff-rel* for which the THEMATIC:EFFECT is constrained to be a *ch-poss-rel* (as defined in example (2.47)). Here the entity which is affected by the causation and undergoes a change of possession, corresponding to the second NP on the SUBCAT list of *pay*, is restricted to be an entity of type *money*. The first NP on the SUBCAT list is tied to the causer (ACTor), and the third NP is associated with the location at which the money ends up. The EXTERNAL feature is not specified in the lexical entry for the verb, since this situational information comes from elements in a sentence other than those indicated on the SUBCAT list of the verb.

2.6 Nominal semantics

Pustejovsky (1991, 1995a) has argued that there are certain core lexical semantic elements of a nominal which must be represented in its lexical entry in order to achieve a more generative description of lexical sense derivation. Pustejovsky proposes that these core elements can be captured by several roles in a structure called *qualia structure*, based on Aristotle's notion of modes of explanation which drive our basic understanding of an object or a relation in the world, which specify four essential aspects of a word's meaning:

- CONSTITUTIVE: the relation between the denotation of the word and its constituent parts. *(generally used for concrete objects)*
  - Material
  - Weight
  - Parts and component elements
- FORMAL: that which distinguishes the denotation of the word within a larger domain.
  - Orientation
  - Magnitude
  - Shape
  - Dimensionality
  - Color
  - Position
- TELIC: the purpose and function of the denotation of the word.
  - Purpose that an agent has in performing an act
  - Built-in function or aim that specifies certain activities
- AGENTIVE: factors involved in how the denotation of the word came into being; its origin.
  - Creator
  - Artifact
  - Natural Kind
  - Causal Chain

The two components of qualia structure which are most directly relevant to the phenomena in this thesis are the TELIC and AGENTIVE roles. I will interpret the TELIC role as indicating a particular event involving the artifact which most directly expresses the use to which the artifact is stereotypically/conventionally put. For example, liquids which are typically ingested by humans, such as *milk*, *beer*, and *water* will be specified for a *drinking* event in their TELIC role. The value of this role will be formalised as a semantic relation from the sort

hierarchy, for which the argument role played by the denotation of the noun in question is indicated via structure sharing with the noun's index. The AGENTIVE role will also be filled by a semantic relation from the sort hierarchy, this time indicating an event which caused the creation of the denotation of the noun. For a *cake*, for instance, this role might be given the value of *bake-rel* and the index of the noun will be associated with the UND argument of this relation.

Nouns will also be added to the ontology captured in the sort hierarchy, in order to capture relations between entities and to make use of inheritance to efficiently characterise related semantic types. In the hierarchy for nominal objects I will follow Copestake (1992) in attempting to capture *general* information that is true about individual nominal types. As a result, much of the information associated with the nominal types will be inherited by default — specific instances of a nominal type may not be accurately characterised in terms of the information it receives by default through the specifications in the hierarchy. However, the representation of this information should allow verbs to specify selectional restrictions to a certain extent, and to account for co-compositional behaviour.

In order to allow for defaults and default inheritance, I assume an order independent default unification for typed feature structures as defined by Lascarides *et al* (1996). Typed feature structures, as defined by the sort hierarchy, can include default information (which is marked as being default) which can be inherited by their subsorts or overridden with sort-specific values. An example of the application of this default inheritance to TELIC role values from Lascarides *et al* (1996) and adapted to the representation formalism introduced here, is shown in Figure 2.15.

Defeasibility is indicated with a slash notation: values are of the form *indefeasible/defeasible*, which is abbreviated in the figure to */defeasible* where the indefeasible value is completely general. An operation called Persistent Default Unification (PDU) is defined to capture the default inheritance appropriately.

Figure 2.15 reflects several features of the representation of nominal semantics I assume.

- Nouns are defined in terms of relation sorts. These are meant to be viewed as capturing linguistically-relevant lexical semantics of the associated nouns. Intuitively, some of these sorts correspond to particular concepts in our world knowledge. So for example the sort *book* would have an associated concept *book* in world knowledge. However, the sorts are part of a *linguistic* ontology and as such there may be specifically linguistic features and the hierarchical structure may not reflect every aspect of relations between entities in the world.

- Nominal relations are defined for a feature QUALIA which is of type *qualia*. The *qualia* type is in turn defined for features corresponding to the four qualia structure roles identified by Pustejovsky (1991, 1995a).

- Nouns are defined for an index, analogous to the INDEX of *nom-obj*s in HPSG (Section 2.3.1). It is labelled INST (instantiated) and is of the type *nom-ind*, defined in the previous section.

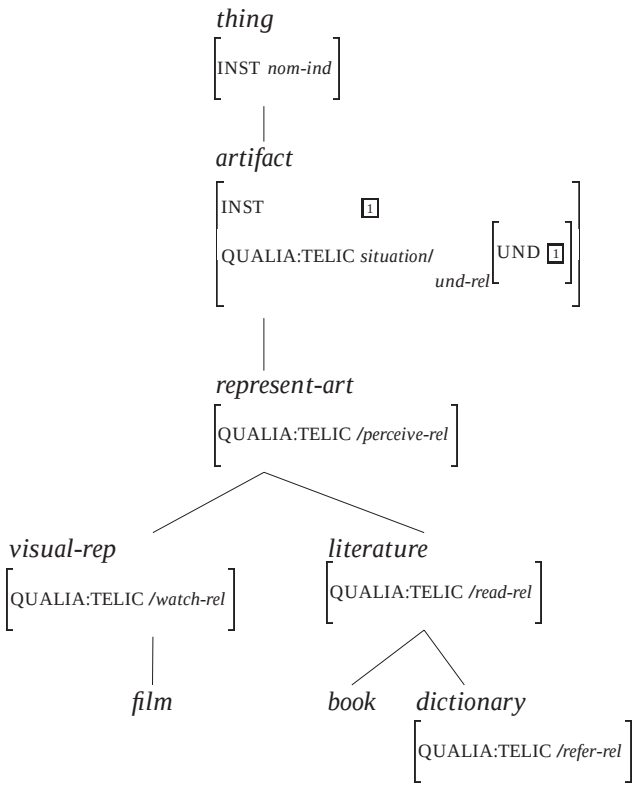

Figure 2.15: Inheritance of the Telic roles of artifacts (Lascarides *et al* 1996)

The figure is intended to convey that the TELIC value of an artifact's qualia should always be subsumed by a *situation*, and by default will be a relation in which the artifact is an UNDergoer. More specifically, something characterised by a *represent-art* will by default have a TELIC role of type *perceive-rel*. This relation will inherit from *artifact* the default structure sharing between the INST and the UND value. Further down the hierarchy, the subtype *dictionary* of the type *literature* will override its supertype's default TELIC value with a more specific default, while *book* doesn't. So a *book* will have a TELIC value of *reading* and a *dictionary* will have a TELIC value of *referring_to*.

2.7 Conclusions

In this chapter I have discussed the need for lexical semantic representation, and the particular form which that representation should take. The conclusions to be drawn from this discussion can be summarised as follows:

- Lexical decomposition is important for capturing the interaction between syntactic structure and meaning, and for capturing generalisations over the syntactic behaviour and interpretation of related words. The identification of lexical semantic components should derive from examination of the interaction between syntactic structure and meaning as reflected in sentences exhibiting minimal variations.
- The examination of linguistic data can lead to development of a highly constrained framework.
- An ontological structure is required which reflects linguistic (as opposed to world knowledge) generalisations, laying the foundation for a modular treatment of language processing in which language-specific knowledge is isolated from general reasoning.
- There must exist a logical interpretation of the representation which allows for certain inferences to be drawn on the basis of a consistent structure.

I reviewed two proposals for lexical semantic representation, Jackendoff (1983, 1990) and Davis (1995). I argued that Jackendoff's system captures the appropriate level of decomposition but that Davis' architecture captures relationships between representations in a much more natural and efficient manner, particularly in light of current developments of constraint-based grammars. I therefore introduced a representation which is based on Jackendoff's Conceptual Semantics in spirit, but on Davis' architecture in the formal details. This representation is fully compatible with the details of HPSG, overriding only the kind of relations which are assumed to reflect the semantics of words and phrases. I expanded Davis' representation on the basis of Jackendoff's proposals to allow for the representation of situation-level interpretations of full sentences, and to accommodate default information associated with individual words. The subsequent chapters of this thesis will build on this representation to address the problems of sense shifts, lexical semantic interactions, and polysemy.

three Prepositional Phrases and Verb Semantics

3.1 Introduction

Prepositional phrases (PPs) which modify verb phrases contribute semantic content which is consistent (unchanging) across uses and identifiable. Their interpretation is, however, constrained by the lexical semantic structure of the modified phrase. In this chapter, I will investigate the treatment of the semantic contribution of prepositional phrases in HPSG, with particular reference to dative PPs. Dative PPs are those PPs which participate in the dative alternation. There are two basic kinds: *to*-datives, which appear in the two alternate forms shown in (3.1), and *for*-datives, as shown in (3.2).

(3.1) a. John gave a book to Mary.

b. John gave Mary a book.

(3.2) a. John baked a cake for Mary.

b. John baked Mary a cake.

Such PPs have traditionally been treated as having little semantic content independent of the verbs with which they appear. I will argue that the consistency of their semantic content suggest that a compositional analysis of their contribution is more appropriate. I will show that the overall interpretation of a sentence containing a VP-modifying PP can only be determined by an interaction of the semantics of the VP and the semantics of PP, and certain principles governing this interaction. I will end the chapter by introducing a general account of the semantic function of PPs, and by extension also of other entities which can function as semantic obliques in a sentence, which takes into consideration the different ways that they can interact with the semantic content of the sentence in which they appear.

Dative prepositional phrases provide an interesting testing ground for investigating various syntactic and semantic properties of prepositional phrases because they have tradi-

tionally been viewed (e.g. Jackendoff 1990, Verspoor 1994) as elements which are subcategorised by the verbs with which they appear and as having very little meaning independent of those verbs. The prepositions are viewed as markers for particular semantic roles. This view derives from analogous phenomena in other languages, such as German, in which these semantic roles are introduced by (dative) case-marking on the nouns which fill the same position as the PPs in English.

More generally, a distinction has traditionally been drawn between prepositional phrases which behave as verbal complements, i.e. which participate in the main relation expressed by the verb (like the dative PPs), and those which are adjuncts, introducing information which helps to contextualise or situate more precisely an eventuality expressed by a sentence. I introduced this distinction in Chapter 2, Section 2.5.1, and incorporated a representational distinction between internal and external semantics to accommodate the integration of the semantics of the different kinds of modifying phrases (via internal predication or external predication). An example of the distinction can be found by comparing (3.1a) with (3.3).

(3.3) John ran a mile in the park.

The PP *to Mary* in (3.1a) introduces an entity which plays a central role in the semantics of *giving* — the role of recipient. This PP is therefore considered to be a (subcategorised) complement of the verb. The PP *in the park* in (3.3), in contrast, does not introduce a central entity in the semantic relation expressed by the main verb. Instead it simply provides location information about the event expressed in the main clause, and is treated as an (unsubcategorised) adjunct.

It is arguable that the class of verbal complements is divisible further. Gawron (1986) identifies *argument PPs*, in which a preposition marks a verbal argument, and *co-predicating PPs*, which introduce a new argument and thereby extend the main verbal relation. Every use of a preposition, however, contributes the same lexical content. It is only the manner in which that lexical content interacts with the content of the modified phrase which varies. This analysis provides the basis of a compositional treatment of the semantics of prepositional phrases.

In this chapter, I build on this insight to present a treatment of prepositional phrases which also includes a three-way basic distinction between PP types. But, in contrast to Gawron, I look in further detail at the range of semantic contribution of adjuncts. The distinction in types which I propose differs from that of Gawron (1986) in that it is driven by syntactic considerations as well as semantic ones. Hence I assume that there are (1) *complement PPs* in which the preposition marks a verbal argument and is subcategorised by the verb, (2) *pseudo-complement PPs*, which share certain syntactic properties with complements yet are not introduced by subcategorisation and which can either mark a verbal argument or introduce a new argument and extend the verbal relation, and (3) *adjunct PPs* which introduce external predication. The pseudo-complements therefore cross-cut the traditional adjunct vs. complement distinction, and subsume some of Gawron's argument PPs and some of his co-predicating PPs. The syntactic distinction will be shown to be important for accounting for syntactic optionality of argument PPs without explicitly encoding such optionality in lexical entries.

I introduce mechanisms within HPSG for licensing pseudo-complement and adjunct PPs which accounts for the syntactic properties of co-predicators in contrast to adjuncts as well as their distinct compositional behaviour. In this treatment I, like Gawron, assume a single lexical entry for each preposition. I further rely on the mechanisms licensing the introduction of non-subcategorised PPs for controlling the appropriate semantic interaction. The range of interactions will be seen to depend on the type of the preposition and the lexical semantics of both the PP and the modified phrase.

The chapter begins with an investigation of the status of *for*- and *to*-dative PPs in Section 3.2 — does the syntactic behaviour of these PPs as verbal complements warrant a treatment of them as essentially semantically void? This will be shown to miss a generalisation about their semantic contribution. I will also show that certain of these PPs must be construed as entities at the level of pseudo-complements, interacting with the main verb's internal semantics.

The integration of the semantics of pseudo-complements with the semantics of the modified entities can be treated in a manner analogous to the treatment for pure adjuncts. However, in Section 3.3 the internal vs. external predication which distinguishes them will be explicitly accommodated in the framework controlling the semantic integration.

Section 3.4 will discuss issues related to the treatment of prepositional phrases in HPSG, including general properties of adjuncts which must be accommodated and the existing HPSG treatments of adjuncts. The existing approaches will be shown to fail in adequately handling certain syntactic and semantic properties of adjuncts. In particular, interactions between surface order and semantic precedence are stumbling blocks for those approaches.

Section 3.5 will propose a framework based on work by Kasper (1993) and van Noord and Bouma (1994) which supports integration of the semantic contribution of all PP types. The semantic representation introduced in Chapter 2 provides the foundation for the treatment of both pseudo-complements and adjuncts within the same system. The treatment of the semantic integration is handled via rules which will be explicitly stated and shown to provide a more satisfactory handling of surface order / semantic precedence interactions as well as other properties of the adjuncts. The framework will also be shown to accommodate a treatment of the phenomenon of the dative alternation.

Finally, we will see a brief example of how world knowledge can influence the grammaticality of a sentence. If world knowledge cannot support an interpretation which stems from lexical semantic composition of the components of a sentence, then that sentence is infelicitous.

3.2 Adjuncts or Complements?

The traditional treatments of dative PPs force the semantic contribution of these PPs to be specified explicitly in the lexical entry for each verb which can appear with a dative PP, essentially treating the preposition as a marker for a particular verbal relation and a particular semantic role, with no independent contribution. These approaches seem to demand the syntactic subcategorisation of dative PPs, even in cases in which they are syntactically optional such as (3.2), since they model the dative alternation as a permutation of a verb's subcategorisation list. This ignores generalisations over the contribution of the PPs, in that the PPs seem

to add similar information regardless of which specific verb they appear with. Thus in each of (3.2a) and (3.4a), the PP *for Mary* specifies who benefits from (some aspect of)[1] the event described in the remainder of the sentence.

(3.4) a. John made a drawing for Mary.

b. John made Mary a drawing.

Likewise, the contribution of the inner NP in the alternate in (3.2b) of (3.2a) and the alternate in (3.4b) of (3.4a) can be identified as specifying who receives benefit from the outer NP (*a cake* and *a drawing*, respectively).

A general question arises from the observation of such generalisations — what is the status of these PP elements? Should they be treated as subcategorised-for complements or as adjuncts which make an independent, identifiable, semantic contribution across verbal heads? This will be investigated below through a series of standard syntactic and semantic tests for argument structure. The two types of dative PPs, *to*-PPs and *for*-PPs, will be contrasted in this investigation. The analysis will show that *for*-dative PPs and certain *to*-dative PPs behave as complements syntactically while behaving as adjuncts semantically. The remaining *to*-dative PPs behave as complements both syntactically and semantically.

3.2.1 Syntactic Tests for Argument Structure

The "do so" Test

The standard syntactic (structural) test for argument structure might be called the "do so" test. In X-bar theory terms, a complement is seen as combining with a lexical category to form an intermediate phrasal category while adjuncts combine with an intermediate phrasal category to produce the same category. The claim is that a full V-bar level constituent can be replaced by "do so". In the case of dative sentences, if the constituent {verb NP} in isolation (i.e. without the PP element) can be replaced by "do so", this indicates that the prepositional phrase is acting as a V-bar adjunct, because {verb NP} is construed as a V-bar constituent in isolation. If only the full constituent {verb NP PP} can be replaced by "do so", the PP must be construed as a complement.

to-datives

As noted in Jackendoff (1990), there seem to be two types of verb classes which can appear with *to*-datives. The first type are verbs for which the PP is a complement, while the PP is an adjunct for verbs of the second type. The "do so" data in (3.5)-(3.12) suggest that *give* and *tell* belong to the first subclass (despite the optionality of the *to*-PP with *tell*), while *send* and *kick* belong to the second.

[1] This qualification derives from the fact that the *for*-PP in these sentences is actually ambiguous between pseudo-complement and adjunctive readings. So the sentences can either be interpreted as indicating that Mary benefited from the event as a whole or that Mary benefited specifically from the cake / drawing. This issue will be addressed in Section 3.3.

(3.5) a. Adam gave a book to Debbie and Brian also did so.
b. *Adam gave a book to Debbie and Brian did so to Susan.

(3.6) a. Adam told a story to Debbie and Brian also did so.
b. *Adam told a story to Debbie and Brian did so to Susan.

(3.7) a. Adam gave a book to Debbie in the library.
b. *Adam gave a book in the library.

(3.8) a. Adam told a story to the children in the bedroom.
b. Adam told a story in the bedroom.

(3.9) a. Sam sent a letter to Bill and Mark also did so.
b. Sam sent a letter to Bill and Mark did (so) to Susan.

(3.10) a. Sam kicked a ball to Bill and Mark also did so.
b. Sam kicked a ball to Bill and Mark did (so) to Susan.

(3.11) a. Sam sent a letter at the post office.
b. Sam sent a letter to Bill at the post office.

(3.12) a. Sam kicked a ball in the park.
b. Sam kicked a ball to Bill in the park.

for-datives

The application of this test to *for*-datives, as shown in (3.13)-(3.14), provides evidence that these prepositional phrases should be treated as adjuncts since "do so" can replace either the full {verb NP PP} constituent or just the {verb NP} constituent.

(3.13) a. Adam baked a cake for Debbie and Brian also did so.
b. Adam baked a cake for Debbie and Brian did so for Susan.

(3.14) a. Adam sang a song for Debbie and Brian also did so.
b. Adam sang a song for Debbie and Brian did so for Susan.

Iterability test

Pollard and Sag (P&S 1987) discuss the complement vs. adjunct distinction, reviewing several syntactic and semantic tests which generally capture usage distinctions between the two types of constituents. One of the syntactic tests is the iterability test. In general, several instances of the same adjunct type can combine with the same head, as shown in (3.15).

(3.15) Kim and Sandy met in Baltimore in the Hyatt hotel in the lobby.
[P&S 1987, (257a)]

Complements, on the other hand, cannot be iterated. Thus in (3.16)-(3.18) the prepositional phrases seem to be complements rather than adjuncts.

(3.16) *Adam gave a book to Debbie to Frank.

(3.17) *Adam told a story to the kids to the adults.

(3.18) *Adam sent a letter to Mary to Diane.

By the same logic, however, it would appear that the prepositional phrases in (3.19)-(3.21) are complements as well, in contrast to the results of the "do so" test reported above.

(3.19) *Sam kicked a ball to Bill to Frank.

(3.20) *Sam sent a letter to Bill to Frank.

(3.21) *Adam baked a cake for Debbie for Susan.[2]

The problem here is that the iterability criterion presupposes a particular semantic relation between the iterated constituents. Adjuncts are not in general iterable if the semantic contribution of each adjunct contradicts the semantic contribution of a previous adjunct. Consider (3.22a) in contrast to (3.22b): the first sentence is ungrammatical while the second is not.

(3.22) a. *Sam kicked a ball at 10 o'clock at 8 o'clock.

b. Sam kicked a ball in the morning at 10 o'clock.

The difference here has nothing to do with a difference in status between the PPs, but rather with the fact that certain semantic roles can be multiply specified via a particular semantic relation. In this example, we see that temporal adjuncts can only be iterated if the information conveyed by a given adjunct is contained in the information conveyed by previous adjuncts. One point in time does not contain another and so the adjuncts in (3.22a) are contradictory, but a point of time is contained in a span of time and so the adjuncts in (3.22b) are considered together to make precise the time at which the event occurred.

So adjuncts can really only be iterated if the semantic (meaning) contribution each makes is in a relation of containment to the previous adjuncts. Thus in (3.15) above, the adjuncts can be iterated because each one can be interpreted as being contained within the location specified by the previous adjunct, making more precise the locative information, rather than providing an overriding semantic contribution. In (3.16)-(3.21), the prepositional phrase specifies the (intended) recipient of some object. The containment relation does not apply to distinct recipients and therefore these adjuncts are incompatible with iteration. This analysis is confirmed by the data in (3.23) and (3.24), which contrast with (3.15) and (3.20) respectively. Sentence (3.23) is ungrammatical because Chicago cannot be contained within Baltimore, while (3.24) is grammatical because the head office of the Times is contained within New York.[3]

[2]This sentence is okay, however, on an interpretation in which the entire action of *Adam baking a cake for Debbie* has been performed for Susan's benefit. See Section 3.3.

[3]Thanks to Janet Hitzeman for the suggestion of this data.

(3.23) *Kim and Sandy met in Baltimore in Chicago.

(3.24) Adam sent a letter to New York to the head office of the Times.

Furthermore, some of the examples Pollard and Sag provide of adjunct iteration rely on pragmatic factors and do not seem to be wholly grammatical. For example, sentence (3.25), introduced as grammatical and felicitous by Pollard and Sag, in my opinion is not entirely felicitous, and certainly can only be interpreted with the two prepositional phrases as adjuncts if the comma indicates a conjunction such as "and".

(3.25) Heather opened the rusty lock with a key, with a pair of pliers. [P&S 1987, (257e)]

Applying this conjunctive interpretation requirement to the ungrammatical sentences above improves their acceptability, as shown in (3.26)-(3.29). These sentences seem to display ellipsis, rather than providing a sense of the underlying argument structure. The analysis of such sentences must be discourse-based, rather than purely based on the syntax/semantics interface, in order to appropriately identify the ellided meaning. In these examples the ellision is intra-sentential, but it is still governed by discourse factors controlling reference resolution.

(3.26) Adam gave a book to Debbie and to Frank.

(3.27) Sam kicked a ball to Bill and to Frank.

(3.28) Sam sent a letter to Bill and to Frank.

(3.29) Adam baked a cake for Debbie and for Susan.

We must conclude on the basis of the examples above that the possibility of iteration of PPs is not a reliable indicator of argument structure, and in fact cannot be viewed as purely syntactic since the phenomenon of iterability seems to interact with semantic and discourse-level factors. The evidence it provides for treating both types of dative prepositional phrases as complements rather than adjuncts will not be taken as definitive.

Relative Order

Pollard and Sag (1987) also point out that in English adjuncts tend to be ordered after complements, suggesting that prepositional phrases which are required to precede other kinds of adjuncts (i.e. which cannot be flexibly interspersed with other adjuncts) are actually complements.

The data in (3.30)-(3.35) suggest that the dative prepositional phrases should be treated as complements according to the relative order diagnostic. These PPs must appear before any other adjunctive phrases.

(3.30) a. Adam gave a book to Debbie in the library.

b. *Adam gave a book in the library to Debbie.[4]

[4]Note that the structure I intend for this sentence is *[Adam gave [a book] [in the library] [to Debbie]] and not [Adam gave [a book [in the library]] [to Debbie]], in which the PP is a noun phrase modifier and which is grammatical.

(3.31) a. Adam told a story to the children in the bedroom.
b. *Adam told a story in the bedroom to the children.

(3.32) a. Adam sent a letter to Mary at the post office.
b. *Adam sent a letter at the post office to Mary.

(3.33) a. Adam kicked a ball to Mary in the park.
b. *Adam kicked a ball in the park to Mary.

(3.34) a. Adam baked a cake for Mary in the kitchen.
b. ?Adam baked a cake in the kitchen for Mary.

(3.35) a. Adam sang a song for Mary in the pub.
b. ?Adam sang a song in the pub for Mary.

Complement-Internal Gaps

Some adjuncts appear to be extraction islands, as shown in (3.36), while unbounded dependencies into complements are generally possible, as shown in (3.37). This provides a further syntactic test for the adjunct / complement distinction.

(3.36) *Which endangered species did Sandy meet someone fond of __? [P&S 1987, (260c)]

(3.37) Which endangered species did Kim impress you as being most fond of __?
[P&S 1987, (261c)]

The data in (3.38)-(3.40) therefore suggest that the dative prepositional phrases are complements rather than adjuncts, since the objects of the PPs can be extracted.

(3.38) Whom did Adam give a book to __?

(3.39) Whom did Adam kick a ball to __?

(3.40) Whom did Adam bake the cake for __?

Pollard and Sag acknowledge, however, that certain adjunct types do appear to sanction internal gaps, as shown in (3.41)-(3.42). One would not want to assume that *sleep* and *die* subcategorise for *without*-PPs, as these PPs do not seem to contribute an argument which plays a role in the main verbal relation.

(3.41) This is the blanket that Rebecca refuses to sleep without __. [P&S 1987, (264a)]

(3.42) Which symphony did Schubert die without finishing __? [P&S 1987, (264c)]

Furthermore, one of the examples they give of an ungrammatical sentence with an adjunct-internal gap, shown in (3.43a), seems to become more acceptable with a different adjunct, as in (3.43b). The change involves replacing the "motivational" adjunct with a *for*-PP (which again one would not want to assume without question is a subcategorised complement), suggesting that this type of PP allows internal gaps.

(3.43) a. *Which famous professor did Kim climb K-2 without oxygen in order to impress __? [P&S 1987, (260b)]

b. Which famous professor did Kim climb K-2 without oxygen for __?

Due to the problems with this test, it is very likely incorrect to assume that the data in (3.38)-(3.40) necessarily indicate that the dative PPs should be treated as complements.

3.2.2 Semantic Tests for Argument Structure

Several semantic differences between complements and adjuncts have also been proposed. I review them in this section.

Constancy of Semantic Contribution

Pollard and Sag (1987) discuss the semantic basis for the distinction between arguments and adjuncts as follows:

> In general, a given adjunct can co-occur with a relatively broad range of heads while seeming to make a more-or-less uniform contribution to semantic content across that range. A given optional complement, by contrast, is typically limited in its distribution to co-occurrence with a small (and often semantically restricted) class of heads (possibly even a single item); in addition, the semantic contribution of the complement is idiosyncratically dependent on the head. (p. 136)

The *to*- and *for*-datives seem to have a constant semantic contribution across the verbal heads with which they appear. In general terms, the *to*-datives seem to indicate the intended recipient of some object, and the *for*-datives seem to indicate the intended beneficiary of something. This semantic uniformity constitutes evidence in support of their treatment as adjuncts. This unchanging contribution was observed by Jackendoff (1990) and formalised in terms of "adjunct rules" which identify the semantic contribution of particular elements in certain syntactic constructions and indicate how this contribution is to be integrated into the semantic representation for the overall construction. The contribution can also be observed in the similarity of the core structures of dative verbs as presented in Verspoor (1994).

The dative PPs, however, appear with only a restricted set of verbs which undergo dative alternation (see Levin (1993) for an outline of the semantic classes of verbs relevant to this alternation). These can be divided into subsets of semantically related verbs. As such, the conclusion from their limited distribution might be that these PPs are complements.

If one adds the criterion that the contribution of an adjunct to the semantic content of a phrase should not simply be the filling of some role in the head's relation, however, one finds that the dative PPs cannot be uniformly classified. There are some "true" *to*-dative verbs, such as *give* and *tell*, for which the semantic information contributed by the dative PP is directly relevant to the meaning of the verb, fills a particular role and would therefore be construed as a "true" argument of the verb (Gawron 1986's *argument* PPs). On the other hand, there are "adjunct" *to*-dative verbs, such as *kick*, for which the semantics of the dative PP provides

additional information about the situation being described by the sentence, adding a semantic role instead of filling a pre-existing role in the verb's relation (Gawron's *co-predicating* PPs). We will see that this is so on the basis of the verbal argument structures below.

Entailment Tests

A good indicator of a verb's argument structure is the entailments of sentences containing the verb. For example, the optional prepositional phrases in (3.44a) seem to be optional complements of the verb *complain* rather than adjuncts due to the entailments displayed in (3.44b). (Examples from Wechsler 1991.)

(3.44) a. John complained (to Mary) (about the heat).

b. John complained.
$\models \exists x, y \mid$ John complained to *x* about *y*.

In contrast, adjunctive prepositional phrases do not result in such entailment patterns, as shown in (3.45). These entailment patterns indicate whether or not particular semantic information is directly relevant to the meaning of the verb. Information which is not directly relevant should be treated as an adjunct rather than an argument.

(3.45) a. John sang (to Mary) (about his homeland).

b. John sang.
$\not\models \exists x \mid$ John sang to *x*.
$\not\models \exists y \mid$ John sang about *y*.

Considering the application of this test to *to*-datives, we find the entailment patterns in (3.46). The data supports Jackendoff's assertion (see Section 3.2.1) that some *to*-datives require the PP as an argument, while for others it is an adjunct. Particularly interesting is the contrast between (3.46b) and (3.46c). The notion of a recipient is more central to the notion of *sending* than to *kicking*, since *send* necessarily involves an (intended) transfer, while *kick* does not. Thus there are semantic differences in the relationship these verbs have to the *to*-PP which were not reflected in the syntactic tests for argument structure, where these two verbs behave similarly with respect to the attachment of the PPs.

(3.46) a. Adam told a story.
$\models \exists x \mid$ Adam told a story to *x*.

b. Sam sent a letter.
$\models \exists x \mid$ Sam sent a letter to *x*.

c. Sam kicked a ball.
$\not\models \exists x \mid$ Sam kicked a ball to *x*.

Applying the test to *for*-datives, we have the entailment patterns in (3.47). The information added by the PPs thus seems to be adjunctive.

(3.47) a. Adam baked a cake.
$\not\models \exists x \mid$ Adam baked a cake for *x*.

b. Adam sang a song.
$\not\models \exists x \mid$ Adam sang a song for *x*.

However, the semantic intuitions about the adjunctive nature of *for*-datives become a bit confused when we consider the data in (3.48) (from Haegeman 1991).

(3.48) a. Hercule bought a detective story for Jane.

b. Hercule bought Jane a detective story.

c. Hercule bought a detective story.

In an unmarked context, (3.48c) is taken to mean that Hercule bought the story for himself. This suggests that the semantic relation of *buying* includes the person for whom the bought object is intended. Thus, (3.48c) $\models \exists x \mid$ Hercule bought a detective story for *x*. This would seem to argue against the treatment of the *for*-PP as an adjunct for the verb *buy*, but the semantic intuitions about *for*-datives remain valid in the general case.

3.2.3 Conclusions

These syntactic and semantic tests for investigating the status of the dative PPs have given conflicting evidence, making it difficult to convincingly establish PPs as either complements or adjuncts. We can, however, identify the following properties:

- The semantic contribution of dative PPs is consistent across verbal heads, regardless of whether this contribution is redundant with respect to the entailments of the main verbal relation (as for *give*) or whether this contribution extends the main verbal relation (as for *kick*).
- Dative PPs tend to behave as complements rather than adjuncts syntactically, even if they are optional.

The most influential argument for the treatment of certain dative PPs as adjuncts rather than complements is that they seem to have a constant semantic contribution across all applications. The fact that these PPs often appear to behave as complements syntactically is overshadowed by the semantic generalisations which can be made by treating them as adjuncts. It does not seem to make sense to treat these PPs as idiosyncratically contributing semantic information to the heads they modify (or not contributing any semantics) when this semantic contribution is so consistent and predictable across word classes. A way of acknowledging their independent semantic contribution while accounting for the syntactic facts is needed. To do this, I appeal to Gawron's (1986) arguments that prepositions have consistent lexical content but that they can modify a verbal head in different ways. I will, however, classify the prepositions according to both syntactic and semantic properties while Gawron's classification is purely semantic.

All of the dative PPs are instances of something which I will call a *pseudo-complement*. This is an element which often behaves syntactically as a complement but which behaves semantically as an adjunct. The information expressed by the pseudo-complement cannot always be

PP Type	Syntax	Semantics
complement	required	mark a verbal argument
pseudo-complement	optional	mark a verbal argument or extend a verbal relation (internal predication)
adjunct	optional	augment the situational description expressed in a sentence (external predication)

Table 3.1: Summary of the properties of different PP types

logically inferred from the use of the verb, but is somehow "closer" to the meaning expressed by the verb than true adjuncts (in Gawron's terms, they are internal predicators). This idea will be developed further in the section which follows.

3.3 Pseudo-Complements

3.3.1 Definition and Relation to adjuncts

It is possible to define a pseudo-complement precisely in terms of its relationship to the semantics of the verbs it modifies. Specifically, a pseudo-complement is an element with an independent semantic contribution involving a semantic argument of the verb. In contrast, adjuncts are elements with an independent semantic contribution involving the full event described by the verb and its semantic arguments. Pseudo-complements and adjuncts are both syntactically optional modifiers of the verbs with which they appear, in opposition to complements which are syntactically required (e.g. subcategorised). Pseudo-complements may introduce a semantically entailed argument of a verb or extend the verbal relation, while the complement PPs are restricted to introducing a semantically entailed argument (such as the *to*-PP obligatory complement of *give*) or providing an idiosyncratic semantic contribution to the verb meaning (e.g. the arbitrariness of the PP complements of the semantically similar verbs *charge (with), blame (for), accuse (of)* noted by Wechsler 1995). Adjuncts do not affect the verbal relation expressed by the verb. These distinctions are summarised in Table 3.1.

The parallel sentences in (3.49), on their most salient interpretations, exhibit the pseudo-complement/adjunct distinction clearly: in (3.49a) the PP *with brochures* is interpreted as indicating what John put into the envelopes and as such introduces a relationship between the envelopes and the brochures, while in (3.49b) the PP *with Sarah* introduces someone who accompanied John in the *stuffing envelopes* event, introducing a relationship between Sarah and the full event, not the envelopes. The preposition *with* exhibits an ambiguity which is resolved by the way it predicates what it modifies (where an infelicitous predication would be ruled out via pragmatic reasoning).

(3.49) a. John stuffed the envelopes with brochures.

b. John stuffed the envelopes with Sarah.

Thus the semantics of a pseudo-complement preposition specifies a relation between an element *within* the semantics of the verb it modifies and the object of the preposition (internal predication), while the semantics of an adjunct specifies some operation on the full event conveyed by the sentence, minus the adjunct (external predication).

This definition differs from Gawron's definition of co-predicators in that I do not assume that a pseudo-complement will always introduce a relation which is not already a component of the lexical semantics of the main verb. It is, on the other hand, simply an internal predicator which is not strictly subcategorised by the main verb.

Consider the sentences in (3.50)-(3.52).

(3.50) a. John sang a song about his homeland.

b. John sang a song for Mary.

c. John sang a song in the park.

d. John sang a song at noon.

e. John sang a song { about his homeland / for Mary } { in the park. / at noon. }

f. John sang a song { in the park / at noon } { *about his homeland. / ?for Mary. }

(3.51) a. Sam kicked a ball to Bill.

b. Sam kicked a ball to Bill in the park.

c. *Sam kicked a ball in the park to Bill.

(3.52) a. *John ran a marathon about his homeland.

b. John ran a marathon for Mary.

c. John ran a marathon in the park.

d. John ran a marathon at noon.

e. John ran a marathon for Mary { in the park. / at noon. }

f. John ran a marathon { in the park / at noon } for Mary.

None of the PPs in the above sentences contains information which is entailed by the main verb's semantics. However, in (3.50) there is a clear difference between the PPs in the (a,b) sentences and the (c,d) sentences. The PP in (3.50a) expresses a property of the *song* which is sung by John, while the PPs in the (c,d) sentences provide information about the situation described by the main predicate of the sentence. Likewise, in (3.51) the PP *to Bill* specifies a particular goal relation between *Bill* and *the ball* rather than a relation between *Bill* and the full situation expressed by the main predicate of the sentence.

It could be argued that the PP in (3.50a) is a structural modifier of the NP rather than the VP, such that *a song about his homeland* forms a single constituent. This would explain the grammaticality of (3.53a).

(3.53) a. A song about his homeland was sung by John.

b. A song was sung by John about his homeland.

c. A song was sung by the choir about freedom.

d. John sang about his homeland.

It seems, however, that the analysis in which the PP modifies the VP constituent must also be available, as shown by the grammaticality of the sentences in (3.53b)-(3.53c). These sentences appear to be licensed semantically — there is an argument of *sing* which is embedded into the semantics of the verb (singing entails singing *something*, even if that something is an unnamed tune; that is, it involves producing sound which is normally called a song) and this argument is available as the element within the verbal semantic representation which can be picked out for the relation contributed by a pseudo-complement. Furthermore, the existence of sentences such as (3.53d), in which there is no explicit NP to which the PP could be attached, provides evidence that the PP can be viewed as specifying a relation involving an argument internal to the verb — an argument which is unexpressed in this case but still entailed by the verb and therefore a part of the verbal semantic argument structure.[5] The event of *singing* is not *about John's homeland*, but rather *what* John was singing.

Sentence (3.50b) is ambiguous between two interpretations — one in which the PP behaves as a pseudo-complement and a second in which it behaves as a true adjunct. On the pseudo-complement interpretation, the PP expresses that the song itself is for Mary's benefit, while on the adjunct interpretation it expresses that the entire activity of singing is for Mary's benefit. This ambiguity is more marked in the case of (3.54), where the pseudo-complement interpretation would be preferred if, for example, Mary were a teacher who will receive the essay, while the adjunct interpretation would be preferred if Mary were ill and asked John to do her homework of writing an essay for her.

(3.54) John wrote an essay for Mary.

This ambiguity provides an another example of the influence of pragmatic coherence on interpretation — both interpretations might be feasible independent of a context, but only one will make sense in a particular context.

The availability of both of these interpretations implies that both a pseudo-complement and an adjunct can appear in the same sentence. Not only is this evidenced by (3.50e), but more interestingly by (3.55), which can only be interpreted as indicating that the song was for Mary's benefit and that the entire activity was done for Bill's benefit. If this sentence is difficult to interpret, imagine a context, for example, in which Bill and Mary are unable to

[5] Note that this sentence differs from sentences which superficially resemble it, e.g. *John worried about his homeland*, in which no verb-internal argument to be modified by a pseudo-complement clearly exists. The difference is that an act of worrying entails a topic about which one worries, while an act of singing does not entail a song topic. That is, *John worried* $\models \exists x | $*John worried about x*, while *John sang* $\not\models \exists x |$ *John sang about x*. Thus the fact that the object of the event in the *worry* sentence is not clear does not discredit a pseudo-complement analysis of verbs like *sing* which do have a semantically entailed product. The PP appearing with *worry* would be licensed through compatibility between the semantic relation type of *worry* and the semantics of the *about* PP.

celebrate their anniversary together because they are living in different places, so Bill asks John to go to where Mary is and sing.

(3.55) John sang a song for Mary for Bill.

There is a syntactic ordering preference for the pseudo-complement PPs to precede the adjunctive PPs, as shown by the contrast between (3.50e) and (3.50f), and between (3.51b) and (3.51c). The interpretation of the *for Mary* version of the sentences in (3.50f) is questionable — it is unclear whether the pseudo-complement interpretation of the PP is available when preceded by another adjunct. It could be postulated on the basis of the contrast in (3.50f) and the sentences in (3.51) that the pseudo-complement interpretation of a PP is only available in immediately post-verbal-complement position, and that therefore *for Mary* in (3.50f) must be interpreted as specifying a relation involving the entire event expressed by the verb. This constraint can be captured in terms of obliqueness — pseudo-complements are semantically less oblique than adjuncts, and less oblique elements precede more oblique elements in English.

The distinction between pseudo-complements and adjuncts leads to an explanation for the ungrammaticality of (3.52a). The PP *about his homeland* can only behave as a pseudo-complement with respect to a verb phrase; it does not provide information which could apply to a full situation. There appear to be certain PPs which can only behave as pseudo-complements and other PPs which can only behave as adjuncts. Furthermore, the semantic contribution of the pseudo-complement must be compatible with the semantics of the modified verb, and there must be appropriate arguments in the relation expressed by the verb available for modification by the pseudo-complement. This will be discussed further in Section 3.3.3. Semantic incompatibilities will rule out (3.52a), because a *move-rel*[6] verb like *run* cannot be extended to a relation involving a topic, and the preposition *about* introduces a topic. This means that the PP in (3.52b) can only be interpreted as a true adjunct, that is that the whole activity of *John running a marathon* was done for Mary. The marathon itself cannot be interpreted as benefitting Mary. This also explains the contrast in acceptability between (3.50f) and (3.52f). There is an obliqueness difference between the PPs in the former on the pseudo-complement interpretation, which prevents the PPs from freely alternating in syntactic order. In contrast, there is no obliqueness difference between the adjuncts in the latter, enabling the PPs to appear in any order.

3.3.2 Pseudo-complement semantics

The *for*-dative pseudo-complement

In the analysis in Verspoor (1994), *for*-dative prepositional phrases are treated as arguments of the verbs with which they appear. Their semantic contribution is therefore directly integrated into the semantics of the verb at the lexical level. The semantic analysis there is based on the discussion of Pinker (1989). The core semantic content of each of the *for*-dative verbs for

[6]See Section 4.7.2 for details of this relation, which is a subtype of *act-rel*.

sentences of the form *NP_1 gimbles NP_2 for NP_3*[7] can be paraphrased as, "NP_1 acts on NP_2 in order for NP_3 to have NP_2". The contribution of the *for*-PP can be identified as the "in order for NP_3 to have NP_2" portion of the paraphrase. In the notation of Verspoor (1994), this is represented as a HAVE event related by a *for_to* subordinating relation to the main event expressed by the semantics of *gimble*.

Jackendoff (1990) argues that this HAVE event doesn't properly capture the semantics contributed by the *for*-PP. He claims that the event is rather forced when applied to certain verbs. For example, when *John sings a song for Mary*, in what sense does Mary *have* the song? Jackendoff therefore suggests that the contribution of the *for*-PP is better described as indicating that the object of the preposition (NP_3) is intended to benefit from the action of the subject (NP_1). The event embedded by the *for_to* relation would more appropriately be as in (3.56b) rather than Pinker's proposal of (3.56a). This event represents "NP_1 affects NP_3 positively", or in other words, "NP_3 is intended to benefit from the actions of NP_1".

(3.56) a. $(\text{HAVE}\ (\text{THING}_3,\ \text{THING}_2))$ [Pinker's proposal]

b. $(\text{AFF}^+\ (\text{THING}_1,\ \text{THING}_3))$ [Jackendoff's proposal][8]

However, Jackendoff's proposal also does not seem to accurately capture the interpretation associated with the PP in this form. The benefit represented in his form is indirect — since NP_1 does not act upon NP_3 directly, what actually is intended to benefit NP_3 remains unclear. In fact, it seems that what is intended to benefit NP_3 directly is NP_2, the object upon which NP_1 acts in order to benefit NP_3. Thus it seems more accurate to represent the semantics of the *for*-dative as a relationship between the referent of NP_2 and the referent of NP_3. Incorporating a *benefit_rel* relation from the representation developed in Chapter 2, I propose that the semantics associated with the *for*-dative preposition should be as indicated in (3.57).

(3.57)
$$\left[\begin{array}{l}\text{FOR\_TO} \\ \quad \textit{benefit\_rel}\left[\begin{array}{l}\text{ACT}\ \textit{nom-ind}_2 \\ \text{UND}\ \textit{nom-ind}_3\end{array}\right]\end{array}\right]$$

It is clear that this representation involves a pseudo-complement interpretation since the ACT argument of the *benefit_rel* relation is also a semantic argument of the verb. In addition, the treatment of the pseudo-complement modification will include incorporating the subordinated *for_to* relation directly into the internal semantics expressed by the verb rather than the external semantics.

In contrast, the *for*-adjunct preposition (as in *John ran a marathon for Mary*) adds the semantic content in (3.58) to the representation of the full situation (to the external semantics). Its definition specifies that the influencing argument of the *benefit_rel* corresponds to the entire situation expressed in the sentence. The object of the preposition is therefore affected positively by the event (referred to via the *sit-ind*, situation index), rather than by a particular semantic element within the event representation.

[7] *Gimbles* is a marker for verbs which can appear in this construction.

[8] AFF^+ is the notation for "affects positively".

(3.58)
$$\begin{bmatrix} \text{FOR\_TO} \\ \textit{benefit\_rel} \begin{bmatrix} \text{ACT } \textit{sit\_ind} \\ \text{UND } \textit{nom-ind}_3 \end{bmatrix} \end{bmatrix}$$

Note that although Jackendoff (1990:195) suggests that the *for*-PP can be given precisely such an event interpretation, he provides no formal mechanism for doing so, or for distinguishing between the two possible interpretations of the *for*-PP. In Jackendoff's approach, the two different readings of the *for*-dative form must fall out of a single representation (that in (3.56b)), which fails to adequately reflect either reading and does not account for the identity of the interpretation of the double object form with only *one* of these readings (that in (3.57)).

The distinction that pseudo-complements pick out a semantic argument from within the verb semantics while adjuncts incorporate the event expressed in the sentence as an argument in the relation they express is thus formalised in the semantics of the two forms associated with *for*. These will, upon formalisation, actually be captured by a single lexical entry for the preposition which can interact with the modified verb phrase in multiple ways. This interaction and the difference in where the semantic contribution is integrated with respect to the verb semantics — that pseudo-complements contribute to the internal semantics of the verb (the verbal relation) while adjuncts contribute to the full situation expressed by a sentence (the external semantics) — will be discussed in more detail in Section 3.5 and handled by the lexical rules which will be introduced there.

We will see that the double object form of the dative sentences is only possible on the pseudo-complement interpretation. This will account for the inability of some verbs to exhibit dative alternation, as shown in (3.59). In (3.59a) only the adjunct interpretation of the PP is possible and hence the double object form is infelicitous. This syntactic structure will be licensed by a lexical rule which will constrain its interpretation to internal predication.

(3.59) a. John extinguished the light for Mary.

b. *John extinguished Mary the light.

The *to*-dative pseudo-complement

In Verspoor (1994), the *to*-dative prepositional phrase is also explicitly specified as an argument of each verb with which it appears. The core semantic content for each of the *to*-dative verbs for sentences of the form *NP*$_1$ *gimbles NP*$_2$ *to NP*$_3$ can be paraphrased as, "NP$_1$ acts on NP$_2$, causing NP$_2$ to go to NP$_3$". The contribution of the *to*-PP can thus be identified as the "causing NP$_2$ to go to NP$_3$" portion of the paraphrase. In the notation of Verspoor (1994), this is represented as a GO event related by an *effect* subordinating relation to the main event expressed by the semantics of *gimble*, as shown in (3.60). The meaning of this preposition is such that no adjunctive interpretation would make sense — it wouldn't make sense for an event to GO somewhere, and thus no adjunct interpretation exists for it.

Dative alternation (*to*-datives)	Benefactive alternation (*for*-datives)
give verbs, verbs of *future having*, *bring*, *take*, *send* verbs, *carry* verbs, *transfer of message* verbs, verbs of *instrument of communication*	*build* verbs, *create* verbs, *prepare* verbs, verbs of *performance*
slide verbs, *drive* verbs, *throw* verbs	*get* verbs

Table 3.2: Alternating classes of *to*- and *for*-datives (from Levin 1993)

(3.60)
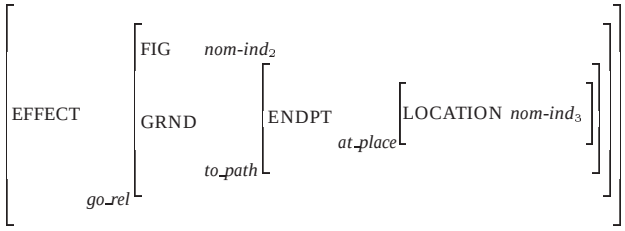

3.3.3 Pseudo-complement lexical semantic compatibility

It has been noted many times in the literature that there seem to exist semantic constraints on the verbs which can appear with dative PPs. Jackendoff (1990:196), for example, proposes that only verbs of *creation* or *preparation* can appear in the double object form of *for*-beneficiary sentences. This follows from the fact that "singing" involves the physical formation of a song, "baking" involves bringing into existence baked goods, etc. and that only these type of verbs appear in this form. This notion of semantic constraints on the double object form of *for*-datives also extends to *to*-datives: only certain verbs can appear in the double object form with a *to*-dative interpretation. Levin (1993) identifies the classes of alternating verbs as summarised in Table 3.2.

What is striking upon observation of that table is that the alternating verbs group into clear sets of semantically related verbs. In fact all of the verbs listed in the table, on their (di)transitive uses, express relations which are subtypes of *cause-eff-rel*. More specifically, the dative-alternation verbs all express the transfer of something somewhere or to someone (e.g. if you *slide* or *throw* or *send* a book, you are causing it to go somewhere), while the benefactive-alternation verbs divide into two sets: Jackendoff's creation and preparation verbs (which correspond to *cause-create-rel* and *cause-ch-st-rel*, respectively); and *get*-type verbs which are transfer verbs like the dative-alternation verbs, except that the location to which the object is transferred is encoded in the verb semantics to be the location of the entity referred to in the subject NP (e.g. *John got a present (for Sally)* where John first acts to transfer the present to his possession).

The conclusion that can be drawn from this is that only verbs expressing specific semantic relations are compatible with PPs expressing particular meanings. This point has also been argued by Wechsler (1995), who suggests that complement PPs express a meaning which must be compatible with the type of the verb. Here I extend this notion to pseudo-complements as well. The constraints on which verbs can be extended by particular pseudo-complements (or

appear with particular complement PPs) can therefore be captured in the semantic relation hierarchy. A possible structure for that hierarchy, as far as *to*- and *for*-dative verbs are concerned, appears in Figure 3.1.

There are several things to note about this hierarchy. First is the use of the *thematic* attribute THEM-ARG (thematic argument). The value of this attribute is constrained at the level of *cause-eff-rel* to unify with the entity which is the UNDergoer of the causation. This constraint is inherited by each subtype at lower levels in the hierarchy, and ensures that the thematic elements of the verbal relation are all relativised to this same entity. The fact that this constraint is relative to a particular relation type is justified by arguments made by Nam (1995) that when certain prepositional phrases modify an n-ary function, they introduce properties restricting the n^{th} argument of that function (the thematic argument in my terms). Thus different verbal relations, with different argument structures, will have different thematic arguments and this can be efficiently captured in the hierarchy.

The second point to note is the use of defaults. Here I am attempting to prevent arbitrary extensions of the THEMATIC attribute associated with a particular semantic relation through grounding of the thematic type, while maintaining generality over the relationships between the verbal relations. Thus it is clear that the relation *transfer-to-ben-rel* (intended to accommodate sentences such as *John gave the presents to Mary for the kids*) is an extension of the relation *transfer-to-rel* and therefore should be (and is) a subtype of this latter relation. These relations both inherit a THEMATIC value of *effect* from *cause-eff-rel*, but in the case of *transfer-to-rel* the THEMATIC value must be restricted to a particular kind of EFFECT (*go-rel*) which should not be extended with any other thematic information. A THEMATIC value of *effect* is, however, compatible with other thematic information and could in theory be extended by pseudo-complements contributing other thematic information. In order to rule this out, the relation *transfer-to-rel* is given a default (node-specific) THEMATIC value of *effect_only*, with which any pseudo-complement PPs must be compatible in order to modify a verb of this type. The subtype *transfer-to-ben-rel*, by similar reasoning, is given a default THEMATIC value of *effect_for-to_only*, while its non-default value (*effect_for-to*) is an extension of the non-default value *effect* inherited from *transfer-to-rel*. If the (indefeasible) THEMATIC value of *transfer-to-rel* had been specified as *effect_only*, it would have been impossible to capture the relationship between it and *transfer-to-ben-rel* due to the incompatibility between *effect_only* and *effect_for-to_only*. The use of defaults enables this intuitive relationship to be captured without losing the constraints which exist for individual verbal relations.

When the semantics of a PP is integrated with the verbal relation it modifies, it must unify with the default THEMATIC value of that relation, or a subtype of that relation with which it is compatible. The effect of this is to ground the PP via unification (i.e. if the PP expresses an *effect* theme and it combines with a *transfer-to-rel* verb, it will be ground to type *effect_only*), and to only allow the PP to combine with verbs of compatible type (so a *for-to* PP might combine with a *transfer-to-rel* verb but the result would be a relation of type *transfer-to-ben-rel*).

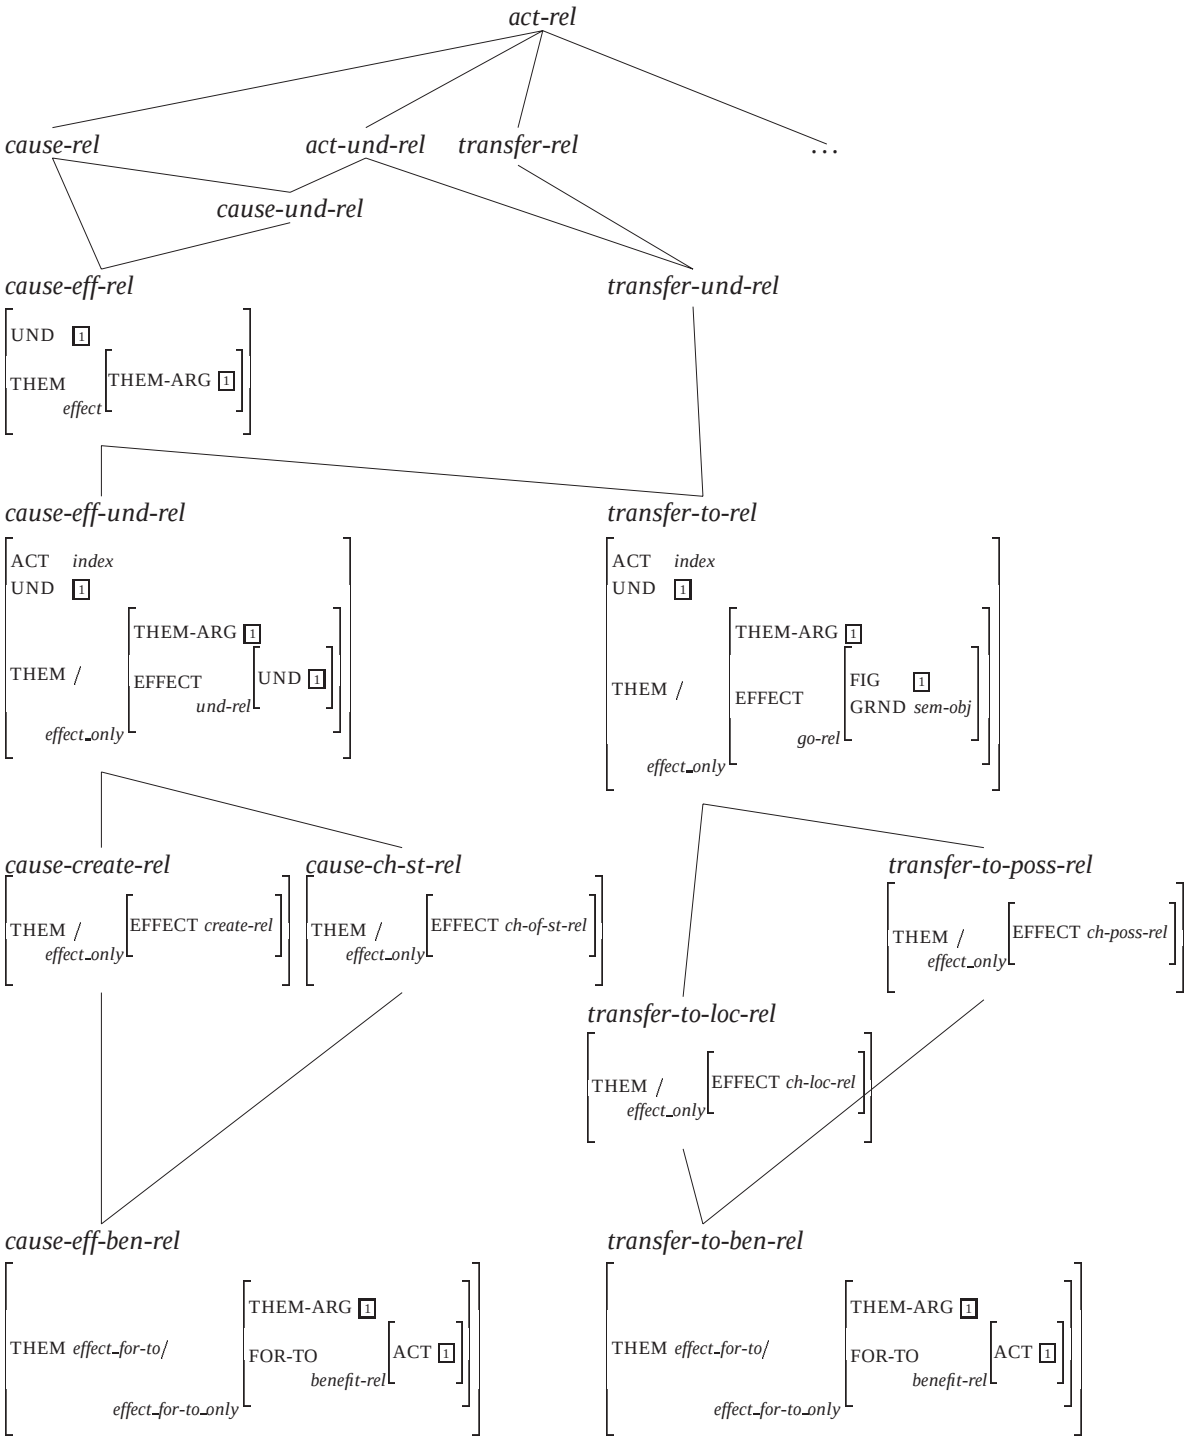

Figure 3.1: Extended segment of the *rel* sort hierarchy (THEM = THEMATIC)

3.3.4 Conclusions

Pseudo-complements are elements very close in nature to true verbal complements. They specify a particular relation between a semantic argument of the verb and the object of the preposition. Their semantics can, however, be treated by the same mechanisms as true adjuncts. Pseudo-complements and adjuncts share the property of specifying a consistent, contentful, and identifiable relation which can be applied across modified heads.

Specification of the type of modification which an individual PP may provide with respect to a head must occur lexically. Three types of PPs will be allowed for in the treatment of PPs to be introduced in Section 3.5 (in particular, in the sort hierarchy) — PPs which can only behave as pseudo-complements, PPs which can only behave as adjuncts, and PPs which are ambiguous between the two.

Prepositions which are strict complements will not be represented as a distinct semantic type in the sort hierarchy. Complement PPs which mark a verbal argument according to the normal interpretation of the head preposition can be treated semantically as pseudo-complements: their semantics, marking a verbal argument, is simply integrated into the verbal relation directly in the lexical entry through unification, rather than through an external combination process as will be defined for standard pseudo-complementation. Thus there will be no semantic (type) distinction between the *to*-PP complements of *give* and *kick*, for example, although they will be licensed in different ways (for *give* in the lexical entry and for *kick* through the mechanisms of pseudo-complementation).

The type of a particular preposition must then be specified in the lexical entry of the preposition. This type will be used as a criteria for determining how semantic integration between the semantics of the PP containing the preposition and the semantics of the modified verb is to occur.[9]

3.4 Adjuncts

3.4.1 Characteristics of adjuncts to be accounted for

There are certain characteristics which pseudo-complements and adjuncts share which must be taken into consideration in any treatment of the semantic contribution of these elements.

Adjuncts have an unchanging semantic contribution

As was discussed in Section 3.2.2, adjuncts have an identifiable, consistent semantic contribution across heads. The implication of this for any treatment of adjuncts is that there should be a single lexical entry which specifies the meaning of the adjunct with respect to a particular type of head. That is, the content of the adjunct combines in a certain general way with the content of the element it modifies (a verb or noun phrase, for example) and this must be specified only once. Since this combination does not change with every type of head, a single

[9] i.e. which lexical rule will apply — see Section 3.5; 3.5.3 in particular.

specification is much more efficient than incorporation into the lexical entry, into the subcat list, of each verbal head with which an adjunct can appear.

In particular, a single lexical entry can only be realised if adjuncts select the types of heads they modify. Were individual heads to idiosyncratically specify the adjuncts with which they can appear, the semantics of the adjunct could conceivably be incorporated with the content of the head in a different way for each head, and in effect the adjunct need not have any independent meaning. Furthermore, this approach would require that the set of adjuncts which could appear with a particular head be specified in advance, at the level of the lexicon, for every individual element in the lexicon which could potentially be modified by an adjunct. This is clearly not a desirable consequence.

An additional semantic argument for the selection of a head by an adjunct is observed by Kasper (1993): "The semantic contribution of a modifier generally must incorporate the semantic contribution of the element that it modifies, whereas the semantic content of the modified element (the syntactic head) does not depend crucially on any of its potential modifiers".

Restrictive, Operator, and Thematic adjuncts

Adjuncts have traditionally been analysed as being of one of two types:[10] *restrictive* adjuncts and *operator* adjuncts. Restrictive adjuncts are adjuncts which "restrict" the value of a particular index representing an object, event, or situation, such as the index for location or time of the event. These adjuncts specify properties to be associated with the indices. Operator adjuncts are adjuncts which take the content of what they modify as an argument in a semantic operation, predicating something of that content. Examples of this type of adjunct include negatives, frequentatives, and duratives.

The distinction between these adjunct types provides an explanation of the differences in the semantics of (3.61a) and (3.61b), containing operator adjuncts, as compared to the lack of semantic difference between (3.62a) and (3.62b), containing restrictive adjuncts. In (3.61a), the twenty minute duration is a property of the event whose frequency is described, while in (3.61b), the 'twice-dailiness' is a property of the event whose duration is described (Pollard and Sag 1987). In both (3.62a) and (3.62b), *in the park* specifies the location of the jogging, and *yesterday* specifies the time of the jogging, regardless of their surface order. They each specify (or *restrict*) properties of the main event described in the sentence, rather than predicating something of an event they receive as an argument, as in (3.61). (Sentences from P&S 1987, (252))

(3.61) a. John jogged for twenty minutes twice a day.

b. John jogged twice a day for twenty years.

(3.62) a. John jogged in the park yesterday.

b. John jogged yesterday in the park.

[10]This discussion of adjunct types is mainly derived from the discussion in Kasper 1993.

In essence, restrictive adjuncts seem to add new information about an index for which the event was previously underspecified (e.g. location) while operator adjuncts take the event as an argument, thereby building up a more complexly structured semantic representation for the sentence.

There is a group of adjuncts which semantically do not clearly fit either of these two types. These adjuncts, like all other adjuncts, add information to the basic event expressed by the verb plus its semantic arguments. However, they do not simply restrict an index specifying something about the situation in which the event occurs or predicate something of that situation. Rather, they relate information via one of a predetermined, limited, set of subordinating relations. They can be viewed as adding a theme to the verb semantics, and thus will be called *thematic* adjuncts. Examples of thematic adjuncts can be found in (3.63)-(3.65). In (3.63), the *because_of*-PP adds information which explains the cause of the situation expressed in the remainder of the sentence. In (3.64), the *with*-PP expresses the means by which the situation expressed in the sentence minus the PP occurred. In (3.65), the *to*-PP expresses a motivation for the situation in the remainder of the sentence.

(3.63) Peter reads well because of the tutoring. [Kasper 1993, (10a)]

(3.64) Peter opened the door with the key.

(3.65) Peter read the book to learn about World War II.

Surface order vs. Semantic precedence

The relative surface order of multiple restrictive adjuncts generally has no effect on their interpretation. The relative semantic scope of multiple operator adjuncts, on the other hand, sometimes does and sometimes does not depend on their relative surface order.[11] Since the order of interpretation of operator adjuncts can affect the overall interpretation of a sentence, it is important to account for interpretation orders which vary from straight surface order, in addition to accounting for interpretation orders which are dependent on surface order.

We saw in (3.62) that the surface order of restrictive adjuncts in some cases has no effect on their interpretation. The sentences in (3.61) showed, on the other hand, that relative surface order can influence the interpretation of the sentence. Contrasting (3.61) with (3.66) indicates that the content of the multiple adverbials can also influence their relative interpretation. Sentences (3.61a) and (3.66a) have the same semantics despite their differences in surface order. Sentence (3.66b) is ungrammatical because the combination dictated by the surface order is temporally impossible — it is not possible to repeat an event which itself lasts twenty years twice within one day.

(3.66) a. John jogged twice a day for twenty minutes.

b. *John jogged for twenty years twice a day.

[11] See Kasper (1993) for a good overview of the cases of interaction among multiple adjuncts.

Clearly there are very complex constraints governing both the semantic composition and the relative surface order of multiple adverbials. A treatment of adjuncts must therefore provide a mechanism for the application of these constraints.

Redundancy constraints

It is important in any treatment of adjuncts to prevent multiple adjuncts from providing information which fills the same role. Sentences such as those presented in (3.67)-(3.69) must be prohibited. For thematic adjuncts, the constraint seems to be that only one thematic adjunct corresponding to a particular subordinating relation is allowed, while for restrictive adjuncts the constraint is that multiple restrictive adjuncts relating to the same index must have values which are related via containment (as discussed in Section 3.2.1).

(3.67) a. Peter reads well because of the tutorials and because of the homework exercises.

b. *Peter reads well because of the tutorials because of the homework exercises.

(3.68) a. Heather opened the rusty lock by oiling it and by applying force. [cf. (3.25)]

b. *Heather opened the rusty lock by oiling it by applying force.

(3.69) a. Sam kicked a ball at 10 o'clock and at 8 o'clock. [cf. (3.22)]

b. *Sam kicked a ball at 10 o'clock at 8 o'clock.

If information is explicitly coordinated through a conjunction or disjunction, it is possible for multiple PPs of the same type to appear in a sentence. Since coordination in effect builds a complex element of the same category as its components, this data can be interpreted as evidence that exactly one PP making a particular type of semantic contribution can appear in a sentence. Although this does not hold for PPs which supply information related by containment, it does hold for most PP types.

Interspersal of adjuncts with complements (Mittelfeld phenomena)

Kasper (1993) presents a detailed analysis of word-order phenomena in the German Mittelfeld: "the part of the German clause between the finite verb (or the beginning of verb final clauses) and the clause final verb or verb cluster, if any." In particular, he observes that the linear order of verb complements and adjuncts within the Mittelfeld is relatively free. Any treatment of adjuncts must therefore be able to account for this interspersal.

3.4.2 The standard HPSG approach

The treatment of adjuncts in Pollard and Sag (1994) centres on the selection of a head by an adjunct. The adjunct specifies the type of head which it modifies via the MOD feature of its SYNSEM:LOC:CATEGORY:HEAD field. Semantic integration is specified in the lexical entry of the adjunct, via structure sharing between a substructure of the head's content and the content of the adjunct. Adjuncts differ from complements in that they have a non-null MOD value,

that they are not subcategorised-for by the element with which they combine, and that they are joined with that element via a different mechanism.

Immediate dominance (ID) schemata govern the permissible configurations of immediate consituency (akin to phrase-structure trees) in HPSG. One such schema creates a *head-adjunct-structure*, combining a head and an adjunct into one structure, and ensuring that the head of the constituent is an element allowed by the MOD feature of the adjunct. The content of the mother in a head-adjunct-structure is required to be token-identical with the content of the adjunct via the Semantics Principle. This guarantees that the appropriately integrated semantics is associated with the phrase as a whole.

The specification of the ID schemata in standard HPSG does not allow for Mittelfeld phenomena. The schemata handling complements require that all complements other than the subject must be combined at once into a phrase. The adjunct attachment schema allows the adjunct to appear immediately before or after the head it selects,[12] or before or after the phrase containing the head and all of its complements, but does not license the appearance of the adjunct within a group of complements.

Other characteristics of adjuncts are handled in this approach, however. A single lexical entry specifies the integration of an adjunct's semantics with the element it modifies. The difference between restrictive and operator adjuncts can be accommodated by variances in the definitions in the CONTEXT field of the adjunct's SYNSEM feature. Redundancy constraints are not explicitly accommodated, but could conceivably be implemented within the MOD feature of an adjunct in terms of restrictions on the modified head. How this implementation would be accomplished is not, however, entirely clear.

Surface order and semantic precedence issues remain a stumbling block for the standard HPSG approach. Since linear precedence constraints (constraints defined in terms of obliqueness which control the surface order of elements relative to one another) apply at the level of individual phrases built by the ID schemata, and only one adjunct at a time can be attached to a head via an ID schema, the order of modification is constrained to surface order.

3.4.3 A "Semantic Obliqueness" hierarchy

Kasper (1993) proposes a treatment of adjuncts aimed specifically at handling Mittelfeld phenomena. He adopts the standard HPSG representation of adjuncts, in that the adjuncts specify the heads they modify via the MOD field and semantic integration occurs through coindexing between parts of the CONTEXTs of the selected head and of the adjuncts themselves.

Kasper makes several relevant semantic assumptions. First, states of affairs (*soa*s) come in two basic kinds: those that are spatio-temporally located (*located-qfsoa*) and those that are not

[12] Note that this in fact does not constrain adjunctive placement enough, improperly allowing lexical heads rather than phrasal heads to be modified by an adjunct. This would therefore not rule out phrases such as **The king in the bath of France* or sentences like **John kicked in the park the ball*. These sentences must be ruled out via the lexical entries of the prepositions which select for nominal / verbal heads: a head with an empty SUBCAT list must be explicitly selected for in the MOD field of the preposition.

(*unlocated-qfsoa*). Second, the NUCLEUS of a state of affairs is split into a primary quantifier-free *soa* (*qfsoa*) and a set of restrictions. Multiple semantic restrictions with respect to the same state of affairs can thereby be specified in the restrictions set. This set plays a role analogous to the RESTRICTIONS feature on referential indices in the semantic content of nominal objects in standard HPSG. Thus adverbials and adnominals can be treated in a parallel manner.

The *head-complement* structure of standard HPSG is extended by Kasper to include an ADJUNCT-DAUGHTERS attribute. This is a list of adjunct signs ordered in terms of a "semantic obliqueness" hierarchy, i.e. from widest to narrowest semantic scope.

To handle the syntax and semantics of adjunction, Kasper splits the MOD field of the adjuncts into two parts: a SYN attribute which indicates the syntactic category of the head with which the adjunct must combine and a SEM attribute specifying the semantic value to which the adjunct is applied. Kasper then specifies an *Adjunct Syntax Principle* requiring the MOD:SYN attribute of all signs on the ADJUNCT-DAUGHTERS list of a *head-complement* structure to be token-identical with the CAT value of the head daughter. Furthermore, his *Adjunct Semantics Principle* forces semantic composition to occur in terms of "semantic obliqueness" order: the element with narrowest scope is applied to the head's semantics, then the element with second-narrowest scope is applied to the resulting semantics, and so on down the list.

The relative surface order of complements and adjuncts would then have to be constrained by separate principles of constituent order which constrain the possible combinations of elements from the ADJUNCT-DAUGHTERS and COMP-DAUGHTERS attributes.

An issue about which Kasper remains vague is how elements are put onto the ADJUNCT-DAUGHTERS list. Apparently the HEAD-COMPLEMENT and HEAD-SUBJECT-COMPLEMENT schemata must be redefined to allow for arbitrary insertion of adjuncts into the ADJUNCT-DAUGHTERS list of the *head-complement* structure. What drives this insertion, however, remains unclear. Some mechanism must exist to identify all adjunctive sentence constituents, evaluate their relative "semantic obliqueness", and insert them into the list.

Since Kasper opts for a semantic obliqueness order on the ADJUNCT-DAUGHTERS list rather than an order reflecting surface order, semantic differences which depend on syntactic order may not be appropriately handled. The adjunct insertion mechanism discussed above must be defined in such a way as to take order effects into account. Furthermore, the mechanism must also provide for adjuncts which are not hierarchically related semantically (as in the case of restrictive adjuncts) so as to avoid analysis redundancies deriving from differences in order on the list.

It is observed by van Noord and Bouma (1994) that Kasper's approach cannot account for interpretation ambiguities in Germanic verb cluster constructions. These ambiguities occur because adjuncts are able to modify any verb within a verb cluster. Thus in the Dutch sentences in (3.70) (From van Noord and Bouma (1994)) the adjuncts (*today*, *with the telescope*) can either be interpreted as having narrow scope and modifying the event introduced by the main verb or as having wide scope and modifying the event introduced by the auxiliary.

(3.70) a. dat Arie vandaag Bob wil slaan
that Arie today Bob wants to hit
that Arie wants to hit Bob today

b. dat Arie Bob de vrouwen met een verrekijker zag bekijken
that Arie Bob the women with a telescope saw look at
that Arie saw Bob looking at the women with a telescope

Under the standard treatment of such clusters within a flat structure, the first auxiliary verb is treated as the head of the structure. Kasper's solution thus dictates that any adjuncts must modify this head rather than an embedded verb, not allowing for any narrow-scope readings.

3.4.4 The lexical rule approach

To solve the problem of accounting for the ambiguity of adjunctive modification in Germanic verb clusters, van Noord and Bouma (1994) propose a solution treating adjunction via a lexical rule. The lexical rule specifies the addition of a single adjunct to the SUBCAT list of a verb. The ambiguity in the verb cluster modification then derives from the possibility of the lexical rule applying to any verb in the cluster. In the narrow scope case the lexical rule applies to the embedded verb, placing the adjuncts on its subcat list. The subcat requirement will then be inherited by the head verb, but the semantics of the adjunct will be incorporated into the semantics of the embedded verb. In the wide scope case the adjunct is simply on the list of the head verb and its semantics applies to the head.

Van Noord and Bouma propose to treat lexical rules as constraints on lexical categories, and to use delayed evaluation techniques. These constraints are implemented as rules which must be satisfied by the lexical entry of a word in a particular category. The constraints are evaluated with respect to the base (or "stem") form of a word in the lexicon. The true lexical entry for the word used in an attempted parse results from evaluation of constraints with respect to the base form.

The delayed evaluation techniques prevent constraints from being evaluated until enough information is available to do so. This means that constraints may actually only be partially evaluated at any step in the application of multiple constraints to a single lexical entry. The benefit of these techniques is that parsing mechanisms can interact with lexical information, allowing constraints from both structural and lexical levels to apply simultaneously as input is processed. This essentially means that sentence-level semantics can be encapsulated in a lexical rule. This accommodates a notion of a construction (e.g. Goldberg 1995) in which a particular syntactic form is associated with a specific interpretation, and provides an elegant way of defining precisely how different parts of a sentence combine to produce an interpretation. We will see additional need for such constructions in Chapter 4, and there is other work in HPSG (Sag in press) which argues for the specification of non-compositional phrasal-level semantics.

The van Noord and Bouma approach accommodates most of the characteristics of adjuncts well. A single lexical entry is necessary for each adjunct, and they allow for both restrictive and operator adjuncts by requiring the appropriate semantic combinations to be specified in the MOD field of the adjunct, following Kasper's (1993) approach. Mittelfeld phenomena are handled by allowing for the insertion of the adjuncts at any point in the verbal subcat list.

It is noted by van Noord and Bouma that their approach is flexible enough to accom-

modate various approaches to the ordering of adverbials on the subcat list. Although the lexical constraint controlling the addition of adjuncts as defined in their paper assumes that the adjuncts are inserted into the subcat list in order of semantic obliqueness (adopting Kasper's idea of semantic combination from narrow to wide scope), there is nothing in their methodology which restricts the definition of the constraint. It is difficult, as mentioned above, to see precisely how syntactic ordering effects could be accommodated in an approach that relies entirely on semantic obliqueness. Changes in the ordering on the subcat list, however, would require radical changes in van Noord and Bouma's definition of the constraint adding adjuncts. In particular, if the ordering on the subcat list were changed to reflect surface order of the adverbials, their recursive approach to semantic composition would no longer suffice. Other mechanisms, analogous to the linear precedence (LP) constraints which are required to handle word order restrictions in their existing approach, would be necessary to control semantic composition. These mechanisms could only be applied to a fully expanded subcat list and thus would prevent semantic content from being truly recursively computed.

Redundancy constraints are a problem in the van Noord and Bouma (1994) approach, as in all other approaches discussed here. They could conceivably be defined in the requirements in the MOD field specifications, but again it is not clear how to do this in a straightforward manner.

3.4.5 Conclusions

None of the existing approaches to the treatment of adjuncts provides a satisfactory framework for explaining surface order and semantic precedence effects. The standard HPSG approach makes no attempt to accommodate these effects whatsoever; the Kasper (1993) and van Noord and Bouma (1994) approaches both rely on unspecified principles for determining semantic precedence, and principles of constituent order to control the surface order of adjuncts. None of the approaches satisfactorily allows for interactions between these various principles. In the section that follows, I will attempt to develop a more satisfactory framework.

3.5 Semantic Integration

The van Noord and Bouma (1994) treatment of adjuncts adopts the positive aspects of Kasper's (1993) treatment, integrating them into a framework which solves several problems with Kasper's original treatment. They accomplish this via a delayed-evaluation lexical rule approach to the incorporation of adjuncts. I will adopt this general approach, but will refine the semantic representation in terms of that introduced in Chapter 2 to show how it can be used to handle the phenomena of redundancy restrictions, adjunct combination restrictions, and pseudo-complementation. Additionally, a more explicit methodology for handling word order and semantic precedence constraints will be introduced.

3.5.1 Adjunct semantics

The form of the lexical entries for adjuncts used in this approach relies on the semantic representation defined in Chapter 2. The MOD field of an adjunct is used as the main vehicle for identification of the type of object which the adjunct can modify. Specification of the type of the adjunct is, however, also crucial. The definition of lexical rules for incorporating different kinds of adjuncts depends upon this specification.

Kasper's (1993) splitting of the MOD field into syntactic and semantic parts is unnecessary in a lexical rule approach. Since the adjuncts are added directly to the SUBCAT list of the element they modify, the lexical rules account for the appropriate structure sharing between the SYNSEM specified in every adjunct's MOD field and the modified "head".[13] In fact, the lexical rule approach to adjunct modification even eliminates the need for HPSG's *Head-Adjunct Schema* (schema 5) and the complex definition of a semantic head needed for the Semantics Principle (Pollard and Sag 1994, p. 322) — adjunctive elements are essentially given the status of subcategorised elements and the differences in how the semantics of the different types of adjuncts interacts with the semantics of the modified phrase are handled directly in the rules. The verb therefore remains the semantic head of the sentence, and all phrases with complements and (possibly) adjuncts are licensed by the *Head-Complement Schema*.

In Section 3.3.4 it was proposed to divide prepositions into three types, reflecting their behaviour as strictly a pseudo-complement, strictly a true adjunct, or a preposition which can behave as both. In fact, more types are necessary, to capture the difference between restrictive, operator, and thematic adjuncts. The relevant piece of the type hierarchy appeared in Chapter 2 in Figure 2.12 (page 53). I will provide representative lexical entries for each of these types.

The pure pseudo-complement *to* is shown in (3.71). This type of preposition adds an EFFECT attribute to the internal semantics associated with a situation (see Section 3.3.2), if compatible with the relation expressed by the modified verb (see Section 3.3.3). The lexical entry need only specify the value of this attribute. Structure sharing between the semantics of the prepositional phrase and the INTERNAL value of the situation is specified in the lexical rule bringing out the pseudo-complement modification. This is handled in this way because this aspect of modification remains constant across prepositions of this type. Constraints on the interaction of the preposition with a particular verb will be governed by the type hierarchy, as outlined in Section 3.3.3.

[13] Note that it is necessary to interpret the MOD field as specifying a particular type of element with which an adjunct can combine, rather than necessitating that the adjunct modify a phrasal head. This is because the head of a phrase may not be the element in the phrase which the adjuncts actually modifies, as was discussed in Section 3.4.3.

(3.71)

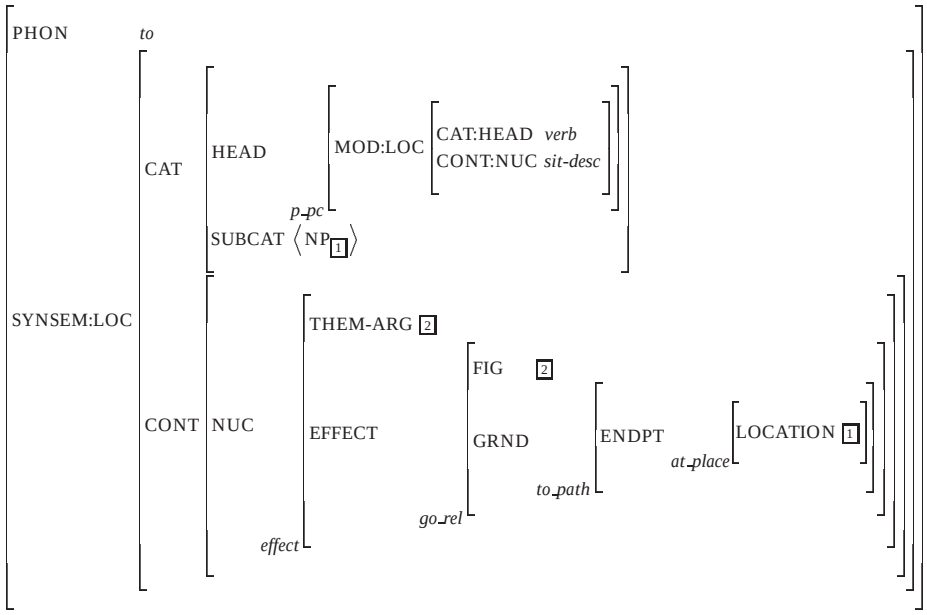

The lexical entry for *by*, a thematic adjunct, is shown in (3.72). The sense of *by* expressed here is that in *John broke the lock by hitting it with a hammer*. This type of preposition adds a thematic element to the external semantics associated with a situation. As above, the lexical entry need only specify this thematic element and the basic structure of the modified entity, as the appropriate structure sharing between the semantics of the modified verb and the semantics of the preposition is accomplished in the lexical rule. In this case, the thematic element expresses the particular means through which the situation is accomplished. It also expresses an additional constraint that the actor of the embedded VP (the complement of the preposition) be token-identical with the actor in the main situation.

(3.72)

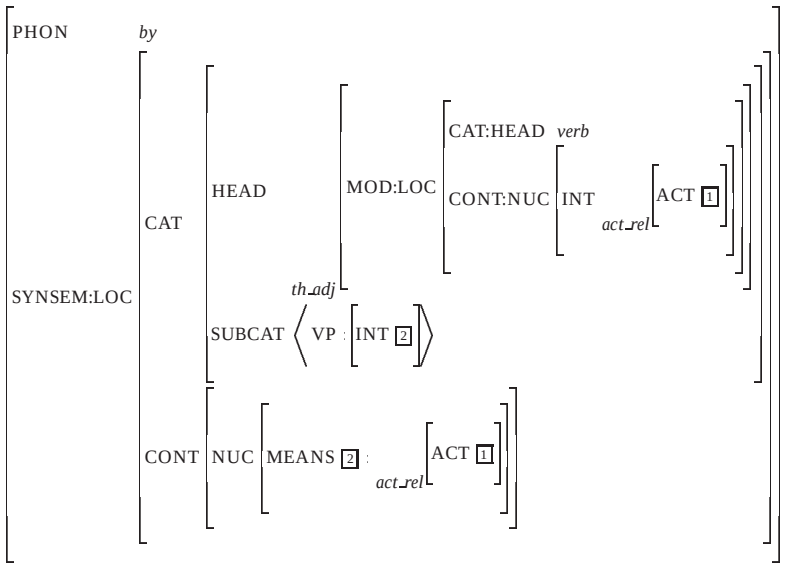

The lexical entry for the beneficiary *for,* which can behave as both a pseudo-complement and an adjunct, is shown in (3.73).

(3.73)

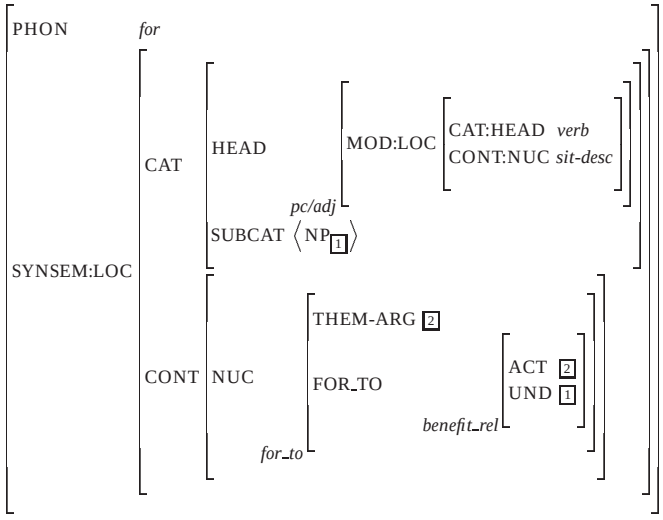

This type of preposition adds a thematic element to either the internal or the external semantics associated with a situation, depending on how it is used in a particular utterance, i.e. whether it is used as a pseudo-complement or adjunct. Again, the lexical entry need only specify the thematic element and the basic structure of the modified entity. Either the lexical rule for pseudo-complements or the lexical rule for thematic adjuncts will be applied to incorporate this adjunct into a sentence. Which rule is used determines where the thematic element is attached, and constraints on the THEM-ARG attribute will result in the appropriate entity playing a role in the thematic situation (see Section 2.5.2). This provides a mechanism for structure-sharing a thematic argument with the appropriate semantic element from the verb semantics, regardless of which particular type of modification occurs on a particular use of the preposition. The lexical rule controlling thematic adjunct attachment essentially enforces the constraint in (3.74), unifying the thematic argument with the index for the situation as a whole, while constraints on the THEM-ARG for pseudo-complement attachment are defined in the semantic relation hierarchy (see Section 3.3.3).

(3.74)

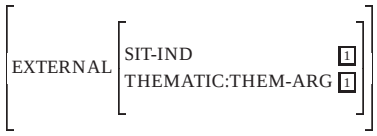

The lexical entry for the restrictive preposition *in*, as in *John ran in the park* is shown in (3.75). This definition follows that of Kasper (1993), specifying the restriction of an index picked out from the verb semantics. The semantic effect of a simple locative adverbial such as *in* is only to add a restriction on the LOCATION index associated with the situation – in this case, the location of the situation is in (*in-place*) *the park*. This restriction is recorded in the RESTR feature associated with the external semantics of the situation. The union of this restriction with any existing restrictions allows for the possibility of multiple restrictive PPs within a single sentence. Note also that this lexical entry utilises the standard HPSG approach to semantic specification, as the SYNSEM:LOC:CONT:NUC attribute of the entry specifies the full *sit-desc* to be associated with the sentence. Thus the lexical rule for restrictive adjuncts will need only to specify the replacement of the semantics associated with the sentence by this *sit-*

desc, in effect allowing this PP to become the semantic head of the sentence. Since this aspect of the semantics will be controlled by the lexical rules, it therefore does not need to be addressed in phrase structure schemata (see Section 3.5.1).

(3.75)

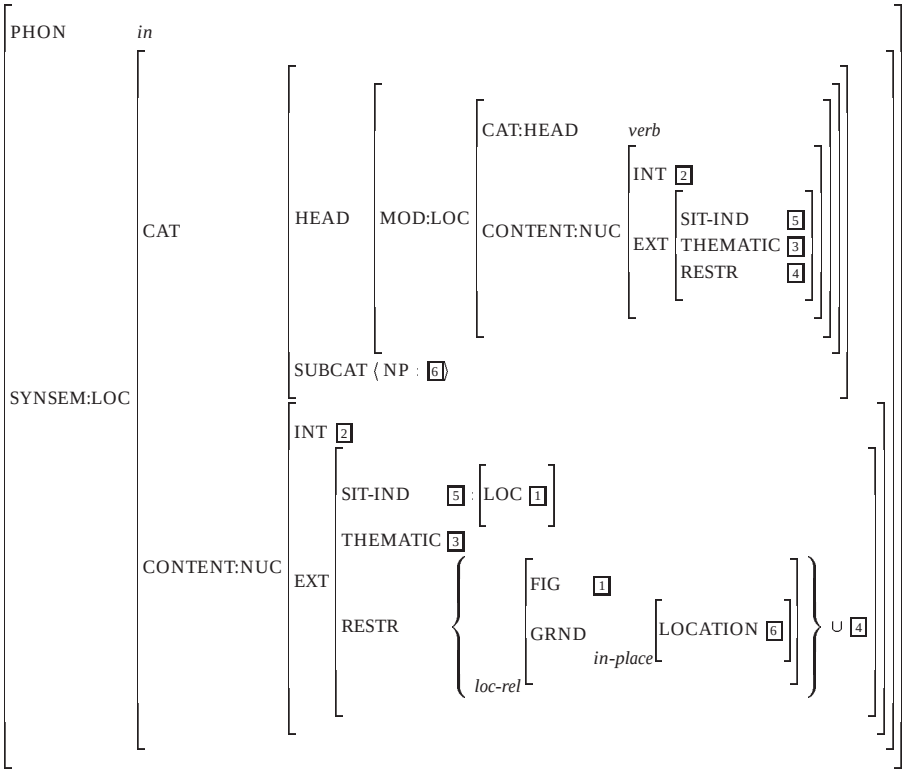

An approximated representation for the durative operator preposition *for*, as in *Peter reads for two hours*, is shown in (3.76). This also follows Kasper's (1993) treatment of operator adverbials. Here the semantic content of the modified VP appears as an argument of the *dur-soa*, reflecting the behaviour of operator adjuncts as adjuncts which predicate something of the content they modify.

(3.76)

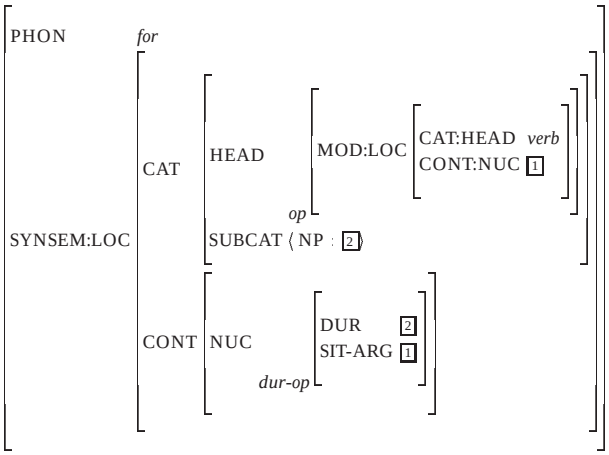

3.5.2 The OP-ADJUNCTS feature

Both the Kasper (1993) and the van Noord and Bouma (1994) approaches generate a list on which adjuncts theoretically appear in order of their semantic obliqueness. Surface order of these adjuncts is then controlled by separate principles of constituent order. The motivation behind building these lists in terms of semantic obliqueness lies in the compositional approach to semantic interpretation in the two approaches.

The problem with these approaches is that they cannot easily account for the interaction between semantic scope of modification and surface order. Furthermore, it is not clear in either approach how or when the relative semantic obliqueness of adjuncts on these lists is determined. In the Kasper (1993) approach, a mechanism must exist which drives the insertion of elements into the ADJ-DTRS list, although it is not explicitly specified. This mechanism must also be responsible for evaluating the relative semantic obliqueness of inserted elements. It is not at all obvious how the surface order of the elements would be taken into account in this evaluation.

The van Noord and Bouma (1994) approach assumes that the parser hypothesises a structure for the SUBCAT list of the head of a phrase which is evaluated against the constraints captured in the lexical rules. The hypothesised SUBCAT list must therefore reflect consultation of linear precedence rules imposed upon the parser. These linear precedence rules must be able to generate a SUBCAT list arranged in terms of semantic obliqueness from the surface order of the elements. Once the SUBCAT list is hypothesised to be a list of elements in a certain order, the lexical rules adding adjuncts to the subcat list act to perform the appropriate semantic integration of the adjuncts into the overall representation of the verb semantics. Because the system treats these rules as constraints to be verified, no mechanisms controlling the relative order of adjuncts on the SUBCAT list need be applied at the level of the rules. These mechanisms would be redundant.

This general approach is quite interesting, and effectively handles the word-order effects on the adjunct semantics if the linear precedence constraints are defined correctly. However, it is difficult to imagine how these constraints would be defined given that they would have to accommodate all variances in surface order among all adjunct types.

The approach presented here will restrict the domain of the constraints controlling semantic obliqueness to operator adjuncts. The constraints only need to take into account the relative semantic order of operator adjuncts, and will thus be easier to define. This restriction is possible since all other types of adjuncts provide information which actually modifies only the main *sit-desc* object associated with a verb. For example, in the sentences in (3.77), the "John-jogged" event is what is located in the park, regardless of the position of the restrictive PP relative to the operator adjuncts. It is not the "twenty-minutes-duration (John-jogged)" event which is located in the park, as would be suggested by (3.77b), or the "twice-daily (twenty-minutes-duration (John-jogged))" event which is located in the park, as suggested by (3.77c). All three of these sentences should have the interpretation "twice-daily (twenty-minutes-duration (in-park (John-jogged)))". Thus the semantic contribution of the restrictive adjunct must be incorporated before the operations specified by the operator adjuncts are

processed.

(3.77) a. John jogged in the park for twenty minutes twice a day.

b. John jogged for twenty minutes in the park twice a day.

c. John jogged for twenty minutes twice a day in the park.

Neither of the van Noord and Bouma (1994) and Kasper (1993) approaches handles these phenomena appropriately. Both approaches will give rise to errors in the semantic representation associated with a sentence containing interspersed operator and other adjuncts — namely that the restrictive or thematic adjuncts will be seen as modifying complex operator SOAs rather than the main SOA expressed by a sentence — because they do not postpone evaluation of operator adjuncts until after the other types.

Because all adjuncts other than operator adjuncts provide information relevant to the main *sit-desc* associated with a verb, the semantic contribution made by these adjuncts can be incorporated into the structure representing the semantics of the situation being modified as soon as they are encountered (i.e. as soon as the adjuncts are inserted into the SUBCAT list of the modified word by a lexical rule). On the other hand, operator adjuncts must always be processed after all other adjuncts, as evidenced by the example above.

To accommodate this difference between operator adjuncts and other adjuncts, I propose to make a distinction between the treatment of operator adjuncts and the treatment of other adjuncts. In the lexical rules controlling the treatment of all types of adjuncts other than operator adjuncts, the semantic contribution of these adjuncts is incorporated into the representation of the semantics of the situation immediately. For operator adjuncts, however, incorporation of their semantic contribution will be postponed until after all adjuncts have been inserted into the SUBCAT list.

As operator adjuncts are added to the SUBCAT list in a lexical rule, they are also added to an OPERATOR-ADJUNCTS (OP-ADJ) list associated with the SYNSEM:CAT of the lexical element whose SUBCAT list is being manipulated. This will be used in the handling of semantic status and surface order interactions. The OP-ADJ list will reflect the operator adjuncts in order of semantic obliqueness, while the SUBCAT list reflects the surface order of all complements and adjuncts.

The approach involves keeping track of both surface order and relative semantic obliqueness of operator adjuncts. Following van Noord and Bouma (1994), the application of the lexical rules will be driven by a structure for the SUBCAT list as proposed by the parser. However, this structure will reflect the natural surface order of the adjuncts rather than incorporating any evaluation of their semantic obliqueness. Thus linear precedence constraints on the parser will simply require that all adjuncts appear after the complements on the SUBCAT list, with the adjuncts in surface order. The evaluation of semantic obliqueness will occur when an operator adjunct is added to the subcat list in a lexical rule. The evaluation function will be given the existing OP-ADJ list and the new element, and then must determine the placement of the new element onto the list. This function will be able to take into consideration the relative surface order of the operator adjuncts, as any adjunct which it is attempting to insert into the OP-ADJ list must appear later in the surface order than any elements already on the list.

After all operator adjuncts have been inserted into the OP-ADJ list, and the semantic contribution of all other adjuncts has been integrated into the semantic representation for the situation as a whole, the semantics of the operator adjuncts can be processed. The OP-ADJ list will contain all of the operator adjuncts, listed from narrowest to widest scope. A function *process_op_adjs* will essentially accomplish what Kasper's (1993) *Adjunct Semantics Principle* does, but only for operator adjuncts: the MOD:LOC:CONT:NUC value of the adjunct of narrowest scope will be made token identical to the *sit-desc* object representing the situation. Then, if there are $n > 1$ elements on the OP-ADJ list, the MOD:LOC:CONT:NUC value of OP-ADJ$_i$ is token-identical with the SYNSEM:LOC:CONT:NUC value of OP-ADJ$_{i-1}$ for all i between 2 and n. The result of this processing is a semantic value which then becomes the semantics associated with the sentence as a whole.

In sum, the approach proposed here differentiates between operator and other adjunct types, integrating the semantics of other adjunct types immediately and postponing the semantic integration of operator adjuncts. This results in an appropriate representation of the semantics of sentences in which adjunct types are interspersed, and reflects the fact that only the semantic obliqueness of operator adjuncts relative to one another (but not to other adjunct types) plays a role in interpretation. The approach also allows the surface order of adjuncts to influence the evaluation of semantic obliqueness in a more straightforward manner by allowing the SUBCAT list to reflect their surface order.

3.5.3 Lexical rules

The approach presented here requires there to be different lexical rules for different types of adjunction. Each rule allows for the integration of semantics and structure-sharing between modifier and modified appropriate to the relevant type of adjunction. The design of the rules essentially follows that of van Noord and Bouma (1994), in that the *add_adj_control* rule builds an output structure based on the input structure it receives, by relying on other rules to modify elements of the structure in appropriate ways.

The basic algorithm for extending the SUBCAT list of the main verb in a sentence can be summarised as follows:

1. Take a verb's SUBCAT list as specified in its lexical entry.
2. Append adjuncts to the end of the SUBCAT list.
   (a) Add at most one pseudo-complement, if semantically compatible with the verbal relation.
   (b) Add any number of thematic adjuncts and integrate their semantics immediately into the verb's *sit-desc*.
   (c) Add any number of restrictive adjuncts and integrate their semantics immediately into the verb's *sit-desc*.
   (d) Add any number of operator adjuncts to the SUBCAT list (in surface order). Also add this adjunct to the *op-adj* list in order of semantic obliqueness relative to other

operator adjuncts already represented on that list as determined by the function *eval_sem_obliqueness* (not worked out below).

3. Process the semantics of any operator adjuncts added (as outlined in the previous section).

The application of the lexical rules depends on dynamic interaction with the parsing mechanisms through delayed evaluation techniques, as outlined in van Noord and Bouma (1994) and briefly introduced in Section 3.4.4. The lexical rules here can be seen as defining *potential* licensed additions to a verb's SUBCAT list, which can be actualised during the processing (parsing) of actual sentences.

The rules are described and presented in detail below. Note that the subsort check needed to prevent redundant thematic PPs is not explicitly represented, nor is the definition of *process_op_adjs*. The definition of the function *eval_sem_obliqueness* is left out, as it is beyond the scope of this chapter to determine precisely what the constraints on relative semantic obliqueness of operator adjuncts are.

Each clause of the function *add_adj* as defined below is a lexical rule which adds just one adjunct, of a particular adjunct type, at a time, and then recursively calls the function *add_adj* to allow the addition of any number of other adjuncts (which might be none). In this way adjuncts of different types can be interspersed with one another.

- Controlling rule: allows for the addition of all adjuncts to the element's SUBCAT list, and the processing of all operator adjuncts. The first argument is the original *synsem* object input, the second argument is the *synsem* object which results after all adjuncts have been added and processed. This rule calls *process_op_adjs*, which is responsible for processing the semantic contribution of the operator adjuncts and integrating it to establish the full semantics expressed in the sentence (*SemanticsOut*).

$$\text{add\_adj\_control}\left(\left[\text{LOC}\left[\begin{array}{ll}\text{CAT} & \left[\begin{array}{l}\text{HEAD Head}\boxed{1}\\ \text{SUBCAT SubcatIn}\boxed{2}\\ \text{OP-ADJ }\langle\rangle\end{array}\right]\\ \text{CONT} & \left[\text{NUC SemanticsIn}\boxed{3}\right]\end{array}\right]\right], \left[\text{LOC}\left[\begin{array}{ll}\text{CAT} & \left[\begin{array}{l}\text{HEAD Head}\boxed{1}\\ \text{SUBCAT SubcatOut}\boxed{4}\\ \text{OP-ADJ Operator-adjs}\boxed{5}\end{array}\right]\\ \text{CONT} & \left[\text{NUC SemanticsOut}\boxed{6}\right]\end{array}\right]\right]\right) :-$$

add_adj_top(Head[1], SubcatIn[2], SubcatOut[4], SemanticsIn[3],
SemanticsMid[7], ⟨⟩, Operator-adjs[5]),
process_op_adjs(SemanticsMid[7], SemanticsOut[6], Operator-adjs[5]).

- Top level rule used in the addition of adjuncts: forces all adjuncts to be added to the SUBCAT list after all complements. This clause applies when there are no more pseudo-complements to be added.

add_adj_top(Head[1], SubcatIn[2], SubcatOut[3], SemanticsIn[4], SemanticsOut[5],
Operator-adjsIn[6], Operator-adjsOut[8]) :–

add_adj(Head[1], AdjunctList[7], SemanticsIn[4], SemanticsOut[5],
Operator-adjsIn[6], Operator-adjsOut[8]),
append(SubcatIn[2], AdjunctList[7], SubcatOut[3]).

- Top level rule which licenses the addition of a pseudo-complement prior to any other adjuncts. The rule specifies the unification of the *thematic* information added by the pseudo-complement with the verb's default internal thematic element. This adds the information contributed by the pseudo-complement into the existing representation of the verb's internal semantics, if it is compatible (see Section 3.3.3). It then makes a recursive call to *add_adj_top* to allow for the addition of subsequent pseudo-complements prior to any other adjuncts.[14]

add_adj_top(Head[1], SubcatIn[2], SubcatOut[3], SemanticsIn[4]: [INT:THEMATIC / [11]],
SemanticsOut[5], Op-adjsIn[8], Op-adjsOut[9]) :-

append(SubcatIn[2], ⟨ [LOC [CAT [HEAD [MOD:LOC:CAT:HEAD [1]] pseudo-comp, SUBCAT ⟨⟩], CONT [NUC thematic[11]]]] ⟩,
SubcatMid[10]),

add_adj_top(Head[1], SubcatMid[10], SubcatOut[3], SemanticsIn[4], SemanticsOut[5],
Operator-adjsIn[8], Operator-adjsOut[9]).

- Base case for adding adjuncts – adds no adjunct, and the output semantics is unified with the input semantics. The argument structure for this function is:
add_adj(Head, Adjuncts, SemanticsIn, SemanticsOut, Operator-adjsIn, Operator-adjsOut).

add_adj(_, ⟨⟩, Semantics[1], Semantics[1], Operator-adjs[2], Operator-adjs[2]).

- Adds a thematic adjunct; specifies the unification of the *thematic* information added by the adjunct with the situational (external) thematic elements, adding the information into

[14] In defining this rule, I have ignored linking issues stemming from the definition of the semantic relation type hierarchy. Unification of the semantics of the preposition with the semantic relation of the verb is not sufficient for ensuring grammatical sentences, due to issues of redundancy as discussed in Section 3.4.1. A pseudo-complement can only either fill in an unfilled argument in a relation (to rule out sentences such as *John baked a cake for Mary for Mary*), or extend the relation to one which has all of its arguments linked either to explicit surface elements or to discourse referents (to prevent the extension of a *transfer-und-rel*, e.g. *John sent the letter* directly to a *transfer-to-ben-rel* without the recipient component of the relation being specified either through the context or through the addition of a "to" pseudo-complement, e.g. to prevent *?John sent the letter for Mary* without an understood recipient). These issues could potentially be resolved through linking constraints on the relations resulting from unification (see Davis 1995) or through discourse interactions, but I leave them for future investigation.

the representation of the situation as a whole. This clause handles both pure thematic adjuncts and the event-modifying use of *pc/adj* PPs.

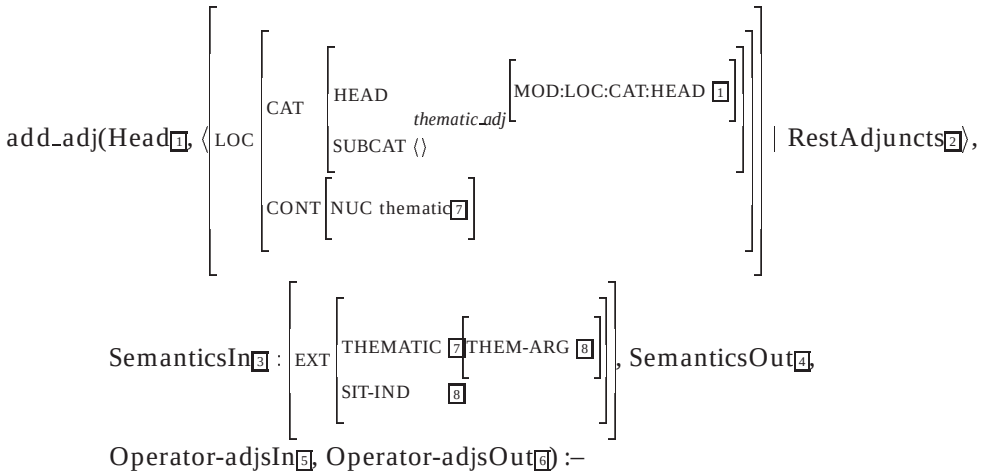

add_adj(Head[1], RestAdjuncts[2], SemanticsIn[3], SemanticsOut[4],
Operator-adjsIn[5], Operator-adjsOut[6]).

- Adds a restrictive adjunct; specifies that the *sit-desc* object found in the SYNSEM:LOC:CONT:NUC field of the adjunct definition becomes the semantics associated with the current situation, on the basis of the assumption that for restrictive adjuncts the change to the semantics of the modified situation is fully specified in the lexical entry.

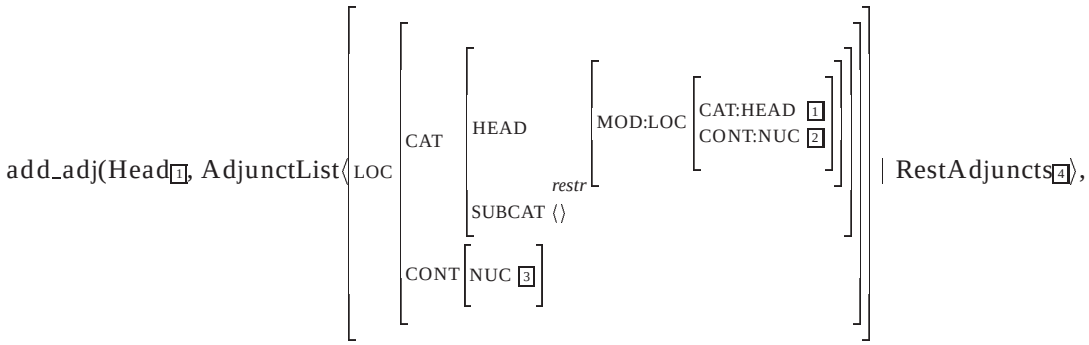

add_adj(Head[1], RestAdjuncts[4], Semantics-of-Adjunct[3], SemanticsOut[5],
Operator-adjsIn[6], Operator-adjsOut[7]).

- Adds an operator adjunct to both the adjuncts list and the operator-adjuncts list. Function *eval_sem_oblique* evaluates the semantic obliqueness of this adjunct with respect to other elements of the operator-adjuncts list and inserts it in the appropriate place ([2] is the *synsem* value associated with the adjunct, [6] is the original OP-ADJ list and [7] is the modified OP-ADJ list). Does not change the semantics associated with the current situation, as the integration of the semantics contributed by the operator adjunct will take place at the end of the addition of adjuncts, in *process_op_adjs*.

add_adj(Head[1], AdjunctList⟨[2] : [LOC:CAT [HEAD [op_adj [MOD:LOC:CAT:HEAD [1]]]] SUBCAT ⟨⟩]⟩ | RestAdjuncts[3]⟩,
SemanticsIn[4], SemanticsOut[5], Operator-adjsIn[6], Operator-ajdsOut[8]) :–

eval_sem_obliq(AddedAdjSynsem[2], Operator-adjsIn[6], Operator-adjsNew[7]),

add_adj(Head[1], RestAdjuncts[3], SemanticsIn[4], SemanticsOut[5],
Operator-adjsNew[7], Operator-adjsOut[8]).

3.5.4 Dative alternation

Within the framework developed in this chapter, dative alternation must be seen as an alternation between two forms of pseudo-complementation. The phenomenon can therefore be captured in terms of lexical rules. The dative form is accounted for straightforwardly by the pseudo-complementation lexical rule introduced in the previous section (Section 3.5.3). The double object form must be allowed by another rule, such as the one specified in (3.78). This rule identifies a pseudo-complement preposition in the lexicon (using the function *get_lex*) which supplies the semantics associated with the NP inserted into the SUBCAT list (the dative NP). This NP can be seen as the object of the missing preposition in the double object form. The lexical rule induces a "focus shift", raising the inserted NP in obliqueness to the level of direct object and pushing the original direct object down to the level of indirect object.

(3.78) a. add_adj_top(Head[4], SubcatIn⟨NP:[1], NP:[2] | RestSubcat[11]⟩,
SubcatOut[6], SemanticsIn[10]: [INT:THEMATIC / [5]], SemanticsOut[7],
Operator-adjsIn[8], Operator-adjsOut[9]) :–
test_rest(RestSubcat[11]),
get_lex([SYNSEM:LOC [CAT [HEAD [pseudo-comp [MOD:LOC:CAT:HEAD [4]]] SUBCAT ⟨NP : [3]⟩] CONT:NUC thematic[5]]]),
add_adj_top(Head[4], ⟨NP:[1], NP:[3], NP:[2]⟩, SubcatOut[6], SemanticsIn[10],
SemanticsOut[7], Operator-adjsIn[8], Operator-adjsOut[9]).

b. test_rest(⟨⟩).
test_rest(⟨PP⟩).

This rule can apply to transitive forms of verbs which have semantics compatible with a pseudo-complement preposition. So, for example, the presence of a lexical entry which licenses *John sang a song* will lead to a lexical entry which licenses *John sang Mary a song* where *sing* is compatible with the benefit expressed by the *for* pseudo-complement. In this case,

RestSubcat is ⟨⟩ and the test *test_rest* succeeds with its first clause. It also can apply to ditransitive forms of verbs with a SUBCAT list ⟨NP, NP, PP⟩ again, if the semantics of the verb is compatible with the semantics of a pseudo-complement preposition. For instance, a lexical entry which licenses *John gave a book to Mary* will lead to one which licenses *John gave Mary a book* by virtue of the presence of the pseudo-complement preposition *to* in lexicon, which has a semantics compatible with the meaning expressed by *give*. In this case the second clause of *test_rest* will succeed. The function *add_adj_top* is called recursively to allow for the addition of other compatible pseudo-complements, e.g. to license *John gave Mary a book for the kids* (interpretation: John gave a book to Mary, and the book is intended to benefit the children).

This approach to the dative alternation links the alternate forms through the semantics associated with the dative (pseudo-complement) preposition — the semantics provided in one case by the preposition is in the other case indicated by the obliqueness of one NP relative to the other. Thus the approach makes a generalisation about the relationship between dative PPs and inner double object NPs. Furthermore, the approach ensures that there is only one available interpretation of the double object form — the pseudo-complement interpretation — even if the "missing" preposition can be interpreted as either a pseudo-complement or an adjunct type preposition. It is also in line with Jackendoff's (1990) analysis in which the double object form only allows an interpretation in which the object of the preposition benefits from the object of the verb, but differs from that work in that here the double object form has an interpretation identical to one of the interpretations of the dative form (see Section 3.3.2).

The rule in (3.78) above is only an example of how the double-object form lexical rule could be defined. In actuality, this rule would likely have to define more complicated modifications of the internal semantic structure expressed by the verb in the alternate form. Several lexical rules of this type may also be necessary, probably involving a more precise definition of the initial internal semantics associated with the verb, to capture different types of semantic alternation between the dative and double object forms.[15] The form of the rules is not critical for the current discussion; the fact that such rules can be defined to account for the dative alternation is important.

The lexical rule approach to the treatment of the semantics of the two forms involved in the dative alternation provides a means of accounting for alternation contrasts previously difficult to explain. Consider the sentences in (3.79)-(3.84). (From Jackendoff 1990, who attributes (3.79)-(3.80) and (3.83)-(3.84) to Jane Grimshaw.)

(3.79) a. John fixed the roof for Mary.

b. *John fixed Mary the roof.

(3.80) a. John fixed a sandwich for Mary.

b. John fixed Mary a sandwich.

[15] See Verspoor (1994) and Pinker (1989) for a fuller discussion of lexical rules used to capture syntactic alternations with corresponding semantic consequences which depend on a verb's semantics. Goldberg (1995) also addresses this issue, discussing examples such as **John blew a kiss to Mary/*John gave a kick to Mary* in contrast to *John blew Mary a kiss/John gave Mary a kick*. These differences in acceptability could stem from slight variation in the meaning of the double object form as compared with the standard dative form.

(3.81) a. Bill removed the garbage for Harold.
b. *Bill removed Harold the garbage.

(3.82) a. John chose a dress for Mary.
b. *John chose Mary a dress.

(3.83) a. Sue poured some cement for Dick.
b. *Sue poured Dick some cement.

(3.84) a. Sue poured some coffee for Dick.
b. Sue poured Dick some coffee.

The contrast between (3.79) and (3.80) stems from differences in the meaning expressed by the verb *fix*. In (3.79), *fix* means *repair*, and is apparently not compatible with the semantics of the pseudo-complement on this interpretation. The only lexical rule which can be used to interpret (3.79a) is the thematic adjunction lexical rule, resulting in an interpretation in which the entire fixing event is done for Mary. No interpretation of (3.79b) is possible because the lexical rule licensing the double object form requires the modified verb to have semantics compatible with the meaning of the pseudo-complement. In (3.80), on the other hand, *fix* is being used to mean *make*, which is compatible with *for* pseudo-complementation (as a *cause-create-rel*), and therefore the double object form lexical rule can apply to provide an interpretation for (3.80b). Likewise, the semantic relations expressed by the verbs in (3.81) and (3.82) are not compatible with pseudo-complements and thus the double object forms involving these verbs are not permitted. Only the thematic adjunct interpretation of the PPs is available.

The contrast between (3.83) and (3.84) must be a result of consultation of world knowledge in the application of the lexical rules. There is no difference in the senses of *pour* expressed in these sentences. On the pseudo-complementation interpretation of these sentences (that is, when the pseudo-complement lexical rule introduces the *for*-phrase), what is being poured is interpreted as affecting Dick directly. While there are clearly several senses in which coffee can benefit Dick (e.g. because it is liquid and humans need liquid to survive; because it is warm; etc.), there is no sense in which the cement in (3.83) can benefit Dick directly, likely because Dick is not intended to receive the cement. The *benefit-rel* which is conveyed by the preposition *for* is an underspecified relation which must be fully specified through pragmatic reasoning (to determine how something benefited), and in the case of (3.83) the relation cannot be specified on the basis of world knowledge. Thus the pseudo-complement interpretation of these sentences is ruled out on the grounds of limitations in the world. Pragmatic reasoning is required to support and augment the underspecified lexically-proposed interpretation.

3.6 Conclusions

In this chapter, I have argued with reference to data from the dative and benefactive constructions in comparison with other adjunctive constructions that there is a three-way distinction in the syntactic and semantic function which a prepositional phrase can play in a sentence.

- ***complement PPs:*** syntactically obligatory PPs which introduce a semantically entailed component of a verbal relation. The contribution of the PP can either be idiosyncratic, or can be the standard semantic contribution of that PP on a pseudo-complement use.
- ***pseudo-complement PPs:*** syntactically optional PPs which either mark an entailed participant in a verbal relation or extend a verbal relation in a way licensed by the verbal relation hierarchy.
- ***adjunct PPs:*** syntactically optional PPs which provide situational information but do not interact with the verbal relation expressed in the main clause.

Making this distinction allowed development of a lexical rule-based treatment of PP (more generally, adjunct) integration driven by the semantic properties of the prepositions and the verbs, as lexically defined. Combinations of VPs and PPs are constrained through the semantic relation hierarchy and the types of individual prepositions, in conjunction with lexical rules which control in very general terms the composition of semantics from the combined phrases.

The approach eliminates lexical representation of syntactically optional verbal "complement" PPs, instead treating these PPs as pseudo-complements and allowing them to be productively licensed on the basis of semantic properties of the verbs. This captures a generalisation about a verb's (potential) syntactic argument structure on the basis of lexical semantics.

The approach explicitly acknowledges the semantic content of prepositional phrases, and shows how this content may interact in various ways with the semantic content expressed by other elements in a sentence. Furthermore, syntactic ordering is taken into consideration in determining how this interaction might occur, thereby ruling out infelicitous readings (e.g. a PP which appears after other adjuncts cannot have a pseudo-complement interpretation). Finally, the lexical semantically derived interpretation of certain constructions was briefly discussed and shown to require support from pragmatic reasoning for complete felicity.

The proposals made in this chapter concerning the treatment of adjuncts go a long way towards appropriately handling the characteristics of adjuncts:

- ***Consistent semantic contribution:*** There is only one lexical entry required in this approach for each meaning associated with an adjunct, even if the adjunct is involved in different types of adjunction.
- ***Restrictive and operator adjuncts:*** Both of these types of adjuncts are accounted for and treated in a way which reflects precisely the type of modification which must be associated with them; namely that restrictive adjuncts directly modify the situation expressed by the verb and that operator adjuncts take a full situation as an argument. A third type of adjunct, *thematic adjuncts*, has also been identified as a type of adjunct which adds information about a situation as a whole.
- ***Surface order vs. semantic obliqueness:*** The interaction between surface order and semantic obliqueness for operator adjuncts is accounted for by maintaining lists which reflect both of these types of information — surface order in the SUBCAT list and semantic obliqueness in the OP-ADJS list.

- ***Redundancy constraints:*** Redundant PPs are avoided through use of a type system which keeps track of modifying information associated with a sentence. The lexical rules would then simply need to include a subsort check to prevent two modifiers of the same type in a sentence.

- ***Mittelfeld phenomena:*** Complements and adjuncts both appear on the SUBCAT list of a head. There is thus nothing structural which prevents these elements from being interspersed. The linear precedence rules must be responsible for determining their allowed relative order.

The advantages of the approach presented in this chapter over the previous approaches from which it is derived can be summarised as the following:

- The division between external and internal semantics allows various types of modification, including types not handled in the previous approaches (thematic adjuncts and pseudo-complements), to be accommodated within the same framework. In particular, the use of a semantic object of type *thematic* common to external and internal semantics provides for a general treatment of prepositions which can behave both as a thematic adjunct and as a pseudo-complement. This treatment can even account for the ambiguity of interpretation found in sentences involving such prepositions.

- Dative alternation can be easily accounted for by defining variants of the basic pseudo-complementation lexical rule. It follows only from specific underlying lexical semantic properties of a verb.

- The interspersal of operator adjuncts with other types of adjuncts does not lead to interpretation errors.

- There is a more straightforward framework in which to account for the interaction between surface order and semantic precedence. The use of delayed evaluation and linear precedence rules which follow surface order allows the context to drive adjunct interpretation.

four Manner of Motion Verbs and Resultatives

4.1 Introduction

In this chapter I will examine in detail the *resultative construction*. This construction is composed of a verb plus its arguments and an additional unsubcategorised phrase (either an adjective phrase [AP] or a prepositional phrase [PP]) which expresses a result state of the event expressed by the verb. This result state is predicated of one of the verbal arguments, or in some cases of an unsubcategorised noun phrase [NP] which is also inserted (4.1d). Examples of the construction appear in (4.1).[1]

(4.1) a. John hammered *the metal* flat.

b. *The river* froze solid.

c. John sneezed <u>the tissue</u> off the table.

d. John laughed <u>himself</u> silly.

Manner of motion verbs are a group of semantically similar verbs which can be defined as follows (Levin 1993:264):

> These verbs describe motion that typically, though not necessarily, involves displacement, but none of them specifies an inherent direction as part of its meaning. All of these verbs have meanings that include a notion of manner or means of motion. They differ from each other in terms of the specific manner or means.

[1] My annotations: italics indicate an argument of the main verb which is also the subject of the resultative predicate; underlining indicates something that is not an argument of the main verb, but which is the subject of the resultative predicate. Additionally, in what follows I use a '*' to indicate ungrammatical sentences, '#' to indicate pragmatically infelicitous sentences, and '?' to indicate sentences of questionable grammaticality.

These verbs can be extended to a *directed motion* use, in which an "inherent direction" is contributed via a path prepositional phrase. Examples of standard and extended uses of some of these verbs appear in (4.2)-(4.3).

(4.2) a. John swam.

b. John swam across the pool.

(4.3) a. John ran.

b. John ran to the school.

In most existing work on resultatives that I am aware of (Levin and Rappaport Hovav 1995, Wechsler 1996, Markantonatou and Sadler 1995, *inter alia* — exceptions are Jackendoff 1990 and Goldberg 1995), manner of motion verbs on a directed motion sense are treated as resultative constructions due to the result-like quality of the directional PP (i.e. the PP seems to introduce a result state: in (4.2b) John is at the other end of the pool, and in (4.3b) John is at school). The syntactic form of the directed motion use of manner of motion verbs, however, differs from the standard form of resultatives for unergative verbs (verbs which have only an underlying subject, including *sneeze* and *laugh* in (4.1)). Instead, they pattern with unaccusative verbs (verbs which have an underlying object but not subject, like *freeze* and *break*: *The river froze*, *The glass broke*) although they are not normally considered to be unaccusative. This difference poses a challenge to uniform explanatory theories of the resultative construction.

In this chapter I will investigate the syntactic and semantic properties of the resultative construction, arguing that they are distinct from the properties associated with manner of motion verbs and that a uniform analysis cannot explain the distinctions. I will look specifically at the unaccusativity-based proposals of Levin and Rappaport Hovav (1995) (henceforth L&RH), and will expose problems with the evidence for their analysis of the manner of motion verbs, and will introduce issues that their account fails to address adequately. In particular I will show that semantic differences between the manner of motion verbs on a directed motion use and other resultatives fail to be captured by a uniform treatment. I will also examine a semantic solution to the modeling of the resultative construction proposed by Wechsler (1996), concluding that his intuitions are generally accurate and moreover point to a non-syntactic explanation for unaccusativity. In contrast to Wechsler, however, I will suggest that the behaviour of verbs in the resultative construction stems from *a combination of the semantics of the construction, lexicalisation of certain forms, and contextual constraints*. Moreover, I will argue for a non-resultative treatment of the manner of motion verbs with a directional phrase complement.

I will end with an alternative explanation of the manner of motion verbs data, and resultative constructions in general, drawing on work done in the framework of Construction Grammar (Goldberg 1995). This proposal will obviate the need for the unaccusative / unergative distinction within the analysis of resultatives and will clearly differentiate the treatment of manner of motion verbs from true resultatives. It will be formalised in terms of the representations outlined in Chapter 2, and with the lexical rule mechanisms described in Chapter 3.

I will also argue throughout the chapter that the data in question cannot be considered in isolation of pragmatic reasoning, and that the felicity of a particular sentence depends on its coherence in the discourse it appears in. I will provide a tentative proposal for how discourse coherence constrains the interpretation of the data. I will show how an acknowledgement of the influence of lexical conventionalisation and pragmatic reasoning can lead to a fuller account of the data.

4.2 Levin and Rappaport Hovav (1995)

4.2.1 The account

Levin and Rappaport Hovav (1995) aim for a unified treatment of resultative constructions, on the basis of the distinction between unaccusative and unergative verbs. Their proposals with respect to resultatives are intended to support their main argument that unaccusativity is syntactically encoded while being semantically motivated.

According to the L&RH approach to resultatives, variations in syntactic behaviour in the resultative construction stem only from underlying differences in the unaccusativity of the verbs participating in the construction and the existence of a *Direct Object Restriction*, summarised in (4.4).

(4.4) **Direct Object Restriction (DOR):** a resultative phrase may be predicated only of an immediately postverbal NP, not of a subject or of an oblique complement.

According to this constraint, unergative (intransitive) verbs with no underlying object and hence no immediately postverbal NP require a "fake reflexive" (an NP co-indexed with the subject) or some other NP to be inserted into postverbal position, to introduce something which the resultative can predicate (4.1d). Transitive verbs, on the other hand, have an overt postverbal NP and unaccusatives have an underlying object (which moves to subject position due to syntactic constraints) and so these verbs have (underlying) direct objects which the resultative can predicate (4.1a-b). The implication of this treatment of resultatives is that any verb which appears in what L&RH term the *unaccusative resultative construction*, in which an intransitive verb is directly followed by a resultative phrase, must be treated as having unaccusative argument structure in order to account for the felicity of the verb in a resultative construction without a postverbal NP and their infelicity in this construction with a postverbal NP. The verbs in (4.5)-(4.7) must therefore all be given an unaccusative analysis.

(4.5) a. *They* slowly swam apart. [L&RH (5.15b)]

b. *They slowly themselves swam apart.

(4.6) a. *The refrigerator* door clicked open. [L&RH (5.27a)]

b. *The refrigerator door clicked itself open.

(4.7) a. *The prisoners* froze to death. [L&RH (2.19b)]

b. *The prisoners froze themselves to death.

The intransitive *freeze* is generally accepted to be unaccusative and therefore this behaviour is explained via the Direct Object Restriction, while *swim* and *click* are not. In the context of change-of-location resultatives, however, the manner of motion and sound emission verbs must be viewed as unaccusative in order for the DOR to hold. This is accomplished in the Levin and Rappaport Hovav (1995) account via a lexical rule applied to verbs of manner of motion (e.g. *swim*) and verbs of sound emission (e.g. *click*), to the exclusion of other verb classes, which adds a directional phrase requirement and shifts the lexical classification of the verb to be unaccusative. There will therefore be two lexical entries for *swim*: $swim_{\text{unerg}}$ and $swim_{\text{unacc}}$ + dirP. Note that the added directional phrase need not be one that makes the denoted eventuality telic, and that thus its addition is a syntactic rather than semantic specification.

The evidence which L&RH present in favor of this unaccusative treatment of verbs of manner of motion and verbs of sound emission on a directed motion use comes in part from auxiliary selection data in Dutch and Italian, but mainly from the English resultative construction itself and from the causative alternation. I will discuss each of these pieces of evidence in turn below.

4.2.2 Problems with the Levin and Rappaport Hovav analysis

Auxiliary Selection

Although L&RH concentrate on providing evidence from English for the shift in unaccusativity of verbs of manner of motion in the presence of a directional phrase, they do suggest that the pattern of auxiliary selection in languages such as Dutch and Italian supports a treatment of verbs of manner of motion as unaccusative on a directed motion use. The evidence they cite is shown in (4.8)-(4.9) below (L&RH :185(13)). These data show that the shift of a manner of motion verb to a directed motion use is accompanied by a change in the selected auxiliary, from *have* to *be*. Levin and Rappaport Hovav assume that a verb's selection of the auxiliary *be* is an indication of its underlying unaccusativity.

(4.8) a. Hij heeft / *is gelopen.
he has/is run

b. Hij is / ?heeft naar huis gelopen.
he is/?has to home run

(4.9) a. Ugo ha / *è corso meglio ieri.
*Ugo has/*is run better yesterday*

b. Ugo è / *ha corso a casa.
*Ugo is/*has run to home*

There has, however, been much discussion on the topic of the relationship between auxiliary selection and unaccusativity. For example, Zaenen (1988, 1993) suggests that auxiliary selection in Dutch has to do with semantic features such as *controllability* and *telicity*, rather than being a reflection of the underlying syntactic argument structure of the verb

phrase. Everaert (1992) proposes that auxiliary selection depends on the theta- or case-assigning properties of a verb, again independent of the argument structure of the verb phrase. He argues that both auxiliary selection and unaccusativity can be defined in terms of these properties, but that there is not necessarily a direct correspondence between them. Furthermore, consider the Italian sentences in (4.10).[2]

(4.10) a. i. L'uomo è/*ha corso all'università.
*The man is/*has run to the university.*

ii. L'uomo ha/*è corso verso l'università.
*The man has/*is run towards the university.*

b. i. L'uomo corre all'università.
The man runs to the university.

ii. *[3]L'uomo cammina all'università.
The man walks to the university.

In Italian, different types of directional phrases result in the selection of distinct auxiliaries (as in (4.10a)), suggesting that the telicity of the directional phrase influences auxiliary selection. This is in contrast to the syntactic specification of the directional phrase as introduced by L&RH, as mentioned earlier. In addition, parallel sentences, containing verbs which are semantically highly similar and identical directional phrases, may differ in acceptability (4.10b) (Di Tomaso, 1996). These data suggest that semantic factors interact with auxiliary selection, causing conclusions about syntactic encoding of unaccusativity drawn solely on the basis of auxiliary selection to be highly questionable.

Although there is no consensus on precisely which factors affect auxiliary selection, the analysis of the relevant data does suggest that while unaccusativity might be sufficient to determine the selection of the *be* auxiliary, it may not be a necessary condition for *be* selection (i.e. all unaccusative verbs may require a *be* auxiliary, while it may not be true that all verbs which select a *be* auxiliary are unaccusative). Thus on the basis of auxiliary selection alone, we should not assume that the uses of manner of motion and sound emission verbs in a directed motion sense are necessarily unaccusative.

Resultative Constructions

The first set of data in English which Levin and Rappaport Hovav (1995) present in favor of an unaccusative analysis of directed motion sentences comes from the resultative construction itself. They contrast manner of motion resultatives which specify a change of location (4.11)-(4.12) with those that specify a change of state (4.13)-(4.14), (L&RH :186-187, (15)-(18); judgements of acceptability noted here are theirs).

(4.11) a. *She* danced/swam free of her captors.

[2]Data from Di Tomaso (1996); see that paper for detailed discussion of Italian verbs of manner of motion and spatial prepositions.

[3]Note here that on the interpretation of the Italian preposition *a* as the locative *at*, this sentence is acceptable.

b. *They* slowly swam apart.

(4.12) a. *She danced / swam herself free of her captors.

b. *They slowly swam themselves apart.

(4.13) a. **She* danced / swam sore.

b. *Don't [*you*] expect to swim sober.

(4.14) a. She danced her feet sore.

b. Don't [you] expect to swim yourself sober!

From the treatment of resultatives in terms of the Direct Object Restriction in (4.4) it follows that the verbs in (4.11) must be given an unaccusative analysis. In these sentences there is no postverbal NP of which the change of location resultative phrase can be predicated. Instead, the surface subjects are interpreted as entities undergoing a change of location. Under the DOR, these surface subjects must then be underlying (deep structure) objects. That is, the verbs in these sentences are unaccusative. In contrast, the required reflexive postverbal elements in (4.14) (cf. (4.13)) suggest that the relevant verbs in these sentences have an unergative analysis and therefore do not provide the resultative phrase with a subject in the underlying argument structure.

The main verbs of the sentences in (4.11)-(4.12) are identical to those in (4.13)-(4.14), yet the former set of examples are felicitous only when they appear without a (reflexive) postverbal element, while the reverse holds for the latter. Comparison of these two sets shows that, in addition to this syntactic difference, the resultative phrases differ: in the first set the resultative phrases express a change of location, while in the second set, the resultative phrases express a change of state. L&RH identify this distinction in resultative phrase type as the explanation for the syntactic differences: manner of motion verbs behave like all other unergative verbs in the resultative construction, except when the resultative phrase expresses a change of location, in which case they behave like unaccusative verbs. L&RH therefore propose that the change of location resultative phrase is added to the subcategorisation frame of manner of motion and sound emission verbs by a special lexical rule which shifts the sense of the verb to a directed motion interpretation and simultaneously changes the argument structure of the verb to be unaccusative. This is in contrast to their treatment of all other resultative constructions, for which the resultative phrases (regardless of their type) are not syntactically licensed by a lexical rule, but instead on the basis of a more general process resulting from the identificaiton of event positions in the argument structures of the verb and the resultative phrase (see Levin and Rappaport Hovav 1995:53).

This analysis accounts for the data above. The sentences in (4.12) are ruled out because the verb is in the context of a directional phrase and so must have unaccusative structure due to the lexical rule which licenses the directional phrase. The single underlying argument position cannot be filled twice. The sentences in (4.13) in contrast do not appear with a directional phrase. The verb in these cases is unergative, and hence the sentences are ruled out because there is neither an underlying object nor an overt postverbal element of which the resultative phrase can be predicated.

Nonetheless, the account has a flaw which can immediately be identified. The extended sense of manner of motion verbs is linked to an explicit directional phrase through the lexical rule, causing the sense shift to be dependent on the presence of this phrase. It is, however, possible for a manner of motion verb to acquire a directed motion sense through context, without an explicit directional phrase appearing as one of its arguments. In (4.15), for example, the sentence *He walked* means *He walked to the store*.

(4.15) a. John had to go to the store to buy some flowers for Mary.

b. He walked, arriving just before it closed.

This example implies that the directional phrase requirement imposed by the lexical rule must be overridable by the discourse context, or that the unergative-unaccusative shift can be induced by discourse context. This fact calls into question the dependency of the sense of the verb on a lexical rule and, as a result, the association of the verb's underlying argument structure with a particular syntactic frame. I will return to this issue in Section 4.3.1.

In addition, there is a problem with using resultative data to support an analysis of the directed manner of motion verbs as unaccusative in that the analysis presupposes a particular treatment of the resultative construction. The effect is that the argumentation concerning resultatives is rather circular. The arguments in favor of the L&RH proposal for the syntax of the resultative construction depend in part on the distinct behaviour of unaccusative and unergative verbs with respect to this construction. In particular, the Direct Object Restriction is justified in large part on the basis of unaccusative data. Then in order to bring the directed manner of motion resultative data in line with the analysis required under this restriction, L&RH must assume that these verbs are unaccusative. So in fact the data do not provide independent evidence for the unaccusativity of these verbs; rather, their unaccusativity follows from the analysis of resultatives L&RH assume. A different treatment of resultatives may not require these verbs to be unaccusative.

Causative Alternation

The second set of data in English presented in favor of an unaccusative analysis comes from the causative alternation. This evidence is independent of their analysis of the resultative construction, and therefore more convincing.

A causative construction is a sentence containing a transitive verb which essentially expresses "cause to V-intransitive", where V-intransitive is the intransitive counterpart of the verb (Levin 1993). Examples of this construction are in (4.16)-(4.17), for which the causative (a) sentence entails the intransitive (b) one, (from Levin 1993:26-32).

(4.16) a. John broke the glass.

b. The glass broke.

(4.17) a. The visitor rang the bell.

b. The bell rang.

Particularly relevant data, involving verbs of manner of motion, are shown in (4.18)-(4.21), (L&RH :188,(19)-(21)).

(4.18) a. The soldiers marched (to the tents).

b. The general marched the soldiers to the tents.

c. ??The general marched the soldiers.

(4.19) a. The horse jumped (over the fence).

b. The rider jumped the horse over the fence.

c. The rider jumped the horse. *[directional phrase understood]*

(4.20) a. The mouse ran (through the maze).

b. We ran the mouse through the maze.

c. *We ran the mouse.

(4.21) a. The tricycle rumbled across the sidewalk.

b. They rumbled the tricycle across the sidewalk.

c. *They rumbled the tricycle.

The examples suggest that a directional phrase must be present or at least understood when the manner of motion and sound emission verbs are used causatively. Thus the verbs differ with respect to their ability to undergo causativisation on their different senses, supporting the Levin and Rappaport Hovav (1995) assertion that the lexical properties of these verbs change when they undergo the sense shift. Specifically, L&RH argue that the pattern in the causativisation data stems from the unaccusative argument structure of the verbs on a directed motion use. This conclusion depends on their view of causativisation, in which an external cause can only be introduced into an argument structure which has an empty external argument position. Since the causative form is felicitous on the directed motion sense of the verbs in (4.18)-(4.21), the data suggest that on this sense the verbs have an unfilled external argument position — that is, they have an unaccusative argument structure — in contrast to the unshifted sense of the verbs which have an unergative argument structure and cannot appear in the causative construction.

The evidence from causativation, however, is weak at best. Levin and Rappaport Hovav (1995:193) acknowledge that causativation examples for the verbs of sound emission are infrequent, and the attested examples they cite (such as (4.21b) above) are not so clearly acceptable.[4] Furthermore, most manner of motion verbs cannot appear in the causative construction, even on a directed motion interpretation, as shown in (4.22). The example (4.20b) shows great variability in acceptability when different objects and directional phrases are chosen, as in (4.23), suggesting that the causative use of *run* may be idiosyncratic or lexicalised with highly specific selectional restrictions and a particular interpretation varying from a true causative meaning. Other examples of "causative" manner of motion verbs, as shown in (4.24), also

[4]The uses could be for stylistic or literary effect.

seem to have idiomatic, specific interpretations.[5] In these cases, the directional phrase seems to be optional as well, again suggesting a lexicalised causative version of the verb rather than a causative form derived from an underlying unaccusative.

(4.22) a. The child limped / hobbled / ambled / meandered / swaggered / sauntered / sashayed / wiggled to the store.

b. #John limped / hobbled / ambled / meandered / swaggered / sauntered / sashayed / wiggled the child to the store.

(4.23) a. We ran the cat out of the house.

b. *We ran the child to the store.

c. *We ran the cat down the road.

d. #We swam the mouse through the water maze.

(4.24) a. We walked the dog (down the road).

b. We walked Bill home.

c. ?We walked the cat across the road.

d. We jumped the horse (over the wall).

e. ??We jumped the cat across the ditch.

The range of acceptability in this data points to the interaction of several knowledge sources in the evaluation of the felicity of instances of the causative construction. While there does appear to be a syntactic process which underlies the generation of causative forms of certain verbs, a purely syntactic treatment is bound to overgenerate and must be highly constrained through non-syntactic means. Constraints seem to derive both from semantics, e.g. from lexicalised (or conventionalised) meanings of particular verbs, such as *walk* and *jump* above, and from pragmatics — the causative interpretation of a verb use must be made sense of in the context of use. We can in fact improve the felicity of certain instances of this construction by placing them in an appropriate context. Consider (4.25). This sentence, independent of a context, seems to make little sense. However, if we place it into the context of the discourse in (4.26) which establishes the causal influence John can have on the child limping to the finish line, it improves in felicity.

(4.25) #John limped the child to the finish line.

(4.26) The children were racing towards the finish line. Sarah tripped over a treebranch and hurt her ankle. Determined not to see his favorite niece lose, John grabbed Sarah's hand, pulled her off the ground, and **limped the child to the finish line**.

The process of causativisation, then, is only semi-productive. Its productivity is restricted via the mutual combination of world knowledge as dictated by lexical semantics and pragmatic

[5]Note that (4.24a) and (4.24b) differ greatly in interpretation. (4.24b) clearly does not have a causative meaning. Rather it has an accompaniment / aiding meaning.

inference. I will return to these issues to investigate the interaction of different knowledge sources in explaining such sentences, and also with respect to the resultative construction, in Section 4.3.1.

The lack of consistency in the causativisation data for verbs of manner of motion and verbs of sound emission means that sweeping claims about the underlying argument structure for these verbs on the directed motion sense are untenable on the basis of this data. This evidence is therefore inconclusive for the L&RH aim of justifying a particular syntactic encoding of the resultative construction.

Conclusions

The influence of semantic factors on auxiliary selection, the circularity of the arguments in favor of an unaccusative treatment of verbs of manner of motion and verbs of sound emission on a directed motion use on the basis of resultative data, and the lack of clear patterns in the causativisation data for these verbs indicate a need to find stronger evidence in favor of a treatment of these verbs as underlyingly unaccusative. These problems do not, however, necessarily mean that L&RH 's analysis of the resultative construction and in particular the argument structure of these verbs is incorrect. They simply indicate that their explanation has not as yet been adequately motivated in the case of these verbs. However, even if we were to accept the evidence L&RH provide for their account of the treatment of verbs of manner of motion, there remain several semantic issues which are inadequately or incorrectly handled in their approach. I turn to those issues in the next section.

4.3 Semantic properties

In this section I will discuss in detail semantic and pragmatic issues which must be taken into consideration in modeling manner of motion verbs and resultatives. These issues will become clear through analysis of the L&RH account.

4.3.1 The meaning of the constructions

By attempting to treat all resultative constructions via a uniform mechanism, Levin and Rappaport Hovav (1995) fail to capture subtle differences in the meaning of directed manner of motion constructions and other resultative constructions. Specifically, they argue that these constructions both express causative changes of state.[6] I will show, however, that a change of state is not necessarily causative and the distinction between directed manner of motion constructions and other resultative constructions is precisely that the former express non-causative changes of state while the latter express causative changes of state.

[6]This point about the CAUSE component of meaning in the unergative resultative construction has also been made by Markantonatou and Sadler (1996:4,7). Note that they do not explicitly address any semantic differences between manner of motion verbs and other verbs in the resultative costruction.

Let us consider two similar examples. (4.27) is a directed manner of motion construction and (4.28) is a change-of-state resultative involving the same verb.

(4.27) John walked to the store.

(4.28) John walked his feet sore.

The L&RH account of the interpretation of the resultative construction is based on compositionality, as in most semantic theories, with the additional assumptions that an eventuality which expresses a process followed by a state is interpreted as an accomplishment (cf. Moens and Steedman 1988), and that an accomplishment always describes causative changes of state (L&RH :54). The meaning of resultative constructions consists of the basic meaning expressed by the main verb, which I will represent in the case of (4.27) as *walk(John)* in line with predicate logic representations, and the meaning expressed by the resultative predicate, approximately *at_the_store(John)*.[7] The process expressed in the main verb is composed with the state expressed by the resultative phrase to generate an interpretation in which the main event causes a change to the state conveyed by the resultative phrase. Sentence (4.27), then, is interpreted as meaning essentially *John's walking causes the state of John being at the store* on the L&RH account (cf. Dowty 1979). Similarly, sentence (4.28) expresses *John's walking causes the state of John's feet being sore*.

One identifiable weakness of this account is the failure to explain the distinction in meaning between motion sentences containing directional phrases which explicitly include the endpoint of the path, and those containing directional phrases which don't.

(4.29) John walked to the school.

(4.30) John walked towards the school / along the river / around the park.

In (4.30), in contrast to (4.29), John does not necessarily reach the school. This sentence can therefore not be considered a true resultative: there is no change of location to a specified point. Rather, it expresses that the activity of walking took place along a particular path. Levin and Rappaport Hovav (1995) do not make any distinction between these two types of directional phrases. A uniform treatment of these directed manner of motion constructions as resultative constructions ignores the possibility of non-resultative interpretations for them. Levin and Rappaport Hovav (1995) fail to explicitly address the issue of how and why certain directional phrases are interpreted as result predicates, while others are not.

There is in addition a distinction in the causative nature of the unergative resultative construction, as in (4.28), in comparison with the unaccusative resultative construction. The former requires a causative interpretation, while the latter explicitly does not. This additional required component of meaning of the unergative resultative construction becomes evident upon examination of cases in which the causation component is incompatible with the

[7] Levin and Rappaport Hovav (1995) do not provide any formal representation of meaning for the resultative constructions. Furthermore, they do not specify how path prepositional phrases, such as *to the store*, are reinterpreted as expressing a result state of being in some location. The discussion here assumes that this issue could be satisfactorily resolved on the L&RH account and that the representations used approximate the meanings L&RH assume.

meaning of the verb, and which are as a result infelicitous. The problem can be seen clearly in comparison with a set of sentences for a non-agentive verb, as in (4.31)-(4.32).

(4.31) The bottle floated to the bridge.

(4.32) ??The bottle floated itself broken.

The interpretation of (4.31) as would be predicted by L&RH is *The bottle's floating caused it to be at the bridge*, and that of (4.32) is *The bottle's floating caused it to be broken*. These two sentences seemingly should be equally felicitous on the given interpretations, but are not. I suggest that this is due to the association of different interpretations with these two constructions: (4.31) actually means *The bottle is at the bridge as a result of it floating* while (4.32) should mean *The bottle caused itself to be broken by floating*. The unergative resultative construction expresses a relation of causation (*cause-eff-rel*) between the event in the main clause (headed by the verb) and the state expressed by the resultative phrase, in which the subject of the main verb must be interpreted as the *causer* (i.e. the ACTor, the initiator of the causation). This is in contrast to the unaccusative resultative construction which does not express a causative relation.[8] Hence (4.31) is felicitous and sentences like (4.33)-(4.34) are interpreted, respectively, as *The river was solid as a result of it freezing* and *The vase was in pieces as a result of it breaking*.

(4.33) The river froze solid.

(4.34) The vase broke into pieces.

A verb which instantiates the unergative resultative construction, then, must be compatible with a causative interpretation. As the interpretation of (4.32) as dictated by the construction is semantically incompatible with the specifics of bottles and floating (how can a bottle floating cause breakage?), the sentence as a whole is infelicitous. On the other hand, the interpretation of (4.28) as *John caused his feet to be sore by (John) walking* is perfectly felicitous. This causation is a necessary component of the meaning of the unergative resultative — if the semantics of a particular verb or verb / argument phrase are incompatible with that meaning, the construction is infelicitous.

It could be argued that what is reflected by the contrast in the data for these two constructions is not a causativity relation but rather a volitionality constraint which applies to the unergative resultative construction but not the unaccusative resultative construction. The non-volitionality of bottles with respect to floating (4.32), etc. could then explain the infelicity of these events in the unergative resultative construction. However, consider the data in (4.35)-(4.37) below.

(4.35) John cried himself to sleep.

(4.36) John laughed tomato soup up his nose.[9]

[8]This difference in meaning is also implicit in the analyses Jackendoff (1990, ch. 10) gives of resultative constructions and of manner of motion verbs.

[9]This example is a modified version of a sentence spoken by Henry Thompson at lunch on 17/4/97: "Make me laugh hot tomato soup up my nose and you'll regret it." Thanks to Claire Grover for spotting it and passing it on.

(4.37) The ball squashed the tin flat.

In (4.35) and (4.36), *John* is not volitional with respect to the crying or the dreaming (if he were volitional, it should make sense to say *#John accidentally cried/laughed* but it does not) yet he does instigate a change of state to the state expressed in the resultative phrase by being the agent of the event which leads to the change of state. Similarly in (4.37) the agent of the main event which causes the change of state is clearly non-volitional as it is inanimate.

The difference in interpretation between these two kinds of constructions becomes even more clear if we consider a verb which seems to be possible in both the unaccusative resultative construction and the unergative resultative construction, as in (4.38).[10] On the L&RH account, such a verb is impossible since the lexical rule forces the manner of motion verb *swing* to be unaccusative in the presence of the directional phrase *apart*, ruling out its use in the unergative resultative construction. Both the unaccusative and unergative resultative sentences, however, seem to be equally felicitous in this example. Note that my judgement on the unergative resultative construction for these verbs differs from that which I assume Levin and Rappaport Hovav would assign, given the * they assign to the sentences I cited in (4.12) above.

(4.38) a. The clowns swung apart.

b. The clowns swung themselves apart.

The interpretations of these sentences are closely related, and in fact the meaning of (4.38a) is entailed by the meaning of (4.38b). The latter conveys the meaning of the former, with the additional suggestion that the clowns are actively swinging in such a way that they end up apart. That is, (4.38b) means *The clowns cause themselves to be apart by swinging*, and (4.38a) simply means *The clowns are apart as a result of their swinging*. Intuitively the unergative resultative construction conveys a certain meaning which the unaccusative resultative construction does not. The former could be used in a context, such as (4.39), in which the clowns are not agentive with respect to the swinging, and can therefore not be construed as *causers*, while the latter cannot.

(4.39) The director pulled the puppet strings and

a. the clowns swung apart.

b. *the clowns swung themselves apart.

Consider also the contrast in (4.40).

(4.40) a. The clowns swung over the net.

b. The clowns swung themselves over the net.

(4.40a) is highly ambiguous due to the nature of *over*: the PP could be expressing either the location of the swinging event, a point on or the direction of the path along which the swinging

[10]Thanks to Joan Maling for the suggestion of this example.

occurs, or a result phrase. On the other hand, (4.40b) is not ambiguous. The PP in that case can only be interpreted as a result phrase.

The treatment of these two constructions as identical semantically fails to capture the distinctions between them. In addition, the meaning of the unergative resultative construction does not seem to lend itself to a solely compositional treatment, since the element of causation conveyed by these contructions cannot be derived solely from the constituents of the sentence (the constituents parallel that in unaccusative resultative constructions, yet the meanings of the two constructions differ).[11] An adequate treatment of the interpretation of these constructions is missing in the Levin and Rappaport Hovav (1995) account of the resultative constructions. I will introduce a possible solution to this problem in Section 4.6, on the basis of Jackendoff (1990)'s adjunct rules and Goldberg (1995)'s analysis of these PPs in Construction Grammar, in which certain syntactic configurations have a particular semantic content not fully predictable from the semantics of their components.

4.3.2 Argument Structure

The shift to an unaccusative argument structure from an unergative argument structure induced by the lexical rule for verbs of manner of motion and verbs of sound emission is highly unintuitive in semantic terms and may have certain negative implications for the thematic structure of the verb.

The L&RH proposal of a shift in argument structure in the presence of a directional phrase means formally that each verb of manner of motion has two distinct entries in the lexicon. The verb *walk*, for example, will be assigned the representations shown in (4.41), with the corresponding surface projection of each argument structure. The verb in (4.41a) is unergative, with a single argument in the external argument position. The verb in (4.41b) is unaccusative, with an empty external argument position, an underlying direct object and a mandatory directional phrase (PATH).

(4.41) a. *walk*(j) John walked.

b. *walk*(__, **j**, [PATH]) John walked to the store.

The lexical rule which creates the entry in (4.41b) from (4.41a) therefore *demotes* the argument which fulfills the role of actor in the event expressed by the verb from an external argument to an internal argument, leaving the external argument position open.

The shift in argument structure is proposed by L&RH to account for certain syntactic phenomena as discussed above, but it may have certain implications for semantic interpretation. The demotion of an external argument to an internal argument position is an operation which has unclear side-effects in terms of thematic structure and lexical semantics. That argument must maintain its thematic role, as the the actor remains an actor even when the sense of the manner of motion verb is extended to a directed motion use. This argument structure demotion is therefore problematic under the assumption of linking rules which

[11]This observation is also made by Markantonatou and Sadler (1996) who argue that "the semantics of the construction is not constructed solely from the meaning of the verb itself" (pg. 5).

map semantic arguments to syntactic argument positions: in these cases we have the same semantic argument being mapped to two different syntactic argument positions. There is no clear semantic difference in the function this argument serves in the two cases which could be a distinguishing factor in the application of the linking rules. There would therefore be ambiguity in the potential argument structures for many verbs as dictated by the linking rules. Levin and Rappoport Hovav avoid these issues by not discussing any aspect of the (lexical) semantic representation of the verbs they are attempting to account for syntactically.

The directed motion sense of manner of motion verbs can be triggered by either an explicit prepositional phrase (4.42) or by the discourse context, as we saw above in (4.15). Furthermore, we noted in the previous section that this directed motion sense does not necessarily entail a change of location to a particular point and is therefore not necessarily strictly resultative.

(4.42) John walked to the store / around in circles / in place.

Because of these semantic properties, it seems rather unintuitive to assume that the underlying representation for a manner of motion verb may differ depending on whether a directional phrase is explicitly present or not. The *directed* manner of motion interpretation one gets of a manner of motion verb when it appears with a directional phrase is not a result of a sense shift in the meaning of the verb. Rather, it is a result of the composition of the meaning of the directional phrase with the meaning of the verb (cf. Dowty 1979; Moens and Steedman 1988; Jackendoff 1983, 1990; Pustejovsky 1995a). Motion verbs do not lexically encode a path, but they are also not incompatible with a path. In fact, it could be argued that due to their nature, motion verbs have a potential path argument. When these verbs appear with a directional prepositional phrase (PP), the PP extends the motion relation and thereby is interpreted as expressing the path which the motion follows. Directional phrases (such as *to the store* or *towards the park* in contrast with *in place*) which encode a path with an endpoint additionally indicate that the motion is following a path to a particular end location. Whether or not this end location is interpreted as the result location depends on the nature of the preposition heading the phrase. Thus all that is required to achieve the correct interpretation of the sentence is the integration of the meaning of the directional phrases with the verb semantics. Whether this meaning comes from context or from overt specification in the sentence is unimportant. The implication of these points is that manner of motion verbs on a directed motion use do not need to be collapsed with resultatives into a single treatment.

A lexical rule which extends the sense of manner of motion verbs in the presence of a directional phrase is unnecessary. The extended sense follows from the merging of lexical semantic properties of the verb with the meaning of the directional phrase. The criticism of the use of such a lexical rule is further strengthened by the observation that the context can trigger a sense shift of a verb, independent of the syntactic frame in which the verb appears. The lexical rule approach is thus semantically overrestrictive. The semantic evidence in the case of manner of motion verbs to a directed motion sense favors a non-lexical treatment of the sense extension.

4.3.3 Idiosyncrasy of Resultatives

The resultative construction is a highly idiosyncratic phenomenon (see also Section 4.2.2). Certain resultative phrases occur with great frequency with a range of verbs, suggesting that the sentences in which they appear are examples of a productive phenomenon. However, these examples are in fact instances of semantically restricted conventionalised phrases which constrain the productivity of the resultative construction. Consider the data in (4.43)-(4.44).

(4.43) a. i. He laughed himself to death.
ii. *He laughed himself dead.

b. i. He laughed himself to sleep.
ii. *He laughed himself sleepy / asleep.

c. i. He laughed himself out of a job.
ii. *He laughed himself jobless / unemployed.
iii. *He laughed himself out of the room / down the hall.

d. i. He laughed himself silly.
ii. He laughed himself faint / dizzy.
iii?? He laughed himself tired.

e. i. They laughed John out of the room.
ii. #They tittered John out of the room.
iii. #They laughed John into the room / down the hall.
iv. #They insulted John out of the room.

(4.44) a. i. He danced himself to fame.
ii. *He danced himself famous.

b. i. He danced his feet sore.
ii. *He danced his feet to soreness.
iii. ?He danced himself sore.
iv. *He danced himself crippled.

On the basis of these examples, it can be concluded that there are clearly specific lexical and semantic constraints on the resultative construction which must be identified and incorporated into any treatment which hopes to achieve a complete model of the constructions. There is no obvious way to incorporate such semantic constraints into the L&RH model given that their treatment concentrates purely on syntactic structure and does not involve lexical semantic features of individual words.

Even minor variations in some component of a felicitous resultative construction results in an infelicitous instance. The variations displayed in (4.43a-b) and (4.44a) are straightforward syntactic substitutions, where the resultative phrases have the same semantics but different syntactic form. Each variation conveys the same the result state. The general treatment Levin and Rappaport Hovav (1995) give of resultatives cannot account for such differences, as they

do not make any distinctions on the basis of the *form* of the resultative phrase. The *type* of the resultative (change of state or change of location) only affects the behaviour of the manner of motion verbs / verbs of sound emission and is therefore only an issue for these verbs. All other resultative constructions are seemingly allowed on their account, regardless of the form or type of the resultative phrase. Hence the approach also cannot account for the infelicity of the semantic substitutions in (4.43c-e) and (4.44b), in which the resultative phrases are exchanged for syntactically identical but semantically distinct, albeit closely related, phrases of the same syntactic type, a verb is interchanged with a semantically related verb, or a reflexive of one sort is exchanged for another reflexive.

The idiosyncratic nature of these constructions suggests that the L&RH view of resultatives as phenomena governed solely by syntactic constraints is inadequate to fully account for their behaviour. However, a purely semantic treatment of resultatives would also not be able to explain these substitutions. Syntactic substitutions with identical resultative interpretations could not be distinguished on a semantic account. Additionally, the semantic distinctions involved in semantic substitutions would need to be extremely fine-grained in order to adequately model the data above.

How can the restrictions on the productivity of this phenomenon be explained? The conventionality of the resultative construction must be acknowledged. Rather than assuming the existence of any systematic semantic constraints which govern contrasts such as those in (4.43c.i) vs. (4.43c.iii) and (4.43e.i) vs. (4.43e.ii), certain uses of the resultative construction seem to be restricted solely on the basis of partially lexicalised instances of the resultative construction, which, like idioms, allow little variation in their component parts. Thus some mechanism for encoding conventional constraints is needed.

Furthermore, pragmatics plays a role in the felicity of resultative constructions. Discourse coherence is essential. This was observed by Simpson (1983), who suggested that only predicates consistent with a change of state can occur in the resultative construction. I suggest that this "consistency" is relative to the discourse context in which the predicate appears. Thus the constrast between example (4.43e.i) and (4.43e.ii) can be explained on the basis of discourse coherence. The discourse must support a link between the cause and the effect expressed in the resultative construction. In the case of (4.43e.iii), there is no obvious way in which laughing at someone can cause that person to go into a room or down a hall, while it is more clear how the laughing can cause that person to go *out* of the room where people are laughing at him. This example also shows clearly discourse coherence constraint interacts with conventionality: there are certain cause-effect links which must be conventionalised in world knowledge.

Lastly, syntactic form must play a role in the interpretation of resultative constructions, as the contrast between the i. and ii. sentences in (4.43a-b) shows. Prepositional phrase and adjective resultative phrases seem to behave differently in the resultative construction, such that semantically equivalent resultative phrases in one syntactic form are felicitous in one sentence and infelicitous in others (also consider (4.44b.i-ii)). Thus reference to the syntactic form of the resultative phrases is also critical to the modeling of the resultative construction.

4.3.4 Conclusions

The Levin and Rappaport Hovav (1995) syntactic treatment of manner of motion and sound emission verbs as resultatives misses a semantic difference between these constructions and forces a change in argument structure which does not match the intuitive meaning conveyed by these verbs and has unclear consequences for linking. Furthermore, the evidence in favor of this account is not strong — the conclusions which can be drawn from the causativisation data are unclear and the semantic issues raised are quite serious.

The discussion in this section has pointed to several factors which need to be taken into account in an adequate model of these constructions: the semantics of the constructions, which may not be capturable in terms of traditional compositional semantic models, the avoidance of lexical rules in the treatment of the extended sense of manner of motion verbs, conventionality in the use of the constructions, discourse coherence constraints, and variations in the syntactic form of the constructions.

4.4 Wechsler

I turn next to a more semantically-based analysis of resultative constructions, proposed by Wechsler (1996), in which the syntactic mechanisms are handled by HPSG (Pollard and Sag 1994) and the behaviour of different verbs with respect to the resultative construction depends on background information encoded in the lexical entries of the verbs.

4.4.1 Wechsler's approach

Wechsler (1996) suggests an analysis of resultatives which does not rely on the syntactic encoding of unaccusativity / unergativity. In fact, he explicitly rejects this encoding as an adequate basis for explaining the distribution of data in the resultative construction. Instead, he draws a distinction between types of resultative constructions which parallels the control / raising distinction for complement-taking verbs (see e.g. Pollard and Sag 1994, ch. 3). Wechsler argues that differences between the semantic restrictions that *control resultatives* impose on their complement resultative phrase and those that are imposed by *raising resultatives* can be used to account for the resultative data. He provides a formalisation of the distinction between these resultative types in HPSG, and shows how his analysis can be used to account for a range of resultative data.

Briefly, *control resultatives* are resultatives in which the resultative phrase is predicated of an argument of the main verb, (4.45), and *raising resultatives* are resultatives in which the resultative phrase is predicated of something which is not an argument of the main verb, (4.46).[12]

(4.45) a. John hammered *the metal* flat.

[12]This distinction corresponds to a binding difference in Jackendoff (1990)'s Resultative Adjunct rule: the patient role in the meaning expressed by the resultative phrase may be bound to the patient of the verb (a 'control' resultative) but need not be (a 'raising' resultative).

b. *Sally* jumped free of the car.

(4.46) a. The joggers ran their Nikes threadbare.

b. The audience laughed the speaker off the stage.

The main distinction between these two types of resultatives according to Wechsler's analysis is that control resultatives place semantic restrictions on their complement resultative phrases, while raising resultatives do not. Specifically, the resultative predicate in a control resultative must express a "canonical or generic result (or intended result) of the action denoted by the verb" (Wechsler 1996, p. 2). The nature of this canonical result is encoded in the lexical semantics of the verb. The resultative phrase is required to unify with the encoded result through the mechanisms of the formalisation. In contrast, raising resultatives do not specify any canonical result or other specification in their lexical semantics and therefore do not constrain the kinds of resultative phrases they can appear with. This distinction is supported on the basis of data such as that in (4.47), which shows control resultatives to be acceptable only with result phrases that are compatible with the canonical result of the action expressed by the main verb, and (4.48), which shows that the result phrases appearing with raising resultatives do not need to have any connection to the meaning of the main verb.

(4.47) a. *Robert* ran clear of the car / *exhausted. [Wechsler 1996, p. 2, (6a)]

b. John hammered *the metal* flat / *safe. [Wechsler 1996, p. 2, (6c)]

(4.48) a. Olof painted himself into a corner. [Wechsler 1996, p. 2, (7b)]

b. We laughed ourselves silly. [Wechsler 1996, p. 2, (7d)]

Wechsler captures his insights in HPSG by utilising the BACKGROUND feature in verbal *synsem*s which encodes pragmatic information relevant to the proper use of the word. This feature has a value which may contain the feature TELOS, which in turn has a value representing the result state or endpoint of the event conveyed by the verb. Control resultatives have the TELOS value specified while raising resultatives do not. Resultatives are licensed via a lexical rule which adds resultative phrases (APs or PPs) to a verb's SUBCATegorisation list and requires the unification of the CONTENT value of the resultative phrase with the TELOS value of the verb. If that TELOS value is unspecified, then no restrictions are placed on the semantic content of the resultative phrase.

Wechsler's use of a unification-based grammar formalism such as HPSG allows him to easily specify interactions between the semantics of the main verb and the semantics of the result phrase in the resultative. In particular, control resultatives can indicate in their lexical semantics precisely what role their arguments must play in the canonical result state they are associated with. This is achieved via structure-sharing between the elements in the core semantics of the verb and elements in the value of the TELOS feature. It ensures that resultative phrases are predicated of the appropriate verbal argument in control resultatives, and enables Wechsler to account for the data in (4.49) (Wechsler 1996, p. 11, (29a-b)), in which the result phrase is predicated of the subject of a transitive verb. These sentences would violate L&RH's Direct Object Restriction in (4.4), as the resultative is not predicated of the direct object as required by the DOR. L&RH would therefore be unable to account for these sentences.

(4.49) a. *The wise men* followed the star out of Bethlehem.

b. *The sailors* rode a breeze clear of the rocks.

Syntactic issues associated with the resultative construction are mainly handled by the mechanisms of HPSG of which Wechsler takes advantage, with the addition of a lexical rule which adds the non-argument NPs which serve as the subjects of resultative predicates in raising resultatives (the *Raising Rule*). This addition, in conjunction with the control / raising resultatives distinction, leads to an account of the appearance of fake reflexives for resultative constructions with many unergative intransitives, (4.50), in contrast with the lack of reflexives for other unergative intransitives, specifically the manner of motion verbs (4.51), which behave as control resultatives.

(4.50) a. * *The dog* barked hoarse. [Wechsler 1996, p. 4, (11a)]

b. The dog barked itself hoarse. [Wechsler 1996, p. 4, (11b)]

(4.51) *John* ran to the store.

In (4.50a), Wechsler argues, hoarseness cannot be a canonical result of barking and thus this result state is compatible only with a raising use of *bark*. In this case, the Raising Rule is triggered and the requirement for a post-verbal NP is added. Semantic considerations will determine that the NP is token-identical with the subject, and syntactic constraints will ensure that it appears as a reflexive. In contrast, all manner of motion verbs optionally encode a location in their TELOS value. Thus the result state in (4.51) is compatible with the canonical result of *running*, and so this is a control resultative which does not trigger the addition of any non-argument noun phrases. Instead the single argument of *run* is structure-shared with the subject of the resultative predicate due to unification between the verb's TELOS value and the CONTENT of the result phrase.

This treatment does not give unaccusativity a role in the modeling of the resultative construction, thereby avoiding many of the pitfalls of the Levin and Rappaport Hovav approach. In particular, the treatment avoids any notion of a lexical shift in the underlying syntactic classification of certain verbs and is therefore not subject to many of the criticisms given in this chapter. There are, however, other difficulties, to which I now turn.

4.4.2 Discussion of the approach

The problem of the 'canonical end state'

Wechsler's treatment of control resultatives assumes the representation of a canonical end state associated with the activity expressed by the verb. This assumption is also made by Markantonatou and Sadler (1996), who suggest that a resultative phrase can be inferred for all verbs which contribute a "theme" argument, and Rappaport Hovav and Levin (1996) who suggest the content of the resultative phrase must be construable as part of the prototypical event described by the variable. The assumption is difficult to implement, however, given the restrictedness and context-dependence of such supposedly 'canonical' end results.

The default end result which one would be inclined to associate with a particular event depends on the participants in that event. Consider the set of sentences in (4.52).

(4.52) a. John hammered the metal. (flat, into a ball)

b. John hammered the porcelain. (to dust, cracked)

c. John hammered the fruit. (to a pulp)

For each different kind of object which John hammers, a different default end result would be inferred. It is true that in each of the above examples that result is a state. It is therefore conceivable that all that should be encoded in the TELOS should be that the result is some state. In that case a pragmatic constraint would clearly have to be recruited to account for object-dependent variations in the construction of felicitous resultative sentences, as shown by the contrast between (4.52) and (4.53).

(4.53) a. ??John hammered the metal to dust.

b. ??John hammered the porcelain to a pulp.

c. ??John hammered the fruit flat.

Even simply encoding that the canonical result for *hammer* is a state would be too restrictive, in light of examples such as (4.54), in which the resultative phrase adds a change of *location* rather than a change of *state*.

(4.54) a. John hammered the nail into the wall.

b. John hammered the board out of the window / off of the slide / onto the ground.

Also consider the sentences in (4.55). (4.55a) primarily specifies the creation of an object (*a portrait of John*) rather than the covering of a surface by paint, in contrast to *John painted the wall*, while (4.55b) is a resultative which specifies not a change of state of the painted object (*a star*), but the location at which the stars are painted. Specifying that the canonical endpoint of the painting event is for the painted object to change state can't reflect the full intended meaning of such sentences.

(4.55) a. I painted a portrait of John.

b. I painted a star on each of my Christmas cards.

Thus the assumption of a canonical end state buys us little in accounting for the restrictedness of the resultative construction: no encoded canonical end state can be made specific enough to account for restrictions which arise from the use of the verb in a specific sentence, nor general enough to account for the full range of possible felicitous uses of the verb in the resultative construction. I will argue in Sections 4.4.3, 4.5, and 4.7.4 that the application of a pragmatic constraint to the construction better models the data.

Furthermore, Wechsler's and Markantonatou and Sadler's motivation for the representation of the canonical result in terms of potential telic interpretations of the verbs is suspect, since it depends on an analysis of aspect driven by lexical representations. Consider Wechsler's (1996, p. 6) discussion of the representation of the word *paint*:

> Not all painting events have a definite end point, so the value for TELOS is optional (indicated by parentheses). But painting events *can* be telic; and if a painting event *is* telic, then the result state is canonically a state involving the state of the PAINTED participant.

Painting events can be both atelic and telic, as is clear from the two variants in (4.56). The first sentence has an atelic main clause (with no inherent endpoint) while the main clause is telic (with an inherent endpoint) in the second, according to the standard Vendlerian tests.

(4.56) a. John painted the wall for two hours.

b. John painted the wall in two hours.

The most salient interpretation for *John painted the wall* seems to be the telic one, but this generalisation does not hold across all verbs which can function as control resultatives. Consider the sentences in (4.57), with my grammatical judgements. The sentence in (4.57c) is from Markantonatou and Sadler (1996, p. 13).

(4.57) a. The blacksmith / child hammered the metal.

b. The blacksmith / child hammered the metal for three days.

c. ?The blacksmith hammered the metal in three days.

d. #The child hammered the metal in three days.

The telic interpretation of (4.57a) is far more marked than the atelic interpretation, and the questionable felicity of the telic variants suggests that the representation of a generally available "canonical end state" for the activity of hammering is overly strong.

The question of how the telic interpretation of such sentences is arrived at must be carefully considered. It is by no means obvious that the (potential) result state should be represented in the lexical entries of these verbs. The aspectual category of a proposition can be affected by adverbials (Moens and Steedman 1988), and even by the kind of objects which serve as a verb's arguments (contrast *John built a house in a day* with **John built houses in a day* (Verkuyl 1989, Krifka 1989). In short, to quote Moens and Steedman (1988, p. 17), "aspectual profiles are properties of sentences *used in a context:* sense-meanings of sentences or verbs in isolation are usually compatible with several (or even all possible) Vendlerian profiles".

It even seems that our world knowledge plays a role in determining the appropriate aspectual category for a proposition. Consider the sentences in (4.57) again. The difference in felicity between (4.57c) and (4.57d) is that for the former we are able to infer an end state for the metal and interpret the "hammering the metal" event as telic on the basis of our knowledge that that part of what a blacksmith does is to hammer metal into particular shapes, i.e. to make tools out of metal. In the case of a child, on the other hand, we don't have any knowledge of a "standard" end state which his / her hammering could lead to and so the sentence is odd. Pragmatic reasoning is again seen to interact with the interpretation process. The assumption of a lexically-encoded telic interpretation for a verb cannot adequately accommodate these pragmatic influences.

It could be argued that the independent motivation for the lexical representation of a canonical end state comes from the semantics of the verbs, that verbs like *laugh* can only appear as a raising resultative because they are never associated with any canonical end result while verbs like *hammer* and *run* can appear as control resultatives because they do have telic interpretations. However, the telic interpretation of *hammer* seems only to come from the fact that it can appear in a resultative construction. Similarly, the telic interpretation of manner of motion verbs like *run* arises only when it appears with telic prepositional phrases or telic arguments (e.g. *run a mile*). Thus it seems that the syntactic constructions in which a verb can appear determine the range of its interpretation. This means that the telic interpretation of these verbs stems from properties of a particular construction and properties of elements like PPs and arguments interacting with the core meaning (semantic relation) expressed by the verbs. It suggests that the syntactic context in which a verb is used can influence its potential interpretation, and that therefore the semantics of a verb in different constructions cannot be considered an independent criteria for determining the range of lexical entries it should have. The notion of canonical end state simply cannot be isolated from the syntactic constructions in which a telic interpretation is licensed.

The problem of the optional TELOS feature

The Wechsler (1996) lexical representations pre-encode the behaviour of particular verbs in the resultative construction. If a verb can appear in a resultative construction in which the resultative is predicated of a verbal argument, then the verb must have a TELOS value specified. There does not seem to be any independent basis proposed by Wechsler for determining when a particular verb has a canonical result specified in its BACKGROUND:TELOS feature.

The pre-encoding becomes clear upon examination of Wechsler's detailed illustration of the behaviour of the verb *run* within his approach (Wechsler 1996, Appendix, p. 14-16). In this analysis the sentence *We ran our Nikes threadbare* is analysed using a variant of the lexical entry for *run* which does not have a TELOS value specified, which then leads to the addition of the non-argument NP to the verb's SUBCAT list via the Raising Rule. The sentence *John ran into a tavern*, in comparison, is analyzed with a lexical entry for *run* for which the TELOS value is specified as a location. So we have one verb which can appear in two different resultative constructions, and as a result we have two variants of its lexical entry (cf. Jackendoff 1990, to appear; Goldberg 1995). A verb like *laugh*, on the other hand, is only ever used as a raising resultative and hence only has one lexical entry, with no TELOS value specified. The absence of an additional entry for *laugh* with a value for the TELOS feature is precluded only on the basis of the lack of examples of this verb in the control resultative construction.

World knowledge of situations involving the relations expressed by particular verbs could be considered an independent basis for encoding a canonical result, but we saw above with reference to the examples (4.52) and (4.53) that these 'canonical' results can be highly context-specific. This will either lead to generation of individual lexical entries for a verb corresponding to distinct contexts, which is heavily redundant and fails to capture any generalisations, or specification of vague canonical results which will result in overgeneration. I therefore see no adequate justification for having both raising resultative and control resultative lexical

entries for a particular verb, independent of that verb's behaviour. This pre-encoding undermines the explanatory adequacy of Wechsler's theory.

The problem of idiosyncrasy

The main issue which Wechsler's approach aims to explain is the conditions under which the different resultative constructions can be used, in particular the semantic restrictedness of result phrases in control resultatives. This, however, is somewhat suspect when the idiosyncrasy of the resultative construction, as discussed in Section 4.3.3, is taken into consideration. We showed above (Section 4.4.2, example (4.54)) that a verb which appears to specify a stative canonical end result can felicitously appear with a locative end result. That example therefore violates the predictions of Wechsler's account. In addition, the differences in acceptability between (4.58a-b) and (4.58c-e) and the contrast in (4.59) suggest that even the result phrase in a raising resultative is not totally unrestricted, as is argued by Wechsler. The conventionalised restrictions on the phenomenon creates difficulties for the general application of the semantic claims about the resultative construction which Wechsler makes. The result phrases in control resultatives do not seem to be consistently predictable on the basis of background knowledge, while those in raising resultatives must be somewhat constrained, at least by pragmatic factors.

(4.58) a. Robert ran himself into the ground.

b. Robert ran his shoes threadbare.

c. *Robert ran himself into the shower.

d. ?Robert ran his shoes to pieces.

e. *Robert ran himself to pieces.[13]

(4.59) a. The audience laughed the speaker out of the room.

b. *The audience laughed the speaker silly / embarrassed.

Semantics

There are several unresolved semantic issues in Wechsler's treatment as well. As discussed in Section 4.3.1, there are subtle differences in the meaning of directed manner of motion constructions and that of other resultative constructions. Furthermore, there are differences in meaning between resultative constructions and the standard syntactic frames in which the verbs appear. In particular, as pointed out by Levin and Rappaport Hovav (1995, p. 54), the resultative construction must incorporate a causal relation. In Wechsler's approach, resultative phrases are added to a verb's SUBCAT list, and the CONTENT of that phrase is unified with the TELOS value, as a result giving the sentence a telic interpretation, but no elements of meaning are added. So the meaning of *We ran our Nikes threadbare* is represented as, essentially, *We ran and at the end our Nikes were threadbare*, rather than *We caused our Nikes to be threadbare by running*. This problem is remedied by Wechsler (p.c.) through reinterpretation of

[13] I am assuming a literal interpretation of *run* here, not a metaphorical one.

the TELOS feature as involving causation. Even when the causative component is incorporated, however, the problem of distinguishing between the semantics of verbs of manner of motion on a directed motion use and the semantics of other resultatives remains, due to the uniform semantic treatment of all resultatives on this account. That is, if I am correct in asserting that the directed manner of motion verbs do not induce causation while standard resultatives do, Wechsler cannot account for the semantic distinction between them.

It is furthermore unclear that all instances of raising resultatives are actually truly raising, in the sense that they do not assign a semantic role to their NP complement. Consider again (4.48), repeated here, and contrast these examples with (4.60).

(4.48) a. Olof painted himself into a corner. [Wechsler 1996, p. 2, (7b)]

b. We laughed ourselves silly. [Wechsler 1996, p. 2, (7d)]

(4.60) a. Olof painted the dog into a corner.

b. *We laughed Billy / the guests silly.

The variant of (4.48a), (4.60a), is acceptable, while the variant of (4.48b), (4.60b), is not. In the latter example, the referential identity of the NP complement and the subject is crucial to the acceptability of the example, and hence there is clearly something controlling the nature of the unsubcategorised complement. I suggest that the form of this resultative is conventionalised, and this conventionalisation contrains the general applicability of the mechanisms proposed by Wechsler.

4.4.3 Conclusions

In spite of my criticisms of the implementational details of Wechsler's (1996) approach to resultatives, in particular of the requirement of extra lexical entries specifically for the purpose of capturing resultative behaviour and the lack of independent justification for the specification of a BACKGROUND:TELOS value, Wechsler's intuitions about the semantic constraints on the different resultative constructions do go much further towards capturing the apparent semantic restrictions on this construction than the Levin and Rappaport Hovav explanation (in fact this is unsurprising given that L&RH focus on syntactic variations rather than semantic restrictions). This seems to be particularly true for the case of control resultatives in which the resultative phrase is predicated of the direct object of a transitive verb. Consider the sentence in (4.47b), repeated here as (4.61a), the further variations of the sentence in (4.61b) (Wechsler 1996, p. 2, (6c)), and sentences (4.62). In each case, the acceptable resultative phrases do seem to make specific some canonical goal of the activity.

(4.61) a. John hammered *the metal* flat / *safe / *red.

b. John hammered *the metal* into a triangle / smooth / shiny / into the ground.

(4.62) a. John chiseled *the ice* smooth / into a bird / onto the floor.

b. John chiseled *the ice* *cold / *shiny.

Wechsler's presentation of the treatment of resultative constructions in terms of the TELOS feature implies that there is a semantic generalisation about the inherent goal-oriented nature of certain events which is derivable from general world knowledge. I have shown in the preceding sections that this generalisation cannot be isolated from the syntactic construction for which it is an issue, and that the semantic restrictedness of the construction cannot be accounted for entirely in terms of this generalisation.

Rather than assuming that these verbs specify a highly specific canonical result (the examples above and in Section 4.4.2 show clearly that this canonical result cannot simply be required to be a *state* as this is not restrictive enough) which cannot independently be motivated, then, it seems to make more sense to argue that the semantic restrictedness hinges in part on the requirement of establishing a causal relation between the main event expressed in the sentence and the result state. This enables us to account for sentences such as (4.61)-(4.62), and would also allow us to account for situations in which pragmatics could ameliorate a particular sentence. For example, in a context such as (4.63), the sentence (4.63c), infelicitous in a context-neutral environment, seems to be completely felicitous. Here the context allows us to establish a natural causal relation between hammering and the metal being safe. I will pursue this idea further in Section 4.7, in light of the approach suggested by Construction Grammar (Goldberg 1995).

(4.63) a. The slide at the park had a section which had come loose.

b. Several children had hurt themselves on the protruding edge.

c. In order to prevent further injuries, John hammered the metal safe.

4.5 Construction Grammar

I have pointed out at several points in this chapter that the resultative construction seems to be governed in part by conventionalised constraints. A framework is needed in which the conventionalised nature of the resultative construction can be captured. To serve this purpose, I will adopt certain principles of the Construction Grammar approach to grammar. These are outlined below with respect to the analysis of the resultative construction proposed by Goldberg (1995). I will in addition assume the existence of standard compositional interpretation mechanisms, which will be applied whenever possible to maintain maximum flexibility and generality in the interpretive process (e.g. to handle standard adverbial modification).

The Construction Grammar approach assumes that form-meaning correspondences are the basic units of language. Within this framework, constructions exist independently of the particular words which instantiate them. Each construction has a specific syntactic configuration which is associated with a specific semantics (and possibly certain pragmatic properties) — each construction specifies the semantic roles of different syntactic positions and relations between the roles, and the semantics of the words which appear in the construction must fuse with the semantics of the construction itself. This approach is similar in spirit to Jackendoff's (1990, 1996) analysis of certain phenomena, including the resultative.

(4.64) a. Caused Motion Construction

e.g. *John sneezed the tissue off the table.*

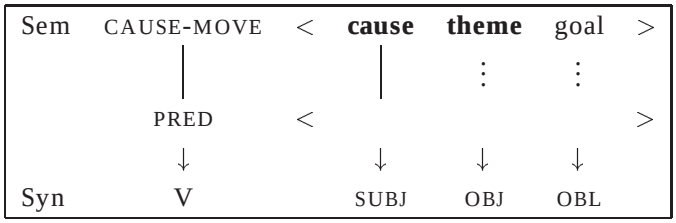

I_S

Intransitive Motion Construction

e.g. *John ran to the store.*

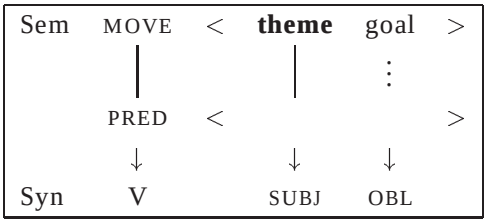

b. Resultative Construction

e.g. *John hammered the metal flat.*

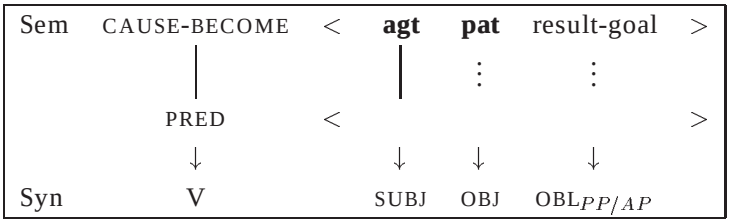

I_S

Intransitive Resultative Construction

e.g. *The river froze solid.*

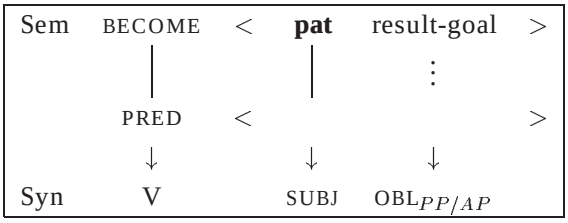

The analysis of the resultative construction given by Goldberg (1995) does not treat all of the constructions discussed elsewhere under this name as uniform. In fact, she isolates four different constructions (cf. Jackendoff 1990), which are divided into two groups: the *Intransitive Motion Construction* which is a subpart (in terms of syntactic and semantic specifications) of the *Caused Motion Construction*, and the *Intransitive Resultative Construction* which is a subpart of the *Resultative Construction*. The four constructions and their specifications are summarised in (4.64), where OBL stands for *oblique*, solid lines between semantic roles and roles in the PREDicate's role array indicate that the semantic role must be fused with an *independently existing* verbal participant role, and roles represented in bold are *profiled* arguments — that is, entities in a verb's semantics that are "obligatorily accessed and function as focal points within the scene, achieving a special degree of prominence (Langacker 1987)" (Goldberg 1995, p. 44). The I_S label on the arrow between constructions indicates that the construction at the head of the arrow is a subpart of the construction at the tail.

The existence of these various constructions is motivated on the basis of data in which the semantic interpretation of a sentence as a whole cannot be attributed to the meaning of the main verb or any composition of the meanings of the sentential components. I will not delve deeply into the arguments Goldberg presents that the meaning of these constructions does not straightforwardly come from the meaning of the words in the constructions (cf. the discussion of the interpretation of manner of motion verbs in Section 4.3.2). I simply cite one of her strongest arguments against the adequacy of strictly lexical compositionality for determining this meaning (Goldberg 1995, p. 154-155):

> [I]t is fallacious to argue that because we may be able to pragmatically infer the meaning of a construction, its existence is therefore predictable rather than conventionalised. Such reasoning is based solely on a model of interpretation, yet we must also account for production.

The relevance of this point is that it may be possible for a purely compositional analysis of the resultative construction(s) to account for the *interpretation* given to of one of these constructions, but it cannot explain how the constructions are licensed, given that the resultative phrase is not normally considered to be subcategorised by the verb (as it does not play a clear role in the verbal relation) yet does not behave as an adjunct (as it influences the "internal" semantic relation expressed in the main clause of the sentence). In fact, this criticism can be applied to the Levin and Rappaport Hovav (1995) analysis of unergative resultatives, as that analysis does not explain *why* it is possible to predicate something of an entity which isn't an argument of the verb. This point poses a further challenge in the attempt to model these constructions: there must be some theory of why particular syntactic configurations are possible. The Construction Grammar approach provides this explanation directly through the assumption of form-meaning pairs which license and constrain particular syntactic patterns.

The interpretations Goldberg associates with the various constructions should be clear from the representations she postulates for them, as presented in (4.64). They can be summarised as in (4.65) (compare with the meanings identified in Section 4.3.1 and summarised below in (4.78)).

(4.65) a. Caused Motion: $NP_{subject}$ causes NP_{object} to move to Result Location

b. Intransitive Motion: $NP_{subject}$ moves to Result Location

c. Resultative: $NP_{subject}$ causes NP_{object} to become Result State

d. Intransitive Resultative: $NP_{subject}$ becomes Result State

Each different construction, then, has a meaning independent of the others and therefore different thematic constraints. Furthermore, each has a meaning independent of the words which instantiate it. Instantiation-specific meanings result from the merging of the meaning of the construction with the meaning of the main verb, which is seen as an instance of the main semantic relation (CAUSE-MOVE, MOVE, etc.) or as specifying the means by which that semantic relation is accomplished. For example, in *John hammered the metal flat*, *hammer* specifies the means by which the CAUSE-BECOME relation is achieved and so the meaning of this instantation of the resultative construction is *John caused the metal to become flat by means of hammering*.

Another property of this account is that certain semantic restrictions on the instantiation of these constructions can be shown to fall out from the thematic constraints encoded in the construction definitions. The sentences in (4.66) are ruled out because the object of the verb must be construable as a patient in this construction, but cannot be construed as such in those sentences. The objects fail the traditional test for patienthood (*What Sue did to the monster was watch it/*What happened to the book was Sue kept it) and so the sentences are not compatible with Goldberg's Resultative Construction.

(4.66) a. *Sue watched the monster to sleep.
on interp.: Sue caused the monster to fall asleep by watching it

b. *Sue kept the book dirty.
on interp: Sue caused the book to become dirty by keeping it

Furthermore, the restriction that among the only unergative verbs which can appear in the "unaccusative resultative construction" (using the label L&RH give to the particular syntactic construction of NP+Verb+ResP) are the manner of motion verbs falls out from the definition of the Intransitive Motion construction: the predicate must be an instance or means of MOVE-ing, which manner of motion verbs clearly are and very few other verbs are. This leaves, however, the verbs of sound emission which can appear in this construction unaccounted for.

Goldberg (1995) also identifies general semantic constraints which help to account for the apparent idiosyncrasy of these constructions. These are summarised (4.67) below.

(4.67) a. Constraints on the Caused Motion construction

1. The causer must be an agent or natural force (i.e. not an instrument).
2. The caused motion along the path must be directly caused.[14]

b. Constraints on the Resultative construction

[14]See Goldberg (1995) for further details about what it means for something to be directly caused. The notion is relevant here to the extent that pragmatics can play a role in determining direct causation.

1. The agent must be animate.
2. The change of state must occur simultaneously with the endpoint of the action denoted by the verb.
3. The resultative phrase codes a clearly delimited endpoint (is an end-of-scale adjective).
4. The resultative phrase cannot be a deverbal adjective.

These constraints restrict the instantiation of the constructions. Thus the constraints in (4.67a) can be shown to account for the differences in (4.68)-(4.70); in particular via the agent / natural force constraint (4.68) and the constraint that the motion must be directly caused by the action (the motion in (4.69b) is secondary to the shooting and there is no obvious causation between laughing and causing someone to get into his car). Similarly the constraints in (4.67b) can account for the differences in (4.71)-(4.72): (4.71) due to the animate agent constraint and (4.72) due to the end-of-scale constraint.

(4.68) a. The sound waves blasted the dust off the table.

b. *The loudspeaker blasted the dust off the table.

(4.69) a. Pat shot the bullet across the room.

b. *Pat shot Sam across the room. *(unacceptable on the interpretation that Pat shot Sam and the bullet forced him across the room)*

(4.70) a. They laughed the poor guy off the stage / out of the auditorium.

b. ?They laughed the poor guy into his car.

(4.71) a. She slept herself sober.

b. *The feather tickled her silly.

(4.72) a. He drank himself drunk / sick / dead.

b. ?He drank himself a little sick.

c. *He drank himself funny / happy.

Several of these constraints, however, could be explained by more general pragmatic principles which govern the use of these constructions. I would like to suggest that rather than assuming constraints on the semantics of the construction itself, it is preferable to assume that these constraints apply at the pragmatic level, to determine the felicity of the utterance within a discourse context, as the constraints derive from general pragmatic principles rather than construction-specific properties. I will explain how in what follows.

As suggested in Sections 4.3.3 and 4.4.3, the causation expressed in the resultative constructions must be coherent within the discourse context. The relationship between the causing eventuality and the caused eventuality must be clear from either world knowledge or discourse context. I refer the reader to example (4.63) and the preceding discussion, and the analysis of that example found in Section 4.7. Pragmatic reasoning there is seen to be triggered by requirements derived from the (fixed) interpretation dictated by the construction.

Furthermore, the constraints as expressed by Goldberg are defeasible. Consider for example (4.73), in which the instrument of the causation is specified rather than the agent, in apparent violation of the constraint in (4.67a.1). Relevant world knowledge or an appropriate discourse can seemingly license violation of Goldberg's proposed constraints. This defeasibility is a typical property of pragmatic constraints, in contrast to semantic constraints which should not be dependent on pragmatic reasoning.

(4.73) a. The feather excited her into a frenzy.

b. The work pushed him to the brink of insanity.

These observations then lead to a reinterpretation of the constraints on these constructions in terms of pragmatic principles. Grice's Maxim of Quantity (Grice 1975), for example, provides the basis for the direct causation constraint in (4.67a.2). This Maxim states *(i) make your contribution as informative as is required for the current purposes of the exchange* and *(ii) do not make your contribution more informative than is required*, that, for example, a speaker will not use a construction whose primary semantic content is caused motion along a path if he wishes to focus on an aspect of an event other than caused motion. Verbs which express a primary effect other than motion, like the sense of *shoot* in (4.69b), would therefore not be used in this construction. Thus I would argue that the direct causation constraint is actually nothing more than a restatement of this pragmatic principle, specified to the case of the Caused Motion construction. The other constraints in (4.67) (with the possible exception of (4.67b.4) which seems to be an idiosyncratic syntactic constraint) can similarly be argued to follow from the need to firmly establish the causal relation during pragmatic reasoning.

Although the Construction Grammar approach is highly appealing, not all syntactic patterns are suitable for treatment via constructions (see Jackendoff 1996 for a summary). Instances where the meaning of a sentence can be accounted for entirely in terms of lexical composition and where its syntax is licensed by other generative aspects of the grammar should probably not be handled via constructions. There certainly seem to be relevant instances in which a constructional approach can be overrestrictive, or at least misses some more general syntactic or semantic process underlying the generation of a particular phrase. Consider the sentences in (4.74), for example.

(4.74) a. John danced mazurkas across the room.

b. John walked the dog to the store.

c. John swam laps to exhaustion.

d. The children played leapfrog across the park.

These sentences correspond approximately to Goldberg's Intransitive Motion construction in terms of their basic meaning, yet do not fit this construction syntactically due to the presence of a verbal noun phrase complement. That noun phrase complement makes the activity expressed in the verb more specific — John didn't just dance in haphazard fashion across the room, he danced mazurkas — but the basic interpretation of motion along a path is preserved. Syntactically, these examples fit the Caused Motion Construction, but the noun phrase object

can hardly be construed as a theme, since the noun phrase object in each instance is not the entity involved in the motion expressed by the verbal predicate (*mazurkas* and *leapfrog* do not change position), so this construction cannot be applied either. Goldberg would be forced to assume a constructional variant of the Intransitive Motion Construction which licenses the noun phrase complement. This, however, seems unnecessary when a straightforward compositional analysis can account for both types of sentences, as will be shown in Section 4.7.2.

Other examples which Goldberg's analysis would have difficulty accounting for are found in (4.75)-(4.76).

(4.75) a. The pebbles rolled smooth.

b. John ran to exhaustion.

(4.76) *Mildred exercised into the room.

In (4.75), a manner of motion verb appears with a non-locative resultative phrase. The interpretation is not an instance of the Intransitive Motion construction (CAUSE-MOVE), but rather the Intransitive Resultative construction (CAUSE-BECOME). Because of Goldberg's requirement of compatibility between the predicate expressed by the verb and that expressed by the construction, however, it is impossible for a motion verb to be licensed by the Intransitive Resultative construction. (4.76), in contrast, would be compatible with the Intransitive Motion Construction since it contains a MOVE verb and therefore would incorrectly be predicted to be grammatical. Again the problem hinges on the compatibility between the predicate expressed by the verb and that specified by the construcition.

4.6 Criteria which the solution must satisfy

As has become apparent through the preceding discussion, any solution to the problem of modeling the resultative construction and the use of manner of motion verbs must meet several criteria. These can be summarised as follows:

- **Syntax:** What licenses the construction must be clearly identified, since their generation must be accounted for as well as their interpretation.
- **Syntactic constraints:** The restriction of the syntactic form of the "unaccusative resultative construction" (Levin and Rappaport Hovav 1995) to only certain classes of unergative verbs, specifically the verbs of manner of motion and verbs of sound emission, must be accounted for.
- **Semantics:** The appropriate semantics of each of the variations of the resultative construction must be captured.
- **Conventionality and Pragmatic felicity:** The semantic restrictedness of the phenomena must be captured (i.e. an explanation must be provided for why certain result phrases cannot combine with certain verb phrases to produce a felicitous resultative, and why semantically similar items cannot always be felicitously substituted for one another in the resultative constructions).

We will discuss each of these in more detail in the subsequent sections, pointing out where the previous proposals fall short of the criteria. In what follows, we isolate two main categories of resultatives: what Levin and Rappaport Hovav (1995) call the unaccusative resultative construction, in which an intransitive verb is directly followed by a resultative phrase (see Section 4.2), and the unergative resultative construction, in which a resultative phrase appears after, and is predicated of, a post-verbal NP (I will continue to use the L&RH terminology for these constructions for ease of reference to the syntactic structures, although I do not take on the notion implicit in these terms that unaccusativity should play a part in the analysis). The basic syntactic structures are therefore as shown in (4.77), where $NP_{controlled}$ is the noun phrase of which the resultative phrase (ResP, either a PP or an AP) is predicated. Each of these construction categories is further subdivided in order to identify more specific properties of the sentences which instantiate them (cf. differences identified in Jackendoff 1990).

(4.77) a. **$NP_{\{subject,controlled\}}$ V ResP**

i. V is a true unaccusative (e.g. freeze)

ii. V is one of the manner of motion verbs or a verb of sound emission

b. **$NP_{subject}$ V $NP_{controlled}$ ResP**

i. $NP_{controlled}$ is a fake reflexive, coindexed with the subject $NP_{subject}$, and V is an unergative intransitive.

ii. $NP_{controlled}$ is the object of a transitive verb

iii. $NP_{controlled}$ is neither subcategorised by V nor a fake reflexive

4.6.1 What licenses the constructions?

As pointed out by Goldberg (1995) and introduced in Section 4.5, an analysis of the resultative construction is incomplete without some account of the existence of the construction. Thus analyses which define the syntactic and semantic relationships between the constituents of the contruction yet avoid the issue of why those relationships might be possible are inadequate.

Goldberg's Construction Grammar directly addresses this issue by defining constructions as components of grammar in their own right — certain syntactic constructions, such as the resultative construction, exist because there is a form-meaning pair in the grammar which licenses the construction. Similarly, the Wechsler (1996) approach licenses the construction via a lexical rule which adds the resultative phrases and ensures appropriate semantic integration, and the more specific Raising Rule which also adds a complement NP to the subcategorisation lists of the verbs which appear in the construction.

The Levin and Rappaport Hovav (1995) analysis, in contrast, provides a post-hoc explanation of the construction without explaining why it occurs in the first place. They suggest (p. 55) that the resultative phrase is licensed by virtue of the identification of an event position in the argument structure of the verb and an event position in the argument structure of the head of the resultative phrase. How this identification could also account for the addition of a post-verbal NP, however, is unclear — these NPs seem to be required on the L&RH analysis only due to the semantic role they play in the resultative construction. According to the Change of

State Linking Rule (Levin and Rappaport Hovav (1995) p. 51), these NPs must refer to an entity undergoing a change of state, but, the verbs in this construction do not independently encode a change of state and so the post-verbal NPs cannot independently be licensed via this linking rule. These NPs only play this semantic role by virtue of the interpretation they receive *in the context of the resultative construction*. Thus one of the core syntactic properties of the resultative construction cannot be accounted for in terms of independent syntactic or semantic factors.

4.6.2 Why are only two classes of unergative verbs different from all the others?

Any account of the resultative construction must provide some explanation of the distinct behaviour of the verbs of manner of motion and the verbs of sound emission in what appears to be the unaccusative resultative construction, to the exclusion of all other unergative verbs.

Levin and Rappaport Hovav (1995) are forced to assume a lexical rule which applies only to these verb classes, without any explanation for why the lexical rule is only relevant for these particular classes. Their approach therefore does not satisfy this criterion.

Wechsler (1996) does not directly address the behaviour of verbs of manner of motion and verbs of sound emission, but relies on the mechanisms of his theory to capture the facts. In particular, each verb of manner of motion and verb of sound emission could be associated with a lexical entry in which a locative canonical result for the event is encoded in the BACKGROUND:TELOS feature, and in which the located entity is structure-shared with the subject of the sentence (i.e. the entity undergoing motion or the entity emitting sound, respectively). The structure-sharing is accounted for in Wechsler's analysis by a constraint for English which requires a resultative to be predicated of the AFFECTED THEME — the argument which undergoes a change of state or location in consequence of the event described by the main verb. This constraint ensures that the resultative phrase is predicated of the subject of the verbs of manner of motion, the verbs of sound emission, and the unaccusative verbs and the object of other transitive verbs, without reference to any syntactic properties of the sentence (in contrast to Levin and Rappaport Hovav's DOR in (4.4)). Thus the answer to the question posed in this section heading on Wechsler's account is simply that these verbs are control resultatives whose affected theme corresponds to subject position. This explanation of course hinges on the control/raising distinction Wechsler draws and the association of a canonical end result with the verbs of manner of motion and verbs of sound emission. It is, however, far from clear whether the latter assumption is valid (see Section 4.4.2), so while the account implicitly meets this criterion, it may do so on the basis of problematic analyses.

By identifying four different constructions and explicitly defining the thematic relations expressed in each one, Goldberg (1995) does not have to isolate verbs of manner of motion from other unergative verbs. The definitions of the constructions are simply such that verbs other than the manner of motion verbs are thematically compatible only with the Intransitive motion construction since the construction requires the verbs to be an instance or means of MOVE-ing. She does not, however, address verbs of sound emission, which would require some additional mechanism such as a lexical rule controlling a sense shift of these verbs, as

these verbs are not compatible with the MOVE predicate.

4.6.3 Semantics of the constructions

In this section I introduce the semantics I associate with the various types of sentences under consideration. As discussed in Section 4.3.1, most of these constructions share a basic interpretation of (4.78a), while the meaning of the verbs of manner of motion, verbs of sound emission, and unaccusative verbs, when combined with a resultative phrase differs from this. In particular, the verbs of manner of motion resultative sentences have the interpretation (4.78b), the verbs of sound emission resultative sentences have the interpretation (4.78c), and unaccusatives have the interpretation in (4.78d).

(4.78) a. **$NP_{subject}$ cause $NP_{controlled}$ to be in Result State by $NP_{subject}$ V-ing**

i. John hammered the metal flat.
John caused the metal to be flat by John hammering (it).

ii. John sneezed the tissue off the table.
John caused the tissue to be off the table by John sneezing.

b. **$NP_{\{subject,controlled\}}$ is in Result Location as a result of $NP_{\{subject,controlled\}}$ V-ing**

i. John ran to the store.
John is at the store as a result of John running.

ii. The bottle floated to the bridge.
The bottle is at the bridge as a result of it floating.

c. **$NP_{\{subject,controlled\}}$ is in Result Location and $NP_{\{subject,controlled\}}$ V-ed (emitted the sound V) as a part of its motion there.**

i. The truck rumbled across the intersection.
The truck is across the interesection and rumbled as a part of moving there.

ii. John wheezed across the road.
John is across the road and he wheezed as a part of moving there.

d. **$NP_{\{subject,controlled\}}$ is in Result State as a result of $NP_{\{subject,controlled\}}$ V-ing**

i. The river froze solid.
The river is solid as a result of it freezing.

ii. The snow melted to liquid.
The river is liquid as a result of it melting.

Neither of the Levin and Rappaport Hovav (1995) and Wechsler (1996) accounts captures these interpretation differences, as both accounts treat all of the constructions in (4.77) as semantically uniform. Although Wechsler does distinguish between two types of resultatives (i.e. control and raising), this distinction occurs along a dimension orthogonal to the contrast between verbs of manner of motion / sound emission constructions and the other resultatives and therefore has no basis from which to account for the interpretation differences. In contrast, Goldberg (1995) is able to capture the interpretation differences explicitly in the definitions of

her various constructions. Only the verbs of sound emission are not explicitly addressed in her analysis.

4.6.4 Semantic restrictedness

Section 4.3.3 introduced data which reflect the restricted nature of the resultative phenomenon. This is not a fully productive phenomenon, in which any resultative phrase can be predicated of any noun phrase. This fact must be accounted for in any treatment of resultatives.

Levin and Rappaport Hovav (1995) do not address the issue of semantic restrictions on the resultative phrase at all, as they aim to give a syntactic account of resultatives. Their approach does not, however, exclude an interaction with pragmatics. In addition, the lack of productivity in this phenomenon creates particular problems for Wechsler's (1996) treatment, since he claims that raising resultatives impose no semantic constraints on the resultative phrases whatsoever. Goldberg (1995), on the other hand, approaches the issue by identifying semantic constraints which apply to the constructions she proposes, and assuming lexicalisation of highly idiosyncratic instances of the constructions. The problem of fully accounting for infelicitous substitutions remains, however.

4.7 The proposal

I propose a model of the data introduced in this chapter which has the following characteristics:

- **A non-uniform treatment of the range of data.**
  Constructions with the syntactic configuration identified in (4.77a) are treated differently than those in the syntactic configuration in (4.77b). Only the latter fall under the heading *Resultatives*.

- **Syntactic constraints and semantics: Resultatives derive from fixed constructions.**
  Following Goldberg (1995), Resultatives will be licensed by a lexicalised construction with specific syntactic and semantic characteristics.

- **Conventionalisation: certain instantiations of the Resultative construction will be assumed to be lexicalised.**
  Again following Goldberg (1995), this property seems to be required to account for certain idiosyncrasies in the range of data allowed.

- **Pragmatics: plays an important role in determining the felicity of an instantation of the Resultative construction.**
  As suggested in Section 4.4.3, and indeed by Goldberg's (1995) direct-causation restriction (see Section 4.5), some of the apparent idiosyncrasy of Resultatives can be accounted for by requiring there to be a natural causal relation between the event expressed by the main verb and the result state provided by the resultative phrase.

- **Compositionality: non-Resultatives acquire their meaning through composition.** The constructions in the syntactic configuration identified in (4.77a) will be assumed to be licensed via the normal process of adjunction and shown to be interpreted compositionally, given a certain perspective on the possible behaviour of adjunct phrases (to be outlined below).

- **Sense shifts:** verbs of sound emission receive their motion sense from a lexical rule.

4.7.1 Resultatives

I begin with outlining the treatment of Resultatives. In this class, I include only a subset of the examples treated as resultative by Levin and Rappaport Hovav (1995) and Wechsler (1996), yet conflate Goldberg's (1995) Caused Motion and Resultative constructions. Specifically, *Resultatives* are defined as sentences of the form [$NP_{subject}$ V $NP_{controlled}$ ResP] which have the associated meaning [$NP_{subject}$ cause $NP_{controlled}$ to be in Result State by $NP_{subject}$ V-ing]. This category then leaves aside the directed motion use of verbs of manner of motion and verbs of sound emission, as well as changes of state expressed with unaccusative verbs. We will return to the treatment of these in Section 4.7.2.

I propose to treat Resultatives in terms of a Construction Grammar construction, shown in (4.79). This construction captures the basic syntactic form of the relevant sentences. The basic semantics of these sentences is also reflected in this construction: $NP_{subject}$ causes NP_{object} to change to the result location or state expressed in ResP. The main verb expresses how the causation is achieved, and its meaning is integrated into the semantics of the construction as a whole via merging between PRED and CAUSE-EFFECT.

(4.79)

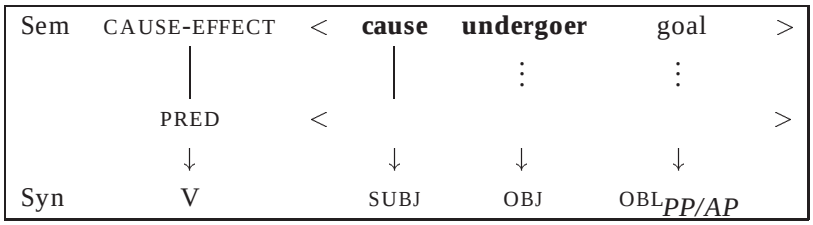

Motivation for the existence of a construction dictating the syntax and semantics of Resultatives comes from (a) the fixed syntactic nature of these sentences in violation of the structure dictated by normal verbal subcategorisation specifications and (b) the semantic similarities among sentences of this type, which cannot clearly be attributed to an interaction between the semantics of the participating components.

There seem to be no valid instances of this construction which vary from the syntactic pattern identified above (also argued by Jackendoff 1990). The position of the resultative phrase is fixed relative to adjuncts, as shown by (4.80)-(4.81) (assuming readings in which the locational prepositional phrases are verbal modifiers rather than noun modifiers). A proposal like Wechsler (1996), in which the resultative phrase is treated on par with adjuncts and licensed by the same lexical rule, cannot account for this restriction on the relative order of the different types of phrases. Similarly, sentences such as (4.82a-b) which correspond roughly to the resultative construction yet express that the subject rather than the object undergoes a

change, are infelicitous. Sentence (4.82c) in which the event which caused the death is phrasal (*drank beer*) is also not grammatical.

(4.80) a. John hammered the metal flat in the workshop.
b. *John hammered the metal in the workshop flat.
c. *John hammered in the workshop the metal flat.

(4.81) a. John laughed himself silly on Saturday evening.
b. *John laughed himself on Saturday evening silly.
c. *John laughed on Saturday evening himself silly.

(4.82) a. *John played (cards) broke.
b. *John drank beer to death.
c. *John drank beer himself to death.

Furthermore, the appearance of unsubcategorised post-verbal elements in this construction must be licensed, and in fact required, by some mechanism. We saw above (Section 4.6.1) that the Levin and Rappaport Hovav (1995) attempt to explain this appearance purely in terms of syntactic constraints and linking rules is inadequate. Goldberg (1995) and Wechsler (1996) both require these elements more or less by stipulation, but the Goldberg account motivates it more convincingly in terms of the semantic properties of the construction and within the context of a general grammatical framework which has other instances of this sort.[15] Thus I choose to follow Goldberg's approach.

The consistency of the interpretations assigned to this construction, despite the varying relations between the verbal head and the other constituents of the sentence as outlined in (4.77b), is striking. There are instances of each form of the construction — varying in relations between the components and the type of resultative phrase — which all convey essentially the same meaning. I summarise these in (4.83).

(4.83) a. $NP_{controlled}$ is a fake reflexive, coindexed with the subject $NP_{subject}$ (V intransitive)
i. John cried himself to sleep. [PP]
John caused himself to be asleep by John crying.
ii. John laughed himself sore. [AP]
John caused himself to be sore by John laughing.
b. $NP_{controlled}$ is the object of a transitive verb
i. John heated *the water* to boiling. [PP]
John caused the water to be boiling by John heating it.
ii. John hammered *the metal* flat. [AP]
John caused the metal to be flat by John hammering it.
c. $NP_{controlled}$ is neither subcategorised by V nor a fake reflexive (V intransitive)

[15]This motivation is also present in Jackendoff's (1990, to appear) work.

i. John sneezed the tissue off the table. [PP]
John caused the tissue to be off the table by John sneezing.

ii. John ran his Nikes threadbare. [AP]
John caused his Nikes to be threadbare by John running.

As pointed out in Section 4.3.1, Resultatives express causative changes of state. Levin and Rappaport Hovav (1995) suggest that the change in telicity from a process to an accomplishment which occurs in the interpretation of events expressed as Resultatives is sufficient to account for the element of causation. It is clearly insufficient, however, as suggested by the contrast with directed uses of the manner of motion verbs which are not interpreted as incorporating this element of causation despite a shift in telicity. The causation component of the interpretation therefore cannot be accounted for by any straightforward compositional mechanism, since the components of the Resultative construction are essentially the same as the directed manner of motion construction — both constructions involve a result state or location being predicated of the referent of a noun phrase. The semantics of the resultative phrase does not directly integrate into the semantics of the remainder of the sentence: a meaning which does not come from either of these parts arises, seemingly without explanation. The explanation under the current proposal (following Goldberg 1995), then, is that this meaning comes from the semantics of the construction in which the words appear.

In proposing this construction, I have conflated two of the constructions which Goldberg (1995) puts forth. This single construction can be used to account for both her Causative Motion and Resultative constructions. This is because the difference between the two Goldberg constructions can be accounted for almost entirely by virtue of the kinds of resultative phrases which are used. If the resultative phrase expresses a location or locative path, then the Resultative is interpreted as a CAUSE-MOVE construction, while if the resultative phrase expresses a state or a path to a state, then the Resultative is interpreted as a CAUSE-BECOME construction. Furthermore, many verbs which can appear in one type of Resultative can appear in the other, as shown in (4.84).

(4.84) a. i. John danced himself across the room.
ii. John danced himself to fame.
iii. John danced his feet sore.

b. i. The critics laughed the show out of town.
ii. John laughed himself out of a job.
iii. John laughed himself sore.

Those that cannot appear in both types can be ruled out on the basis of the pragmatic constraint alluded to in Sections 4.4.3 and 4.5 which requires a coherent causal relation to exist between the main event expressed in the sentence and the result state. So, for example, sentences such as (4.85) are infelicitous because it is difficult to imagine a context in which John hammering the metal or wiping the table could cause the metal or table, respectively, to move. It seems that if such a context could be found, the sentences should be acceptable, as suggested by (4.86) (Jackendoff, p.c.).

(4.85) a. #John hammered the metal down the hall.

b. #John wiped the table across the room.

(4.86) a. John hammered the metal into the hose.

b. John pounded the table down the hall.

These facts suggest that the CAUSE-MOVE and CAUSE-BECOME constructions are really just variations of the same CAUSE-EFFECT construction, and that the construction acquires a more specific interpretation through integration of the meaning of the result phrase. This integration occurs at the time of the merger of the semantics of the construction with the semantics of the main verb and the other components of the sentence. I will outline the process by which this occurs in the following section.

There are, in addition, thematic differences between Goldberg's Resultative and Caused Motion Constructions which I ignore in my conflation of the two. Goldberg treats the post-verbal NP as a *theme* in the Caused Motion Construction, but as a *patient* in the Resultative Construction. I treat the post-verbal NP as simply an *undergoer* in each case (note that this also follows from the Jackendoff (1990) account), (a) to bring the analysis in line with the representation proposed in Chapter 2, (b) to avoid well-known problems with identifying thematic roles, and (c) because the requirement that the post-verbal NP be a patient seems to be an instance of the pragmatic constraint of coherence of the causal relation rather than being an independent constraint. As evidence of the latter point, minimal variants of the examples in (4.66) shown in (4.87), in which the direct objects would certainly not be classified as patients under the standard analysis, seem to me to be pragmatically infelicitous rather than ungrammatical, in that they could be felicitous in a context establishing the appropriate causal relation.

(4.87) a. #Sue watched the monster crazy.
on interp.: Sue caused the monster to be crazy by watching it

b. #Sue loved the book tattered.
on interp: Sue caused the book to become tattered by loving it

Formalisation in HPSG

The formalisation presented here relies on the semantic representation introduced in Chapter 2 and the techniques for implementing lexical rules outlined in Chapter 3.

Constructions, as form-meaning pairs at the phrasal rather than lexical level, do not have any exact correlates in the standard form of Head-driven Phrase Structure Grammar (HPSG) as presented in Pollard and Sag (1994). However, the problem of licensing and integrating non-subcategorised elements (adjuncts in particular) into the representation of a phrase has been addressed by work such as van Noord and Bouma (1994) , and adopted in Chapter 3. Here I will also adopt the basic mechanism proposed in that work: lexical rules define different types of adjunction, altering the SUBCAT list of a verb and the CONTENT associated with that verb as appropriate for each particular type of adjunction. Thus we can have a lexical rule specific to the Resultative Construction, which adds the noun phrase object if necessary and

adds the resultative phrase, and modifies the representation of the semantics associated with the verb phrase as a whole. This approach avoids changes to the fundamental constructs of HPSG which would otherwise be required to capture the phrasal properties of constructions.

As a brief reminder, the lexical rules are formalised in terms of recursive constraints on lexical categories which are processed using delayed evaluation techniques. Information provided by a parser then triggers the application of rules, thus allowing the linguistic context to influence the lexical entry associated with a particular use of a word. This conception of the lexical rules thus allows them to capture constructions in the Goldberg (1995) sense — the verb itself does not undergo a meaning shift or acquire a permanent additional sense encoded in the lexicon; instead a use of the verb in a particular syntactic frame triggers the application of a rule which specifies the appropriate meaning of that use.[16]

This lexical rule will have two syntactic variations (one for each of the intransitive and transitive base syntactic forms which a verb might have before a resultative phrase is added), and three subcases which specify distinct methods of semantic integration: one subcase in which the resultative phrase is an Adjectival Phrase and two subcases for the Prepositional Phrases, one for path PPs and the other for locative PPs. The rule assumes in its present form that the semantic relations between a verb and its arguments remain the same in the resultative construction as they are in a non-resultative use of the verb.[17] This assumption means that

[16]Markantonatou and Sadler (1996) argue against the use of a lexical rule for the sense extension at issue here, instead proposing a monotonic inheritance network of semantic relation types which provides the basis for sense extension without recourse to addition of lexical entries. The lexical rule approach I have advocated also does not involve proliferation of lexical entries, and is in some sense equivalent to the inheritance-based approach (the lexical rule defines potential extensions of the verb meaning, as does the inheritance network). Some motivation for the Markantonatou and Sadler approach comes from the data presented in Section 4.4.2 in example (4.57), repeated below for convenience, and similar examples in (4.88):

(4.57) b. The blacksmith hammered the metal for three days.

c. ?The blacksmith hammered the metal in three days.

(4.88) a. John wiped the table for an hour / in an hour.

b. John painted the picture for an hour / in an hour.

The verbal predicates in these sentences are argued by Markantonatou and Sadler to be ambiguous between a resultative and non-resultative reading and hence possible with both a telic and an atelic aspect. But the telic variants of these sentences, (4.57c) and the *in an hour* variants of (4.88), do not seem to have the standard resultative interpretation, as they lack the element of causation standard for their resultative counterparts (e.g. *John wiped the table in an hour* doesn't seem to mean *John caused the table to be clean in an hour by wiping it* but rather *John completed the activity of wiping the table in an hour*), and many of these telic variants are very odd (4.57c). Furthermore, as discussed in Section 4.4.2, telicity shifts may be triggered by world knowledge or adverbial adjunction, and are morely the result of general non-lexical processes which are independent of the resultative construction (Verkuyl 1989, Krifka 1989, Moens and Steedman 1988). Lastly, I do not agree with the notion of lexical selection of the resultative phrase embedded by the matrix verb in this analysis given the discussion in Section 4.4.2.

I therefore believe that this is not a convincing argument in favor of the inheritance-based approach for the particular problem of resultatives and in fact may lead to inappropriate resultative interpretations of such uses of transitive verbs as given here. In contrast, I do believe it is appropriate for the manner of motion verbs, for encoding potential extensions of that meaning, as we will see in the next section. This is because these extensions are specific to a particular semantic class of verbs and do not involve a change in the core meaning expressed but rather merely an augmentation of that core meaning.

[17]The constructions as proposed by Goldberg (1995) do not make this assumption since the constructions

sentences such as the (a) sentences in (4.89)-(4.90) are not licensed by this construction. In these sentences the relationship between the verb and the direct object differs from the standard relationship as evidenced by the (b) sentences. (Sentences from Goldberg 1995, p. 154.)

(4.89) a. Sam sawed / tore / hacked / ripped a piece off the block.

b. Sam sawed / tore / hacked / ripped the block.

(4.90) a. Sam rinsed / cleaned the soap out of her eyes.

b. Sam rinsed / cleaned her eyes.

However, L&RH (1995) argue that these cases are not instances of the resultative construction, instead resulting from an alternate projection of the arguments of verbs of removal into the syntax. I accept their argumentation and do not perceive these as counterexamples for my proposal. Other apparent counterexamples, e.g. (4.91), stem from independently allowed intransitive variants of the matrix verb (L&RH 1995). In these cases the construction adds the direct object to the intransitive variant.

(4.91) a. Fred drank the teapot dry. [L&RH (1995), 65:(73)]

b. Fred cooked the stove black. [Jackendoff (1990), 227:(38b)]

The distinction between path and locative PP types for the purposes of the lexical rule definitions exists because of a fact about the resultative construction identified by Goldberg (1995): the resultative construction allows PPs which cannot express a path independently of this construction to appear as the resultative phrase, with a path interpretation. So a locative preposition like *inside* can acquire a path interpretation in the resultative construction but cannot be used as a path in other contexts, as suggested by (4.92). As a result, locative prepositions are coerced into a path interpretation by the lexical rule. Specifically, the location they express will be incorporated as the PLACE argument of a path. This coercion follows from a relationship of *endpoint focus* (Brugman 1988) which exists between the meaning of the locative term and its directional interpretation — the location expressed in the locative PP is the endpoint of a path.

(4.92) a. John squeezed the rubber ball inside the jar.

b. *Inside the room he ran, quick as lightning.
on directed motion interp: He ran into the room quickly

There is an additional constraint on this path interpretation, in that the resultative phrases in this construction are interpreted as specifying a change of state / location with a clear end point. This is unlike the path PPs we find with manner of motion and sound emission verbs since they do not necessarily convey a change of location to a specific location. Thus we find a contrast in (4.93), but not in (4.94).

themselves specify the thematic roles of the arguments in the construction. I have chosen not to do this because (a) the thematic specifications Goldberg proposes seem to be too rigid (see example (4.87) and preceding discussion) and (b) the cases for which the semantic relations between a verb and its arguments appear to change are not justifiably treated as instances of the resultative construction (to be discussed presently).

(4.93) a. John sneezed the tissue to the other side of the room.

b. *John sneezed the tissue along the wall.

(4.94) a. John ran to the other side of the room.

b. John ran along the wall/towards the store.

So all paths in the resultative construction must be instances of a *to* path rather than a *toward* or *from* path: that is, the path must lead to a specific endpoint. Locative prepositional phrases are straightforwardly added as the PLACE argument of a *to* path.

The lexical rule will perform the following functions:

1. **Intransitive verbs:** Adds both a noun phrase and a resultative phrase to the SUBCAT list of the verb.
2. **Transitive verbs:** The rule adds only a resultative phrase to the SUBCAT list of the verb.
3. Cases for the resultative phrases:
   (a) The resultative phrase is an Adjectival Phrase.
   (b) The resultative phrase is a Prepositional Phrase.
      i. The Prepositional Phrase head is a path preposition.
      ii. The Prepositional Phrase head is a place preposition.

The semantics given for the verb phrase corresponding to the construction will be virtually identical in each (sub)case. Differences will exist only in whether components of this meaning come from subcategorised or unsubcategorised elements, and in how precisely the resultative phrase is integrated. The lexical rules can be defined as in (4.95), building on the rules previously defined in Section 3.5.3.

The top level rule controls the addition of the resultative phrase and the object NP if necessary, prior to the addition of any adjuncts to the verb's subcategorisation list. The three clauses of *add_ResP* handle the semantic integration of the three types of resultative phrases, embedding the semantic relation normally expressed by the main verb as a MEANS component of the *cause-eff-means* relation which is returned as the internal semantics to be associated with the Resultative Construction. This relation essentially conveys, for the structure {**NP**$_1$ **V NP**$_2$ **ResP**}, the expected semantics *NP*$_1$ *acts on NP*$_2$*, causing NP*$_2$ *to* {*go to some location expressed in ResP/change to some state expressed in ResP*} *by means of V-ing*.

(4.95)
```
add_adj_top(Head[4], SubcatIn[11], SubcatOut[7], SemanticsIn[10]: [INTERNAL Situation[5]],
        SemanticsOut[7], Operator-adjsIn[8], Operator-adjsOut[9]) :–
    test_subcat(SubcatIn[11], SubcatMid[12]),
    add_ResP(SubcatMid[12], SubcatMid2[13], Situation[5], SemanticsMid[14]),
    add_adj(Head[4], AdjunctList[15], SemanticsMid[14], SemanticsOut[7],
        Operator-adjsIn[8], Operator-adjsOut[9]),
    append(SubcatMid2[13], AdjunctList[15], SubcatOut[7]).
```

a. Add an NP to an intransitive Subcat frame;
A transitive Subcat frame remains unchanged.
test_subcat($\langle NP_{\boxed{1}} \rangle$, $\langle NP_{\boxed{1}}, NP_{nom\text{-}ind} \rangle$).
test_subcat($\langle NP_{\boxed{1}}, NP_{\boxed{2}} \rangle$, $\langle NP_{\boxed{1}}, NP_{\boxed{2}} \rangle$).

b. Add different kinds of Resultative Phrases

i. Case one: Path PP

add_ResP($\langle NP_{\boxed{1}}, NP_{\boxed{2}} \rangle$, $\langle NP_{\boxed{1}}, NP_{\boxed{2}}, PP : \boxed{30} : \left[\text{EFFECT} \left[\textit{go-rel} \left[\text{GRND } \textit{path} \right] \right] \right] \rangle$,

$$\text{Situation}_{\boxed{5}}, \left[\begin{array}{l} \text{INTERNAL} \left[\begin{array}{ll} \text{EV-IND} & \boxed{31} \\ \text{ACT} & \boxed{1} \\ \text{UND} & \boxed{2} \\ \text{THEMATIC} & \boxed{30} : \left[\begin{array}{ll} \text{THEM-ARG} & \boxed{2} \\ \text{EFFECT} & \left[\textit{go-rel} \left[\begin{array}{ll} \text{FIG} & \boxed{2} \\ \text{GRND} & \textit{to-path} \end{array} \right] \right] \\ \text{MEANS} & \textit{Situation}_{\boxed{5}} \end{array} \right] \end{array} \right] \textit{cause-eff-means} \\ \text{EXTERNAL} \left[\text{SIT-IND } \boxed{31} \right] \end{array} \right].$$

ii. Case two: Locative PP (Coerce to Path interpretation)

add_ResP($\langle NP_{\boxed{1}}, NP_{\boxed{2}} \rangle$, $\langle NP_{\boxed{1}}, NP_{\boxed{2}}, PP : \boxed{40} : \left[\text{EXT:RESTR} \left\{ \left[\textit{loc-rel} \left[\text{GRND } \boxed{42} : \textit{place} \right] \right] \right\} \cup \textit{set-psoa} \right] \rangle$,

$$\text{Situation}_{\boxed{5}}, \left[\begin{array}{l} \text{INTERNAL} \left[\begin{array}{ll} \text{EV-IND} & \boxed{41} \\ \text{ACT} & \boxed{1} \\ \text{UND} & \boxed{2} \\ \text{THEMATIC} & \left[\begin{array}{ll} \text{THEM-ARG} & \boxed{2} \\ \text{EFFECT} & \left[\textit{go-rel} \left[\begin{array}{ll} \text{FIG} & \boxed{2} \\ \text{GRND} & \textit{to-path} \left[\text{ENDPT } \boxed{42} \right] \end{array} \right] \right] \\ \text{MEANS} & \textit{Situation}_{\boxed{5}} \end{array} \right] \end{array} \right] \textit{cause-eff-means} \\ \text{EXTERNAL} \left[\text{SIT-IND } \boxed{41} \right] \end{array} \right].$$

iii. Case three: Adjective Phrase

add_ResP(⟨NP[1], NP[2]⟩, ⟨NP[1], NP[2], AP : [50] : [be-rel UND [2]]⟩,

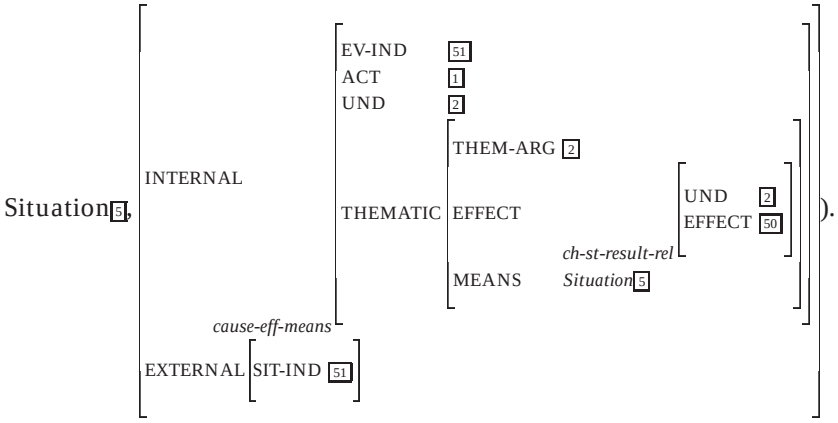

These rules depend on various representational assumptions which build on the representation introduced in Chapter 2. The lexical semantic structure for path and locative prepositions is expected to follow the examples given in Section 3.5.1 (3.71) and (3.75). The semantics of adjectives are assumed to be subtypes of a subtype of *und-rel* and *state*, *be-rel*, which predicate a property of an UNDergoer. The lexical entry for the adjective *flat*, for example, is found in (4.96) and has a semantics of *flat-rel*. This is essentially equivalent to the logical form *flat(x)*.

(4.96)

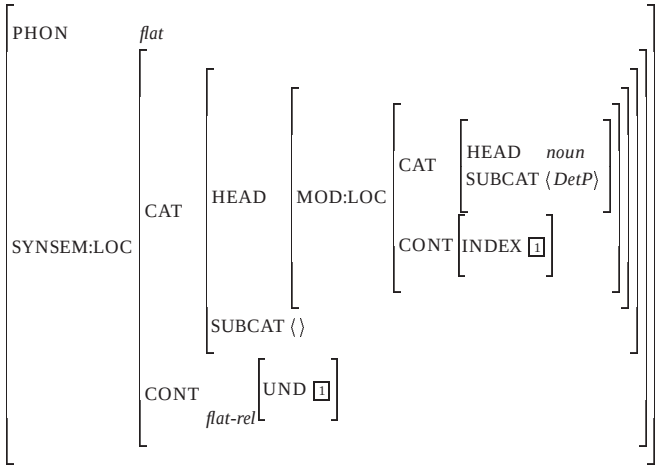

The resultative interpretations which stem from these representations follow from inferencing over the EFFECT values: if something goes to somewhere (*go-rel* plus *to-path*) or changes to a specified state (*ch-st-result-rel*), that thing can be inferred to be at that place or in that state.

Ruling out ungrammatical instances

The proposal made above introduces a mechanism for licensing resultative constructions and accounting for their semantics. As with any account of the resultative construction, this account must also include an explanation of the idiosyncrasy of the phenomenon. The representation utilised provides some of the necessary constraints, restricting path PPs to be a

certain type and coercing locative PPs in a particular way. An additional constraint will be imposed by pragmatics, as will be discussed in Section 4.7.4. There will still, however, be much idiosyncrasy left unaccounted for, specifically cases in which syntactic and semantic substitutions lead to infelicity (see Section 4.3.3).

I mentioned in Section 4.3.3 that some mechanism for encoding conventional constraints is needed in order to capture the idiosyncrasy associated with resultative construction. The framework provided by the Construction Grammar approach to this phenomenon points to a solution. Constructions can be viewed as specifying form-meaning pairs which are part of a language. However, constructions are, like many other linguistic phenomena, governed by conventional usage.[18]

Goldberg (1995, p. 192) suggests there can be lexicalised instances of Constructions — instances in which the form and the meaning specified in the Construction will be preserved, but which will define more precisely the particular words or class of words which can appear. These instances are analogous to idioms in that the structure, meaning and the lexical items which instantiate the construction are fixed, but they differ from idioms because they derive from a more generally available form-meaning pair.

The Resultative Construction can therefore be viewed as defining a semi-productive sense extension mechanism: verbs can be used in sentences with the form and meaning reflected in the Construction (as long as their semantics are compatible with the Construction and any other constraints are satisfied), but some of these sense extensions become lexicalised due to the conventions of language use in particular linguistic communities. In some cases these lexicalised instances are truly idiomatic. They may acquire a meaning which cannot entirely be predicted on the basis of the Construction: Does *John laughed himself silly* really mean that John became *silly* as a result of laughing? In *You scared the daylights out of me*, what are *the daylights* (Jackendoff to appear)? Once there is a lexicalised instance of a construction involving a particular verb, it becomes difficult to use that verb in a different instance of the same construction because that use would conflict with the conventionalised form. Hence the oddness of examples like *John laughed himself tired*.

I suggest that the idiosyncrasy of the resultative construction is a reflection of the high degree of conventionalisation governing the construction. Speakers prefer instances of this construction which conform to their lexicalised instances. Many, if not most, occurrences of this construction which speakers use and come across reflect a lexicalised form. This fact accounts for the idiosyncrasy — variations from lexicalised forms are in theory perfectly acceptable as they can be licensed by the existence of the Construction in the grammar but in practice they are viewed as anomalous or ungrammatical because they don't conform to the "standard" forms in use. Variations (i.e. syntactic and semantic substitutions) of entirely novel instances of the construction are therefore tolerated much more easily than variations of highly colloquial instances. So I can easily accept all variants in (4.97), but the variants of the colloquial (4.98a-b) in (4.98c-d) are less felicitous.

[18] See Chapter 5 for a discussion of the conventional nature of logical metonymy constructions, another example of the role of conventionality.

(4.97) a. Sue brushed her hair smooth.

b. Sue brushed her hair shiny / straight / flat / out of her eyes / . . .

(4.98) a. Sue cried herself to sleep.

b. Sue cried herself sick.

c. ?Sue cried herself asleep.

d. *Sue cried herself sleepy.

This conventionalisation is itself a result of the fact that the resultative constructions are licensed by a form-meaning pair in the grammar. These constructions are very different from constructions which are interpreted strictly compositionally, in that they don't result from generative mechanisms in the grammar and in that their meaning is essentially fixed. These properties indicate that Constructions are "special" in grammatical terms — that is, they do not follow from the normal principles of grammar. That they are subject to a much higher degree of conventionalisation than other constructions seems natural given that their existence in the grammar can be viewed as a result of conventionalisation.

This perspective on the Resultative Construction is in line with observations made about the semi-productivity of many other lexical processes e.g. the generation of denominal verbs (Jackendoff to appear), and blocking by exceptional forms in sense extensions (Briscoe *et al.* 1995, Copestake and Briscoe 1995). The property of semi-productivity has been argued to require lexicalisation of the forms output by the lexical rule. Furthermore, *preemption by synonymy* (Copestake 1995), in which an extended meaning will not be conventionalised if a common synonym exists, has been shown to be overridable in context in that a blocked form can be interpreted. Both of these characteristics surface in resultative constructions, indicating that it is a generative process constrained by conventionality.

In sum, I assume that there are lexicalised instances of this Construction represented in the lexicon, and that these instances limit the acceptability of sentences which vary from the conventionalised pattern.

4.7.2 The other constructions

Unaccusative verbs, manner of motion verbs, and verbs of sound emission clearly differ from the above category of Resultatives. Firstly, the syntactic structure of goal-oriented constructions involving these verbs is different from that for Resultatives, in that it does not require a post-verbal NP of which the goal phrase is predicated, and in the possibility of variability of this structure, as shown by the contrast between "standard" examples such as (4.99) and the sentences in (4.49) and (4.74), repeated here as (4.100) and (4.102), respectively, for convenience. Notice that the sentences (4.100)-(4.101) have a structure which parallels that of the resultative construction as defined above, but the result phrase is predicated of the noun phrase in subject position rather than that in object position.

(4.99) a. John danced across the room.

b. The girls giggled down the hallway.

(4.100) a. The wise men followed the star out of Bethlehem.

b. The sailors rode a breeze clear of the rocks.

(4.101) I love you to distraction.[19]

(4.102) a. John danced mazurkas across the room.

b. John walked the dog to the store.

c. John swam laps to exhaustion.

d. The children played leapfrog across the park.

Secondly, the meaning of these constructions differs enormously from that of Resultatives: (a) there is no element of causation in the meaning, as has been discussed in Section 4.3.1 and above, (b) the goal phrase is restricted by the semantics of the verb which it appears with, as will become more clear below, (c) the shift of a manner of motion verb to a directed motion use can be induced by context (as argued in Section 4.2.2), and (d) there is not necessarily a change of location or state to an explicit endpoint. Whether such a change is inferred depends on the semantics of the goal phrase alone.

I propose a compositional account in which the goal phrases are treated as *pseudo-complements* for the treatment of these cases (see Chapter 3).[20] In particular, the manner of motion verbs will be assumed to encode the potential for a path argument in their lexical semantics and the unaccusative verbs a specific final state.

Sentences such as those in (4.99) are actually ambiguous due to the availability of modification by both pseudo-complements and 'standard' adjuncts: for (4.99a) the reading in which the prepositional phrase is a pseudo-complement has the interpretation *John is across the room as a result of John dancing*, while the reading in which it is an adjunct has the meaning *John is located across the room, and he is dancing*.[21] Similarly, it is possible to have a sentence which contains both a pseudo-complement and a true adjunctive use of the same preposition, as in (4.103a).[22] Notice that the syntactic position of a pseudo-complement is restricted to immediately follow the verb's subcategorised elements. So (4.103b) is not possible on the interpretation on which the PP *in the store* specifies the path of the running rather than the location and *in Washington* specifies a location of the running. The sentences in (4.104) also reflect this syntactic constraint.

(4.103) a. John ran in the store in Washington.

b. *John ran in Washington in the store.

(4.104) a. *John ran in the park to the store.

b. ??John ran in twenty minutes to the store.

[19] I owe this example to a BBC presentation of Anne Brontë's The Tenant of Wildfell Hall.

[20] In Chapter 3, the term *pseudo-complement* is primarily used to refer to certain dative prepositional phrases, such as *for Mary* in *John sang a song for Mary*

[21] Whether one of these sentences is ambiguous depends on whether the preposition heading the goal phrase can behave strictly as an adjunct, strictly as a pseudo-complement, or as both. If it can behave as both, the sentence will be ambiguous.

[22] I assume in these examples that the prepositional phrases modify the main verb rather than a noun phrase.

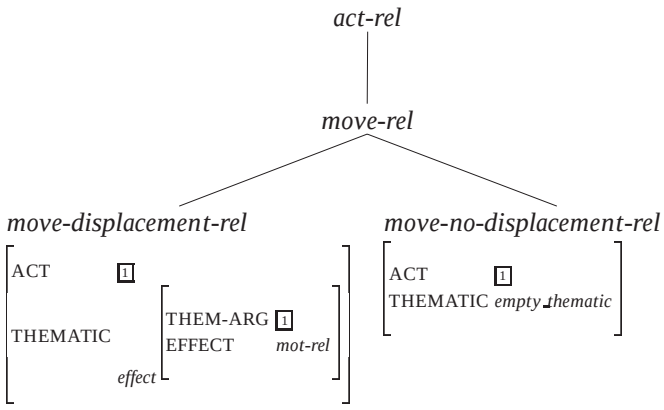

Figure 4.1: Sort hierarchy for Motion relations

Verb-modifying adjective phrases inherently behave as pseudo-complements. They convey a state or property which is relevant to some entity, not an event as a whole, and as such can only be perceived as modifying an argument internal to the verb semantics. That their behaviour parallels that of certain goal phrases therefore becomes clear.

Manner of Motion verbs

Since pseudo-complements interact with verb-internal semantics, their meaning must mesh with the meaning conveyed by the verb. As I suggested in Section 4.3.2, motion verbs can be extended to a relation which includes a path. This potential path is specified by the pseudo-complements. They express a relation between the entity in motion (as indicated by the argument structure of the verb) and a location. This fact alone can account for the limited number of AP modifiers which can appear with verbs of manner of motion. These APs must be construed as expressing locations, not states, as shown in (4.105).

(4.105) a. *Robert ran exhausted.

b. Robert jumped free.

Specifically, each manner of motion verb is associated with the relation *move-rel*, or one of its subtypes, which is a subtype of *act-rel* expressing that the ACTor moves in some way. This relation is defined in the semantic relation hierarchy to have subtypes *move-displacement-rel* and *move-no-displacement-rel*, as shown in Figure 4.1, indicating the possibility of the sense extension of these kinds of verbs. Any pseudo-complement which is to be integrated into the lexical semantics of these verbs, then, must be compatible with the *mot-rel* effect. Path prepositions introduce path phrases which are directly compatible with *move-rel* verbs, by extending them to type *move-displacement-rel*. Phrases headed by locative prepositions can easily be construed as a path: the location they specify is taken to be the endpoint of a path via the coercive relationship of endpoint focus (Brugman 1988), mentioned in the previous section. Adjectival phrases (such as *exhausted* above), however, which cannot be construed as paths are therefore incompatible with the lexical semantics of manner of motion verbs and so

cannot modify these verbs. I will not detail these coercions here, but assume that they will function analogously to the Pustejovsky (1995a) generative type coercions.

Evidence in favor of the lexical encoding of a potential path argument comes from the possibility of contextualisation of this argument: as discussed in Section 4.2.2, and evidenced by example (4.15), the path of motion can be introduced by context. Under the current proposals, this does not induce or require a sense shift of the verb, the context simply instantiates a potential semantic argument and refines the relation expressed in a sentence.

The pseudo-complements are licensed by lexical rules as described in Section 3.3 of Chapter 3, adding a PP to the SUBCAT list of a verb and unifying the semantics appropriately, if compatible. This unification, however, depends on pragmatic reasoning, as argued above, and as such the unification is marked in the lexical entry but validated in the pragmatic component. Here are some example feature structures.

(4.106) The lexical entry for the standard use of *run*.

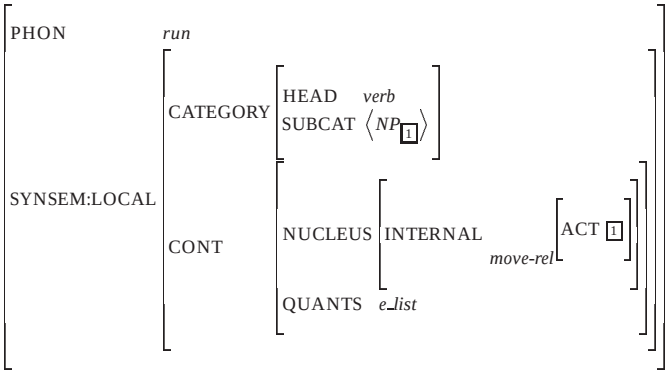

(4.107) The lexical entry for the use of *run* with a PP pseudo-complement. The semantic contribution of the PP must unify (either directly or after coercion to a path interpretation) with the verb semantics, as indicated by the structure sharing between the PP and the internal thematic structure of the verb, represented by 2.

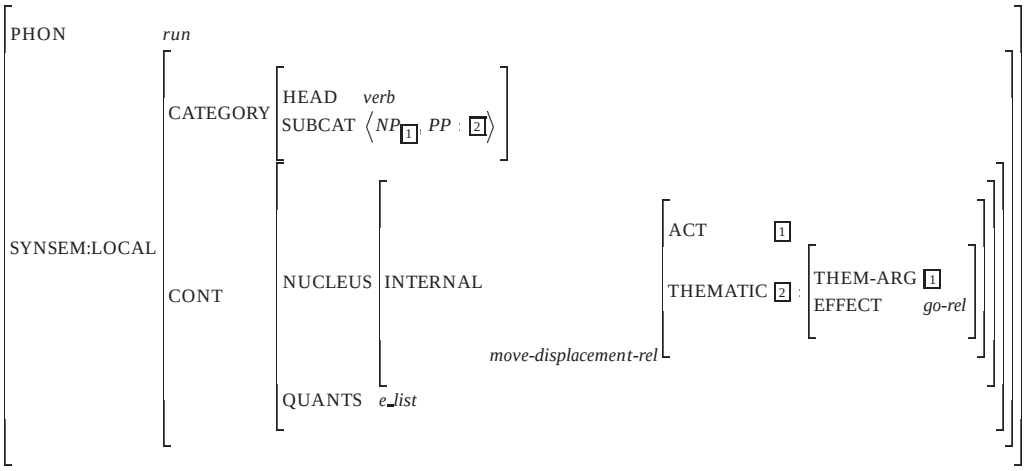

(4.108)

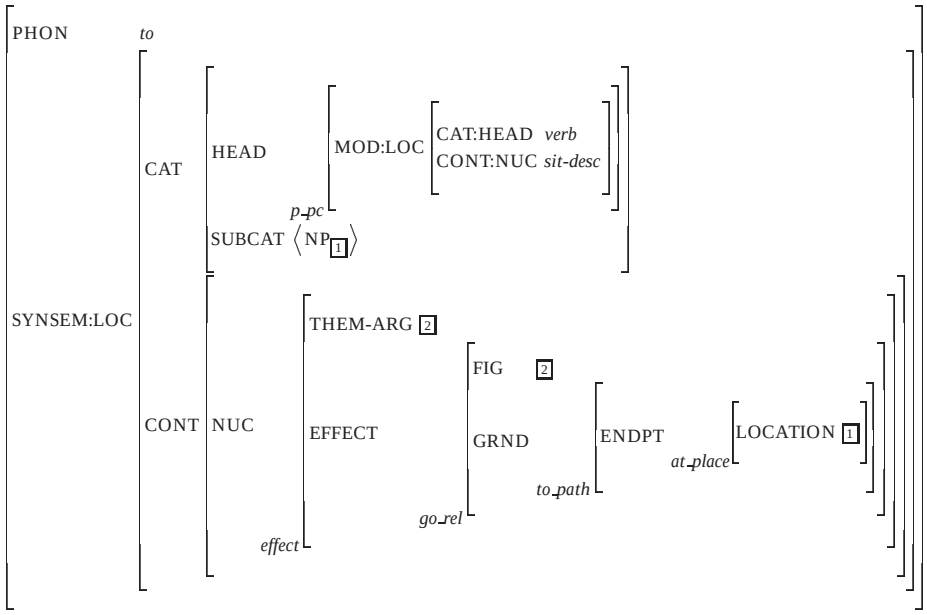

The entry in (4.106) could be combined with a PP headed by *to* whose lexical entry is shown in (4.108) (identical to the entry introduced in Section 3.5.1). The lexical rule controlling this combination would produce the entry in (4.107), which specifies unification of the internal verb semantics and the semantics of the PP. This lexical rule would therefore cause the *path* argument of the verb to be specified to the value expressed by the PP, and would prevent the addition of a PP to the subcat list of a verb which does not express a *go-rel* effect.

This analysis also accounts for the data cited by Wechsler (1996), introduced in (4.49) and repeated here as (4.109).

(4.109) a. *The wise men* followed the star out of Bethlehem.

b. *The sailors* rode a breeze clear of the rocks.

The verbs *follow* and *ride* each subcategorise for a direct object. In addition, since they are motion verbs, they are represented as *move-rel*s. The path is introduced via pseudo-complementation by the path PPs, which extends the relation to *move-eff-rel*.

This analysis finally also accounts (in contrast to Goldberg 1995) for the infelicity of certain movement verbs in a directed motion context: e.g. (4.76), repeated here as (4.110).

(4.110) *Mildred exercised into the room.

Exercise is a verb which expresses movement, but not movement along a path. It differs from other motion verbs in that it does not allow a potential path argument in its lexical semantics, in that it has semantics of type *move-no-displacement-rel*. Pseudo-complementation cannot apply, only standard adjunction, therefore accounting for the unambiguity of (4.111), for which the prepositional phrase can only be construed as situating the event in space, not as expressing the direction of the motion.

(4.111) Mildred exercised in the room.

Unaccusatives

Similarly, the requirement that the meaning of a pseudo-complement must mesh with the meaning conveyed by the verb it modifies can account for the contrasts in (4.112).

(4.112) a. The river froze solid / *icy.

b. The toast burned black / *hard.

Such unaccusative verbs all inherently express changes of state — e.g. if the river froze, then the river wasn't frozen beforehand but is now. The meaning of *to freeze* could be captured as shown in (4.113), expressing a change of state to a state of being frozen (*ch-st-result-rel* is a subtype of *ch-st-rel* which licenses a specific stative effect).

(4.113) The lexical entry for the standard use of *freeze*.

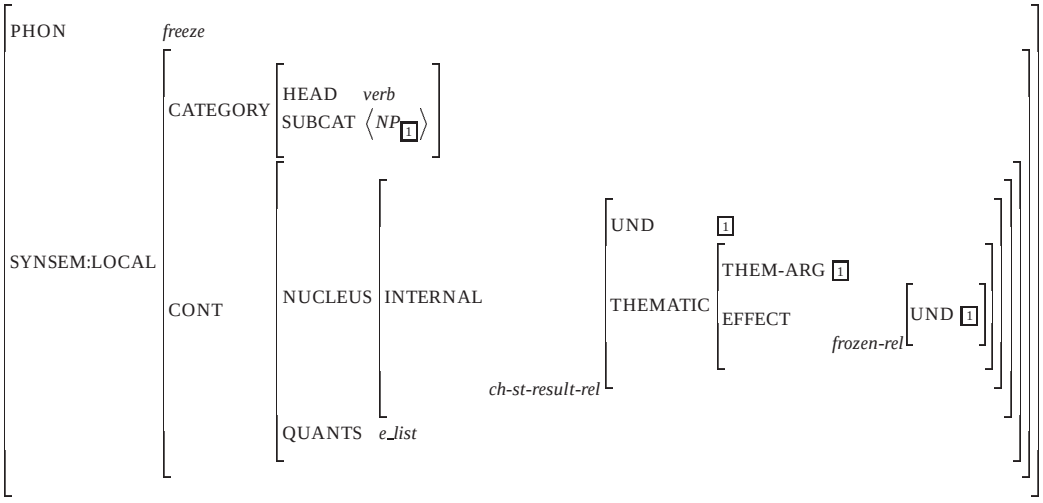

Any AP which modifies such a change of state verb must express a property which unifies with the final state expressed by the verb (e.g. *frozen-rel* in the case of *freeze*). So the entry in (4.113) could be combined with an entry analogous to that introduced in (4.96) on page 148. A pseudo-complement lexical rule for APs would pull out the semantic contribution of the AP (the first element of the RESTR set) and require it to unify with the state expressed in the verb semantics, thereby ruling out APs which bear no relation to the semantics expressed independently by the verb.

(4.114) The lexical entry for the use of *freeze* with an AP pseudo-complement. The state expressed by the AP, represented by 2, must unify with the state expressed internally in the verb semantics (in the thematic effect).

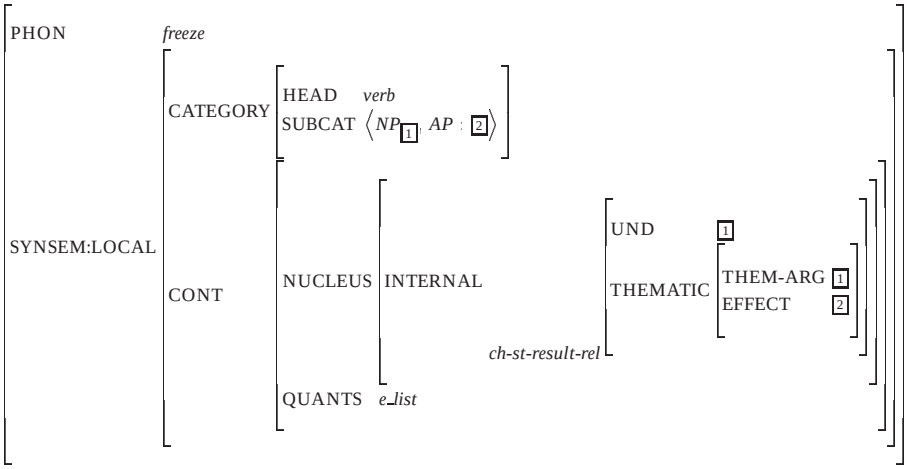

In particular, it seems that these APs are constrained by pragmatic factors to express a state more specific than the property encoded in the lexical semantics of the verb.[23] This constraint is met when an AP expresses how the state expressed by the verb is manifested in the referent of the subject. So, when water is frozen it is solid, and when toast is burned it is black, etc. In fact, the state expressed by the verb seemingly needs to be entailed by the state expressed in the adjective phrase, relative to the referent of the subject (by default, a solid river must be frozen and a black toast must be burned while an icy river may not be entirely frozen and a hard toast may not be burned).[24] World knowledge therefore seems to interact with the ability of an AP to act as a pseudo-complement in these cases. Goldberg's (1995) account cannot explain these differences as her Intransitive Resultative Construction does not require the result phrase to fill a role in the basic predicate expressed by the verb.

Verbs of Sound Emission

In Section 4.6.3 I suggested that the semantics of the verbs of sound emission in the syntactic form (4.115) is as stated in (4.116).

(4.115) **NP_1 V ResP**

(4.116) **NP_1 is in Result Location and NP_1 V-ed (emitted the sound V) as a part of its motion there.**

a. The truck rumbled across the intersection.
The truck is across the interesection and rumbled as a part of moving there.

[23] Grice's Maxim of Quantity (Grice 1975) states that a statement should provide as much information as necessary for adequate interpretation of the meaning, and no more. In this context, an AP which expresses exactly the information already expressed by the semantics of the verb would be redundant and therefore would violate this maxim. Hence the oddness of a sentence such as *The ice melted liquid.*

[24] This is similar to a phenomenon discussed in Jackendoff (1990) in which a Theme internal to the verb semantics can be made explicit via a PP headed by *with*. In that case, the *with*-adjunct (pseudo-complement in my terminology) is subject to a non-redundancy condition: the noun phrase in the PP must be distinct from the implicit value of the Theme as encoded in the lexical semantics.

b. John wheezed across the road.
John is across the road and he wheezed as a part of moving there.

Motion, however, is not a standard component of the semantics of these verbs and so the question of how this interpretation is achieved must be addressed.

I propose that there exists a lexical rule which transforms an intransitive verb of sound emission to a verb which subcategorises for a path PP and means, roughly, that an entity moves along the path and the sound expressed by the verb is emitted as a result of that motion. The evidence for the existence of an independent sense of these verbs on this use, in contrast to the manner of motion verbs for which the directed motion use is merely a precisification[25] of the basic meaning of the verb, is as follows:

- The primary entailment of the sentences in which the verb of sound emission appears with a path PP is of motion rather than sound emission, with the sound emission secondary. This is in contrast to the intransitive use of these verbs for which sound emission is primary, and also to the manner of motion verbs for which the motion is primary for both their uses.

- There is a specific constraint on the verbs of sound emission in this construction: there must be some 'essential connection' between the motion and the sound emission. As Levin and Rappaport Hovav (1995, p. 191) suggest, "the sound must be emitted as a necessary concomitant of motion". Sentences such as those in (4.117) are less felicitous than those in (4.116) above because the sound emission does not result from the motion of the entity. These sentences could express that the sound accompanies the motion, but in isolation of a specific context do not establish that the sound is emitted through the motion (rather than independently) and thus do not satisfy the semantic constraint.

  (4.117) a. ?The car honked across the intersection.
  b. ?John whistled across the road.

  This constraint stems from a particular semantic relation between the input and output forms of the verb, as specified in the lexical rule.

- This is a productive process which also applies to verbs of light emission, as in (4.118), with parallel semantics and semantic constraint (constrast (4.118b) with (4.118c)).

  (4.118) a. The meteor flared across the sky.
  b. The firefly flickered across the field.
  c. ?The flashlight flickered across the field.

- There is no causative relation between the sound emission and the motion (the sound emission does not cause the motion) and so these sentences cannot be instances of the resultative construction.

[25]This is a term borrowed from Manfred Pinkal's (1995) discussion of underspecification, meaning essentially "made more precise".

So the syntax and semantics of these verbs on a directed motion use derives from an additional lexical entry for the verbs, specifying an obligatory PP complement and the appropriate semantics generated by a lexical rule. No additional mechanisms are necessary.

Conclusions

The explanation of the modification of manner of motion and unaccusative verbs via pseudo-complementation clearly bears a certain resemblance to Wechsler's account of the resultative construction, in that they both assume something like an "inherent endpoint" for certain verbs which must unify with the semantics of the phrase which aims to specify this endpoint. However, on my account this "endpoint" is directly licensed by lexical semantic structure, as reflected in the semantic relation type hierarchy, rather than background information governing verb use and as such captures constraints on this kind of sense extension in a semantically motivated manner. This approach does not have the stipulative nature of Wechsler's account as it isolates these manner of motion and sound emission cases from true Resultatives. Furthermore it accommodates "endpoints" which aren't really endpoints at all, in that manner of motion verbs licenses *go-rel*s with general PATH arguments rather than requiring a locative endpoint.

4.7.3 Causativisation revisited

The evidence for the treatment of directed manner of motion verbs as underlyingly unaccusative hinges largely on the causative alternation data in English, as I discussed in Section 4.2.2. In that section, I pointed out the lack of consistency in the causativisation data, indicating that causativisation as conceived of by Levin and Rappaport Hovav (1995) cannot be productively applied to the manner of motion verbs. I suggested there that causative uses of particular verbs of manner of motion are lexicalised. In this section, I would like to informally discuss the possibility of an alternate conception of causativisation which may account for the semi-productivity of the causativisation process with respect to (directed) manner of motion verbs and the requirement of the directional phrase in this context and doesn't require that these verbs have an unaccusative argument structure or that the causative form be lexically encoded for all of these verbs.

As introduced in Section 4.2.2, L&RH propose that causativisation occurs when an empty external argument position is filled by an external cause, hence requiring that any verb undergoing causativisation have an empty external argument position. For manner of motion verbs, their causative forms are generated via a two step process: first the manner of motion verbs are shifted to a directed manner of motion sense, which simultaneously gives them the required empty external argument position and adds a directional phrase, and second an external cause is inserted into the external argument position. Thus from *John walked* we generate *Billy walked John to the store* via *John walked to the store*, and **Billy walked John* will be ruled out.

In lexical semantic work (e.g. Jackendoff 1990), causativisation has been treated as a process which embeds an event representation inside of a causation function (*cause*([CAUSER], [CAUSED-EVENT])). The corresponding change in syntactic form for the causativised verbs

results from different linking from the CAUSE function to surface form than from a simple event representation. The causativisation of *John walked to the store*, represented as in (4.119a), thus generates the representation in (4.119b) which maps to *Billy walked John to the store*.

(4.119) a. walk(j, [*to_the_store*])

b. cause(b, walk(j, [*to_the_store*]))

This causativisation process could conceivably be applied to any caused event (subject of course to the condition that the causation is felicitous given the discourse or world knowledge); the difficulty then is to explain the differences with repect to causativisation between manner of motion verbs and these same verbs on a directed motion interpretation, noted by L&RH and shown in examples (4.18) and (4.20), repeated below for convenience. This data suggests that manner of motion verbs are felicitous in a causative construction only on their directed motion use.

(4.18) a. The soldiers marched (to the tents).

b. The general marched the soldiers to the tents.

c. ??The general marched the soldiers.

(4.20) a. The mouse ran (through the maze).

b. We ran the mouse through the maze.

c. *We ran the mouse.

A possible explanation for the necessity of an explicit path in the event representation embedded by the CAUSE function in order to generate a felicitous causitivised form (that is, to rule out **Billy walked John*) is simply that there exists a constraint on the process of causativisation which requires explicit reference to the end state or location which is being caused. This constraint can be justified through consideration of the function of the causativisation process: its purpose is to express the cause of some change of state or location by something or someone. If the end state or location is not explicitly mentioned in the representation of the event which is caused, then what is caused remains unclear and the meaning of the causativised form is underspecified and therefore infelicitous. The constraint would in effect restrict caused events to be telic events.

This account treats causativisation as a semantic process rather than a syntactic one, with the appropriate syntactic changes being affected via the linking between semantics and syntax. It accounts for the directional phrase requirement for causativised forms of manner of motion verbs via a semantic constraint on causativisation rather than the underlying syntactic argument structure of the verb to be causativised. The inconsistency of the application of this process to verbs of manner of motion still needs to be investigated, but it is likely that the infelicitous examples (such as (4.22b), repeated here) are a result of further semantic constraints on causativisation.

(4.22b) #John limped / hobbled / ambled / meandered / swaggered / sauntered / sashayed / wiggled the child to the store.

The causativisation at issue here is very different from that expressed by the resultative construction. Although the surface forms are parallel, the semantics are very different. Compare (4.120a) with (4.120b).

(4.120) a. Billy rolled the ball down the incline.
$\not\models$ Billy rolled
$\models$ The ball rolled
cause(b, roll(*the ball*))

b. Billy sneezed the tissue off the table.
$\models$ Billy sneezed
$\not\models$ The tissue sneezed
cause(b, be(*the tissue*, *off the table*), means(sneeze(b)))

The caused event bears a different relation to the causer in each case. In the resultatives (e.g. (4.120b)), the subject of the sentence is both causer and participant in the main event expressed by the verb, while in the 'true' causatives (e.g. (4.120a)) the subject is a causer but otherwise plays no role in the event expressed by the verb. In the case of manner of motion verbs, the causer is perceived as having control over both the caused motion and the manner of motion that in the 'true' causative case, but only the caused motion itself, and not the type of that motion, in the resultative case. Consider (4.121). In (4.121a) John explicitly causes the motion of the horse in a jumping manner, while (4.121b) conveys only that the letter changed location, and what caused that change in location (John jumping) but not the manner in which the motion occurs (i.e. the letter does not jump as it moves to the post office).

(4.121) a. John jumped the horse over the fence.
John caused the horse to go over the fence in a jumping manner (means of causation unspecified)

b. John jumped the letter to the post office.
John caused the letter to be at the post office by John jumping (manner of motion of the letter unspecified)

Thus we are dealing with distinct phenomena, which share certain similarities, but which have differing interpretations and which are subject to different semantic constraints. I suggest that differences in the felicity of particular verbs in instantiations of these constructions (e.g. the manner of motion verbs in the casuative form vs. in the Resultative construction) stem not from underlying differences in the unaccusativity of the verbs in the various instantiations, but from these different semantic constraints.

4.7.4 Idiosyncrasy and Pragmatics

I have mentioned several times in this chapter the need to bring pragmatic felicity into the discussion of the resultative construction and the restrictions which there appear to be on its use. It seems clear from the discussion in Sections 4.4.3 and 4.5 that a pragmatic constraint requiring a causal relation to be established between the main event expressed in the resultative construction and the result state must exist. This *pragmatic coherence constraint* alone

can account for much of the apparent idiosyncrasy of this construction, such as the examples repeated below.

(4.43e) i. They laughed John out of the room.

ii. #They laughed John into the room / down the hall.

iii. #They insulted John out of the room.

(4.61a) John hammered *the metal* flat / *safe / *red.

(4.71b) *The feather tickled her silly.

(4.73a) The feather excited her into a frenzy.

I argued in Section 4.5 that the constraints on the resultative contruction derive from general pragmatic principles. Grice's Maxim of Relevance (Grice 1975) is a heuristic governing discourse coherence which suggests that each constituent in a discourse must be as relevant as possible to the current discourse context. In the case of the resultative construction, this heuristic can be taken to mean that a rhetorical connection must exist between the elements in the construction in order for them to be felicitously combined. In this section I will demonstrate that this maxim can be viewed not just as a constraint on attachment of discourse constituents within a discourse but also on the coherence of the semantics expressed by particular constructions, i.e. within a single constituent, and that the recruitment of this maxim to apply intra-sententially can help to explain the impact of pragmatics on the felicity of resultative constructions. This approach has been demonstrated as a useful strategy for modeling the interaction of pragmatic reasoning with lexical semantics in the cases of co-predication and coordination (Lascarides *et al.* 1996).

As the resultative construction essentially conveys a cause and effect relation between the event expressed by the main predicate and the resultative phrase, it must be possible to establish a rhetorical connection between the two components in order for the cause-effect relation to be coherent. The relevant rhetorical relation in this case is *Result* (Hobbs 1985, Polanyi 1985). This causal link must be established on the basis of world knowledge inferences and reasoning about information in the discourse. Sentences like (4.43e.ii.)-(4.43e.iii.) are therefore normally ruled out because the Result relation between the cause and the effect in each case (e.g. some people laughing and John going into a room) cannot easily be supported. I argue that in a highly specific context which establishes the Result relation, these sentences would be entirely felicitous as they would satisfy the Gricean Maxim of Relevance (Grice 1975).

Consider again the discourse presented in Section 4.4.3:

(4.63) a. The slide at the park had a section which had come loose.

b. Several children had hurt themselves on the protruding edge.

c. In order to prevent further injuries, John hammered the metal safe.

Under normal circumstances, hammering metal does not cause the metal to become safe. The context preceding (4.63c) in this discourse, however, establishes the basis for the causal relation — the metal is unsafe because an edge is protruding, so it follows that an event which stops

the edge from protruding (e.g. hammering it) will cause the metal to be safe. Once this basis is established, the constraint on the resultative construction is met and the sentence will be judged felicitous in context.

Let us formalise this reasoning. Following Asher (1993b) and Lascarides and Asher (1993), the Result relation is constrained by the Axiom stated below, where e_x represents the main eventuality described in the constituent x. This Axiom states that if the Result relation holds between two discourse constituents, then there must be a causal relationship between the eventualities they describe. In addition I assume a rule specific to the Resultative Construction, `Resultative Link`, which formalises the pragmatic constraints on the semantics of that construction.

- `Axiom on Result`:
  $\Box(\textit{Result}(\alpha, \beta) \rightarrow \textit{cause}(e_\alpha, e_\beta))$
- `Resultative Link`:
  $\Box \left(\left[\text{THEMATIC} \left[\begin{array}{l} \text{EFFECT } e_\beta \\ \text{MEANS } e_\alpha \end{array} \right] \right] \rightarrow \textit{Result}(\alpha, \beta) \right)$

So the Resultative Construction requires the Result rhetorical relation to be established, which in turn requires a causal relation between the two events participating in the construction. If this causal relation isn't already known on the basis of previous processing, it must be accommodated.

A distinctive feature of the theory of discourse semantics used by Lascarides and Asher (1991, 1993) (Discourse in Commonsense Entailment, or DICE, in combination with Segmented Discourse Representation Theory, SDRT (Asher 1993b)) is that it models this accommodation process: if two constituents are attached together with a rhetorical relation R and R's coherence constraints aren't met by the constituents already, then the required semantic content can be added in a constrained manner on the basis of world knowledge (through a process known as SDRS *Update* (Asher and Lascarides 1996, Copestake and Lascarides 1997)). Intuitively, this models the way speakers expect hearers to fill in gaps when they interpret discourse. In DICE, rules reflecting pragmatic maxims and capturing the influence of world knowledge, syntax and semantics on interpretation are used to compute rhetorical relations. Default knowledge is captured in formulas of the form $A > B$ which mean approximately "if A, then normally B", which allows nonmonotonic inference to take place. Details of the logic underlying such formulas can be found in Asher and Morreau (1991) and Asher (1993a).

In the context under analysis, DICE would have to reason on the basis of world knowledge about what kinds of things people do to make edges safe to accommodate a causal relation as required by the Resultative Construction (cf. Asher and Lascarides 1996). The context in (4.63a)-(4.63b) establishes that there is an unsafe protruding edge on a slide, and resolution of the definite description *the metal* (on the basis of knowledge such as `Slide Material` below) will result in an initial interpretation for (4.63c) of *John hammered the unsafe protruding metal slide edge safe.* To accommodate the *cause* relation, additional world knowledge must be

recruited and added to the discourse structure. Here, the knowledge that a non-protruding edge is normally safe (e.g. `Safe Edge`) and that hammering an edge can result in it being non-protruding (e.g. `Making Protruding Edges Non-protruding`) can be added to the discourse, which together would be enough to support the causal relation between hammering a protruding edge and it being safe.

- `Slide Material`
  $\textit{slide}(x) \wedge \textit{material}(x, z) > \textit{metal}(z)$
- `Safe Edge`
  $\textit{edge}(y) \wedge \neg\textit{protruding}(e, y) > \textit{safe}(e, y)$
- `Making Protruding Edges Non-Protruding`
  $\textit{edge}(y) \wedge \textit{protruding}(e, y) \wedge \textit{hammer}(e) \wedge \textit{patient}(e, y) \wedge \textit{holds}(e, t) >$
  $\textit{state}(e') \wedge \textit{patient}(e', y) \wedge \textit{holds}(e', t') \wedge \neg\textit{protruding}(e', z) \wedge t \prec t'$

These pragmatic constraints can also work to rule out examples which based on general world knowledge might be acceptable but which don't make sense in context. That is, if discourse information suggests that a causal link is not valid, the coherence constraints on Result will not be met (discourse information will conflict with world knowledge in this case but discourse information is more specific and therefore "wins" over world knowledge — see Lascarides and Copestake 1995) and so the sentence will be judged infelicitous in context.

4.8 Conclusions

This chapter has reviewed various proposals for the modeling of the resultative construction and suggested criteria which any model must meet in order to adequately account for the relevant data. I offered a hybrid model, combining insights from the different proposals and from related work, to account for the full range of data.

I have argued, in direct contrast to the proposal put forth by Levin and Rappaport Hovav (1995), that there is no underlying syntactic difference between the verbs that appear in the different kinds of "resultative" constructions. I showed that the evidence presented by L&RH in favor of the syntactic difference is weak. Furthermore, their account requires postulation of a lexical rule which applies only to two verb classes and does not follow from any general principles. The syntactic shift and corresponding semantic shift induced by the lexical rule requires addition of an explicit directional phrase, but I have shown that discourse context can trigger at least the semantic shift and that therefore this rule is overrestrictive. The assumed change in underlying argument structure also potentially creates problems for other components of grammar such as linking rules.

This evidence therefore calls into question the L&RH analysis of manner of motion verbs on a directed motion use. More strongly, the idiosyncrasy of the resultative data and the differing interpretations associated with the various forms of the resultative construction suggests that a purely syntactic explanation is impossible. I showed that semantic differences

exist between the different forms and that pragmatic coherence is a critical factor in determining their felicity. A single syntactic constraint is too restrictive, treating as ungrammatical examples which are merely infelicitous independent of a context which establishes the appropriate semantic relations, and grouping together constructions which actually display entirely distinct syntactic and semantic behaviour.

Instead, I argued that the differing behaviours which verbs display with respect to these constructions follow from the specific semantic properties of the various construction types and the lexical semantics of the verbs themselves. I motivated this through an overview of the interpretations associated with various forms of the construction, a discussion of the syntactic flexibility of those forms, and an analysis of the lexical semantics of both the verbs and the "resultative" phrases appearing in the different forms.

The proposals in this chapter draw on an insight identified by Wechsler (1996) that certain verbs encode an inherent endpoint in their lexical semantics. Although I do not support Wechsler's contention that there are verbs for which this endpoint is optional background information, preferring a more general explanation for these cases in terms of lexicalisation and pragmatic constraints on the construction, I showed that the behaviour of certain classes of verbs, specifically manner of motion verbs and unaccusative verbs, does stem from the semantics they express.

The work by Goldberg (1995) provides a fundamental contribution to the proposals here: her introduction of independent form-meaning pairs to a grammar allows phenomena which have no clear generative basis to be captured. This approach seems to derive from an acknowledgement of the highly conventional nature of language use and the largely fixed nature of the conventions. Acceptance of this insight from Construction Grammar led directly to a view of resultative constructions in terms of a form-meaning pair. This view is supported by the fixed syntactic and semantic behaviour of the construction, as discussed in Section 4.7.2.

In Section 4.6 I identified criteria which any model of the resultative construction must satisfy. I review them here, pointing out how the proposals in this chapter address each one.

- **Syntax:** Two basic mechanisms are used to account for the syntactic behaviour of the constructions examined. Both mechanisms have been shown elsewhere (Goldberg 1995 and in Chapter 3, respectively) to be useful in accounting for various phenomena. Within the HPSG formalisation presented, both are controlled by general lexical rules which may apply across the lexicon.
  - A Construction licenses (true) Resultatives.
  - Pseudo-complementation, a process similar to adjunction but involving entities internal to verb semantics, licenses other goal phrases.
- **Syntactic constraints:** The restriction of certain syntactic configurations to particular verb classes results from the underlying lexical semantics of these verb classes rather than from stipulation of the unique behaviour of these verbs. It does not stem from underlying syntactic differences.

- **Sense shifts:** Just like the syntax, the semantics of the two main classes of data examined is explained via distinct mechanisms.
  - A Construction specifies the interpretation of (true) Resultatives.
  - The lexical semantics of the verb is combined compositionally with the meaning of a modifying pseudo-complement to specify the interpretation of the sentence as a whole.
- **Conventionality and Pragmatic felicity:** The semantic idiosyncrasy demonstrated by the Resultatives is accounted for through two means.
  - Pragmatic coherence: I argued that much of the idiosyncrasy derives from pragmatic infelicity rather than true ungrammaticality.
  - Conventionality: some Resultatives are fixed, lexicalised, instantiations of the general pattern captured in the Construction which behave in certain respects like idioms.

The analysis proposed, in sum, relies on generative mechanisms which have been shown to have wider applicability within grammar, as constrained by conventionalisation of language use and general pragmatic principles. I have taken advantage, where possible, of the compositional nature of much of language yet allowed for the apparent existence of non-compositional units within language. In doing so, I have been able to provide a model of the resultative construction and similar constructions which accounts for the syntactic and semantic differences between them and some of their idiosyncrasy. I have not had to involve the unaccusative / unergative distinction or assume semantic constraints specific to the constructions, but have shown how their behaviour follows from principles more generally applicable to the lexicon and discourse construction.

five Logical Metonymy

5.1 Introduction

One of the stated aims of this thesis is to examine linguistic data which seem to be explicable only in terms of an interaction between syntactic, semantic, and pragmatic processes. I will argue here that the *logical metonymy* data must be addressed through this interactivity, because the interpretations of the verb phrases in this data can vary with the syntactic frame and with the discourse context in which they appear.

I begin by precisely defining the data of interest. Consider the following sentences:

(5.1) a. John began reading / to read the book.

b. John began the book.

(5.2) a. John enjoyed reading the book.

b. John enjoyed the book.

In both of these sets of sentences, the two sentences express the same meaning although in the (b) sentences no reading event is explicitly mentioned.[1] The (b) sentences exemplify logical metonymy, for which more meaning arises than is directly attributable to the sentential components. That is, the interpretation attributed to these sentences includes an event which is not explicitly introduced, in contrast to the (a) sentences in which the event is specified in the VP complement. This phenomenon involves the use of a noun phrase to suggest an event associated with that noun phrase — it is *metonymy* (e.g. Nunberg 1978) in the sense that one phrase is used in place of another (the noun phrase in place of the full eventive verb phrase) and, under the analysis I will adopt, *logical* in that it is triggered by type requirements

[1]For the purposes of this chapter, I am ignoring subtle differences which might exist between the use of infinitival VP complements and progressive (-ing) forms of the VP complement of aspectual verbs. See Freed (1979) for discussion of this issue.

which a verb places onto its arguments (in the examples above, the verb requires an eventive complement). It occurs for certain verbs which have alternate syntactic complement forms with only a single (default) semantic interpretation.

We will see in this chapter that the characteristics of logical metonymy can only be captured through an interaction of syntax, semantics, and contextual influence. This interactivity will be shown to have implications for the representation of lexical information, and the balance between the lexicon and pragmatic reasoning in interpretation. I will present extensive corpus evidence for a conventional, lexicalised basis for the interpretation of logical metonymies, and will argue that a purely pragmatic treatment of this phenomenon (e.g. Hobbs *et al.* 1993) cannot accurately predict grammaticality judgements on the data.

The verbs which will be examined in this chapter are of interest because they can appear in various syntactic constructions which ultimately may be assigned the same interpretation, while each construction is subject to specific constraints. These constraints are either syntagmatic in nature, or reflect idiosyncrasies in the interaction of word meaning with pragmatic reasoning. These facts highlight the importance of the lexicon as a repository of knowledge about verb use and their interpretation in distinct linguistic contexts. In this chapter I will identify the different constructions and introduce data which exemplifies a range of syntactic, lexical semantic, and discourse-level constraints on the constructions. I will argue for an analysis of the phenomena which hinges on lexical semantic information. I will formalise my proposals through the introduction of lexical entries which build on the representation developed in Chapter 2, and the account of the felicitous and infelicitous data will exploit the interaction between discourse reasoning and the information in such lexical entries.

The so-called *eventive verbs* (aspectual verbs like *begin* and *finish* and verbs like *enjoy*, which all require eventive complements) can take on a huge range of meanings in the logical metonymy construction. In each of the sentences in (5.3), *began* takes on different possible meanings which vary with the particular noun phrase which is the complement of the main verb. Similar patterns hold for the other eventive verbs.

(5.3) a. John began the book. *(began = began reading; began writing)*

b. John began the sandwich. *(began = began eating; began making)*

c. John began the beer. *(began = began drinking; began brewing)*

d. John began the cigarette. *(began = began smoking)*

5.2 Analysis of the meaning variation

There are various conceivable analyses of the meaning variation exemplified by the eventive verbs in this data.

- **Sense enumeration:** One could explain this data by assuming that each of the eventive verbs is associated with a list of senses which corresponds to all the possible meanings that verb can take on in the logical metonymy construction. So *begin* might have senses *begin reading, begin eating, begin smoking*, etc. licensed lexically. This approach, however, has several flaws:

1. There is nothing in this analysis which explains the complement-dependent nature of the meaning variation. Thus (5.3a) could just as easily be assigned the interpretation *John began smoking/eating/drinking the book* as the more standard *John began reading the book*. This analysis would lead to overgeneration of interpretations for individual instances of the logical metonymy construction. For example, it could result in the incorrect prediction that the meanings of all the sentences in (5.4) involve a reading event on the basis of the *begin reading* sense associated with *begin*.

   (5.4) a. John began the coffee. *(began drinking)*
   b. John and Mary began Scrabble. *(began playing)*
   c. John began the computer program. *(began writing/running)*

   This problem could be avoided to a certain degree through the addition of selectional restrictions specifying the required semantic type of the NP complement of *begin* on each of its senses, but this misses an important generalisation about the relationship between the sense of begin and the semantics of the complement NP and *why* a particular sense of begin can combine with a particular kind of complement. For example, coffee is normally drunk, which is why it combines with the *begin drinking* sense of begin. This relationship would not be recorded in such an account.

2. It is very inefficient to attempt to list in advance all of the possible meanings the eventive verbs can take on in every possible context and in combination with every possible complement NP, and probably impossible given the creativity of language use. Thus there would likely be undergeneration of the possible interpretations theoretically available for the construction.

3. This analysis is straightforward at the lexical level, but increases the load on the pragmatic component of interpretation dramatically. The problem of identifying the particular sense of the eventive verb intended in context is not aided by restrictions at the lexical semantic level, beyond any selectional restrictions which might be introduced. Selectional restrictions are very rarely hard and fast (we will see below for example that the *begin eating* sense of *begin* can combine with *book* under certain circumstances, which would not be allowed if this sense were restricted to combining with *food*-type objects) and unlikely to be restrictive enough (selectional restrictions will likely only be restrictions on very general semantic types, so for example the *begin drinking* sense might require a complement of type *liquid*, but then this would license a sentence like **John/The car began the petrol* or *John began the river*), and therefore of questionable use for accurately modeling this phenomenon. In the absence of any restrictions, a logical metonymy will be associated with a very long list of possible interpretations, one of which must be chosen through pragmatic reasoning (or some other mechanism for word sense disambiguation — see Section 6.4). The difficulty of this task is proportional to the number of choices which are available, in that each alternative must be evaluated and the most context-compatible option identified, and is thus arduous under the circumstances

deriving from this analysis.

- **Meaning postulates:** An alternative approach is to specify meaning postulates which explicitly encode the interpretation which is to be given to particular logical metonymies. For example, there might be meaning postulates like those in (5.5).

  (5.5) a. *begin the book* $\rightarrow$ *begin reading the book.*

  b. *enjoy the beer* $\rightarrow$ *enjoy drinking the book.*

  These give precise conversions of the underspecified interpretation of the logical metonymies into fully specified interpretations. Meaning postulates, however, suffer from many of the same problems as the sense enumeration technique.

  1. Advance specification of meaning postulates means that no contextual influence on the interpretation of logical metonymies can be accommodated under this analysis, and that novel instances of logical metonymies cannot be given any interpretation.
  2. Some mechanism would again be needed for identifying the precise interpretation of a logical metonymy in a case where multiple "translations" are defined by meaning postulates. The sense disambiguation problem is much more constrained in this case as compared to the sense enumeration analysis, as there should only be a few meaning postulates for any given logical metonymy, but it must still be addressed.
  3. In the absence of the use of types within the meaning postulates, there would be much redundancy, in that the interpretations given to sets of logical metonymies varying in the eventive verb but not in the noun phrase complement should involve the same event (e.g. *begin/finish/enjoy the book* = *begin/finish/enjoy reading the book*).
  4. This analysis generally fails to make any generalisations about how the missing or underspecified event in the logical metonymy might derive from knowledge of the complement noun phrase.

- **Pragmatic extension:** It is possible to imagine a treatment of this data which relies solely on pragmatic reasoning for resolving the meaning underspecification, along the lines of Hobbs *et al.* (1993), which uses inferences drawn from world knowledge or context to elaborate (make more explicit) the logical form associated with a particular syntactic structure. On such an account, logical metonymies would be interpreted on the basis of knowledge of the object referred to by the complement NP of the eventive verb. The search for this knowledge would be triggered by the semantic incompatibility between the eventive argument the eventive verbs seek and the object denoted by the NP. The pragmatic component would identify some event which is highly associated with the object, and suggest that event as the "missing" event. For example, *reading* is a prototypical event in which *a book* participates, so through pragmatic reasoning *John began a book* is likely to mean *John began reading a book.*

  The main problem with such a treatment is that not every event which is commonly associated with an object leads to a felicitous interpretation of a logical metonymy

involving that object. There seem to be various linguistic constraints governing the phenomenon which cannot be adequately captured under this analysis. So the suggested interpretations of the following metonymies do not seem to be possible, despite the apparent relevance of the event to the object.

(5.6) a. *John began a chair. *(began sitting in/on)*
b. *John began a keyboard. *(began typing on)*
c. *John began the tunnel. *(began driving through)*
d. *John began the trees. *(began growing, planting, watering)*

I will discuss the pragmatic solution further in Section 5.5.4. There we will see that a purely pragmatic approach would fail to constrain the possible interpretations of logical metonymies and does not allow for the incorporation of lexically specified restrictions. It therefore would fail to account for the full range of data. Consider the data in (5.7) and the discourse in (5.8), for example.

(5.7) a. John enjoyed the book.
b. John enjoyed reading the book.
c. John enjoyed writing the book.
d. John enjoyed eating the book.

(5.8) a. John will be audited by the tax service, so he has been destroying things which might incriminate him. He has destroyed the files and the computer disks.
b. *He will begin the books tomorrow. *(destroying)*
c. He will begin on the books tomorrow. *(destroying)*

If we consider the sentence (5.7a) independent of a context we are likely to interpret it as (5.7b) or possibly (5.7c). If we then insert the prior context *John is my pet goat; he loves eating things* we suddenly prefer the interpretation in (5.7d). In this case context is able to influence the interpretation of the metonymic construction.

In (5.8), in contrast, we have two cases in which a metonymic construction needs to be interpreted in a way which is at variance with the default interpretation, due to contextual requirements. The sentence *He will begin the books tomorrow* would by default mean *He will begin reading the books tomorrow*, while *He will begin on the books tomorrow* does not have a strong default interpretation. The context clearly cues the interpretation *He will begin destroying the books tomorrow* for both sentences. The sentence (5.8b), however, is infelicitous in the context while (5.8c) is felicitous in the discourse and has the expected interpretation. These contrasts suggest a difference in the behaviour of *begin* and that of *begin on* in metonymic constructions, and unexpected differences in the way particular instances of metonymy interact with the discourse context. There is no way for these differences to be explained on the basis of pragmatics alone, as the difference between the phrases is largely syntactic. Later in this chapter I will present a treatment of lexical differences between *begin* and *begin on* which is able to account for these contrasts.

- **Co-composition:** Similar to the pragmatic account, analyses of logical metonymy have been proposed (e.g. Pustejvosky 1991, 1995a) which rely on lexical encoding of "core aspects" of the meaning of words, and which locate the meaning underspecification of logical metonymies at the level of combination of the eventive verb with the complement noun rather than at the level of the verb itself. So the individual words in the logical metonymy do not have underpecified meanings; it is their combination which requires semantic resolution. The difference between these accounts and the pragmatic account is that they assume lexical specification of certain knowledge, which is then available for interaction with syntactic and / or semantic constraints on the combination of particular words into particular structures. We will see that this fact enables a more thorough analysis of the range of grammatical logical metonymy data.

5.3 The range of the data

I have introduced various grammatical and ungrammatical examples of logical metonymy above. A closer look at the details of the proposed co-composition analyses of logical metonymy will reveal more data which must be taken into consideration.

Much discussion in recent lexical semantic literature has focused on logical metonymy (Pustejovsky 1991, 1995a), linguistic constraints on the phenomenon (Godard & Jayez 1993, Pustejovsky & Bouillon 1995), and the influence of discourse context on the interpretation of these metonymies (Briscoe *et al.* 1990, Lascarides & Copestake 1995). All of this work uses the co-composition analysis as its starting point. The systematic syntactic ambiguity of the eventive verbs is handled on one version by postulating a single logical form for the verb and triggering the operation of *type coercion* to shift the syntactic type of the complement to a type which is compatible with the verb (Pustejovsky 1995a, Pustejovsky & Bouillon 1995). This process is needed because aspectual verbs require eventive complements. Noun phrase complements which refer to objects (like *book* in the above examples) cannot be construed directly as events and must be converted somehow into an event-referring phrase in order to be compatible with the aspectual verbs. Another approach is to postulate multiple logical forms for the eventive verbs (Godard & Jayez 1993, Copestake & Briscoe 1995), each of which is responsible for appropriately combining with a distinct type of complement. In the latter case, type coercion is internal to the verb semantics, in that the lexical entry for a verb specifies precisely how the interpretation of a full VP is to be established on the basis of the particular complement type specified in the lexical entry. The result under both of these approaches is that the logical form for each combination of the verb plus complements will be identical.

Within this work, there are differences in the status of lexical information. Analyses following Pustejovsky (1995a) do not address the issue of contextual influence on logical metonymy interpretation and therefore cannot account for the acceptability of the interpretation of (5.7a) as (5.2d). The lexicon is the sole source of information for the interpretation of logical metonymies under these accounts. Briscoe *et al.* (1990) and subsequent work (e.g. Lascarides & Copestake 1995) argue for the influence of pragmatics on logical metonymy interpretation and treat lexical information as *default* information about the interpretation. A *lexical*

default is lexical information which appears as default in the logical form corresponding to the semantics of a sentence or phrase (this notion is formalised in Lascarides *et al.* 1996). This distinction in the status of lexical information is data-driven, deriving from examples such as the "goat" example. Under this perspective, the lexicon can suggest an interpretation, but that interpretation can be overridden by reasoning derived from the discourse context. These approaches formally address the interaction between lexical semantics and discourse context, but do not account for cases in which context cannot override defaults (e.g. (5.8)) or the full range of constraints on logical metonymy.

The coercion which must occur to get the appropriate readings of the logical metonymies clearly requires more than a simple conversion of the type of the NP argument. The missing element of meaning, the event which serves as an argument to the eventive verb, must be introduced. In the approaches under consideration, this eventive element is suggested by lexical information associated with the object denoted by the complement noun phrase of the eventive verb. Pustejovsky (1991) proposes that the element comes from one of the roles in the noun's lexical semantic structure, the *qualia structure*, which represents the defining attributes of an object. Type coercion looks (in the first instance, for the accounts which also take into account pragmatic influence) to the qualia structure for something of the type required by the verb.

The existing co-composition analyses of logical metonymy assume that a full qualia structure is represented in each noun's lexical semantics, that is a completely specified representation of the core components of the noun's meaning. These components serve as a source for eventualities involving the noun which can be accessed in metonymic interpretation.

5.3.1 Word knowledge represented in Qualia Structure

On the view put forth by Pustejovsky (1991, 1995a), the core lexical semantic components of a word's meaning are captured by several roles in a structure called *qualia structure*. Four essential aspects of a word's meaning are assumed to be represented:

- CONSTITUTIVE: the relation between the denotation of the word and its constituent parts. *(generally used for concrete objects)*
- FORMAL: that which distinguishes the denotation of the word within a larger domain.
- TELIC: the purpose and function of the denotation of the word.
- AGENTIVE: factors involved in how the denotation of the word came into being; its origin.

The roles in qualia structure relevant to logical metonymy are the telic and agentive roles, as these two roles can specify eventualities involving the entity which is the denotation of the noun in the NP complement of an eventive verb. In particular, the telic role captures the activity which is normally performed on / with the entity and the agentive role reflects the activity which brings the entity into existence.

In the examples (5.1)-(5.2) introduced at the start of this chapter, the qualia structure associated with the noun *book* thus must provide events which can serve as the appropriate interpretation of the logical metonymy sentences. The TELIC role of *book* would be *reading*, while the AGENTIVE role would be *writing*. Either of these two events could be picked out by the verb to provide the missing event needed for interpretation in the coercion process.

The existing approaches assume that all concrete nouns always have both the telic and agentive roles specified, and thus the eventualities specified there are always available to the process of type coercion when establishing the interpretation of a logical metonymy structure. No other interpretations will be available (if the discourse context is not take into consideration), correctly ruling out the specified interpretations of the sentences (5.9), since the desired eventualities do not fill a role in the qualia structure of the nominal objects.

(5.9) John began the stone *(*moving)* / the book *(*destroying)* / the desert *(*crossing)*

5.3.2 Problematic data

There exist data which pose a challenge to the co-composition, type-coercive, accounts outlined above. The interpretation of the logical metonymy construction seems to be severely limited in context-neutral environments (i.e. in isolation of a context which might suggest an event in which the entity denoted in the complement NP is participating), as shown by the questionable interpretations of (5.10)-(5.13).

(5.10) John began the telephone. *(making?? calling with??)*

(5.11) John began the violin. *(making?? studying?)*

(5.12) John enjoyed the bicycle. *(riding?? having?? receiving the gift of??)*

(5.13) John enjoyed the piano. *(playing?? its beauty??)*

On the assumption that all artifacts have eventualities associated with them through the qualia structure, the logical metonymy construction should be possible with every artifact and should have clear interpretations. The data above suggests that this is not true, indicating that the assumption that qualia structure can always supply an eventuality for the interpretation of a logical metonymy construction is too strong.

Furthermore, there are interpretations of metonymies predicted on the basis of the eventualities which seem to satisfy roles in qualia structure which are actually ungrammatical, such as those in (5.14).

(5.14) John began the highway *(*driving on)* / the dictionary *(*consulting)*

As a result of similar observations, it has been argued by Godard and Jayez (1993) and Pustejovsky and Bouillon (1995) that the process of type coercion must be constrained. Their proposals will be discussed below in Section 5.4.1. These approaches, and the approaches deriving from Briscoe *et al* (1990), seem largely to ignore the role of conventionality in interpreting logical metonymies. Analysis of a wider range of data than previously considered,

derived from corpora, reveals more limited possibilities for the application of qualia structure-driven type coercion in explaining logical metonymy than predicted by the highly productive generative accounts proposed, particularly for the aspectual verbs. The existing proposed constraints on metonymy, within approaches that assume full specification of qualia structure for all nouns, don't rule out all implausible metonymies while ruling out some plausible ones. In response to this, in this chapter I will reject many of the proposed constraints on type coercion and offer an account in which logical metonymy depends on lexically-specified conventional information: i.e. properties specified in the semantic component of lexical items will determine whether and how those items can participate in metonymic constructions. In particular, I will show on the basis of the corpus data that not every noun has a telic role specified in its qualia structure, and that metonymies involving the two verbs *begin* and *enjoy* display differing behaviour with respect to the ability of context to influence their interpretation. The more general implication of this analysis is that while understanding of language may proceed in general according to certain rules or processes of interpretation, it is ultimately governed by lexical idiosyncrasies and conventionalised linguistic knowledge about individual words.

On its own, a qualia structure-driven approach to non-eventive NP type coercion is inadequate for explaining the data previously introduced in this chapter. It cannot explain variances in the acceptability of sentences and in the availability of default interpretations in context-neutral environments. Nor can it account for the ability of pragmatics to influence interpretation. We will see that it also cannot account for differences in acceptability which derive from changes in the type of the NP complement, such as those caused by pluralisation. Type coercion in this framework must therefore be constrained in some way.

On the other hand, it is clear that some aspects of logical metonymy should be conventionalised via some mechanism such as qualia structure, rather than relegated to pragmatic control. Qualia structure is a good candidate for capturing the information relevant to metonymy, but it must be utilised within a more constrained lexical system than in current approaches.

A purely pragmatic approach would fail to constrain the possible interpretations of logical metonymies, while a purely lexical approach fails to accommodate the potential contextual influence on these interpretations (as required to explain the possibility of (5.7a) above being interpreted as (5.7d)). A combination of the two approaches is necessary to explain the range of logical metonymy data, a point which echoes observations made by Briscoe *et al.* (1990) and Lascarides & Copestake (1995). In this chapter, I will argue that there exist lexical constraints on the potential contextual influence which have not previously been formalised. I will show that information derived in the lexicon can be used to constrain the possible interpretations of a phrase in such a way that even a strong context cannot override the lexical specifications. Furthermore, I will argue that these specifications could not be relegated to the pragmatic component without a reduction in the generality of the treatment of logical metonymy phenomena. Thus, lexical specification of conventions are necessary, and the pragmatic component must be able to utilise the information coming from the lexicon in the appropriate ways.

5.4 An Analysis

5.4.1 Proposed constraints on Logical Metonymy

Godard & Jayez (1993)

Godard & Jayez (1993) focus on constraints on the process of type coercion itself, rather than identification of the missing event introduced as a result of this process. They argue for an approach to type coercion in which differences in the coercive properties of aspectual verbs such as *begin* and *stop* (motivated with reference to French data but which transfer for the most part directly to English) are specified lexically, and in which coercion does not require type-shifting the NP complement, but is incorporated into a richly structured semantic representation for the aspectual verb. Their approach depends on the existence of candidate events for the missing event prior to the type coercion process, and as such does not provide a solution to the problem of constructing an interpretation for logical metonymies, under examination in this chapter.

Godard & Jayez do provide, however, a proposal for the constraints which the French *commencer* (=*begin*) imposes on its NP complement. These are:

(i) The complement must be bounded.[2] *(whether a non-eventive NP or an event)*

(ii) The reconstructed event is an event in which the objected denoted by the NP is controlled by the entity denoted by the subject of *commencer*. (That is, the subject of the main verb is the entity which "triggers and causally maintains" (Godard & Jayez 1993:172) the event into which the complement object is coerced.)

(iii) The reconstructed event should be a type of *modification*.

They use these constraints as the explanation for a range of excluded type coercions. The first constraint (i) explains the contrast in (5.15) in that "some cheese" refers to an indefinite entity — i.e. an unbounded entity — while "the cheese" refers to a definite, discrete entity. It would, however, incorrectly predict that (5.15c) is ungrammatical. The boundedness constraint does appear to apply accurately to the logical metonymy construction, but not to the fully expanded {begin + VP} construction. The second constraint correctly rules out the sentences in (5.16). The third attempts to explain why the sentences in (5.17) are ruled out, expressing an intuition that the object in the NP usually comes into being, is consumed, or undergoes a change of state – hence, it is modified in some way. These sentences do not express such a modification.

(5.15) a. Jean a commencé le fromage.
John began the cheese.

b. *Jean a commencé du fromage.
John began some cheese. [Godard & Jayez 1993:(24)]

[2]Boundedness is given a technical definitions by Godard & Jayez in terms of Krifka's model of aspectual predicates, essentially corresponding to the atelic / telic distinction. Bounded events or objects are viewed as having a set terminal point, while amorphous events / objects do not.

c. Jean a commencé a manger du fromage.

(5.16) a. *A ce moment Jean a commencé un grand mépris pour les politiciens.
At that moment John began a great contempt for politicians. [Godard & Jayez (31)]

b. *L'acide a commencé la destruction du marbre.
The acid began the destruction of the marble. [Godard & Jayez 1993:(35)]

(5.17) a. *Jean a commencé la pierre.
John began (moving) the stone.

b. *Jean a commencé le désert de Gobi.
John began (to go through) the Gobi desert.

I will generalise from these constraints in my proposals for the logical metonymy constructions, adopting (i) and (ii) as these constraints appear to accurately reflect this data, but replacing (iii) with more specific claims about the reconstructed events.

Let us consider briefly the G&J (1993) proposal that the complement of *begin* must be bounded. While this is clearly not the case on a use of *begin* with a VP complement, it does seem to be true when *begin* appears with an NP complement. Thus we have the contrasts in (5.18). Both (a) and (b) are felicitous and involve a finite set of books. While (c) also involves a finite set of books, it seems to be slightly less good, probably due to pragmatic inferences which specify that it is impossible for one person to read multiple books simultaneously and therefore impossible to begin reading multiple books in a single event. Both (d) and (e) involve unbounded sets; in (d) the reading would have to be a generic one (e.g. *John began reading books when he was five years old*) and therefore does not refer to any specific set of books, and (e) also does not refer to a specific set of books. The constraint proposed by G&J – that the NP complement must itself be bounded – accurately captures these data. P&B attempt to capture these same data by pointing out that in (d) the missing event is an activity. As we saw above, however, aspectual restriction does not seem to capture the metonymy data entirely accurately in general. In fact, it cannot rule out (e) either – in this case although the set of books in question is not specified, it is finite. Thus the event of *reading some books* is still a transition as it cannot go on indefinitely and cannot be ruled out with the aspectual constraint.

(5.18) a. They / John began the book.

b. They began the books.

c. ?John began the books.

d. *They / John began books.

e. *They / John began some books.

(5.19) *John began books / houses / sandwiches.

The "bounded NP" constraint, then, also rules out the data in (5.22).

There does, however, seem to be an aspectual constraint which governs the behaviour of {*begin* + *NP*}. The missing event in the metonymy seems to be constrained to be a single

specific durative event, that is, not iterative or generic. This can be argued on the basis of (5.20) and the contrast between (5.18b) and (5.18c).

(5.20) *John began the dictionary. *(consulting)*

In (5.20), *consulting the dictionary* is a transition. However, it is a point-like transition and a single point-like transition is unlikely to be *begun*. As discussed by Freed (1979), *begin* is one of a class of aspectual verbs which presuppose or entail certain temporal facts about the events named in their complements. She argues that *begin* picks out a particular temporal segment of its complement events. This segment includes part of what she calls the *nucleus* of an event, which is an activity which occurs during a stretch of time. It therefore makes sense that *begin* requires its complements to have some duration. For the instances that don't, either an iterative (in which a series of single events of the same type occurs continuously within a given stretch of time) or a generic (in which a series of single events of the same type occurs at different times) interpretation is inferred. Thus, when we say *John began sneezing*, we generally infer that John entered a time period in which he sneezed multiple times. Similarly, (5.20) would have an iterative interpretation. In contrast, (5.18c) is likely to require a generic interpretation since pragmatics suggests that John cannot begin reading multiple books simultaneously, and so we must interpret the metonymic event as a series of events occuring at different times. (5.18b) has neither an iterative or generic interpretation, because it does not refer to a series of events, it rather specifies quantification over single *begin reading the book* events. By ruling out iterative and generic metonymic events, we rule out (5.20) and a generic interpretation of (5.18c). The latter would only be allowed on a reading with quantification over single events, as in (5.18b). The pragmatic oddity of this reading accounts for the oddity of the sentence.

The Godard & Jayez constraint (iii) does not seem to adequately capture the logical metonymy data. In one of the most frequently cited logical metonymies and my first example, *John began the book*, reading surely does not involve a change of state or modification of the book. Godard & Jayez attempt to justify this by claiming that reading a book involves the imposition of an "informational layer" on the book, derived from interpretation of the text, which constitutes "informational modification". I find this explanation somewhat forced; in any case, the range of corpus data which will be presented below (Section 5.5.2) indicate that there is not a unifying semantic type (e.g. modification) of the derived events in logical metonymy interpretations.

Pustejovsky & Bouillon

Their approach

Pustejovsky & Bouillon (1995) use Godard & Jayez (1993)'s analysis to develop constraints on type coercion in terms of the aspectual properties of the reconstructed event and the raising/control properties of the eventive/coercing verbs. Their analysis attempts to explain the logical metonymy data in terms of the linguistic criterion of aspectual class rather than the vague semantic notion of modification adopted by Godard & Jayez. In doing so, Pustejovsky & Bouillon are able to give a uniform treatment across the different syntactic frames of

eventive verbs, as the constraints they propose are claimed to hold for all uses of these verbs, not just the metonymic uses. I will show below, however, that their analysis cannot account for the full range of logical metonymy data and is too strong in its claim of application across all syntactic frames in which the eventive verbs appear.

The Pustejovsky & Bouillon account relies on a structured representation of events, in which subevents are represented and the "focus" of the event is marked as the *head* of the event structure. The event structure provides "a configuration where events are not only ordered by temporal precedence, but also by relative prominence" (Pustejovsky & Bouillon 1995, p. 141). Pustejovsky & Bouillon provide examples of headed transitions, proposing that *build* is a *left-headed event structure* since the event-focus is on the first subevent, the activity of building, while *arrive* is a *right-headed event structure* as the focus is on the actual arrival rather than the pre-arrival process, and *break* is lexically unspecified with respect to headedness. The relevant event structures are shown in (5.21), with the asterisk indicating the headed sub-event. Left-headed structures correspond to accomplishments, while right-headed structures correspond to achievements.

(5.21)

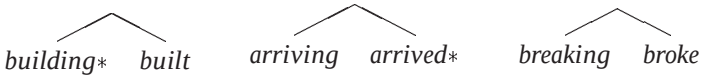

Another element of the Pustejovsky & Bouillon (1995) proposal is the distinction between control and raising aspectual verbs. Using tests proposed by Perlmutter (1970), they argue that *begin* is ambiguous between a control use and a raising use. They build on the insight provided by Godard & Jayez that logical metonymies are restricted to control uses of *begin*.

The constraints on aspectual / eventive verbs proposed by Pustejovsky & Bouillon make use of these two elements, and can be summarised as follows:

- Control use of the verb: the complement of an aspectual / eventive verb on a control use must be a TRANSITION, in particular a *left-headed* transition.
- Raising use of the verb: the complement can have any event type, but no coercion of an NP argument is possible.[3]
- Coercion depends on the availability of an event in qualia structure. If the desired interpretation of a logical metonymy involves an event which is not specified in either the telic or agentive role of the complement noun's qualia structure, that interpretation will be impossible. Note that this is a reflection of Pustejovsky's entirely lexically-based analysis (see Section 5.3), and ignores the ability of the discourse context to facilitate novel coercions.

Sentences like (5.15b) are ruled out under these constraints because *eating some cheese* is an activity, not a left-headed transition. Likewise, the plural sentences in (5.22) (repeated from

[3]This criterion in fact complements the constraint expressed by Godard & Jayez that the object denoted by the NP must be controlled by the subject of *begin* – if the subject of *begin* is not a controller because *begin* is being used as a raising verb, then clearly the NP cannot be controlled by the subject.

(5.19)) are ruled out because *reading books*, *building houses*, and *making sandwiches* are activities. The intended interpretations specified for the sentences in (5.23) are also ruled out because the reconstructed event has the wrong aspectual type for a control reading — the events are activities, rather than transitions.

(5.22) *John began books / houses / sandwiches.

(5.23) John began the highway. *(*driving on)* / dictionary. *(*consulting)*

Sentences like those in (5.16) are ruled out because these are raising uses of *begin*, so NP coercion is impossible. Sentences like those in (5.17) are ruled out simply because the appropriate values are not available in the qualia structures for *stone* and *desert*.

Problems with the approach

The details of the Pustejovsky & Bouillon account of aspectuality in relation to qualia structure are rather vague and would need to be fleshed out for a full evaluation of their approach. They provide no clue as to what event structure is specified for verbs in the lexicon, or how the event structure is built up as a verb is combined with an argument.[4] In the sample representation they offer for *book* (Pustejovsky & Bouillon 1995, p. 151), the events in the qualia structure, *read* and *write*, are specified in the lexical entry as transition events, and no indication is given of how the aspectual type of these events might be shifted (e.g. to account for the aspectual differences between *reading a book* and *reading books*). Their formalisation therefore ignores the influence of the form of the noun on the aspectual nature of the event which may be associated with the NP, although they assume that the aspectual class of an eventuality can be correctly identified in their theoretical discussion. This lack of clarity means that the event headedness of particular eventualities, a critical part of the Pustejovsky & Bouillon approach, is difficult to establish with certainty.

The definition of event headedness is also a problem simply because it is often not clear what the internal structure of an event is or how to establish on which subevent the "focus" is. *Reading a book* is a left-headed transition on the Pustejovsky & Bouillon account, suggesting that it is composed of two subevents: the *reading of the book* event, and the *have read the book* state. The focus of the predicate *read*, then, is on the first of these, the process of reading. *Arriving* is a right-headed transition on their account. Presumably, the two subevents of *arriving* are the *doing something to bring on an arrival* event and the *arrival* event. The focus here is on the actual arrival. But what about *consulting a dictionary*? Analogous to the two events described, one would postulate two subevents for this event: the *consulting of the dictionary* event, and the *have consulted the dictionary* event. Is this a left-headed or a right-headed event, or possibly something which is neutral between the two? The problem with relying on event headedness as a criterion is that it depends on intuitions about event structure which in many

[4]For example, such an event structure can be built up in terms of the transition network ontology proposed by Moens and Steedman (1988) for capturing the influence of tenses, adverbials, and argument type on the aspect of an eventuality. How this network might interact with the Pustejovsky and Bouillon lexical representations is unclear, since they depend on phrasal or sentential-level properties of a sentence referring to an event.

cases are difficult to determine, and difficult to isolate from the context in which the event is described. Pustejovsky and Bouillon do not attempt to define clear tests for determining event structure linguistically, let alone defining an automatic mechanism for determining event structure in the syntax-semantic processor. A fundamental question emerges: if event-headedness is essentially a representation for a verb's aspectual class, what does it buy us over other (simpler) aspectual representations? The answer to this question remains unclear to me.

Even if we assume that somehow the formalism works out such that the eventuality types are appropriately represented, the constraint that *begin* may only appear with left-headed transitions on a control reading is too strong. Each of the sentences in (5.24)-(5.25) have a structure reflecting a clear control use of *begin* (on the basis of Perlmutter (1970)'s tests). In (5.24), we find in each case a control use of *begin* with an eventive complement which is an activity. In (5.25), we again have control uses of *begin*, this time with eventive complements which are achievements, i.e. right-headed transitions.

(5.24) a. Mary forced John to begin building houses.

b. Mary forced John to begin running.

c. Mary forced John to begin discussions.

(5.25) a. Mary forced John to begin leaving the pub.

b. Mary forced John to begin destroying the book.

Pustejovsky & Bouillon base their claim on examples such as those in (5.26).

(5.26) a. *John began to find the house key.

b. ?John began to arrive.

(5.27) a. Mary forced John to leave the pub.

b. Mary forced John to destroy the book.

c. *Mary forced John to find the house key.

d. *Mary forced John to arrive.

However, these sentences pair a control use of *begin* with an event for which the subject cannot be a controller. This distinction is clearly exhibited by the contrasts in (5.27). Thus the sentences in (5.26) are actually instances which require a raising interpretation of *begin*. For discussion of the constraints on this interpretation, see Section 5.5.2.

The Pustejovsky & Bouillon constraint that type coercion is not possible with a raising construction is supported by observations by Jacobson (1990) as reported in Pollard & Sag (1994:141-142), which suggest that raising verbs never allow NPs instead of their VP complements. This constraint is therefore an instance of a more general constraint on raising verbs. The constraint on the complement of aspectual verbs on control uses, however, does not seem to be supported by the data. An alternative explanation for the type coercion phenomena must be sought.

5.4.2 Problematic data remaining

The Pustejovsky & Bouillon (1995) aspectual and qualia-derived restrictions on the control interpretation of *begin* correctly rule out a large number of metonymic sentences, but the sentences in (5.28) would (incorrectly) not be ruled out (overgeneration), since the telic events (in italics in the example) are transitions. In contrast, sentences (5.29), conveying activities, would be ruled out although they are felicitous (undergeneration). More than a simple aspectual restriction is necessary to explain the data.

(5.28) a. *John began the film. *(watching)*

b. *John began the door. *(opening, walking through)*

c. *John began the nails. *(hammering in)*

(5.29) a. John began daycare at his mom's work. *(attending)*

b. John began acupuncture in April and homeopathy in August. *(undergoing)*

c. John began the violin when he was five. *(playing)*

d. The two women began the serious business of the day, gossiping. *(doing)*

e. The two older girls had already begun boarding school. *(attending)*.

5.5 An Alternative

In this section I will isolate uses of eventive verbs in which the verbs take a full VP complement from the metonymic uses in which they take an NP complement. By assuming distinct lexical entries and therefore distinct logical forms for the different uses of these verbs, I follow the proposals of Godard & Jayez (1993) and Copestake & Briscoe (1995) rather than the type-shifting proposals of Pustejovsky (1995a). This analysis is supported by anaphora, relativisation, and coordination data, outlined by Godard & Jayez (1993:169-170) and Copestake & Briscoe (1995:32-33). For example, the coordination of predicates which require complements of different types is possible, as shown in (5.30).

(5.30) a. John picked up and finished his beer.

b. John ate and enjoyed the caviar.

c. Sam wrote but later regretted the article.

In these examples, the first verbal predicate requires a physical object complement, while the second requires an eventive complement, yet there is no problem with their coordination. This suggests that the type of the NP complement cannot itself be shifted. A parallel example from the relativisation data appears in (5.31).

(5.31) a. *John began a book which will take two hours.

b. John began the reading of a book which will take two hours.

Here the relative clauses must be relative to an event. If the type of *a book* in (5.31a) were shifted to the eventive interpretation made explicit in (5.31b), then this example would be grammatical. It is not, however, further evidence that the type of the NP complement is not shifted.

As Copestake & Briscoe (1995:52) point out, however, there is some data which cannot be accounted for under an analysis in which the particular lexical entry for a verb controls the coercion of complement type, shown in (5.32a-b).

(5.32) a. I enjoy films and mending antique clocks.

b. We found Sam swimming the channel, which he enjoys more than golf.

c. *John began the beer and watching the television.

d. *John finished composing his symphony and his novel.

It is difficult to explain these examples under the assumption that one lexical entry for *enjoy* combines with NPs and another with VPs, because there is one instance of the verb which appears to combine happily with both. They argue, however, that these examples occur less frequently and are more marked than the examples in (5.30). Furthermore, the data in (5.32c-d) shows that other eventive verbs do not allow this kind of coordination. I will therefore follow Copestake & Briscoe in preferring the polymorphic analysis of the eventive verbs.

Given the isolation of VP-complement and NP-complement forms of the eventive verbs as a starting point, I will investigate the two uses individually. I will propose simple aspectual constraints on the VP-complement forms of the eventive verbs, but the logical metonymy cases will be shown to be far more complex. No aspectual constraints will suffice for explaining this data; instead I will argue for an analysis which depends on lexical conventions.

My proposals differ from previous proposals in that I do not assume that logical metonymy is a phenomenon which can be fully explained in terms of a generative process governed by general linguistic constraints. My perspective on the logical metonymy data is that if such linguistic constraints exist, they are not straightforwardly identifiable and are much more complex than a simple aspectual or semantic type restriction. My research has not led to the identification of any such constraints. Above we saw that for each of the proposed constraints on logical metonymy, there is a set of examples which violates it. Furthermore, I found that the range of logical metonymies which occurs in a corpus is severely limited, strongly suggesting that this phenomenon is heavily influenced by conventional usage. As a result, I prefer an analysis of the data which derives from the interaction of lexical information about word use with pragmatic reasoning rather than a highly productive process which is constrained by some (arbitrarily complex) set of linguistic restrictions.

5.5.1 {eventive verb + VP}

I propose that the behaviour of an eventive verb with a VP complement can be summarised as follows:

- **Raising Interpretation:**

- The VP complement must be an uncontrolled event.
- The VP complement cannot be an achievement, but can be of any other eventive type.

- **Control Interpretation**:
  - The VP complement must be an event controlled by the subject NP.
  - The VP complement can be of any eventive type.

The distinction between the complements on the raising and control interpretations, in the latter case that the event must be controlled by the subject NP and in the former that the event must be uncontrolled, follows simply from the definition of raising and control verbs. I postulate two distinct lexical entries for the general form {*eventive verb* + *VP*}, corresponding to the raising and control interpretations. This is supported by co-predication data, as shown by (5.33)-(5.34) in constrast to (5.35).

(5.33) a. *John began to read the book and find coffee stains on it.

b. John began to read the book. He began to find coffee stains on it.

(5.34) a. *John began to dance and to feel ill.

b. John began to dance. He began to feel ill.

(5.35) John began to sing and dance / laugh and cry / jump and scream.

These examples show that the raising and control uses of *begin* correspond to distinct lexical entries: (5.33)-(5.34) evidence that it is impossible for a compound VP in which each of the conjunct VPs demand a different interpretation of *begin* to appear with a single instance of *begin*, in contrast to (5.35) in which each of the conjunct VPs requires the same logical interpretation of *begin* (and can therefore co-predicate of the same lexical entry). Therefore it would not be enough to claim that *begin* is ambiguous between the control and raising interpretations on a single use.

The other claims, that the VP complement of *begin* on a control interpretation is unconstrained while it is constrained on a raising interpretation not to be an achievement, are evidenced by the data in (5.36). The judgement of controlled vs. uncontrolled here is made on the basis of Perlmutter's test frame *Mary forced John to X* – if this is an infelicitous sentence, e.g. **Mary forced John to feel ill*, then the event described by *X* is uncontrolled, otherwise, e.g. *Mary forced John to leave the party*, the event is controlled.

(5.36)	a.	The acid began corroding the marble.	[uncontrolled process]
	b.	The guests began arriving.	[uncontrolled process]
	c.	John began to find fleas on his dog.	[uncontrolled process]
	d.	John began to feel ill.	[uncontrolled state]
	e.	*John began arriving.	[uncontrolled achievement]
	f.	*John began to find a flea on his dog.	[uncontrolled achievement]

g. John began to leave the party. [controlled achievement]

h. John began to dance. [controlled process]

i. John began to read the book. [controlled accomplishment]

The aspectual restrictions proposed, and supported by the data, thus conflict with the Pustejovsky & Bouillon claim that *begin* must have a complement of a left-headed transition on a control use and can have a complement of any aspectual type on a raising use. This proposal reflects the data more accurately and remedies the problems raised by the Pustejovsky & Bouillon analysis of sentences such as (5.26) in Section 5.4.1 above.

5.5.2 {eventive verb + NP} and Corpus-supported conventions

To gain a clear view of the possible range of logical metonymy data, it is necessary to look to corpora of spoken and written text. This investigation will give a sense of how widespread the use of logical metonymy is, and how far qualia structure can go towards predicting the interpretations that the examples are given in context. I have therefore consulted both the Lancaster-Oslo/Bergen (LOB) Corpus and the British National Corpus (BNC) for relevant data.[5] These are corpora of British English, the LOB incorporating 500 written text samples of about 2,000 words each, the BNC consisting of over 100 million words of a wide variety of written text (90%) and transcripts of spoken language (10%). I focused mainly on logical metonymies with *begin*, but also looked at those for *finish*.

The Analysis: Methodology and Assumptions

The first phase of the corpus analysis involved extracting all sentences containing *begin* and *finish* in any of their inflected forms. I did this in the case of the LOB simply by utilising regular expression matching commands standard under UNIX. The BNC demanded a more advanced mechanism due to its size and structure, and hence I used the CORSET corpus search toolkit (Corley 1996) . This tool enabled extraction of only those sentences which do not have an explicit VP complement. Further narrowing of the sentences to be examined was accomplished via regular expression matching, picking out only those sentences in which the aspectual verbs were followed by a noun phrase. At the end of this phase, I had a large collection of sentences which were potentially metonymic, as each sentence contained an aspectual verb followed by a non-VP element.

The second phase involved identifying which of the collected sentences were actually metonymic. As neither of the two corpora (or indeed any corpus to which I have access) contains any kind of semantic tagging, much of this work had to be done by hand. The following cases were eliminated from the sets of sentences:

- Sentences in which *beginning* or *finish* are used as a noun, e.g. *From the very beginning the Section Office asked groups for help* and *The finish of the wood* . . . .

[5] See `http://www.hd.uib.no/cd-info.html` for additional information about the LOB and `http://info.ox.ac.uk:80/bnc/` for more information about the BNC.

- Sentences in which the noun phrase following the aspectual verb was not a complement of the verb, e.g. *Even as the show begins several key members appear to be drunk.*
- Sentences in which the aspectual verb appears as part of a larger phrase which seems to impose different interpretation constraints on the phrase than on the metonymic constructions. These include *begin X with, begin X by, finish X with, finish X off, finish X by* (*begin/finish the chapter with the words "[…]", begin/finish the book by giving you an example, finish things off*).
- Sentences containing different senses of the aspectual verbs: the sense of *begin* meaning *found* or *launch* (*I began my business in 1983*), the sense meaning (approximately) *to initiate* (*She began a reassuring smile, Thence begins the ecological richness of Africa, Then began the notion/habit/ritual of […]*); the sense of *finish* meaning *use up* (*finish the ammunition, finish the toilet paper*), and the sense meaning *end* (*The poses finishing the musical phrase*).
- Sentences in which the noun phrase complement is **eventive**, that is directly expressing an event such that no metonymic construction is necessary. This case includes deverbal nouns (*begin a look, begin the cut, begin the inspection*) and other instances such as *begin the game* and *begin a diet*. In these cases, no type coercion is necessary to satisfy the requirements of the aspectual verbs.
- Sentences in which the noun phrase complement is **temporal**, that is referring to something with temporal extent (*begin a relationship, begin the first term of school*). Again, no type coercion is necessary in these cases: the aspectual verbs pick out a particular region of the temporal extent associated with these nouns.
- Certain sentences containing **"event-objects"** as complements: dual nature NPs which seem to have a natural interpretation as an event, but which can also be referred to as an object. These can either be interpreted as events directly or, with certain restrictions, metonymically on an object interpretation, as shown in (5.37). Only the metonymic uses, determined by reading the sentences in context, were included in the analysis.

(5.37) a. John began the speech / lesson. — classified as eventive
b. John began the speech / lesson. (*writing*) — metonymic
c. *John began the speech / lesson. *(hearing, listening to)* — failed metonymy
d. John began the lessons. (*taking*) — metonymic

The third and final phase was to read through each of the metonymic sentences in context in order to determine the interpretation intended by the speaker / author. In addition, the whole procedure was repeated to find metonymic instances of the phrase *begin on* followed by a noun phrase for purposes of comparison with *begin*. The results of these analyses will be introduced in the next section. They provide the basis for the alternate conception of logical metonymy to be put forth in section 5.5.2.

		Begin	Finish	Begin on
Agentive				
	write WRITTEN_OBJECT	34 = 20.7%	44 = 13.8%	0 = 0%
	say sentence	4 = 2.4%	7 = 2.2%	0 = 0%
	have family	1 = 0.6%	0 = 0%	0 = 0%
	OTHER AGENTIVE	26 = 15.9%	43 = 13.5%	4 = 16%
Subtotal Agentive		65 = 39.6%	94 = 29.5%	4 = 16%
Telic				
	eat FOOD / MEAL	11 = 6.7%	87 = 27.3%	1 = 4%
	drink LIQUID	0 = 0%	55 = 17.2%	0 = 0%
	tell STORY	20 = 12.2%	7 = 2.2%	0 = 0%
	attend school / classes	4 = 2.4%	22 = 6.9%	0 = 0%
	sing SONG	9 = 5.5%	6 = 1.9%	0 = 0%
	play MUSIC	7.5 = 4.6%	2 = 0.6%	1 = 4%
	read WRITTEN_OBJECT	2.5 = 1.5%	10 = 3.1%	0 = 0%
	serve (jail) sentence	8 = 4.9%	0 = 0%	0 = 0%
	smoke cigarette	0 = 0%	5 = 1.6%	0 = 0%
	do business	6 = 3.7%	1 = 0.3%	0 = 0%
	take MEDICINE / TREATMENT	5 = 3.0%	1 = 0.3%	2 = 8%
	do DEGREE	1 = 0.6%	6 = 1.9%	0 = 0%
	do homework	0 = 0%	5 = 1.6%	0 = 0%
	play GAME	4 = 2.4%	0 = 0%	0 = 0%
	argue CASE	3 = 1.8%	0 = 0%	0 = 0%
	present EVIDENCE	1 = 0.6%	2 = 0.6%	0 = 0%
	OTHER TELIC	9 = 5.5%	2 = 0.6%	1 = 4%
Subtotal Telic		91 = 55.4%	211 = 66.1%	5 = 20%
Context		8 = 4.9%	14 = 4.4%	16 = 64%

Table 5.1: Distribution of Logical Metonymies for *begin*, *finish*, and *begin on* in the BNC

The Data

The number of occurrences of logical metonymies involving *begin* in the corpora is very low, suggesting that this is not a pervasive phenomenon, at least for this particular verb. Only three relevant instances of *begin* metonymies were found in the LOB, and only 164 were found in the BNC. More significant are the relative figures: there are 40,407 sentences containing verbal forms of *begin* in the BNC. After phase 1 of the analysis of the BNC 4,470 sentences containing *begin* followed by a noun phrase remained as potentially metonymic. Only 3.67% of these were actually metonymic — 0.41% of all sentences containing verbal forms of *begin*. The LOB figures are even lower. Out of 172 sentences containing *begin* followed by a noun phrase in the LOB, only 1.7% of these were metonymic.

For *finish*, the frequency of the phenomenon is much higher. Out of 11,072 uses of finish as a verb, there were 2,799 occurrences of *finish* followed by a noun phrase (25.3%). Out of a sample of 940 of these which were analysed, more than a third (319) were metonymic. The distribution of the metonymy types for both of these aspectual verbs as compared to that for *begin on* is, however, interesting to look at. This information is summarised in table 5.1.[6]

This table shows the number and proportion, relative to the whole set of metonymies for the particular aspectual verb in question (as indicated by the column heading), of each different metonymic interpretation found in the corpus. They are divided into **Agentive**, **Telic**, and **Context**, indicating where the missing event needed in the interpretation seems to come from, assuming the existence of qualia structure. The noun phrase following the eventive verb in the interpretations listed either corresponds to the noun phrase actually found in the corpus examples (those in regular font) or to a generalisation over the types of noun phrases which were found on that interpretation in the corpus (those in small capitals). So, for example, the value 34 in the first row of agentive interpretations for *begin* indicates that 34 instances of *begin* WRITTEN_OBJECT (such as *begin the book* and *begin the diary*) occurred in the BNC on a *writing* interpretation. The OTHER AGENTIVE category includes all examples specifying the creating of an object in the way appropriate to that object, such as *begin (digging) the tunnel* and *begin (painting) the portrait*, while the OTHER TELIC category groups together very low frequency occurrences of other interpretations which seem to come from the telic role of the noun's qualia structure rather than from context.

It is clear from the table that *begin* and *finish* metonymies tend to be interpreted on the basis of information in the qualia structure of their complement noun rather than through the influence of context, as is evident from the high total percentages (95% and 95.6%, respectively) of *begin* and *finish* metonymies which are resolved on the basis of qualia information. Context plays a much larger role for the interpretation of *begin on* (accounting for the interpretation of 64% of the metonymies), suggesting a fundamental difference between the aspectual verbs and *begin on* with respect to the interaction of the verb with the meaning of the complement noun. This will have implications for the treatment of these verbs as proposed in Section 5.5.3.

An inspection of table 5.1 clearly slows that the range of possible metonymies is actually quite small. If we consider the size of the BNC, the fact that telic metonymies occur for only approximately 20 different categories of nouns is striking. The BNC contains a varied collection of texts, covering many topics and therefore certainly referring to far more than 20 categories of nouns. We would expect a far wider range of metonymies on the assumption of a fully specified qualia structure for all nouns: in terms of the Pustejovksy qualia-based theory, every noun should be able to appear in the logical metonymy construction. Although a corpus cannot provide negative examples of the use of the construction — that is, ungrammatical or infelicitous examples — and therefore does not give conclusive evidence in favour of the restrictiveness of telic role-centred interpretation of metonymies and does not allow us to develop a theory solely on the basis of this corpus data, the data retrieved in this

[6] A tally of one half was allocated to each of two possible options if the intended interpretation of the metonymy was not entirely clear from the context. For example, *begin the psalms* could mean *begin reading the psalms* or *begin singing the psalms*.

analysis do suggest that the use of the construction is limited to only a few specific cases. In addition, the use of the results of linguistic analysis (e.g. judgements on infelicitous data) allows us to overcome this "zero-data" problem (see Chapter 6 for some discussion of this) by providing evidence of ungrammatical cases and enabling the development of a theory about this phenomenon. Agentive role-centred metonymies occur for a wider range of objects, which share the property of being artifacts or at least having a clear way in which they come into existence. The information in the agentive role is therefore much more uniform across objects than that in the telic role: all agentive events are creation events. These facts will form the basis of our proposals about logical metonymy in the next section.

Conclusions about Logical Metonymy

The data introduced in the previous section strongly suggests that the metonymic construction is only used with the aspectual verbs *begin* and *finish* if the intended event is conventionally associated with the noun phrase. I propose, therefore, that the use of the metonymic construction with aspectual verbs is restricted to either agentive events or **conventionalised** telic events associated with the complement noun phrase. This follows from the low overall frequency of metonymic constructions and the restricted range of their application. The influence of conventionality is also supported by observations made by Lascarides and Copestake (1995) of the oddity of certain logical metonymies which would not be predicted if the telic role were derived from real world knowledge of an entity. For (5.38), imagine that the doorstop referred to is a book.

(5.38) ?John enjoyed the doorstop.

Despite the knowledge that the doorstop is a book, this sentence is strange. Lascarides and Copestake therefore also argue that logical metonymy is partially conventionalised, stemming from lexical specification of word use. Next I will attempt to establish a clearer formalisation of that conventionalisation.

Metonymic sentences with *begin* and *finish* seem largely acceptable on a "coming into existence" interpretation — i.e. one in which the reconstructed event comes from the agentive role of the qualia structure — particularly if context aids in arriving at such an interpretation.

(5.39) John began the book *(writing)* / house *(building)* / dictionary *(compiling)* / cigarette *(rolling)* / highway *(constructing)* / door *(making)* / car *(building)* / cake *(baking)*

Thus the data in (5.39) are certainly possible, although some may require contextual reinforcement.[7] This conclusion is supported by the high proportion of agentive metonymies in the corpus, as reflected by the Agentive section of Table 5.1, which shows that 39.6% of the

[7]For example, in (5.40), the addition of context in the form of a more explicit subject in (b), for which certain world knowledge exists, gives (b) a clear interpretation which is not so readily available in (a).

(5.40) a. John began the door. *(??)*

b. The carpenter began the door. *(making)*

begin metonymies and 29.5% of the *finish* metonymies are resolved with respect to the agentive role of the complement NP. A possible explanation of the general availability of the agentive role in metonymies is due to the semantic uniformity of this role across objects, as mentioned above.

Interpretations other than the coming-into-existence (agentive) interpretation seem to be largely impossible, except in the restricted cases indicated by the Telic section of Table 5.1, exemplified by the selection of sentences in (5.41)-(5.47).

(5.41) John began the sandwich. *(eating)*

(5.42) John began the beer. *(drinking)*

(5.43) John began the story. *(telling)*

(5.44) John began the book. *(reading/writing)*

(5.45) John began the cigarette. *(smoking)*

(5.46) John began the solo. *(playing)*

(5.47) John began the song. *(singing)*

These cases are highly conventionalised. That is, they occur relatively frequently in English, are highly idiosyncratic, and are likely learned as individual instances of a more general logical metonymy construction. Their usage is triggered by lexical specification of the event conventionally associated with a particular noun on a metonymic usage. As mentioned in Chapter 1, I assume that there is justification for postulating the lexicalisation of knowledge only when a particular linguistic phenomenon cannot be explained purely in terms of world knowledge. This applies to qualia structure: although general knowledge and experience of objects provide potential values for their telic roles, this information is only lexicalised when it is relevant to some linguistic phenomenon. In the case of logical metonymies, the distinction between the felicitous sentences in (5.41)-(5.47) and the infelicitous sentences in (5.28) cannot be explained solely on the basis of world knowledge of the relevant objects, since each of the events which are intended to be coerced from the complement noun are commonly associated with the denoted object and there is thus no pragmatic distinction to be made. Nor, as argued above, can this distinction be made on the basis of Pustejovsky & Bouillon's aspectual restrictions or Godard & Jayez's semantic constraints on the general metonymic process. Instead, accurate modeling of the data can only be achieved through *selective lexicalisation* — lexicalisation only of highly conventionalised events which can be accessed in the metonymic process. Events which are not available in metonymic coercions form part of world knowledge of an object, but must not be lexicalised.

Thus although feasible metonymic interpretations generally correspond to one of Pustejovsky's agentive or telic roles (i.e. contextual information does not often override the core semantic information found in the qualia structure), these do not always provide a feasible metonymic interpretation, as in the sentences in (5.28). The infelicity of the specified interpretations of this data can, however, be explained via the assumption that for these cases, there is no event specified in the telic role of the nouns in the NP.

Further evidence in support of this analysis comes from comparison of the sentences in (5.41)-(5.47) to ones containing semantically related nouns, as shown in (5.48)-(5.51).

(5.48) ?John began the apple. *(eating)*

(5.49) ?John began the orange juice. *(drinking)*

(5.50) ?John began the flyer. *(reading)*

(5.51) ?John began the pipe. *(smoking)*

Since it is reasonable to assume that the nouns in (5.48)-(5.51) would have the same qualia values in their lexical semantic structure as their counterparts in (5.41)-(5.45), the difference in acceptability between these sentences apparently cannot be explained without recourse to idiosyncratic lexical properties of individual words.

The approach outlined above is preferable to one which assumes that all nouns have a telic role specified and then rules some out with various constraints, since telic-role metonymy seems to be possible in only very limited instances, on highly conventionalised uses. However, even an account which assumes that such lexical defaults are optional and may or may not be specified for any given object must provide a justification for the absence or presence of that default on some basis independent of the metonymy data for individual verbs; otherwise the default simply serves as a marker for whether or not metonymy can occur when a particular verb appears with a particular object. We can provide this independent justification on the basis of aspectual verbs other than *begin* and of "metonymic adjectives" – i.e. adjectives which modify events rather than or in addition to objects, which therefore may require type coercion of the object.

The presence or absence of a particular metonymy holds across all eventive verbs, as shown in (5.52). Similarly, the application of metonymic adjectives in (5.53) display the same patterns.

(5.52) a. John began / finished / started / enjoyed the book. *(writing/reading)*

b. John began / finished / started / enjoyed the nails. *(making/*hammering in)*

(5.53) a. i. John read a long book. *(physical length large/num pages large/long to read)*

ii. This is a fast book. *(fast to write/fast to read)*

b. i. John hammered in the long nails. *(physical length large/*long to hammer in)*

ii. This is a fast nail. *(?fast to make/*fast to hammer in)*

There is additional evidence from data involving temporal prepositions, which also require eventive (and hence temporal) arguments. Consider the sentences in (5.54), which display the same patterns of interpretation as the previous examples.

(5.54) a. After five books, John was considered an expert. *(writing, reading)*

b. After fifteen nails, John needed a break. *(making, ?hammering in)*

These examples show that the presence or absence of a telic role for a particular object is consistent across metonymic constructions.[8] However, there are certain inconsistencies in the data. For example, a sentence such as (5.55a) seems to have a natural interpretation with a temporal preposition and in the *enjoy* metonymy which is not available in the *begin* metonymy.

(5.55) a. After three stories, the children wanted to go outside. *(hearing, listening to)*

b. The children began the stories. *(*hearing, *listening to)*

c. The children enjoyed the stories. *(hearing, listening to)*

Such examples will be considered below in the discussion of *enjoy*. The infelicity of (5.55b) will be shown to follow from an additional constraint which *begin* imposes on its NP complement in the metonymy construction.

My proposal is therefore that for certain artifacts a telic event is conventionalised (whether there is some reason for this conventionalisation other than frequency is presently unknown) and therefore specified lexically, available for access in metonymic constructions. For other artifacts such a telic event is not conventionalised or specified lexically and is not available in metonymic interpretation. Thus the nouns in (5.28) will not have telic roles specified, while the nouns in (5.29) and (5.41)-(5.47) will.

{**begin on + NP**}

In contrast, the metonymic phrase {begin on + NP} generally seems to serve as a syntactic marker for pragmatic interpretation. Consider the examples from the BNC in (5.56).

(5.56) a. He parted Jean-Paul's thighs and began on his legs. *(massaging)*

b. Perhaps we began on Elizabeth Bowen's The Death of the Heart. *(discussing the book)*

c. Sometimes I will begin on a central motif such as a tree trunk right in the middle of a blank sheet. *(drawing)*

d. Then she began on me. *(undressing)*

The intended interpretations of the {begin on + NP} phrases indicated in italics clearly can only be established in context. Therefore, I propose that this structure indicates only that something is being done with the NP object, leaving a more specific interpretation to be established using contextual information. It does not need to look to the qualia structure for a default interpretation, as context will provide the interpretation. This proposal is supported

[8]Note, however, that there exist analyses of other phenomena in terms of qualia structure which depend on the telic role to be filled in in cases for which I would claim it isn't. Johnston and Busa (1996), for example, argue for an analysis of the interpretation of nominal compounds which depends on information in the telic role. So the interpretation of *bread knife* as a knife for cutting bread stems from a cutting event represented in the telic role of *knife*. I would clearly not want to allow the coercion of *John began the knife* to *John began cutting the knife*. This particular example could be explained by a restriction that the events coerced from an object must involve that object as UNDergoer (as is in fact represented in the lexical entries for *begin* presented in Section 5.5.3), but it is unclear how conflicts arising from similar examples could be reconciled. This remains a point for future investigation.

by the high percentage of cases of this construction in the BNC which could only be interpreted on the basis of the surrounding context, as indicated in Table 5.1. Any apparently default interpretations of {begin on + NP} sentences (such as a default interpretation of *John began on the book* as *John began reading the book*) should not be considered a reflection of lexical defaults, but rather of the most likely pragmatic interpretation of such sentences in isolation of a defining context.

Metonymic "enjoy"

The numerical data I have introduced have not included an investigation of the use of the verb *enjoy*, due mainly to the high frequency of this verb in the corpus and the lack of effective semantically based extraction mechanisms. However, for completeness and since some discussion in the literature on logical metonymy (e.g. Lascarides & Copestake 1995) concentrates on *enjoy*, the occurrences of {enjoy + NP} were extracted and briefly examined. The conclusions drawn from this examination are that the use of *enjoy* is less constrained than that of *begin* and the other aspectual verbs.

Examples involving *enjoy* which seem perfectly felicitous can be contrasted with parallel examples involving *begin*, as in (5.57)-(5.58).

(5.57) a. John enjoyed the symphony. *(listening to)*

b. *John began the symphony. *(listening to)*

(5.58) a. John enjoyed the film. *(watching)*

b. *John began the film. *(watching)*

(5.59) a. The goat enjoyed the book. *(eating)*

b. *The goat began the book. *(eating)*

Examples (5.57) and (5.58) show that there are interpretations available for metonymic *enjoy* which are not available for metonymic *begin*. Example (5.59) suggests that the default interpretations for *enjoy* are more easily overrideable by preferences deriving in the discourse context than those for *begin*.

However, Lascarides and Copestake (1995) point out that the acceptability of certain *enjoy* sentences cannot be improved with the addition of context, similar to the *begin* sentences, as I mentioned in Section 5.5.2. Some examples are found in (5.60). They argue that this kind of data supports the hypothesis that the process of type coercion is partially conventionalised, motivating a lexical encoding of exceptions.

(5.60) a. ?John enjoyed the path / tunnel / Gobi desert / dictionary / door.

b. ?John enjoyed the doorstop.

c. ?John enjoyed Route 280.

The range of metonymic facts for *enjoy* therefore differs from that of the aspectual verbs, and must be explained within the framework I have suggested for logical metonymy on the

basis of the *begin* data. This implies that cases in which *enjoy* does not pattern with *begin*, such as (5.57)-(5.59), must be explained.

Note first that the majority of cases which are infelicitous for *begin* are also somewhat difficult to interpret with *enjoy* (e.g. (5.60)). These are instances for which I would argue that the telic role of the complement noun phrases is not lexically specified. In this way, just as for the corresponding *begin* sentences, the oddity of this data is explained. There will be no default interpretation for these sentences proposed by the lexicon, hence their interpretation will depend on pragmatic reasoning. So the interpretation of these examples would likely be eased with reference to a rich context supporting a particular interpretation; if the oddity does not go away for particular instances I postulate that this results from the complexity of the pragmatic processing necessary to establish the interpretation. I will not, however, investigate this hypothesis.

The slight difference in the grammaticality judgements of the corresponding *begin* and *enjoy* sentences in (5.60) (* vs. ?), and the differing judgements of the examples in (5.59) stem from the same source, which I will elaborate in Section 5.5.4. I will suggest there that the lexically proposed interpretations of the *begin* sentences (including those instances where there is no lexically proposed interpretation because there is no value specified in the qualia structure of the complement NP) are indefeasible, that is that they cannot be overridden by contextual or pragmatic information. This is an idiosyncrasy of the behaviour of the aspectual verbs in metonymic environments. The verb *enjoy* is not subject to the same constraint and can therefore be influenced by pragmatic reasoning. The sentences in (5.60) are odd because they have no lexically-derived interpretation (on the assumption that the telic role is not represented in the qualia structure of the complement nouns) and possibly because it is difficult to establish a "neutral" context to enable their interpretation,[9] but they are not ungrammatical because they do not violate any linguistic or discourse constraints. The sentence (5.59a) is fully felicitous because the subject NP *the goat* provides enough contextual information to enable a pragmatically induced interpretation which overrides the lexical default of *reading*.

Examples (5.57)-(5.58) remain as problematic. Notice that both of these examples involve "event-objects". Recall that these are NPs which have a natural interpretation as an event (e.g. the performance of the symphony, the playing of the film), but which can also be referred to as an object. I have not been able to come up with any unrelated (i.e. not involving an event-object) examples in which there is such a clear grammatical difference between *begin* and *enjoy* and for which that difference cannot be explained on the basis of pragmatic reasoning. This suggests that there is something unusual about metonymies with this type of NP.

[9]Due to the assumption of the conventional nature of qualia structure, I must depend on a pragmatic account of *enjoy* for those cases in which it doesn't pattern with *begin*. In this section I address most of those cases, but still must account for apparent "default" interpretations of *enjoy* metonymies which do not have complement NPs with a conventionalised telic role. These "default" interpretations must be seen as stemming from pragmatic defaults rather than lexical defaults, that is from reasoning about the most likely interpretation of these sentences given world knowledge. So a sentence like *John enjoyed the pipe* will be assigned, based on lexical processing, an underspecified interpretation of *John enjoyed doing something with the pipe*. This interpretation would then be made more specific on the basis of knowledge about pipes. A "default" interpretation would come from knowledge of the most common activity that is performed with a pipe, e.g. smoking.

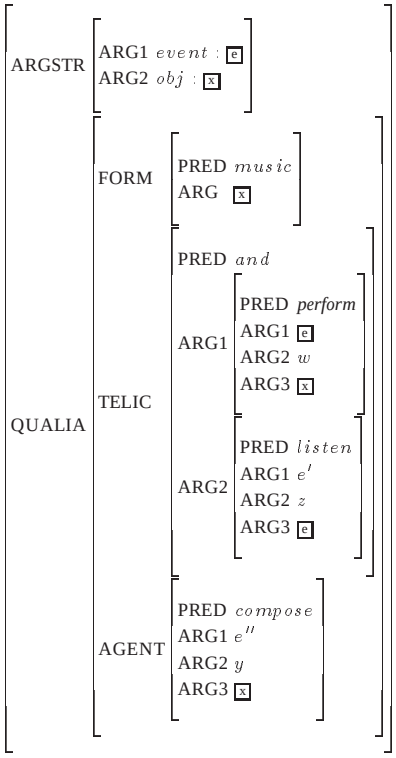

Figure 5.1: The Pustejovsky & Bouillon (1995) proposal for the lexical entry for *sonata*

Pustejovsky & Bouillon (1995) propose that the structure of the event-object *sonata* should be as in Figure 5.1, in their representational notation. In this structure, both the event aspect and the object aspect of the event-object are represented, and the telic role has a compound value which relates the two aspects. This lexical entry indicates that the object *sonata* is performed, and this performance is the event *sonata,* whose function is to be listened to. The structures for *symphony* and *film* can be considered analogous.

The sentence (5.57a), then, does not convey *John enjoyed watching the film-object* (which could be paraphrased as *John liked staring at the film reel*), but rather *John enjoyed watching the film-event* or *John enjoyed watching the playing of the film-object*. The contrast between (5.57a) and (5.57b) can be explained if *begin* does not allow the event coerced from the qualia structure of its complement NP to itself involve an event. The difference between these examples therefore indicates a semantic difference between *enjoy* and *begin* which must be lexically specified.[10] Specifically, it appears that the NP complement of the coercing form *begin* must be constrained to refer to an object, while the complement of the coercing form of *enjoy* is not subject to that

[10]In my representation in Chapter 2, I have not allowed for compound events within the qualia structure of a noun and so this precise structure could not be implemented directly under the proposals there. A similar effect would be achieved by allowing eventive nouns to have qualia structure. There would thus be two lexical entries for the event-objects, one corresponding to the event and the other to the object. The latter would have a telic role of playing / performing while the former would have a telic role of watching / listening. The coercion of the eventive form of the NP could proceed in precisely the same manner as with the object forms.

constraint.[11].This difference will be reflected in the lexical entries proposed in Section 5.5.3. The sentence (5.57a) can now be explained by assuming that the metonymic coercion occurs relative to the event interpretation of the event-object *symphony*, which will not be possible for *begin* since it must coerce relative to the object interpretation.

Returning to the data above in (5.55), then, it appears that temporal prepositions pattern with *enjoy* in allowing coercion relative to an event with elaborated qualia structure. Therefore, (5.55a) can be interpreted as *After listening to the telling of three stories, the children wanted to go outside.* The event of *telling* is in the telic role of the object interpretation of *story*, and the event *listening* is in the telic role of the event interpretation. The coercion in this example occurs relative to the event interpretation.

5.5.3 The Representation

To sum up, the behaviour of the eventive verbs in logical metonymy constructions can be described as follows:

I. Aspectual verbs are control verbs in the logical metonymy constructions.

II. The complement NP in the logical metonymy construction must refer to a bounded entity (from Godard & Jayez).

III. The complement NP of an aspectual verb in the logical metonymy construction must refer to an object.

IV. Default interpretations for the logical metonymy construction stem from either the agentive role of the complement NP's qualia structure, or, if present, the telic role of that qualia structure. These default interpretations thereby derive only from lexical specification of information in qualia structure for particular NPs. For some eventive verbs (e.g. *begin on*), no default interpretation is derived from lexical information. This difference must be lexically captured.

On the basis of the above observations, and building on the representation established in Chapter 2, I will propose two lexical entries for coercing uses of *begin* in this section. Other aspectual verbs will be represented analogously. A representation for nouns as in Figure 5.2 is assumed, in which qualia structure is distinct from other nominal semantics (following Lascarides & Copestake 1995). Values are specified as default, necessary in an inheritance framework when subtypes of a lexical type require distinct qualia values, via the slash notation introduced in Section 2.6: information to the left of the slash (if any) is indefeasible and that to the right is defeasible. I use these typed default feature structures (TDFSs) to allow for the treatment of the influence of pragmatics, to be discussed further in Section 5.5.4, because I need the lexical defaults to persist into the logical form of the interpretation of a sentence so

[11]I have not attempted to explain the origin of this constraint. It may result from the fact that *begin* must be a control verb in the metonymic construction, in contrast to *enjoy* (cf. the infelicity of *Mary forced John to enjoy reading the book). Whether an explanation can be found on this basis remains a question for future work.

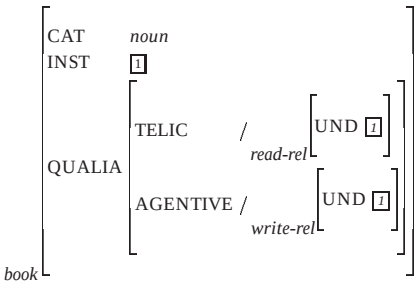

Figure 5.2: Part of the lexical entry for *book*, specifying default telic and agentive events

they are available for interaction with pragmatic reasoning and can be overridden through that reasoning. I further assume that the roles of qualia structure representations for different NP forms are "computed" appropriately using the qualia structure of the head noun.

I also assume a type hierarchy for *begin* which corresponds to that proposed in Copestake & Briscoe (1995) for *enjoy*: a general *begin* type has subtypes corresponding to its different complement forms, each of which has a specific semantic definition. The uses of *begin* in various constructions must be distinguished from one another, as has been discussed previously in this chapter, because each incorporates the semantic contribution of its complement in a different way. Thus a polymorphic treatment of metonymic verbs has been adopted, to allow for different constraints to be implemented for each use of *begin*.

The missing event in a *begin* logical metonymy can either come from the agentive or telic roles, as specified in Figure 5.3. The two lexical entries for the coercing forms of *begin* define the behaviour of this verb on metonymic uses.

I argued in the previous section that the interpretation for the missing event in a {*begin on + NP*} construction is filled in by context. As a result of the context-dependent properties of this form, no defaults / conventions specified in the lexicon are accessed via unification for this phrase; pragmatics must provide all the clues to an interpretation. The lexical entry in Figure 5.4 captures the proposal for the behaviour of *begin* in this construction.

The lexical entry for the coercing form of *enjoy* is shown in Figure 5.5. When it is combined with the representation of an NP argument which has a telic role specified, the default unification framework ensures that the default information specified in the qualia structure for the NP remains as default in the representation of the metonymic event. Thus the metonymic event predicted for *enjoy the book* would be *act_on_pred/read*. This allows for the possibility that *read* can be overridden pragmatically.

If no telic role is specified for an NP, as will be the case for many potential arguments of *enjoy*, we must account for the fact that when *enjoy* is combined with such an NP a default interpretation is still arrived at. I propose that the information which comes from the lexicon into the pragmatic component will be as represented in Figure 5.5 except that the metonymic event will simply correspond to *act_on_pred*, without any default event coming from the NP. Any default interpretation for an {*enjoy + NP*} construction with these NPs is then determined in the pragmatic component, by world knowledge. Clearly world knowledge includes infor-

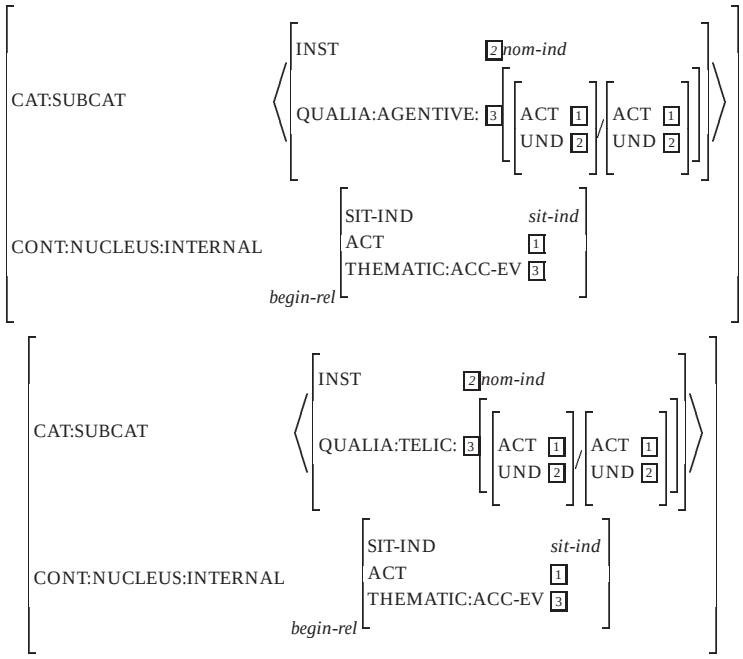

Figure 5.3: Coercing forms of *begin*: One entry picks out agentive role, the other telic role

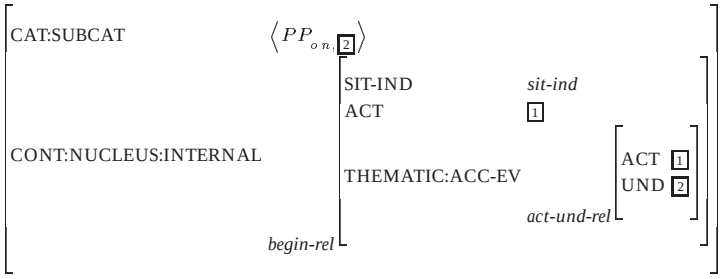

Figure 5.4: Lexical entry for *begin* in a {*begin on* + *NP*} construction

mation about what kinds of events are associated with particular objects. Some of these events will be more strongly associated with the objects than others. It is these strongly associated events which will provide the default interpretation for the *act_on_pred* metonymic event.

5.5.4 Logical Metonymy and Pragmatics

It is clear from the previous discussion that context can sometimes influence the interpretation of metonymic constructions. Although this does not occur frequently for *begin* and *finish*, it does for *begin on* and also *enjoy*. The interaction between logical metonymy and pragmatics must therefore be considered.

Lascarides & Copestake (1995) (L&C) extend the lexical approaches to logical metonymy to develop a system which takes into account the influence of context on the reconstruction of an event in a coercive environment. They utilise the idea that lexical defaults — defaults specified

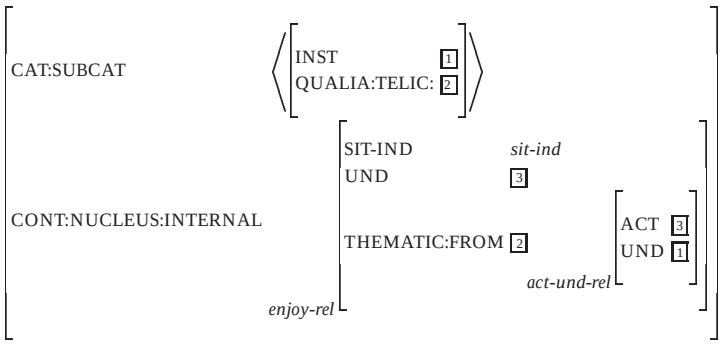

Figure 5.5: Lexical entry for *enjoy* on a coercing use

in qualia structure — persist as default beyond the lexicon into the pragmatic component, and are therefore used in pragmatic reasoning,

L&C formalise their approach in a unification-based framework, with a theory of lexical structure in which the lexical entries are built up from TDFSs so that non-default and default information is demarcated, and which utilises persistent default unification (PDU) (Lascarides *et al* 1996) in order to allow defaults to persist into logical form. The lexicon is hierarchically ordered in the manner adopted in Chapter 2 (as described in e.g. Copestake (1992)). They adopt Pustejovsky's notion of qualia structure, but augment it with the notion of defeasibility. So the values in qualia structure provide lexical information which is treated as default. For example, the telic role of *book* is *read* by default. Other aspects of lexical representation follow the lexical representation language (LRL) (Copestake 1993).

In the pragmatic component DICE (Lascarides and Asher 1991), L&C propose two axioms, i) `Defaults Survive:` lexical generalisations normally apply in a discourse context and ii) `Discourse Wins:` conflicting discourse information wins over lexical defaults. The lexicon links to a pragmatic component via these axioms to interact with discourse context. The axioms together can be used to explain when (5.61a) entails (5.61b) rather than (5.61c).

(5.61) a. My goat eats anything. He really enjoyed your book.

b. The goat enjoyed eating your book.

c. The goat enjoyed reading your book.

Since L&C build on the Pustejovsky approach to logical metonymy, they adopt the assumption of full representation of qualia structure and rely on constraints applied prior to pragmatic processing to rule out implausible metonymies. The approach therefore suffers from the same over- and under-generation as purely lexical approaches under this assumption (see Section 5.4.2). However, it is fully compatible with the view on what information is lexically specified as presented previously in this chapter. I will outline how integration of the underspecified qualia structure with the L&C model of pragmatic reasoning with defaults can account for constraints on the interaction of logical metonymy with discourse context in the discussion which follows.

It is sometimes assumed that with a sufficiently strong context any interpretation of a semantically underspecified sentence can be coerced, as is suggested by the example (5.61). The L&C work, for example, starts from the assumption that contextual information can always override lexical defaults in a logical metonymy. Even when context is added to aid in the interpretation of metonymic constructions, however, there are cases in which the constructions seem impossible despite the availability of an eventive interpretation of the noun phrase from context, as suggested Section 5.2. The example (5.8) presented there is repeated here as (5.62). (5.62b) is infelicitous as a follow-up to (5.62a), despite the strong pragmatic clues to interpret *begin* as *begin destroying*. A construction like (5.62c), however, is grammatical and *begin on* is interpreted as *begin destroying*. A similar pattern holds in (5.63).

(5.62) a. John will be audited by the tax service, so he has been destroying things which might incriminate him. He has destroyed the files and the computer disks.

b. *He will begin the books tomorrow. *(destroying)*

c. He will begin on the books tomorrow. *(destroying)*

(5.63) a. My goat went nuts last night.

b. He ate everything in his cage.

c. i. *He began your book at 9pm.

ii. He began on your book at 9pm.

iii. He particularly enjoyed your book.

The discourse in (5.63) exemplifies a contrast between the behaviour of *begin* (or any aspectual verb which may be substituted) and other eventive verbs. The sentence (5.63c(i)) is infelicitous as a continuation of the discourse (5.63a,b), while the sentences (5.63c(ii-iii)) are not. This parallels the distinction in (5.62). Assuming that default interpretations for *begin (on) your book* and *enjoy your book* are predicted from the lexicon, the examples suggest that the default interpretation of *begin+NP* cannot be overridden by contextual cues for its interpretation, while the default interpretations of *begin on+NP* and *enjoy+NP* can be. In formal terms, *begin* can be considered to force the conversion of lexical defaults into indefeasible information at the interface with logical form, which cannot be overridden regardless of strong contextual clues against the default interpretation. In this situation, a conflict in pragmatics between the interpretation proposed on the basis of lexical information and that suggested by context can only serve to judge the sentence as infelicitous, because it won't connect coherently to the discourse content. I will show how the interaction of lexical information with the pragmatic component can be used to account for incoherent discourses such as that in (5.62a,b) below.

An additional description must therefore be added to the four provided at the start of Section 5.5.3:

V. Eventive verbs differ in their interaction with context: some allow the context to override lexically derived interpretations (e.g. *enjoy*), while others do not (e.g. *begin*). The behaviour of particular eventive verbs with respect to contextual influence must form part of a speaker's linguistic knowledge.

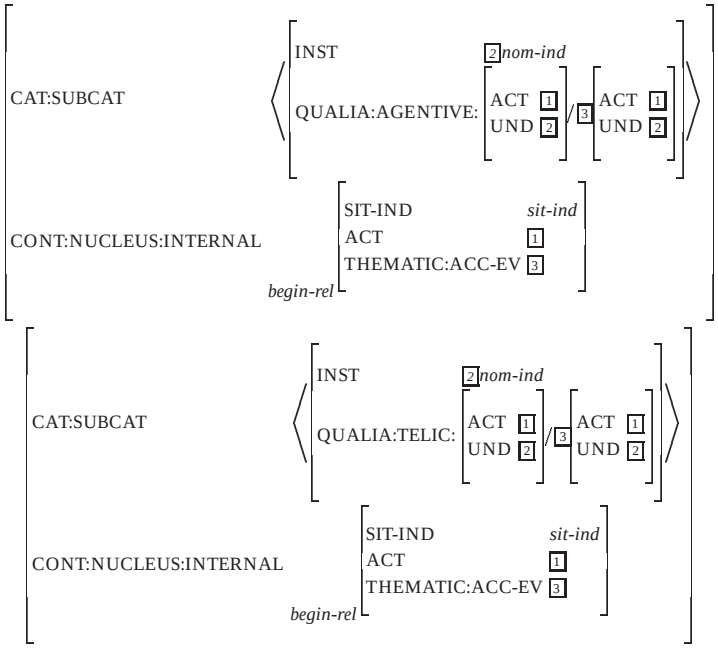

Figure 5.6: Modified coercing forms of *begin*: One entry picks out agentive role, the other telic role

So I must also somehow ensure that defaults specified in the qualia structure of a *begin* complement become indefeasible information in the syntagmatic representation of the sentence. To achieve this, I will modify slightly the lexical entries for the coercing forms of *begin* presented in Figure 5.3, by forcing unification between the event argument position in the logical form of *begin* and the default event represented in the telic or agentive role of the complement noun's qualia structure. The modified lexical entries are shown in Figure 5.6. They differ from the original lexical entries in that only the defeasible portion of the TDFS in the complement's qualia structure is carried into the logical form for the sentence as a whole, which means that the event argument of *begin* will no longer reflect any defaults — it will simply be the non-default feature structure associated with the defeasible portion of the source TDFS. In the original lexical entries, the full TDFS in the complement's qualia structure was carried into the logical form, so both the default interpretation and its defeasible nature was explicitly represented. For the revised lexical entries the default interpretation is represented, but it is no longer treated as defeasible and so cannot be overridden through pragmatic reasoning. The pragmatic reasoning process from which the infelicity of sentence (5.62b) in context is concluded, on the basis of the lexical entries proposed here, is outlined below.

The Analysis

DICE is a theory which allows us to compute rhetorical links between segments of discourse on the basis of the speaker's background semantic and pragmatic knowledge. Discourse representations produced by DICE are in the form of segmented DRSs (SDRSs) (Asher 1993b), in

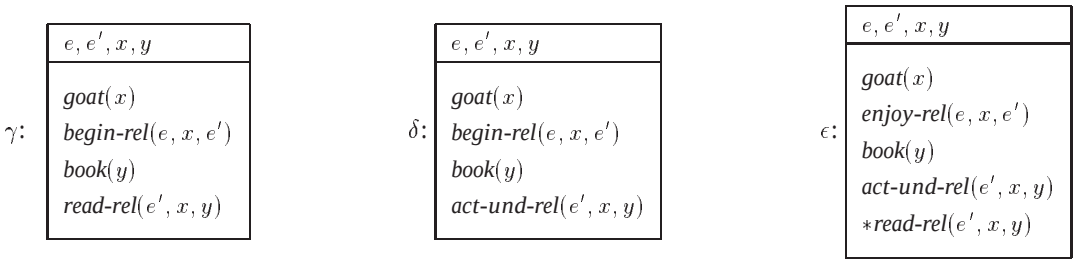

Figure 5.7: Logical forms associated with the three discourse continuations in (5.63c).

which discourses are represented as DRSs plus discourse relations. These discourse relations act as constraints on discourse coherence by constraining the semantic content of the DRSs they connect. For example, if the discourse relation *Narration*, which conveys that one constituent of the discourse is a consequent of a previous constituent, is to hold, it must be possible to compute a common topic between the two constituents. This explains the incoherence of the discourse in (5.64) — there is no topic shared between the two sentences in the discourse.

(5.64) *Max came in. Mary's hair is black.

In this section, I will not give a formal account of the analysis of the discourses in (5.63), as the precise details of the formalism are less important than the describiption of how lexical information might interact with discourse-level reasoning. Thus, I will instead give an informal outline of the analysis. The reader is referred to Lascarides and Asher (1991,1993), Lascarides and Copestake (1995), and Lascarides, Copestake and Briscoe (1996) for details of the formal application of DICE.

I label the DRSs representing (5.63a,b) as α, β and the DRSs representing (5.63c(i-iii)) as γ, δ, ϵ respectively in Figure 5.7. The DRS γ will represent the indefeasible interpretation *He began reading your book at 9pm* for (5.63c(i)), because the sentence has the form *begin+NP* and so the event *read* is picked out from the default portion of the telic role in qualia structure through the lexical entry of coercing *begin* and represented in the logical form as indefeasible. In contrast, δ will represent the interpretation *He began doing something with your book at 9pm* (there is no default interpretation for what the event done with the book is) and ϵ represents the defeasible interpretation *He particularly enjoyed reading your book.* The * marks the fact that the telic event is a default specification. It has remained default after lexicosyntactic processing due to the nature of *enjoy+NP*. The process by which these interpretations are determined will be described in detail below.

Let us consider how the rules apply to the discourse (5.63a,b,c(i)). β must be attached to α. Assuming that DICE calculates that *eating everything in his cage* is a subtype of *going nuts*, we can attach α to β with the discourse relation *Elaboration*, which conveys, as intuitions would dictate, that the event in β is a part of the event in α. Now we must incorporate γ into the SDRS *Elaboration*(α, β). We can attempt to attach γ to either α or β with a discourse relation.

Attaching γ directly to α does not seem to be possible. Intuitively, there is no generalisation between α:*going nuts* and γ:*beginning to read your book*, and so there does not seem to be a

topic common to the two structures which is consistent with world knowledge about *going nuts*. Attaching γ to α with *Narration* is therefore not possible. Furthermore, the latter event is not a subtype of the former, and so *Elaboration* is not possible either. Indeed, the rules in DICE will compute that no discourse relation can adequately relate them on the basis of semantic and pragmatic knowledge. Attaching γ to β also fails in a similar manner. Thus, there is no way to make sense of (5.63c(i)) given the preceding context and the indefeasible interpretation resulting from lexicosyntactic processing, as (5.63a,b,c(i)) is predicted in DICE to be an incoherent discourse.

We can contrast this with the discourses (5.63a,b,c(ii)) and (5.63a,b,c(iii)). In both discourses, we have the same first step as above: β is attached to α with the relation *Elaboration*. Subsequently, the DRSs for the continuation of the discourses, δ and ϵ respectively, must be attached. This is done in both cases according to the analysis in L&C (1995) (for details see that paper). Since δ provides only an underspecified interpretation, the context serves to specify *act-on-pred*(e', x, y) to *eat*(e', x, y). This occurs as a result of the constraints imposed by the rhetorical relations; the strongest coherence for the discourse results when δ is in a clear rhetorical relation to the previous discourse. Here, *doing something with your book* has no clear relation to *eating everything in his cage* whereas *eating your book* is a subtype of that event and so the former event is specified to the latter. Similarly, for ϵ the context overrides the default interpretation *read*(e', x, y) with *eat*(e', x, y), via the axiom `Discourse Wins` combined with the preference for a strongly coherent discourse. In this case, *eating your book* is in a much stronger rhetorical relation to *eating everything in his cage* than *reading your book* is and so the former is preferred. For each of these discourses, then, an interpretation of the continuing sentence is established which would allow DICE to attach the sentence to the discourse via a clear rhetorical relation. This explains the felicity of these discourses in contrast to the discourse (5.63a,b,c(i)).

Infeasibility of a purely pragmatic explanation

I have shown that lexical specification of defaults combined with syntactic control over the persistence of these defaults into the pragmatic component facilitates interpretation of logical metonymies in a discourse context. A purely pragmatic explanation (e.g. Hobbs *et al* 1993), i.e. an explanation not relying on lexicosyntactic factors but only on world knowledge and contextual influences, of such data would fail to account for the incoherence of the discourses (5.8a,b) and (5.63a,b,c(i)) in contrast to the coherence of the discourses (5.8a,c) and (5.63a,b,c(ii)).

The semantic representations for *begin the book* and *begin on/with the book* in such an approach would be identical — they would both correspond to the logical form δ in Figure 5.7, as there would be no lexical specification of default interpretations assumed. There would therefore be no basis for distinguishing their behaviour with respect to the influence of discourse, even if we wanted to assume that the ability of discourse information to influence the interpretation of a particular construction were specified pragmatically. We probably do not wish to make such an assumption in any case since the primary determining factor of this behaviour seems to be syntagmatic knowledge of a specific construction.

Furthermore, the specification of default interpretations of logical metonymies in the pragmatic component would result in a great loss of generality, because information such as the relationship between possible default interpretations and the semantics of the noun in the NP complement (i.e. the fact that default interpretations correspond to the telic or agentive roles) could not be captured in any straightforward way, and also because there are several different types of metonymies (not only *verb+NP* but also *adjective+noun*) which display the same interpretation patterns. A pragmatic approach would be forced to specify the default interpretation of each individual logical metonymy in an *ad hoc* manner.

A purely pragmatic analysis would also fail to accommodate lexical idiosyncrasies and the conventionalisation I have argued for throughout this chapter. How would the oddity of the *began reading* interpretation for *John began the flyer* be explained in light of its clear availability for *John began the book*? Surely the reading activity with respect to magazines is no less common than with respect to books (that is, reading must be the most common activity performed on both flyers and books). There is therefore no pragmatic basis for making this distinction. In fact, none of the data in (5.28), repeated here, could be accounted for on a pragmatic account since all of the events specified in the intended interpretations are commonly performed on the objects denoted by the NPs.

(5.28) a. *John began the film. *(watching)*

b. *John began the door. *(opening, walking through)*

c. *John began the nails. *(hammering in)*

Choice of eventive structure

In a given discourse, the speaker's choice of *begin*, *begin on*, or a full {*begin + VP*} when expressing the beginning of an event associated with a non-eventive NP can be argued to interact with pragmatic principles, in particular, the Gricean maxim of quantity (Grice 1975). This maxim says, in essence, that a speaker should not use more words than necessary to convey a particular meaning. Given the analysis above, this maxim would ensure that {*begin + NP*} will be chosen in all cases for which this choice does not result in pragmatic conflict, as the lexical entry for this form is the most specific of any *begin* entry.[12]

Thus {*begin + NP*} will be chosen when the event associated with the NP which is to be conveyed fulfills the agentive or telic role of the NP, regardless of the salience of the event in the discourse context. In contrast, {*begin on + NP*} will be selected in cases for which the event to be conveyed is not the agentive or telic role of the NP, but rather some other event which is clear from context. If the speaker chooses to convey an agentive or telic event with the {*begin on*} construction, the sentence will still be felicitous if the context suggests the appropriate event, but it is a poor choice of construction given that a more specific form exists which will provide the listener with better clues as to the missing event.

[12]The fact that there is also one less word may also cause preference for {*begin + NP*} over {*begin on + NP*}, but this is unlikely to make much difference.

The {*begin* + *VP*} form will only be chosen in cases in which the event to be conveyed is neither in the qualia of the NP nor entirely clear from context. The lexical entry for *begin* in this form provides no clue whatsoever as to the missing event, and if this event is not clear from context it must be explicitly supplied by the speaker.

These suggestions can account for the distribution of these structures in the corpus: the metonymic uses depend on highly conventionalised lexical information and have highly specific interpretations and as such will only be useful in a limited range of contexts, while the {begin on + NP} form requires a very clearly defined context. Both of these situations must occur far less frequently than somewhat vague contexts in which the intended interpretation of a sentence should be made explicit, given the indeterminacies of linguistic discourse (e.g. the reliance by a speaker on a model of the listener's knowledge which might very well be inaccurate). Therefore the relative infrequency of the coercive constructions relative to the more explicit constructions is a result of the varied and underspecified nature of our linguistic environment.

5.6 Conclusions

A lexically-driven approach to logical metonymy allows predictions about the range of interpretations for these constructions and the defeasibility or indefeasibility of those interpretations to be captured in a general way. The definition of a pragmatic component which has access to this lexical information is critical to the modeling of the behaviour of logical metonymies in discourse contexts. I have shown that the infelicity of certain logical metonymy constructions in some discourses depends on the non-default nature of the lexicosyntactically determined interpretation for such constructions. When a non-default interpretation for a sentence cannot be coherently tied in to the discourse in which the sentence appears, discourse information cannot override that interpretation with a more coherent one, and so the sentence is judged infelicitous in that discourse — the discourse as a whole is weak. This work emphasises the complex nature of the interaction between lexicosyntactic and pragmatic processing; discourse-level analysis is often constrained by lexical properties of the constituents of the discourse.

I have proposed a treatment of logical metonymy which depends on lexical specification of the behaviour of nouns and verbs interacting in this construction; aspectual verbs such as *begin* treat default information coming from the qualia of noneventive NP complements as indefeasible, other eventive verbs leave lexical defaults in their NP complements as defaults, and noneventive NPs will only have a telic event specified lexically if this event has been conventionalised. This conventionalisation has been supported by corpus data and shown to hold across different metonymic constructions, suggesting that it is reasonable to assume lexical specifications reflecting it. The basis of this conventionalisation, however, remains unclear and must be investigated in future work. More generally, an investigation of the motivations for qualia structure seems necessary at this juncture, including a theory of how qualia structure is acquired when learning a language and what dictates the inclusion of information in qualia structure in the lexicon. The current situation demands analysing corpora for specific

phenomena to determine for which verbs the telic role should be specified. Clearly if the proposals here are to be generally useful in a computational framework, such that lexical structures can be built in a principled way, a motivated explanation of the conventions is required. These issues will be taken up in the next chapter.

The general importance of the discussion in this chapter is that it has provided additional evidence of the complex interactivity of syntactic, semantic, and pragmatic knowledge. Certain structures have underspecified meanings which must be resolved through a combination of lexically specified facts, compositionality, and general reasoning processes. None of these knowledge sources can be ignored in modeling the interpretation process.

six Computational Issues

6.1 Introduction

Natural language processing (NLP) systems vary in their goals, and as such vary in what they require from the lexicon. The computational lexicon is the fundamental repository of information about the primary component of language, i.e. words, and therefore critical for systems which aim to handle some aspect of natural language. Two key issues for the lexicon in NLP tasks are lexical **representation** and lexical **acquisition**. This chapter will consider the computational issues resulting from these two lexical aspects, primarily focusing on the challenges meaning ambiguity poses. Polysemy can have a significant impact on the performance of various NLP systems, given the potential unboundedness of sense variation. This is, however, task dependent as some NLP tasks will not require detailed sense discrimination and can effectively make use of lexical entries for which the meaning of a word is underspecified (Kilgarriff 1992, 1997a). Underspecification in the lexicon must therefore be balanced with the needs of the task.

For tasks which require rigorous meaning interpretation, the choice of lexical representation must be carefully considered. It is important to decide on a representation which will not only provide the basis for effective sense disambiguation, but which will also be easy to maintain and update. This means that redundancy must be avoided and generalisations must be captured as far as possible. The most straightforward means of achieving these goals is through the use of an inheritance hierarchy which organises lexical items into groups and allows information relevant for a set of items to be stated once, at a node superior to all the items of the set. When adding a new lexical entry given such a structure, one need only state its type and the properties which are peculiar to that particular word (relative to its type). Indeed, this fact has informed my representational choices in this thesis.

The issue of the acquisition of lexical representations is critical for NLP. A common criticism of NLP systems is that they are often "toy" implementations which would not scale up well to the task of general language understanding of unrestricted texts. This is clearly true in systems

which rely on hand-coding of the lexicon due to the arduous nature of increasing lexical coverage. This factor has led to the search for automated techniques for lexicon acquisition. Initial attempts in this direction were made using Machine Readable Dictionaries (MRDs), and more recently corpora are being consulted as a basis for lexicon development. I will give an overview of various research of both of these types, and will argue that neither results in an adequate representation. I will instead argue for an approach to computational lexicon acquisition which combines top-down linguistic design and bottom-up corpus evidence, and will discuss what a corpus would need to look like in order to give us the relevant information. Although such a corpus is not available today, the desiderata I will put forth point to future research in corpus design and establish a framework in which the lexical representations I utilise in this thesis could in future serve as a basis for larger-scale NLP systems.

6.2 What we want from a Computational Lexicon

The lexicon is the starting point of any NLP system. It must contain information about every potential word form which the system might come across, in order to guide processing. There are several levels of sophistication of (potential) NLP systems, and each of these levels requires lexical information of varying detail. I will attempt to identify different NLP tasks, their lexical needs, and the impact of ambiguity on them. I will consider the tasks of shallow and deep parsing, information retrieval, machine translation, natural language understanding, and natural language generation.

- **Shallow Parsing**

  By shallow parsing, I refer to basic part of speech tagging. This kind of parsing involves identifying the syntactic categories of words in a text. The task requires a lexicon which specifies the various morphological forms of every word, and their associated syntactic category. Semantic information is not necessary since the parser need not discriminate between word forms within the same syntactic category which have different meanings. A single morphological form may, however, very likely occur with various syntactic categories and some mechanism is necessary to tag the forms appropriately in the context of a sentence.

  Most techniques for achieving syntactic category disambiguation are statistical. They hinge on the frequency of a word occurring as a particular part of speech (derived from corpus analysis), in combination with the probability of the word being a particular part of speech given the probability of the categories of words around it. Such techniques can achieve remarkable accuracy (in the 96-97% range). Recent experience with the CLAWS4 tagger designed for the British National Corpus[1] (Leech *et al.* 1994, Leech 1997) suggests, however, that the accuracy can only be improved through the addition of non-probabilistic rules which define highly specific contexts in which a certain part of speech is more likely for a particular (syntactically) ambiguous lexical form. These rules capture

[1] Information about the British National Corpus is available on the web at `http://info.ox.ac.uk:80/bnc/`.

lexical regularities, such as idioms, naming expressions, and noun compounds (Leech 1997). Achieving a few extra percentage points of accuracy requires investing a lot of effort into defining specific rules, but clearly the lexicon can play an important role in achieving near-perfect syntactic category tagging.

- **Deep Parsing**

  Deep parsing, on the other hand, requires identification not only of the syntactic categories of words in a text, but also of phrases and the structural relations between them. There currently exist several different approaches to the task of deep parsing, broadly grouped into theoretically-based approaches and statistical approaches.

  Systems in the first group utilise grammar rules, generally derived from a theory of linguistic structure, to guide the processing of an input string. For this approach the lexicon must contain information in addition to morphological form and syntactic category. Subcategorisation information is generally required to help identify groups of words which form a phrase. This information is also useful for circumscribing the complexity of the disambiguation task, in that a word might be associated with different subcategorisation frames on its different meanings, and so the syntactic context can constrain which lexical entry is appropriate.

  The grammar rules will take on different forms depending on the theoretical framework of the specific grammar. In unification-based grammars, for example, grammar rules make extensive use of features (and their values) represented in the lexical entries of words, including such features as case, gender, and tense. The rules themselves will dictate how the feature structures representing different kinds of words can be combined into feature structures representing phrases, and then how phrases can be combined. Parsing strategies with a grammar can proceed either top-down, i.e. by predicting the category of word to look for according to the grammar rules, or bottom-up, i.e. allowing the lexical information associated with the current word to drive the choice of rule to attempt to instantiate. Carroll (1994) points out that bottom-up parsing is more attractive for natural language grammars, given that highly specific syntactic (and semantic) information can be represented in the lexical entries for particular words, which will severely restrict the number of potential rules applicable at any point in the parsing. Some bottom-up parsers are also augmented with top-down prediction to improve performance, resulting in a hybrid strategy. That the richer the lexical information available to the grammar, the fewer the ambiguous parses that it will generate is, however, clear. The utility of increased lexical representation is shown through the success of lexicalist grammar frameworks like HPSG for computation.[2] This reduction in the number of ambiguous parses, however, comes with the cost of an increase in complexity during parsing — if there are more lexical entries for a particular word, more ambiguities will have to be entertained during parsing. The problem of discriminating senses / lexical entries must therefore be balanced with the reduction of structural ambiguity.

[2] A description of various implementation efforts within the HPSG framework can be found at `http://ling.ohio-state.edu/HPSG/Implementation.html`.

Statistical approaches come in a variety of forms, ranging from hybrid models in which grammars (such as those described above) are augmented with rule application probabilities derived from a corpus (e.g. Carroll and Briscoe 1992, Brew 1995, Abney 1996) to parsers which utilise no grammar rules whatsoever, but derive a probabilistic model of syntactic derivation on the basis of a corpus of labelled syntactic tree structures which is analysed for the occurrence frequencies of tree fragments (e.g. the data-oriented processing approach of Bod and Scha 1996). Both of these kinds of approaches are aimed at identifying the *most probable* parse of a sentence, given that a single sentence can often be associated with multiple syntactic structures. The former type require an amount of lexical information which is dictated by the specificity of the grammar rules (as above), while the latter type require only a finite set of terminal nodes for the syntactic trees, corresponding to word forms.

The questions for this latter type, however, are what guides the initial syntactic analysis of the corpus of syntactic trees from which the statistical model is derived, and how are the lexical categories for the word forms in the corpus determined (e.g. what is the bootstrapping mechanism for acquisition of such a model of language structure?). This approach will also have difficulties coping with unknown words in the text to be parsed, and the ambiguities derived from words which appear in multiple syntactic categories and words with senses which appear in distinct syntactic environments may result in the creation of inaccurate models. How these issues are to be addressed still needs to be investigated; it is likely that the addition of semantic annotations to the syntax trees in the corpus will aid the further disambiguation of the parsing process (Bod and Scha 1996).

Increased lexical information and constraints on semantic composition derived from a theoretical framework are likely to improve the accuracy of such statistical models, as they suffer from the problem of unseen (or zero) data. The issue of how to treat an unknown word arises from the fact that statistical training sets can never encompass every possible word. If the system assigns zero probability to unknown words, it will rule out words which might be licensed linguistically. If it instead assigns a small probability to unknown words, it will allow words which are linguistically impossible. The key to addressing this is to combine the statistical models with a linguistic framework which provides some predictive mechanisms for word formation and sense extensions, as well as constraints on these mechanisms (e.g. as proposed by Briscoe and Copestake 1996).

- **Information Retrieval (IR)**

  The main goal of information retrieval systems is to produce a ranked list of documents in response to a query. One approach to this problem is to index documents based on the words they contain, and retrieve them on the basis of the words that occur in the query (Krovetz 1991). Under the most linguistically naïve version of this approach, the IR system does not even require a basic lexicon, needing at most a list of words (e.g. closed class words) which should not be considered in indexing documents during

pre-processing. Processing proceeds through the application of statistical techniques to unstructured texts. A lexicon with information about syntactic category and word synonyms can, however, help to reduce the number of spurious retrievals which occur (Salton 1968).

Krovetz (1991) highlights lexical ambiguity as problematic for IR: many documents will be retrieved in response to a query which are not relevant, because the sense of a word intended in the query differs from the sense of the same word in the document context. He proposes to combat this problem by making much more elaborate use of the lexicon, and indexing documents according to word *senses* rather than words and to group synonym lists (which are necessary to identify relevant documents which may not contain the exact words in the query) according to word senses. He notes, however, that improvements can be gained without identifying the single correct sense of a word, but instead by ruling out as many *incorrect* senses as possible. This means that even coarse sense disambiguation (at the level of homonymy) can be useful for IR; we will see below (Section 6.4) that this can be achieved through statistical measures of semantic relatedness of the context to a word sense.

Unfortunately, subsequent investigation into the performance benefits from adding word sense disambiguation to IR systems reveals varied judgements on its utility. Kilgarriff (1997a) summarises reports of improved performance resulting from word sense disambiguation as ranging from 2% to 14% in IR systems. Kilgarriff quotes Karen Sparck Jones as saying:

> The fact is, that in relation to IR research as a whole, there's rather little work on WSD just precisely because, as you remark, mechanisms that are independently desirable for improving query quality, namely (in general) adding more query terms to increase the number of conjoint matches, also incidentally achieve disambiguation.

In other words, the ambiguity of word meaning is only an issue for very short queries, which do not provide any disambiguating terms, but not for longer queries in which the combination of multiple terms specifies the desired document context more precisely. Given these facts, the lexicon may not need to be complexly structured for the task of IR. Yet it may turn out that more lexical information could ultimately improve performance; after all, the texts which are to be retrieved are written in natural language and a more sophisticated interpretation would very likely lead to more relevant retrieval. The potential impact on the IR performance will need to be balanced with the complexity of increasing linguistic analysis of the texts.

- **Machine Translation (MT)**
  Machine Translation, if not treated via statistics, is a task which requires highly sophisticated linguistic processing, and correspondingly sophisticated lexica. Translating a sentence from one language to another cannot proceed on the basis of word-by-word analysis (which in itself is a process subject to the difficulties created by lexical ambigu-

ities), because there isn't a one-to-one mapping between words in the source and target languages. Instead, it must involve at a minimum accurate deep parsing and strategies for converting a syntactic structure in the source language to a corresponding structure in the target language, subject to constraints imposed by properties of the particular words in the sentence. For example, subcategorisation requirements of verbs may differ between languages and this affects structural correspondences.

Furthermore, identification of the context is critical for accurate translation of words in the source language which correspond to several different possible words in the target language,[3] and of words ambiguous in the source language for which each sense corresponds to a different word in the target language. However, translation equivalents which are associated with the same set of senses in both source and target language, i.e. which are ambiguous in the same ways in both languages, do not need to be disambiguated for the purposes of translation, since the system does not face a correspondence problem in these cases. So the disambiguation problem for Machine Translation applications is not as serious as for general natural language understanding, but word senses must be taken into consideration in the design of the multilingual lexicon.

- **Natural Language Understanding (NLU)**

  Many NLP applications require deeper understanding of natural language texts than can be achieved through keyword matching or syntactic analysis. Text summarisation and information extraction systems, natural language interfaces (e.g. for tutoring systems or to databases), and interactive software must all be capable of identifying the basic meaning of natural language input deeply enough to respond appropriately. "Understanding" or "interpretation" in these systems corresponds to drawing relevant inferences from the user input and acting on those inferences in particular ways (where an appropriate response depends on the nature of the specific task).

  These systems vary in the specificity of the domains to which they are applied, but it seems clear that the more general the domain, the more important is the capacity of the system to construct an interpretation of the text, relevant to user needs and knowledge, from which inferences can be drawn. This follows from the observation that language use in specific domains is more restricted in terms of both vocabulary and syntactic structure than in general language. Applications working within specific domains can therefore employ heuristics geared towards language use specific to the domain. More general understanding, on the other hand, cannot be done without making use of general linguistic principles and accurate identification of intended word meaning. Lexical structure needs in particular to reflect semantic relations such as hyperonymy / hyponymy (i.e. the hierarchical relations between words) and

[3] A word in the source language which corresponds to multiple words in the target language need not be considered ambiguous in the source language, only underspecified with respect to some feature which the target language discriminates. Colour words provide a good example of this phenomenon, since languages differ as to how specific their colour words are. In English we divide dark colours into black, brown, blue, and green, while in other languages all these colours might be lumped together under one word essentially conveying the meaning *dark colour*.

synonymy/antonymy in order to capture generalisations about the ways in which similar words can be used.

Kilgarriff (1997a) claims that word sense ambiguity is not much of a problem for natural language understanding systems, given that such systems are domain specific and generally have domain models. He says, "It is generally necessary to have a detailed knowledge of the word senses that are in the domain, so the knowledge to disambiguate will often be available in the domain model even where it has not explicitly been added for disambiguation purposes." I find his conclusion unfounded, however, since regardless of where the information needed to disambiguate comes from, disambiguation must happen and can prove to involve arbitrarily complex reasoning. The more lexical structure that exists, the less reasoning is necessary in the domain model or the pragmatic component of the system. A balance must be found between the load which disambiguation imposes on pragmatics and the complexity of the structure of the lexicon. Furthermore, in dialogue systems, the system has virtually no control over what the user says and a domain-restricted lexicon will almost certainly not be able to properly interpet the full range of user input.

Those working on NLU applications may find that a hybrid approach combining statistical techniques with linguistic analyses is the most effective and efficient way of addressing the demands of general language understanding. This is in fact what I will argue for in my discussion of lexical acquisition (Section 6.6). In such an approach the lexicon is likely to play a key role in processing, as the repository of word- and phrase-specific morphological, syntactic, and semantic information. Interpretation of a sentence cannot proceed without both an accurate representation of its syntactic relationships, derived through parsing,[4] and mechanisms for combining the meaning contribution of the words and phrases in the sentence. The semantic contribution of words needs to be identified from the lexicon; to a certain degree the combination of word meanings into a coherent structure can be governed by the specifications in the lexicon as well. The problem of word sense disambiguation must be addressed from a task-specific perspective: how precise must the determination of the intended word sense be in order to generate an appropriate response? The solution must also take into consideration the distribution of load on various modules of the NLU system: how much disambiguation should occur at the level of selecting an appropriate lexical entry, and how much should it rely on discourse, world knowledge, or domain-level processing to refine the sense associated with an (underspecified) lexical entry? The answer to this question will stem from the needs of the task and the framework which drives the linguistic processing.

- **Natural Language Generation (NLG)**
  The problem of generating natural language also can demand highly sophisticated linguistic knowledge of any type of NLP system, as a generation system must not only

[4]This should actually read *mainly* derived through parsing, since syntactic relationships can depend on the semantics of the sentence, such as in the case of prepositional phrase attachment ambiguities.

be capable of constructing an interpretation from a text, but it must produce output which is syntactically and semantically well-formed and contextually relevant. NLG work faces issues of discourse coherence much more overtly than most existing work in NLU due to its emphasis on generating texts consisting of multi-sentence paragraphs or multi-paragraph documents (Dale and Milosavljevic 1996).[5] Furthermore, NLG applications may have a goal of generating different texts for different users and on different occasions, depending on the state of the user model the system maintains. Two essential functions must be performed by any NLG application: *text planning*, the generation of a text plan from an underlying representation, and *linguistic realisation*, realising the text plan in sentences and paragraphs.

The solution for linguistic realisation will need to rely heavily on a rich lexicon of syntactic and semantic information, governed by constraints on the combination of words into sentences and discourse coherence factors. Statistical techniques cannot suffice for generating sentences which are not only plausible syntactically, but contextually relevant, because they capture relationships between words and phrases without having a notion of the underlying structure of those relationships, or the motivations for highlighting such a relationship. Only a solid linguistic framework can guide such processing adequately.

The problem of polysemy must also be carefully addressed within NLG systems, since the range of meanings a word can take on in specific syntactic environments is directly relevant to the problem of forming a grammatically correct and easily understandable sentence using that word. If an NLG system attempts to circumscribe the problem of lexical choice by having a sparse, domain-specific lexicon, it may inadvertently produce something which a user might interpret in a different way or as ambiguous. The particular meaning actually conveyed by a word in context must be no more and no less informative than the sense intended by the NLG system in order to avoid discrepancies between what the user understands and what the system "believes" it has told the user. To tackle these issues, the system must incorporate mechanisms for sense disambiguation which model that of the users, and must therefore look to linguistic theory for insights into the most effective lexical structure and lexical reasoning techniques governing word sense disambiguation.

Statistical information may prove useful in NLG to a certain extent, as an addition to linguistically derived mechanisms. One can imagine a lexicon augmented with probabilities indicating the frequency with which a word appears with a particular meaning; extremely low probability word senses should then generally be dispreferred in generation. This information could thus influence the way the lexicon is used in addressing the polysemy problem in NLG tasks.

[5] However, NLU systems face more difficulty than NLG systems when dealing with multi-sentence texts because of the pragmatic reasoning involved with interpreting multiple sentences, particularly for resolving anaphors and establishing the connections between a set of sentences. NLG systems can avoid these issues by having pre-defined text plans and canned text.

The previous discussion has highlighted the differing needs for lexical representation among various NLP systems. This representation can range from a very shallow list of morphological forms to a highly structured and fine-grained lexicon which derives from linguistic theory. I emphasised the task-dependency of the choice of lexical representation and structure — the increased level of understanding which is necessary in an NLP system, the more important issues stemming from polysemy become and the more attention must be paid to the lexicon. I will review the problems underlying the represention of polysemy in Section 6.3, and will consider existing approaches to word sense disambiguation in Section 6.4.

How the lexicon can be built up and how it can be modified and updated when necessary, however, must be considered in the design of any computational lexicon. I will therefore discuss the problem of lexical acquisition in Section 6.5.

6.3 Polysemy

The fundamental problem posed by polysemy for computation is that particular words can take on an almost indefinite number of subtle meaning variations. These variations can stem from regular sense extension processes, be induced by contextual factors, or result from metaphorical or metonymic extensions. The lexicon is not fixed; word use is flexible and generative. This runs counter to the greatest part of NLP systems, which require every entry in the lexicon to be specified in advance and all potentially relevant sense distinctions to be permanently encoded. We saw in the previous section that for many applications a fixed view of the lexicon is adequate, but as more sophisticated processing of natural language becomes necessary the effectiveness of this approach to lexicon design quickly diminishes. It simply does not reflect the creativity which pervades word use. Let us consider some of the reasons why.

6.3.1 Homonymy vs. Polysemy

It could be argued for many NLP applications that a coarse differentiation between the different meanings associated with a single lexical form is adequate to establish a basic interpretation which can guide subsequent processing. This might be the case for some MT systems or NLU systems which only need to establish the general context of discourse. This coarse differentiation might be at the level of homonyms, leaving polysemous words to be associated with an underspecified lexical semantics which is never made fully precise. So what is the homonymy-polysemy distinction and on what grounds is it made? I will show that the distinction is not clear and therefore not a useful basis for deciding what words/senses to include in the lexicon.

A word with (at least) two entirely distinct meanings yet sharing a lexical form is said to be *homonymous* (e.g. *mogul*, an emperor, or *mogul*, a bump on a ski piste), while a word with several related senses is said to be *polysemous* (e.g. *mouth*, an organ of the body, the entrance of a cave, etc.) (Lyons 1977). While these definitions are intuitively clear, it has been pointed out many times in the literature on lexical semantics that a clear operational distinction between

homonymy and polysemy is lacking. I will review some of the criticisms below, but will begin by introducing an example which emphasises the difficulties of establishing criteria for distinguishing between homonymy and polysemy.

One of the most commonly cited examples of a homonymous word is *bank*, which has a *financial institution* sense and an *edge of a river* sense. These senses seem clearly unrelated, and the fact that they are associated with the same word form seems purely accidental. However, historical linguistics research on Italian has revealed that at some point in the development of the Italian language, these two senses of *bank* actually coincided by virtue of the fact that bankers (lenders of money) sat on the riverbanks while doing their business. So going to the financial institution meant going to the edge of the river, hence to the *bank*. Thus a connection between the two modern senses of *bank* can be established. The relationship between these two senses should presumably not be considered strong enough to establish a relation between them, and therefore to consider *bank* polysemous rather than homonymous, but what criteria for polysemy do they violate? On what criteria do we decide that senses are related?

Lyons

Lyons (1977) devotes a section to the discussion of the homonymy-polysemy distinction (p. 550-569). He identifies the following criteria as those traditionally applied in making the distinction:

1. etymological information — homonymous lexemes "should be known to have developed from what were formally distinct lexemes in some earlier stage of the language"

2. relatedness of meaning — homonymous lexemes have unconnected meanings

Lyons correctly criticises the first criterion as being irrelevant to the synchronic analysis of language, since native speakers are largely unaware of the etymology of the words of their language yet they are able to assign meanings to them. Notice that this criterion will not exclude *bank* as discussed above from being polysemous, given that the two senses historically derive from a single lexeme and in fact a single sense. The second he identifies as an important consideration, but points out that relatedness of meaning is a subjective measure for which intuitions may vary among individuals.

Lyons considers two alternatives to circumventing the homonymy-polysemy issue:

1. Maximise homonymy — associate every meaning of a word with a distinct lexeme. Lyons shows that this will lead to considerable redundancy in the lexicon, as much morphological, syntactic, and even semantic information will be repeated in the lexical entries for the distinct lexemes. However, this redundancy can be greatly reduced given current inheritance-based approaches to lexicon construction. More problematic is the observation that this approach depends on the ability to spell out in advance all of the possible senses in which a word will be used. Lyons suggests that sense distinctions can be "multiplied indefinitely" (1977:554) and that therefore this tack is hopeless. It will never be possible to decide in advance the full range of possible senses a particular

word might be associated with, and furthermore it makes the computational task of selecting the appropriate lexeme daunting given the number of lexemes which might be associated with a particular word form (see Section 6.4 below).

2. Maximise polysemy — adopt the notion that no two lexemes can be entirely distinct when they are syntactically equivalent and when the set of word forms they are associated with are identical. On this view, there are only various kinds of partial homonymy (i.e. when there exist syntactic differences among uses of a word). This removes the vague concept of "semantic relatedness" from the lexicon. However, it would result in an extremely underspecified lexicon from which very little information about the meaning of words could be gleaned. It suffers from the problem of an inability to explain the intuitions that underly the notion of homonymy, and, more relevant to computation, from a complete inability to identify the normal context of use of a particular word and no basis for establishing synonym classes or other semantically-based groupings. How any useful interpretation could be accomplished without some sense differentiation is difficult to see.

Cruse

Cruse (1986) distinguishes *lexemes* from *lexical units*. The former are "the items listed in the lexicon, or 'ideal dictionary' of a language." A lexeme corresponds to a particular word or word form, and can be associated with indefinitely many senses. The latter are "form-meaning complexes with (relatively) stable and discrete semantic properties" (p. 49), and the meaning component is called a *sense*, corresponding to the intuitive notion of sense I have been using. So *bank* is a lexeme, while *bank–financial institution* and *bank–edge of a river* are lexical units.

In discussing the semantic contribution of a word to a sentence, Cruse differentiates between *contextual selection* of a sense and *contextual modulation* of a sense. Selection refers to activation of a particular sense of an ambiguous word form due to the context, while modulation refers to variation induced by the context in terms of emphasising or de-emphasising various aspects of the sense. The sentences in (6.1) exemplify modulation in that the two sentences highlight different parts of the car (the engine and the body, respectively) rather than requiring that *car* refer to different entities in each case. The sentences in (6.2), on the other hand, are instances of contextual selection: a different sense of *light* is selected in each case.

(6.1) a. The car needs servicing.

b. The car needs washing.

(6.2) a. The room was painted in *light* colours. *(cf. dark/*heavy)*

b. Arthur has rather a *light* teaching load. *(cf. heavy/*dark)*

Contextual selection corresponds to the kind of word sense disambiguation undertaken in most NLP systems — there is a pre-existing set of senses for a lexeme (word) and the relevant sense must be identified based on the context. Contextual modulation refers to one creative

aspect of language use, and points to the need for complex representation of knowledge about a word. This is because multiple aspects of a word can be activated simultaneously: *The car needs servicing and washing* is completely felicitous despite the fact that different facets of the car are referred to by each of the verbs (cf. *??Arthur has light teaching loads and rooms in his house* which indicates that multiple senses of *light* cannot be active at once). This kind of knowledge is directly relevant for discourse processing (e.g. anaphor resolution, for example in a discourse such as "*The car needs servicing. It also needs washing.*" where the anaphor refers back to the car as a whole, not just the engine) and generation of coherent sentences / discourses.

Cruse (1986:68) introduces the concept of a *gradient of establishment* of senses. By this he means that a lexical form can be associated with some senses which are potential rather than explicitly represented in the lexicon. Context can stimulate rules which generate an appropriate sense. This idea serves as the foundation of recent work on the Generative Lexicon (Pustejovsky 1991, 1995a: see below). Cruse furthermore argues for the existence of *sense-spectra*, in which the senses of a lexical form lie along a continuum, with no clear boundaries between them, and in some cases without an encompassing sense. Copestake and Briscoe (1995) provide an example, shown in (6.3).

(6.3) a. That book is full of metaphorical language.

b. That book is full of long sentences.

c. That book is full of spelling mistakes.

d. That book is full of typographic errors.

e. That book has an unreadable font.

f. That book has lots of smudged type.

g. That book is covered with coffee.

As they point out, co-predication of the first and the last properties seems odd (6.4a) while co-predication of adjacent pairs seems natural, e.g. (6.4b).

(6.4) a. ?That book is full of metaphorical language and is covered with coffee, so it's very hard to read.

b. That book is full of typographical errors and has an unreadable font.

Cruse (1986:73) suggests that the description of sense-spectra is problematic since a full sense-spectrum does not function as a single lexical unit. Yet he proposes to treat them as a lexical unit, with recognition of the senses along the continuum as *local senses*.

A *lexeme*, corresponding to a lexical entry, is proposed by Cruse (1986:76) to be a family of lexical units. This family can either correspond to a sense-spectrum, or to a set of senses which can be related to one another via regular lexical semantic relationships (captured by lexical rules). Thus the structure of the lexicon on Cruse's view essentially reflects only productive relationships and groups of senses capturing different aspects of a single entity which cannot be consistently delineated. In proposing this, Cruse focuses on the lexical unit as the primary

semantic unit and on the distinction between lexical units, de-emphasising the importance of the word. As such, he skirts the homonymy-polysemy distinction issue, which involves the relationship between lexemes and lexical units. Although the difference between contextual selection and contextual modulation points to phenomena affected by the distinction, he does not propose clear criteria for establishing the distinction. So the computational lexicographer is left with yet another reason for making the distinction but still no basis for making it.

Kilgarriff

Adam Kilgarriff (1992) devotes his thesis to a discussion of polysemy. The essential conclusion which he draws, on the basis of considerations of the traditional distinction drawn between homonymy and polysemy and of investigation of lexicographic techniques for delimiting dictionary senses, is that "Polysemy does not form any kind of 'natural kind'" (Kilgarriff 1992:4). Instead, polysemy describes a "crossroads" between homonymy, alternations, collocations and analogy based on general knowledge and reasoning. Polysemous words can be characterised by at least one of these four methodologies, and Kilgarriff argues that all four must be allowed for in order to capture the full variety of polysemy.

Kilgarriff observes, however, that collocations and analogy depend on frequency information and are subject to contextual variation while description of homonymy and alternations relies on rules, and that joining the two approaches involves augmenting formal lexical structure with frequency data. This view is supported by the results of my investigation of logical metonymy (Chapter 5), which suggests that conventionality plays an important role in predicting language use. The computational lexicon must therefore both reflect linguistic generalisations and provide information on conventional language usage.

The main implication of this work is that polysemy is not a term which can be applied to characterise word senses in an entirely precise way. There cannot be clear-cut tests for identifying polysemy due to its multi-faceted nature. Homonymy is not orthogonal to polysemy, but rather an endpoint of one of the dimensions along which polysemy can be described (fully predictable sense variation — unpredictable sense variation). Furthermore, most words display some variation in the meaning they express and the criteria for pinning down senses are often dependent on questions of frequency and predictability rather than on clearly delineated distinctions. For the purposes of designing a lexicon for an NLU/NLG system, that means distinguishing senses of a word when there are syntactic differences in the way that word is used, and when there are variations in meaning which seem to follow from general, productive relationships.

6.3.2 Perspectives on the representation of polysemy

The conclusion for the computational lexicographer which can be drawn from the previous discussion is that the homonymy-polysemy distinction may not in fact be a useful one for the purpose of designing a lexicon to be used in an NLP task. It does not suffice as a basis for decisions about which words require several lexical entries, and which words can simply be represented by a single, underspecified, lexical entry, because the kinds of varia-

tions in meaning associated with a word do not fall neatly into two groups which might be called "variations which can't be explained semantically" (homonymy) and "variations which seem to have some semantic basis" (polysemy). Given that most NLP systems that require some interpretation capability will have need for making some distinctions between different meanings associated with a word, how might we go about representing these distinctions once they have been decided? The grain of distinctions will, however, have to be chosen with regard to the application of the system, as I suggested in Section 6.2.

Multiple Lexical Entries

The traditional view of lexical organisation, evident both in lexicography and linguistics (e.g. Kempson 1977:79-83), is that distinct word senses should correspond to distinct entries in the lexicon. The arguments in favor of this position hinge on the discreteness of the meaning expressed in a usage of a word — where there is ambiguity, only one sense of a word can be active at any one time. This perspective explicitly denies that the representation of the meaning of a word can be underspecified and modulated (to borrow Cruse's terminology) or specified in context.

Under this view, polysemy is indistinct from homonymy, and so it exemplifies Lyons' "Maximised homonymy" (Lyons 1977; see previous section). It is therefore subject to the criticisms Lyons puts forward: it leads to much lexical redundancy, and it demands the impossible task of enumerating in advance all the senses which might be associated with a lexical form.

Given the productivity and dynamism of language and word meaning, this approach is simply inadequate as a lexical semantic theory, and therefore inadequate for any complex computational language interpretation. It treats the lexicon as independent from grammatical or pragmatic processing; the lexicon is simply a static repository of information which is accessed in, but cannot be influenced by, such processing. It restricts the effect of context on the lexicon to sense selection. Consider example (6.1) above, repeated and expanded here as (6.5), in which different parts of a car can be picked out by different verbs. To handle this on the multiple lexical entry approach, there either need be distinct lexical entries for *car* corresponding to its different parts, which then wouldn't account for a single instance of the word which encompasses two of these lexical entries, as in (6.6), or there could be a single general lexical entry for *car* which was neutral to which part of the car was being referred to, but this would mean that a lot of information about the interpretation of these sentences would be lost at the lexical level. The latter solution would mean that the specific interpretations would have to be left to (computationally expensive and) complex pragmatic reasoning.

(6.5) a. The car needs servicing. *(the engine)*

b. The car needs washing. *(the body)*

c. The car needs rust-proofing. *(the under parts)*

(6.6) The car needs washing and servicing.

Furthermore, to account for examples of *metonymy*, or transfers of reference (Nunberg 1978), for which a commonly cited example is shown in (6.7),

(6.7) The ham sandwich is at table 7.

ham sandwich would have to have a lexically specified interpretation as "person who ordered a ham sandwich". Given the fact that reference transfer is an entirely context-dependent phenomenon, it is difficult to see how it could plausibly be accounted for in terms of a fixed lexicon.

From the perspective of computation, the multiple lexical entry model is quite attractive in that it assumes a finite set of choices from which the NLP system can choose the most appropriate. It is as a result the most widely used model for the computational lexicon to date, and word sense disambiguation research has been framed entirely within this model. This approach will likely suffice for highly-specific domains, in particular when the relevant senses are well-defined with respect to one another, and NLP tasks which do not need to approach human competence in sense discrimination. In fact, highly domain-specific tasks might suffer under the alternative approach of a fully generative lexicon due to the loss of a domain-specific information. However, this could be addressed by adding domain-relative word sense frequencies to the general generative lexicon.

Regular Polysemy

A body of work, beginning most prominently with Apresjan (1974), concentrates on identifying regular shifts in meaning which can occur with particular classes of words. Apresjan (1974:16) defines regular polysemy as follows:

> Polysemy of the word A with the meanings a_i and a_j is called regular if, in the given language, there exists at least one other word B with the meanings b_i and b_j, which are semantically distinguished from each other in exactly the same way as a_i and a_j and if a_i and b_i, a_j and b_j are nonsynonomous.

He suggests that many types of regular polysemy are productive — that if a word has a meaning of a particular type (e.g. the type of a_i and b_i above), it can also be used with the meaning of another type (e.g. the type of a_j and b_j). From this observation the idea of capturing this kind of polysemy in terms of statements (rules) expressing the relationship between semantic types developed.

Taking the systematic relationships identified by Apresjan (1974) and others discussed by Miller (1978), Clark & Clark (1979), Aronoff (1980), and Lehrer (1990) as a starting point, Ostler and Atkins (1991) argue that knowledge of such lexical semantic relationships forms part of linguistic knowledge, and cannot simply be a reflection of regularities in the world or "speakers' free play with analogy" (p. 77), because of the interaction of lexical factors with the application of the relationships (e.g. blocking). They argue for explicit representation of the relationships in terms of *Lexical Implication Rules* (LIRs) which generate derived lexical entries from base lexical entries. In this sort of rule, we find the starting point for a dynamic view of the (computational) lexicon — by making generalisations about potential syntagmatic alternations which a class of words can undergo, one aspect of creativity in language use can be accounted for. For example, if a system captures a rule such as "a countable noun which refers

to an animal can be used as an uncountable noun which refers to the meat from that animal", it would have a basis for distinguishing the senses of *dog* in *John walked his dog* and *We ate dog for dinner last night*, and further would allow for novel instances of this animal→meat rule, such as *I tasted aardvark yesterday*. This sense of *aardvark* cannot be considered established and a part of the fixed lexicon, but can still be interpreted given knowledge of the lexical generalisation.

The use of lexical rules depends on a structured lexicon. That is, it depends on a lexicon in which various generalisations are captured about the semantic classes which words in the lexicon belong to (e.g. animal nouns). This stems from the need to constrain the input to the lexical rules, to restrict their application to those words which participate in the regular relationships. It is difficult to imagine a semantically motivated way of defining the appropriate input to a rule in the absence of represented semantic relationships between words. Lexical rules would be extremely difficult to formulate under an unstructured multiple lexical entry view of the lexicon, in which entries are not grouped into classes.

One of the main issues involved with the use of lexical rules such as these in a computational framework is how to avoid spurious ambiguities which might be generated by these rules. I hinted above at the issue of *blocking*, which occurs when the application of a rule is blocked through the prior existence in the lexicon of a word which is already in the place of the potential output. Two kinds of blocking are *semantic pre-emption*, where the sense to be generated is represented by a previously existing word (consider *beef* which blocks the generation of *cow* on a meat sense), and *lexical pre-emption*, where the word form to be generated already exists in the lexicon with a different sense. These phenomena led Ostler and Atkins (1991) to argue that LIRs must have clear directionality and that blocking phenomena can be formulated as constraints on the application of the LIRs.

The view of the directionality of lexical relationships has, however, been challenged by, among others, Nunberg (1978) and Bredenkamp *et al.* (1996), and this challenge is implicit in the bi-directional representation of lexical rules utilised by Pinker (1989). One central question which arises if directionality of lexical rules is assumed is how to determine which of the senses is the base sense and which is the derived sense. As pointed out by Kilgarriff (1992:89-90), there are cases in which one sense of a systematic relationship is most salient for certain instantiations of it, while the other sense is more salient for other instantiations. An example is found in the tree / wood relationship: the name of a tree can be used to refer to its wood. For *oak*, *ash*, and many other trees / woods, the tree sense is most salient. However, for *teak* and *mahogony*, the wood sense seems to be most salient. Similarly, for *turkey*, the meat sense is more salient than the animal sense, while the reverse is true for *dog*. Kilgarriff addresses this problem by expressing the alternation both at the node for TREE (relating it to WOOD) in his ontology, and at the node for WOOD (relating it to TREE). This is surely unnecessary redundancy. So how, then, could an NLP system instead deal with the problem of blocking?

Copestake and Briscoe (1995) and Copestake (1995) propose adding to lexical entries conditional probabilities that reflect how likely a word is to be used in a specific sense.[6]

[6]A similar proposal has been made for 'standard' dictionaries. Kilgarriff (1997b) reports on how the Longman Dictionary of Contemporary English (3rd edition, 1995) and the Collins COBUILD English Dictionary (New edition,

Thus although a lexical rule will generate a "derived" sense from a "base" sense, different frequencies may be associated with the two senses, and for some word forms (e.g. *teak*) the derived sense will have a higher frequency than the base sense. Under such proposals, novel usages of a word form can be derived through productive application of a lexical rule but the NLP system will have a measure reflecting word usage. Established word senses will be associated with high frequencies, while non-established senses will have low frequencies (and should thus be avoided in, for example, an NLG system). Frequency information could be used to guide parsing, through preference of high-frequency senses. Ambiguous word forms would be parsed initially using the high-frequency sense; low-frequency senses should only be chosen in the case of a syntactic conflict or as a result of subsequent pragmatic processing which determines that it is the "correct" sense.

Furthermore, there are many exceptions to the "constraints" imposed on such lexical rules through blocking (cf. *We had mad cow for dinner last night*) and so application of the rules should not be prevented altogether, even when there are lexical items which might be considered to pre-empt the derived form. Rather, the frequency information would give a clue about how common a particular derived form is, and would obviate the need for hard constraints. It would also guide interpretation, in that the use of an extremely low-frequency form very likely indicates that the speaker chose the word carefully, and wants to convey something additional at the pragmatic level (such as distaste for the meat of a cow). This could be used to trigger pragmatic processing.

In sum, the use of lexical rules in NLP systems can increase the flexibility of the lexicon in those systems. Lexical rules define the space of regular, productive sense extensions which can be used to generate new senses from existing ones. The addition of frequency probabilities to the lexicon can aid in preventing non-established senses from being used in generation, and can influence the interpretation process.

Underspecified Representations

The distinction introduced by Cruse (1986) and described above between *contextual selection* of a sense and *contextual modulation* of a sense has been reiterated in more recent work such as Copestake and Briscoe (1995) and Pustejovsky (1995a). Contextual selection must occur in the case of words which have multiple senses represented in the lexicon, such as those which have two senses related by a regularity which can be captured by a lexical rule. On the other hand, contextual modulation occurs in the case of words which have a single sense which can be contextually specified. For example, the *reel* in *film reel* and *fishing reel* can be defined as "a container artifact with the purpose of (un)winding" (Copestake and Briscoe 1995:18), with the modifiers providing details about what material is (un)wound and thereby making the sense of *reel* more specific. As discussed by Copestake and Briscoe, this modulation involves complex reasoning at the pragmatic level, but depends on a (default) sense for the word which

1995) incorporate word frequency information derived from corpora.

is proposed by the lexicon.[7] The modulation then occurs by specifying an underspecified representation of the word's meaning or by overriding default aspects of the proposed sense. For example, *cloud* refers to a mass of substance which by default is water vapour, but it has extended usages in *dust cloud* and *cloud of smoke* which result from overriding the default to specify that the mass consists of *dust* and *smoke*, respectively, via the modifying phrases. Notice that one would not want to postulate separate lexical entries for these three senses of *cloud*, because the non-default senses can only be triggered in contexts which explicitly refer to the constituency of the cloud, and there are unbounded possibilities for this sense broadening.

One could also conceive of a solution to the sense broadening phenomenon for *cloud* in terms of two lexical entries, one for the standard cloud composed of water vapour, and one which has an underspecified constituency but which requires a complement (adjective or prepositional phrase) which specifies what the cloud is composed of. In the absence of a complement, *cloud* would always be taken to refer to water vapour. This solution has the problem that the constituency of a cloud can also be specified contextually, such that the word *cloud* (without a complement) can in context be taken to refer to a cloud made up of something other than water, as in (6.8).

(6.8) a. A swarm of mosquitoes descended upon Peter's head.

b. The thick cloud obscured his vision and prevented him from sleeping.

If *cloud*, on the other hand, is taken to be ambiguous between the water vapour and the underspecified senses even without a complement, then (a) there is no account of the preference for the water vapour sense in isolation of a specific context, and (b) the load of determining the specific sense is shifted to pragmatics. The pragmatic component must then for every use of *cloud* decide which sense is relevant based on the context. This is a much more difficult task for pragmatics to resolve than checking whether a proposed sense (i.e. the default sense) is coherent in the context. Making "guesses" on the basis of lexical information about the intended sense should be more efficient than computing the sense from scratch in each case. Furthermore, *cloud* is often used in contexts in which there is no prior talk of weather or anything which might prime the *water vapour* sense, yet that sense is (by default) the intended one. The default interpretation of *cloud* would have to be represented in pragmatics in order to accurately model this, which forces conventional information about the use of individual words, in specific languages, to be brought into the pragmatic component. This is linguistic information, not world knowledge, and specifically lexical semantic information. This fact about the word *cloud* is unlikely to have exact correspondences in languages other than English, and means that such a pragmatic model would not be reusable cross-linguistically. This is undesirable, for reasons discussed in the introduction of Chapter 2.

[7] Empirical support for this position is provided by Briscoe *et al.* (1990), who show that logical metonymy is utilised only in instances where either the default interpretation is intended, or in contexts which are rich enough to override the default. It is not used when the context is not rich enough to override the default which would arise from the use of logical metonymy, in favor of a construction which more explicitly provides an interpretation. The data I introduce in Chapter 5 also follow this pattern.

Copestake and Briscoe (1995) refer to the relation between different senses of a word which arise due to contextual modulation as one of *constructional polysemy* — their claim is that the polysemy arises from the packaging together of several aspects of an entity into a single representation. They argue that acceptable instances of co-predication are evidence that a single sense or lexical structure is available for a word. So the non-zeugmatic acceptability of examples such as (6.9) are evidence that the aperture and physical object senses of *door* and the institution and physical object senses of *newspaper* are 'bundled together' into a single sense (Pustejovsky 1995b).

(6.9) a. John painted and walked through the door.

b. John used to work for the newspaper that you are reading.

Pustejovsky (1991, 1995a) has addressed directly the issue of what the underspecified representation of a word might look like, which I have discussed in various points in this thesis (e.g. Section 2.6), in particular with respect to nouns. As a reminder, the representation he proposes includes specification of various core aspects of the denotation of a noun (form, content, agentive, telic roles). Each of these aspects can be picked out in context, but none of them forms an independent sense of the word. These representations therefore allow the behaviour associated with constructional polysemy to be accounted for directly.

As with lexical rules, the incorporation of structured yet underspecified lexical representations and the use of defaults in NLP systems would greatly enhance the flexibility of the lexicon and would help to achieve context-specific interpretation of words. Acknowledging the ways in which words can interact and providing mechanisms for modeling the interactions will lead to more precise representations of the words in context, and therefore will allow more accurate inferences and responses.

Consider again the example of *cloud of mosquitoes*. A system which only had the standard water-vapour sense of *cloud* represented would incorrectly infer from *The cloud of mosquitoes descended upon Peter's head* that there is water vapour around Peter's head and that he becomes wet, and would probably suffer irresolvable conflict when attempting to establish the role of *mosquitoes* in the interpretation. Under these circumstances, a text summarisation system might incorrectly inform a user that this text has to do with weather phenomena rather than a mosquito attack (cf. my earlier remarks). This example suggests that NLP systems can benefit from integrating linguistic insights even in cases where limited interpretation is necessary by allowing for interactions between words which shift their meanings.

Furthermore, the structure in lexical representations can help to define the range of shifts which a particular word can undergo and to constrain the ways in which it can interact with other words. Pustejovsky's core aspects of nominals, for example, establishes several dimensions along which nouns can interact with other words (e.g. verbs, adjectives), but excludes others which aren't represented in the lexical entries. Thus, the representations can help not only to increase the generativity of the computational lexicon, but also to constrain the generativity and avoid generation of spurious interpretations.

6.4 Lexical Disambiguation

In the previous section, I reviewed various perspectives on lexical polysemy. It is clear from that discussion that, as Kilgarriff (1992:4) says, polysemy is not a 'natural kind', but rather that there exist various types of polysemy, and, following from that, various ways of identifying and representing lexical senses. Given that polysemy exists, however, the natural language processing system is left with the task of identifying the particular sense of a word in context — i.e. disambiguation of a lexical item. There is a substantial body of literature on this topic; I will not attempt to give a full overview of the field (for such an overview see Kilgarriff 1992, ch. 2). In this section I will simply review a few of the techniques which have been proposed. It will become clear that existing approaches to lexical disambiguation rely heavily on advance identification of the potential senses which a given word can take and do not allow for contextual modulation of senses, only contextual selection.

The point of departure for most work on word sense disambiguation is the multiple lexical entry view of the lexicon as introduced above: a lexical item is associated with discrete senses identified in advance, and the job of the disambiguation module is to select one of these senses as the meaning intended by the use of a particular word in a particular context. This approach is therefore subject to the criticisms of inadequacy put forth above, in that it ignores potential contextual influences on the precise sense a use of a word has — context can only influence the *selection* of a sense, not the *determination* of a sense.

Wilks (1975) relies on selectional restrictions to drive sense disambiguation — different senses of nouns are specified as having different semantic features and different senses of verbs require arguments with specific semantic features. These features and requirements must be encoded in the lexicon. During processing, the selectional restrictions imposed by the verbs in a sentence interact with the semantic features associated with the nouns in the sentence to identify the intended sense of each word. As pointed out by Kilgarriff (1992), this approach is limited because it does not consider syntax in the disambiguation process, nor is semantic context beyond verb-argument relations taken into consideration. Furthermore, selectional restrictions can be overridden in sufficiently rich discourse contexts.

In direct contrast to Wilks' single-pronged approach to disambiguation is the system developed by Hirst (1987), which allows for the interaction of cues from selectional restrictions, syntax, semantic relations, and the linguistic context to select a word sense. Hirst utilises a marker-passing technique. Markers are passed from lexical entries activated by the sentence to related nodes in the knowledge base, in the spirit of spreading activation in a semantic network. Convergences between specific senses of different words in a sentence are taken to be indicative of selection of those senses in that sentence. This approach follows data from the psycholinguistic literature (e.g. semantic priming experiments) which indicate that the occurrence of a word initially results in the activation of all of its senses before the context establishes the intended sense. Again, this approach assumes that the possible senses for a word are established in advance and can be discriminated in context.

McRoy (1992) develops further the idea that word sense disambiguation must involve information from many sources. She depends on a lexicon in which coarse distinctions among

senses are encoded (in addition to syntactic information), and which is linked to a conceptual hierarchy and information about word frequency and collocations derived from a corpus. The effect of semantic context is modeled via sets of senses, or *clusters*, which group together concepts/senses sharing some concept. There are three kinds of clusters: *categorial clusters* which are sets of senses which share a conceptual parent in the hierarchy, *functional clusters* which share a specified functional relationship to some entity (e.g. a part-whole relationship), and *situational clusters* which groups together senses which tend to occur together in a common setting or event. These clusters are defined in advance in the lexicon of the system. They serve essentially the same purpose as Hirst's semantic network; a cluster containing a sense of a word will be active if it "contains any of the senses under consideration for other words in the current paragraph" (McRoy 1992:20). The process of word sense disambiguation corresponds to sense selection based on context. Senses of all the words in a sentence which are preferred after a preliminary processing phase integrating sense indications from morphology, syntax, and clusters are fed into a parser and then a semantic interpreter which establish the precise relations between the words in the sentence (where these relations come from the clusters rather than through modulation) and leads to final sense selection.

Another mechanism which could be integrated with the McRoy-style approach has been proposed by Agirre and Rigau (1995), using the taxonomy provided by WordNet (Miller 1990)[8] as the basis for a measure of *conceptual distance* of words in a window of context around the word to be disambiguated.[9] The sense selected is the sense (predefined) in the subhierarchy of WordNet for which the highest conceptual density is calculated. This is a more formulaic, probabilistic approach to determining contextually-based sense preferences, but it is similar in spirit to the cluster approach in use by McRoy. Other algorithmic approaches to word sense disambiguation using WordNet and similarity measures are reported in Szpakowicz *et al.* (1997) and Resnik (1995).

Each of the above approaches uses explicitly encoded semantic relationships associated with particular pre-established senses in order to disambiguate a word. More recently, methodologies have been explored which allow a system to *learn* how to disambiguate a word (e.g. Cottrell 1989). For example, Véronis and Ide (1995) construct a neural network on the basis of a machine readable dictionary — input nodes correspond to words, output nodes correspond to senses (words tagged with a sense number as found in the dictionary), and the network is trained on the texts in the dictionary definitions. In this way, the network *acquires* a representation of the significant semantic relations among senses of words, in contrast to systems for which these semantic relations must be represented in symbolic terms (usually through a conceptual hierarchy). However, these approaches depend on a very large initial training set to establish reliable relationships between words. Since a given sense of a word

[8]WordNet is a public-domain lexical knowledge base. It consists of a hierarchy of lexicalised concepts, which correspond to sets of synonymys. It is organised around semantic relations including synonymy and antonymy, hypernymy and hyponymy, and meronymy and holonymy. Its development was based on psycho-linguistic considerations of concept relations.

[9]The idea of using a context window dates back to Lesk (1986), who proposed that the correct sense of a word is the sense with the greatest number of overlaps with senses of other words in a 10-word window of context.

may appear in varous kinds of contexts, the training set must incorporate a wide range of contexts in order to develop a satisfactory predictive model. The development of such a training set would be very time-consuming.

Statistical techniques are currently also being exploited to solve the disambiguation problem (e.g. Yarowski 1995, Fujii *et al.* 1996, Pedersen *et al.* 1997). These approaches generally use features of the surrounding words, such as surface form, part of speech, and morphology, as the basis for classification of different senses of an ambiguous word. A set of disambiguated example sentences, usually derived from a large corpus, drives the development of a classification algorithm. In some cases (e.g. Fujii *et al.* 1996) the features used for classification also include a similarity value of the context words to the ambiguous word, estimated on the basis of relationships drawn from a thesaurus. Guthrie *et al.* (1991) use the definitions in the machine readable version of the Longman Dictionary of Contemporary English (LDOCE) to create word sense neighbourhoods, divided according to the subject codes represented in the MRD, which then are used for a probabilistic account of word sense disambiguation. Similarly, Wilks and Stevenson (1997) utilise the senses discriminated in LDOCE to disambiguate texts, combining part of speech information and a measure of the overlap between the dictionary definition for various senses and the textual context to achieve 86% accuracy in assigning words in their small sample text to the correct homograph (and 57% accuracy in the sense assignment). This result shows that an approach to word sense disambiguation which combines information sources and techniques is promising for disambiguation to pre-encoded senses. The approaches outlined here in some cases depend on a large corpus which needs to be disambiguated in advance and therefore suffer from the same problems as the neural net work approaches. Other approaches proceed by analysing the raw (undisambiguated) corpus, allowing the model to create its own sense distinctions, which might result in ad-hoc senses for a word. The approaches which rely on the lexica in machine readable dictionaries or thesauruses allow sense disambiguation to proceed relative to a well-defined set of senses, but these lexica are expensive to create (Wilks and Stevenson 1997). Ultimately a hybrid semi-interactive approach, combining raw corpus analysis with theoretically-motivated "tweaking" (possibly with respect to an MRD) of the resulting model may prove to give the best results.

A review of the various learning techniques as applied to the sense disambiguation problem appears in Mooney (1996), to which I refer the interested reader. I will not go into further detail here; the general discussion here is enough to support the conclusion that extant automatic approaches to word sense disambiguation depend on pre-encoded distinctions between senses and do not consider the influence of context on determining the precise sense of a word. Although these approaches are adequate for certain NLP tasks which can function solely on the basis of coarse meaning discrimination, they do not suffice as useful techniques under a generative, contextually modulated view of lexical semantics and more general NLU / NLG tasks.

6.5 Acquisition of the Lexicon

Any attempt to represent and disambiguate word senses depends heavily on the representation of lexical semantic information. Many of the early NLP systems relied on hand-coding of the lexicon, but this was quickly realised to be problematic for the development of large-scale systems. Research turned to development of automated techniques for encoding of the lexicon. The initial attempts in this area were made by basing lexica on electronic versions of dictionaries, Machine Readable Dictionaries (MRDs). However, the need for frequency and co-occurrence information as argued for by McRoy (1992) and Copestake and Briscoe (1995) points to the need to augment lexica with information which can be derived only through corpus analysis. In this section, I will review several attempts to extract lexica from each of these sources.

6.5.1 Machine Readable Dictionaries

Dictionaries have long been a source of information about word meaning. The goal of lexicographers is to identify and describe discrete senses associated with a particular lexical form, for each lexical form in the language. As such, they seem a natural resource for development of computational lexica. Hand-coding of lexica, even for highly specific domains, is a daunting task given the number of words in the language and the myriad ways in which each word can be used. Lexicographers invest many years in developing a dictionary, using a depth of experience, so it makes sense to avoid reduplication of work through maximum utilisation of the information that is available in the form of MRDs.

Much effort has been invested in attempting to automatically extract information from MRDs which is then converted or incorporated into lexical entries for an NLP system, and there has been some headway in acquiring lexical information from MRDs. Initial attempts in this area focused on identification of morphological and syntactic information (e.g. Byrd 1983, Boguraev and Briscoe 1987), but more recently there has been research into extraction of lexical semantic information as well. This research has largely concentrated on identifying taxonomies (e.g. Chodorow *et al.* 1985, Lesk 1986, Vossen *et al.* 1989, *head finding* as described in Byrd *et al.* 1987, Calzolari 1991), although some work in identifying related sets of words has also been undertaken (e.g. *filtering* as described in Byrd *et al.* 1987). The identification of taxonomies has proceeded largely on the basis of analyzing dictionary entries for the genus term of a word and thus identifying hypernym relations between entries.

However, several problems with using MRDs are repeatedly identified in the literature. They can be summarised as follows:

- Dictionaries are written with a human user in mind and using natural language; the definitions therefore make full use of the subtleties in natural language and contain many ambiguities which are easily resolved by humans but not by systems attempting to mine information from them.

- Dictionaries are written by lexicographers (although sometimes with reference to a corpus). The divisions and definitions therefore reflect the biases of the human creators.

They are a secondary source of information about language use.

- Word meaning is divided into discrete senses which are treated as independent; relationships between words and word senses are only implicitly represented in the text of the dictionaries. Hierarchical relationships between senses aren't systematically reflected (Atkins 1991). Dictionaries ignore evidence for the synchronic organisation of the lexicon (Beckwith *et al.* 1991). Much domain-specific knowledge of word use is missing.

- The senses captured in dictionaries generally correspond only to established meanings. Furthermore, they are by their nature finite. They can therefore not account for the generative nature of word use, or lexical productivity (Pustejovsky and Boguraev 1994). Word senses are identified in advance, independent of specific contexts. There is therefore no distinction possible between contextual selection of word sense and contextual modulation of sense. Non-established but linguistically licensed senses cannot be predicted from the information in the dictionary.

- MRDs often do not use consistent formats in their definitions; there may be typographical errors; definitions may be circular or have internal inconsistencies (Boguraev and Briscoe 1989b).

Much MRD research (see, *inter alia*, the articles in Boguraev and Briscoe 1989a and Zernik 1991) has focused on the machine readable version of the *Longman Dictionary of Contemporary English* (LDOCE) because it seems to overcome some of the problems with MRDs to a certain degree. This is a result of (1) limiting definition vocabulary to a set of core vocabulary which is used consistently, (2) marking of the words in definitions with a sense number so that the exact intended meaning is specified, and (3) providing "subject" codes which indicate the domain in which a word sense is most likely to appear and some selectional restrictions. However, the latter codes are used inconsistently and are incomplete (Boguraev and Briscoe 1989b:17) and so little use has been made of them.

Work with LDOCE has achieved better results than other MRD research because it does not involve completely free text analysis. Despite this, the amount of semantic information useful for NLP which has been automatically extracted from dictionary definitions is severely limited. The construction of semantic taxonomies is clearly very important for NLP systems, but if these taxonomies are derived from a source which makes artificial divisions between word senses in some cases and conflates word senses which might have linguistically significant differences in other cases (consider the causative and unaccusative forms of a verb like *roll* which in some dictionaries are described under a single sense), their utility for precise interpretation seems questionable. Even extraction of subcategorisation information cannot always proceed systematically from MRDs, so their utility for establishing more complex semantic lexicons is certainly in doubt.

Because dictionaries are written with a human user in mind — even if the vocabulary in the definitions is restricted to a subset of natural language — they leave much of a speaker's knowledge of a word or concept unexpressed, relying on world knowledge, understanding of the general context in which a word appears, and cognitive processes such as analogy to fill in

the gaps. They also do not need to explicitly mention the "linguistically-relevant" components of the meaning of a word. These are identified by a language user on the basis of perceived similarity between words, in ways which do not necessarily correspond to general semantic relatedness but to shared lexical entailments of the kind discussed in Chapter 2. For an NLP system to use and interpret words appropriately, however, these lexical relations must be made explicit in the lexicon. For these reasons, MRDs are not adequate on their own as a source of lexical knowledge for computational systems.

6.5.2 Corpus-based Acquisition

In contrast to dictionaries, text corpora are primary sources of information about language use. They can be analyzed using statistical techniques, to derive information about word frequency, co-occurrence, etc. They support detailed studies of how particular words are used by providing extensive examples of natural language sentences in context (Atkins 1991).

However, corpora suffer from the fact that they consist only of surface words — in many corpora not even syntactic annotations are provided, much less clues about the meaning conveyed by the text. Linguistic pre-processing is necessary to convert corpora into a form useful for lexical acquisition, by identifying parts of speech, syntactic structures, morphologically related words, etc.

Some statistical work has treated MRDs themselves as corpora on which analysis is performed. For example, Tony Plate describes a technique in Wilks *et al.* 1989 for acquiring data on the co-occurrence of pairs of words in LDOCE. The frequency of co-occurrence of two words (or senses as identified in the dictionary) is argued to be a measure of the strength of the semantic relationship between them. Such data could be useful in solving various NLP problems, such as identification of the topic of a discourse or lexical disambiguation (e.g. Guthrie *et al.* 1991, Wilks and Stevenson 1997).[10]

Various statistical tools for gathering information from corpora are described in Church *et al.* (1991) and Church *et al.* (1994). These are the *mutual information* test, which measures the similarity of two words, the *t-test*, which measures the difference between two words, and the *substitutability test* which identifies sets of words with similar distributions. The last test might pick out words which are synonyms, antonyms and co-hyponyms. Church *et al.* argue that these tests can provide different insights into relations between words, but that "great care and skill will be needed in interpreting the salient features of the sets that are identified" (Church *et al.* 1994:174). They view the role of these tools as aiding the lexicographer in his work (and thereby hopefully improving the quality and accuracy of the information in dictionary definitions), rather than as a basis for lexicon acquisition. However, they can also be seen as tools to be used in semi-automatic lexical acquisition, establishing relationships which can be evaluated by a human working with the system to develop a lexicon which reflects the

[10]Since this work was done on LDOCE which is constrained to a limited vocabulary, the co-occurrence data is of restricted size and reflects relations between specific senses. How the technique could scale up to analysis of unrestricted texts is unclear, given that these can contain hundreds of thousands of words which are potentially highly ambiguous.

idiosyncrasies of language use.

Other statistical work (e.g. Bruce and Wiebe 1995), however, views the lexicon as a probabilistic model characterising a rich set of relationships between a large number of variables. These models are developed on the basis of training data, including some data tagged for word senses and syntactic structure, and assume a predefined and finite set of sense distinctions for each word. Under the Bruce and Wiebe analysis, the relationships identified on the basis of the training data are further combined with constraints derived from propositional logic expressions of relationships among word senses, such as those which can be derived from WordNet. Thus they use theoretical knowledge to interconnect the statistical knowledge.

Fukumoto and Tsujii (1995), on the other hand, propose to identify semantic classes of verbs entirely on the basis of statistical clustering, following from the premise that semantically similar words appear in similar contexts. They also argue that polysemous verbs can be recognised by splitting a word cluster into two sets and comparing the semantic deviation of the sets: the distributions of each of the two distinct senses of a word will differ. This definition of polysemy, however, is closer to homonymy and does not allow for subtle gradations between the meanings of polysemous verbs, for which the clusters will differ very little.

In a similar vein, Zhai (1997) attempts to identify *lexical atoms*, or two-word idiomatic phrases with a non-compositional meaning, on the basis of several statistical heuristics which measure compositionality. These heuristics compare co-occurrence frequencies, word associations, and context similarity of the two words in a lexical atom independently and together. The idea is that the meaning of a lexical atom [X Y] is radically different from X or Y independently and that therefore the lexical atom will appear in distinct contexts.

In sum, statistical techniques are useful for measuring various relationships between words in a corpus, and for predicting semantic connections on the basis of frequent co-occurrence of certain words in different contexts. The work of Fukumoto and Tsujii (1995) and Zhai (1997) points towards the usefulness of corpus analysis for identifying entirely distinct uses of a particular word, but such analysis could not easily be extended to automatic discovery of closely related uses of a word or productive lexical rules. No statistical techniques can result in the identification of subtle meaning / usage distinctions. None of these acquisition techniques will result in a generative lexicon which captures the regular relationships between *groups* of words and addresses the productivity of the lexicon, so the NLP systems in which they can be of use is limited. In addition, the observations made by Church *et al.* (1994) and the techniques proposed by Bruce and Wiebe (1995) indicate that the information mined from a corpus can most effectively be applied within a theoretical framework which structures and guides the interpretation of the data.

6.5.3 Prospects for automatic lexicon acquisition

In the literature on lexical acquisition and word sense disambiguation, many doubts are expressed about the usefulness of MRDs and corpora for full-blown acquisition of a computational lexicon. I present a selection of them here, focusing on those which have to do with sense discrimination.

Atkins (1991) asks some fundamental questions about the nature of dictionaries, in comparison to the needs of computational lexica.

> [H]ow much semantic information accurate enough to be useful in a computational lexicon is contained in a dictionary definition written for the human user, who often consciously supplements and corrects what is being read? Is it indeed possible to write dictionary definitions that encapsulate the essential facts about the senses of a word? Can the meaning of a word be divided into discrete senses without distorting it beyond reason? (Atkins 1991:168)

The conclusion from her investigations is that objective evidence supporting the accuracy of dictionary entries does not exist; that there is little systematicity in the sense differentiations made for a word across dictionaries. She argues:

> The traditional dictionary entry is trying to do what the language simply will not allow. Word meaning cannot be sliced up into distinct bundles, labelled (however carefully) and packaged into a dictionary entry which will tell the truth, the whole truth and nothing but the truth. (Atkins 1991:180)

Atkins (1991) and Atkins and Levin (1991) suggest that MRDs can only be effectively utilised in conjunction with a comprehensive theory of the lexicon, which defines 'templates' for lexical entries, establishes what information about particular words might be expected to be found in the dictionary entries for the word, given the identification of the word's semantic class. This approach, then, requires an interaction between the information in the MRD and the structure of the lexicon set out by the theory.

The need for a theory of the lexicon is also implicit in the following statement by Ann Copestake:

> The process of constructing lexical entries [from LDOCE] results in [a Lexical Knowledge Base] which is defective in that it retains the LDOCE sense distinctions. [...] There is little alternative to this currently, given the lack of an adequate theory of sense distinction, although it would clearly be desirable to have a more linguisticaly motivated treatment. (Copestake 1992:136)

The survey of MRD research presented above (Section 6.5.1) showed how narrow the range of semantic information successfully extracted from MRDs to date has been. There are many more relationships among words and word senses which are important for NLP applications than the hyperonymy relation. Synonymy/antonymy and meronymy/holonymy, for example, are important structuring relations (Miller 1990), as is knowledge of regular polysemy. The work of Davis (1995), as discussed in Chapter 2, also strongly indicates that semantic relations at a level deeper than semantic class, the level of semantic roles, is critical for capturing generalisations about language structure. Whether such information exists within MRDs, even implicitly, and whether it could be extracted automatically in a consistent way is not immediately obvious.

For corpus-based lexicon acquisition, the picture is even more bleak. There is no structure in corpora whatsoever, and statistical techniques can only provide very coarse distinctions

in word use. The linguistic information necessary for identification of regular polysemies and for accurate classification of words into taxonomies is simply not available from a corpus. Corpora are best used as a data resource for evidence of how particular words are used, and for identification of collocations and non-lexical units (e.g. idioms), in conjunction with analytical methodology for identifying relationships between uses of words.

The fundamental problem for these automatic techniques is that they depend on pre-existing divisions between word senses which, as we saw in Section 6.3, are not easily justifiable and cannot by their nature be fully identified in isolation of particular contexts. These approaches seemingly deny the creative aspect of language use from the outset and will therefore always fall short of the ultimate goal of identifying the underlying principles of generative language use. NLP systems which require sophisticated language processing demand a framework which will accommodate the flexibility of language use and which will result in fine-grained interpretation. This framework can only come from linguistic theory.

6.6 Linguistic analysis plus Conventionality

Throughout the previous sections, I have emphasised the importance of lexical generativity for advanced interpretation in NLP tasks, and for accommodating meaning ambiguities. Although rich lexical structure is not required for all NLP tasks, it will become increasingly important as NLP systems aim to extend their capabilities to handle more general natural language understanding and generation. We have seen that the design of the lexicon must be informed by linguistic analysis, to identify regularities in sense extensions and to capture syntactic and semantic generalisations associated with particular groups of words.

Following Atkins (1991) and Atkins and Levin (1991), I would like to suggest that an adequate computational lexicon can only be established on the basis of top-down design derived from a linguistic theory in combination with bottom-up information derived from corpora about specific usage of language. Such an approach is supported by evidence that the probabilistic model of Bruce and Wiebe (1995) improves in accuracy when augmented with analytical (theoretically derived) knowledge. The information derived form corpora might include, as suggested by *inter alia* Krovetz (1991), sense frequency information, co-occurrence relations and collocations. It should also include idioms and representation of proper nouns, which establish contexts in which a word can take on non-compositional meanings.

The notion of harnessing linguistically-derived insights to aid lexicon design and automatic lexical acquisition has also been convincingly advocated by Light (1996), who shows that surface cues, such as morphological features of a word, can have consistent correspondences to lexical semantic features associated with that word (or its base form). For example, the prefix *un-* applied to a verb (e.g. *unlatch*, *unhinge*) signals that that verb is a member of the *telic* aspectual class. Such correspondences, once identified on the basis of theoretical research, can be utilised to establish lexical semantic structures for words through corpus analysis. Light demonstrates the utility of morphological cues for identifying a range of lexical semantic properties, ranging from aspectual class to general semantic relation (e.g. *change-of-state-rel*) to antonymy. Corpus analysis driven by surface cue–lexical semantic correspondences can

clearly play a useful role in automatic lexicon acquisition, but it relies on linguistic observations of those correspondences.

The linguistic analysis of logical metonymy in Chapter 5 resulted in identification of certain semantic information which would need to be represented in the lexicon in order to accurately model the conventionality of the phenomenon while still capturing a generalisation about how logical metonymy takes place. To automatically acquire the appropriate representation, corpora would need to be analysed for evidence of specific components of qualia structure. This corpus analysis would very clearly have to be guided by the linguistic theory underlying the explanatory model, including assumptions of the generative devices encoded in the lexicon (e.g. Pustejovsky 1991, 1995a), since the results of the acquisition depend on a particular view of the processes involved in logical metonymy and a particular view of the kind of lexical structure associated with nouns.

Let us consider how the automatic acquisition of the knowledge relevant to logical metonymy might proceed, given the theoretical analysis in Chapter 5 which assumes that logical metonymy always occurs with respect to either the agentive or telic roles of a noun, but that these roles are not represented in the lexical entry of every noun. Although a certain amount of the work of acquiring qualia structure can apparently proceed via automatic means, some of it still must be built up by hand due to the interpretation required to establish whether or not the telic role should be represented for a particular noun, as will be pointed out in step (22c) below.

1. The values of potential agentive and telic roles must be identified for every artifact-referring noun. This involves identifying the verbal relations in which the noun most frequently plays a role. Two particular kinds of verbal relations are most of interest:

   (a) The agentive role of a noun is likely to be the most frequent occurrence of a *creation* verb in which the noun is the created entity. For example, *bake* would be assumed to be the value of the agentive role for *cake* if *bake a/the cake* appears more frequently in the corpus than any other creation activity involving *cake*.

   (b) The telic role of a noun is likely to be the most frequent occurrence of any *non-creation* verb in which the noun plays a non-agentive role. For example, *read* would be assumed to be the value of the telic role for *book* if *read a/the book* appears more frequently in the corpus than any other non-creation activity involving *book*.

2. Instances of logical metonymies must be identified and analysed.

   (a) Pick out instances of an aspectual verb or metonymic adjective followed by a noun (phrase) which does not refer to an event. In the case of aspectual verbs, this process must be restricted to instances in which the noun phrase is a complement of the verb and must therefore occur after deep parsing has established the structure of the sentence in which the verb appears.

   (b) For those nouns which don't participate in logical metonymies in the corpus, propose that their telic role is not accessible to the process of logical metonymy, and that therefore their telic role should not be lexically represented.

(c) For those nouns which participate in logical metonymies in the corpus, attempt to identify whether the logical metonymies are agentive role-centred or telic role-centred, i.e. whether the ellided event is a creation or a non-creation event. How this portion of the analysis could proceed automatically is not clear to me, as it involves extensive context-dependent interpretation and therefore would involve the full power of a NLU system. However, the preceding stages will have identified the relevant set of examples in the corpus, which is likely to be limited to a small set of nouns (as evidenced by the small range of possibilities found in my corpus analysis in Chapter 5). As soon as a single non-creation metonymy involving a certain noun is found, it should be assumed that the telic role for that noun is represented and the next noun can be considered.

3. Add the potential agentive role to the lexical entry for each noun; add the potential telic role to the lexical entry only if there was evidence to do so found in the logical metonymy data.[11]

What do the needs of the process described above tell us about the framework which must already be in place before this specific corpus analysis can proceed?

- An ontology must have been established which divides nouns into very general types, such as nouns which refer to artifacts, nouns which refer to natural entities, nouns which refer to events, etc. Furthermore, the ontology must specify general verbal types, such as the hierarchy in 2.13 on page 54, which can be used to identify particular kinds of relations in which a noun occurs most frequently. These verbal types must be associated (via a linking theory or subcategorisation) with particular syntactic frames. During pre-processing of the corpus through deep parsing, then, the relations can be clearly identified.

- These general noun types must be defined for the relevant qualia structure roles. The agentive role, for example, is relevant to the nouns which refer to artifacts, while it is not for nouns which refer to natural entities. Each role for each of the noun types should be specified with a very general value which can be made specific in the case of individual nouns. So the agentive role of artifact-referring nouns should be associated with a very general agentive value, such as something of the type *creation-rel* identified in Chapter 2.

We also would like to use corpora to identify the frequency with which a certain word undergoes a potential alternation, as suggested by Copestake and Briscoe (1995). For sense

[11] Note that the acquisition procedure outlined here suffers from the zero-data problem, in that there may be potential possible logical metonymies which are not represented in the corpus. This is particularly a problem for logical metonymy since it is such an infrequent phenomenon. A corpus unfortunately also cannot provide negative, i.e. ungrammatical, instances of a phenomenon. The strategy outlined makes use of the positive instances of logical metonymy but treats the lack of an instance of the phenomenon as evidence for a negative instance, which is likely to be too restrictive. There is no obvious way to get around this since this phenomenon does seem to be governed by conventionality and an NLP system will have no way of learning the conventions without reference to a corpus. The best solution is to base the acquisition on as large a corpus as possible.

extensions which have no syntactic reflexes, this is a virtually impossible task, even in corpora that have been processed for syntactic structure. This is because there will be no basis for distinguishing one sense from another in the corpus. However, many sense extensions do have syntactic effects and therefore a parsed corpus can provide the basis for identifying the frequency of some of the different senses of a word.

The addition of rudimentary semantic tagging to the corpus would also aid in calculating the frequency of various sense extensions, particularly if the lexicon is augmented to include certain selectional restrictions. For example, the verb *eat* would likely specify that its *eaten* complement is *foodstuff* or something similar. In the context of *eat*, then, a noun phrase like *the lamb* (e.g. *John ate the lamb*) would be interpreted under its *meat* sense rather than its *animal* sense. This kind of information could guide the identification of a use of a word with a particular sense.

What does the previous discussion tell us about what the corpus needs to look like in order to support the desired processing? Most corpora in existence have at most part-of-speech tagging (e.g. the BNC) resulting from shallow parsing. They can be useful for identifying collocations and general co-occurrence frequencies. However, in order to identify semantic relationships, the corpus must be given more structure. Specifically, I suggest the following desiderata:

- The corpus must have been parsed (deeply), in order to identify the structural relations between elements in sentences in the corpus. This process also can determine, for example, which alternant of an alternating verb is used in a given sentence, or when a verb is being used in the passive form, and which elements of a sentence correspond to syntactic and semantic arguments of the main verb.

- Verbs should be tagged according to their semantic type (which can depend on the results of the parsing above). This information will be derived in the pre-existing ontology for verbs and the associated syntagmatic information represented there.

- Nouns in the corpus must be tagged according to their general semantic type (artifact-referring, event-referring, etc.). Again, this information will depend on the pre-existing ontology for nouns.

In conclusion, the extraction of information useful to advanced NLP tasks from a corpus demands a certain level of linguistic sophistication both from the corpus and from the framework which drives the corpus analysis. This information will ultimately be necessary in order for computational systems to achieve the capability to handle the problems posed by polysemy and the creativity of language use.

seven Conclusions

7.1 Summary

In this thesis I have investigated how lexical semantic structure can be utilised in a contextually-dependent interpretation framework, to handle linguistic phenomena which involve inference of meaning which does not appear to be explicitly present. I have looked at how the meaning of certain prepositions and verbs can vary in particular contexts, and how information represented in the lexicon can lead to models of this variation. I have argued that such models are important for the development of flexible natural language understanding systems, and I have hinted at how that development can be aided by lexical semantic research. To conclude, I would like to highlight the core themes developed in this thesis and provide an overall picture of the resulting framework.

I have proposed a lexical semantic representation which is typed and utilises multiple inheritance, incorporates limited semantic decomposition derived from linguistic research at the syntax/semantics interface (e.g. Jackendoff 1983, 1990; Davis 1995), makes use of non-monotonicity to a limited degree to allow for generalisations to be stated despite the existence of exceptions, and utilises a consistent and constrained structure defined by the lexical hierarchy in which the features are associated with particular entailments. The use of types and multiple inheritance allows the efficient representation of information in common to related words, by minimising redundancy. I have concentrated on verb semantics, but have also adopted the general framework for nominal representation as found in Pustejovsky (1991, 1995a), to the extent that the information encoded there can help to explain the use of certain nouns in particular phenomena. The representation fits in with the spirit of the constraint-based grammatical framework HPSG (Pollard and Sag 1994), although some of the details of that framework have been altered to accommodate more interaction between syntax and semantics at the sentence level. It could fairly straightforwardly be implemented in a feature-structure formalism such as ALE, the attribute logic engine (Carpenter 1993), and integrated with the parsing mechanisms provided there, although the default feature structures may pose

a challenge to that implementation. The result would be a system which would generate representations reflecting the interpretation of a sentence as results from syntactic and lexical semantic processing. These representations might include some default elements, or some underspecification, which would need to be resolved with respect to the discourse context.

The thesis is essentially an extended argument in favor of the interaction of distinct knowledge sources in language interpretation, and a view of the lexicon as repository of specifically linguistic knowledge about a word which can help to account for any idiosyncrasies in its use, and the way that word is used in context. I have provided many arguments for maintaining a distinction between world knowledge and linguistic knowledge, but have suggested ways in which the result of linguistic processing might interact with pragmatic processing.

The use of a rich lexical semantic structure in the development of accounts of several phenomena has been shown to be critical for explaining the range of the data and why particular combinations of words or particular interpretations of a certain combination are impossible or infelicitous in context. We saw in Chapter 5, for example, that certain linguistic phenomena like logical metonymy are governed by conventionality, even though the process of interpreting logical metonymy boils down to inferring content that's not explicitly stated and which appears to stem from world knowledge. That is to say that the interpretation of certain sentences and the ability of certain words to appear in particular constructions cannot be predicted from world knowledge or even contextual influences, but depends on specifically linguistic knowledge about a word which must be lexically captured.

In Chapter 3, we saw that the type system could be constructed in such a way as to allow potentially ambiguous prepositions to only be represented once in the lexicon. The ambiguity therefore did not derive from lexical ambiguity, but rather from lexical underspecification of the kind of entity the preposition could modify and the fact that the semantic contribution of that preposition could be made in several different ways. Such techniques for lexicon construction capture generalisations, since several prepositions might be ambiguous in precisely the same ways. The generalisations are represented implicitly in the inheritance relations between types, and explicitly in the existence of different rules which govern the interaction of elements of a particular type with other elements in a sentence. The combination of underspecified types with which words are associated and rules which effectively specify (or precisify, in Pinkal's (1995) terminology) the type of those words in a particular linguistic context and control their semantic contribution, allows regularities to be efficiently captured and increases the flexibility of the linguistic processor.

Chapter 4 teased apart two phenomena, the use of manner of motion verbs on a directed motion interpretation and resultatives, that have traditionally been lumped together into a single analysis due to their syntactic similarity. Through consideration of the semantics of sample data for each phenomenon and of their interpretation in context, I showed that they are distinct. Differences between the two phenomena which had posed challenges to the uniform analyses were shown to stem from underlying lexical differences in the verbs participating in the different forms, and from the distinct analyses which the phenomena require. This chapter showed how those analyses could be formalised given the lexical semantic struc-

tures and interactivity mechanisms developed in other parts of the thesis. As in Chapter 3, I showed the use of rules governing the interaction of elements at the sentential level to be critical in capturing regularities and explaining interpretations which do not seem to derive directly from the literal meaning of the words in the sentence. The proposal of these rules is in line with research in the framework of Construction Grammar (Fillmore 1988; Goldberg 1995) which suggests that certain syntactic constructions are paired with specific, non-compositional interpretations, but it goes one step further in claiming that the lexical semantic structure of the words in the constructions also play a role in determining the applicability of the rules.

This thesis has contained much linguistic discussion and has examined a wide range of data. This work differs from much linguistic research in that in addition to considering sentences in isolation, I looked at minimal variants of sentences for comparative purposes and at sentences in context. I believe that the results have shown that a complete account of a linguistic phenomenon must involve contextual analysis, in addition to consideration of subtle semantic differences between similar sentences, in order to identify those aspects of interpretation which stem from conventionalisation and word-specific idiosyncrasies, those aspects which stem from regular, productive processes, and those aspects which stem from an interaction of conventionality and regular processes. The information which is to be encoded in the computational lexicon must reflect the results of such detailed analysis.

7.2 Future Research

As discussed in the introduction to this thesis, the pragmatic reasoning framework and discourse model to which I appeal in this thesis is a highly intuitive one. I have not attempted to propose a concrete model of contextually-based interpretation, but have instead concentrated on what kind of lexical semantic structure and linguistic processing might underpin such interpretation, and what kind of representations might interact with pragmatic reasoning. The development of a full-scale model of language understanding, will however, require a formal model of conceptual structure, world knowledge, and how this knowledge interacts with the result of linguistic processing and therefore lexical knowledge. This is a daunting task, given the huge variations in experience (linguistic and otherwise) among speakers, but there do seem to be certain conceptual relations which are primary and consistent (at the word level, these might include synonymy, antonymy, hyponymy, and semantic similarity; at the discourse level, coherence constraints should be represented; world knowledge might include script-like descriptions of common situations). Interpretation can clearly be influenced by all such knowledge, even at the level of individual sentences, but formal models of precisely how this knowledge is structured are the subject of on-going research (e.g. the WordNet project (Miller 1990) and Cycorp's attempts at developing an ontology of world knowledge[1]). Further research is also necessary into precisely how world knowledge and discourse processing interacts with the results of linguistic processing.

In this work, I have not fully explored the exploitation of the entailments associated with

[1]See `http://www.cyc.com` for information on the CYC project.

the attributes in the lexical semantic structures for complex semantic reasoning, although this is clearly an important function served by a consistent representation. These entailments would form part of the pragmatic reasoning system, and could provide a basis for establishing the inference mechanisms which interact with world knowledge.

I also have not fully explored the role of selectional restrictions. There are clearly certain combinations of words which are semantically anomolous due to conflicts in the expectations of semantic type which a verb places on its arguments, or more generally due to our expectations of how certain entities can behave (e.g. *ideas* are not alive and so we do not expect them to be able to *sleep*: *#The ideas are sleeping*). But many of these apparent restrictions can be violated in certain contexts, like metaphors, poetry or animated films (where, for example, goats might read and the interpretation of *The goat began the book* would be entirely felicitous). Investigating the nature of these restrictions and the constraints under which they can be violated will be important for complete models of language understanding, and will again involve examining the interaction between lexical semantic constraints and reasoning at the pragmatic level.

I have considered a few constructions in this thesis whose interpretation depended on lexical semantic structure and the linguistic and discourse context. There are clearly many more which need to be investigated, and one of the most active areas of lexical semantic research is in the analysis of verbal alternations. Many verbs have alternate syntactic forms (in addition to the dative and causative alternations, there are alternations such as the spray / load alternation: *John loaded books onto the truck* and *John loaded the truck with books*), and corresponding slightly varying interpretations or usage constraints. I have not considered exactly how the lexical semantic representations I have utilised in this thesis might be applied to the problem of efficient representation of these alternations, other than for the dative alternation, or the constraints in place which prevent certain verbs from alternating. Given the frequency of these alternations and the amount of discussion of them existing in the literature, this would be a natural next step to take with this framework in place.

Bibliography

Abney, S. (1996). Stochastic attribute-value grammars. Unpublished Manuscript; Available from the cmp-lg archive as cmp-lg/9610003.

Agirre, E. and G. Rigau (1995). A proposal for word sense disambiguation using conceptual distance. In *Proceedings of First International Conference on recent advances in NLP*, Bulgaria.

Apresjan, J. D. (1974). Regular polysemy. *Linguistics 142*, 5–32.

Aronoff, M. (1980). Contextuals. *Language 56*(4), 744–758.

Asher, N. (1993a). Extensions for commonsense entailment. In C. Boutilier and J. Delgrande (Eds.), *IJCAI Worskhop on Conditionals*. Morgan Kaufmann.

Asher, N. (1993b). *Reference to Abstract Objects in Discourse*. Kluwer Academic Publishers.

Asher, N. and A. Lascarides (1995a). Lexical disambiguation in a discourse context. *Journal of Semantics 12*(1), 69–108.

Asher, N. and A. Lascarides (1995b). Metaphor in discourse. In J. Klavans (Ed.), *Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity*, Menlo Park, California, pp. 3–7. 1995 AAAI Symposium: AAAI Press. Technical Report SS-95-01.

Asher, N. and A. Lascarides (1996). Bridging. In *Proceedings of the International Workshop on Semantic Underspecification*, Berlin, October 1996. Available from the Max Plank Institute.

Asher, N. and M. Morreau (1991). Commonsense entailment – a modal theory of nonmonotonic reasoning. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence* (IJCAI-91), Sydney, Australia, pp. 387–392.

Atkins, B. (1991). Building a lexicon: The contribution of lexicography. *International Journal of Lexicography 4*(3), 167–191.

Atkins, B. T. S. and B. Levin (1991). Admitting impediments. In U. Zernik (Ed.), *Lexical Acquisition: Exploiting On-Line Resources to Build a Lexicon*, Chapter 10, pp. 233–262. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. Publishers.

Barwise, J. and J. Perry (1983). *Situations and Attitudes*. The MIT Press.

Beckwith, R., C. Fellbaum, D. Gross, and G. A. Miller (1991). WordNet: A lexical database organized on psycholinguistic principles. In U. Zernik (Ed.), *Lexical Acquisition: Exploiting On-Line Resources to Build a Lexicon*, Chapter 9, pp. 211–232. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. Publishers.

Bod, R. and R. Scha (1996). Data-oriented language processing. An overview. ILLC Technical Report LP-96-13, University of Amsterdam. Available from the cmp-lg archive as cmp-lg/9611003.

Boguraev, B. and T. Briscoe (1987). Large lexicons for natural language processing: Utilising the grammar coding system of LDOCE. *Computational Linguistics 13*(3-4), 203–218.

Boguraev, B. and T. Briscoe (Eds.) (1989a). *Computational Lexicography for NLP*. Longman.

Boguraev, B. and T. Briscoe (1989b). Introduction. In B. Boguraev and T. Briscoe (Eds.), *Computational Lexicography for NLP*, Chapter 1, pp. 1–40. London: Longman.

Bredenkamp, A., S. Markantonatou, and L. Sadler (1996). Lexical rules: What are they? In *Proceedings of COLING 96*, Copenhagen, pp. 163–168.

Brew, C. (1995). Stochastic HPSG. In *Proceedings of 7th Conference of the European Chapter of the Association for Computational Linguistics (EACL 95)*, pp. 83–89. EACL. Available from the cmp-lg archive as cmp-lg/9502022.

Briscoe, E., A. Copestake, and B. Boguraev (1990). Enjoy the paper: lexical semantics via lexicology. In *Proceedings of the 13th International Conference in Computational Linguistics (*COLING-90*)*, Helsinki, pp. 42–47.

Briscoe, T. and A. Copestake (1996). Controlling the application of lexical rules. In *Proceedings of the SIGLEX Workshop on Breadth and Depth of Semantic Lexicons*, Santa Cruz, CA, pp. 7–19. ACL.

Briscoe, T., A. Copestake, and A. Lascarides (1995). Blocking. In P. Saint-Dizier and E. Viegas (Eds.), *Computational Lexical Semantics*, pp. 273–302. Cambridge: Cambridge University Press.

Bruce, R. and J. Wiebe (1995). Towards the acquisition and representation of a broad-coverage lexicon. In J. Klavans (Ed.), *Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity*, Menlo Park, California, pp. 15–20. 1995 AAAI Symposium: AAAI Press. Technical Report SS-95-01.

Brugman, C. (1988). *The Syntax and Semantics of 'have' and its Complements*. Ph. D. thesis, University of California, Berkeley.

Byrd, R. J. (1983). Word formation in natural language processing systems. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, Karlsruhe, Germany, pp. 704–706.

Byrd, R. J., N. Calzolari, M. S. Chodorow, J. L. Klavans, M. S. Neff, and O. A. Rizk (1987). Tools and methods for computational linguistics. *Computational Linguistics 13*(3-4), 219–240.

Calzolari, N. (1991). Acquiring and representing semantic information in a lexical knowledge base. In *Lexical Semantics and Knowledge Representation*, pp. 188–197. Association of Computational Linguistics. SIGLEX workshop.

Carpenter, B. (1993). ALE: The Attribute Logic Engine user's guide. Version Beta.

Carroll, J. (1994). Relating complexity to practical performance in parsing with wide-coverage unification grammars. In *Proceedings of the 32nd Annual Meeting of the Association of Computational Linguistics*, NMSU, Las Cruses, NM, pp. 287–294. ACL.

Carroll, J. and T. Briscoe (1992). Probabilistic normalisation and unpacking of packed parse forests for unification-based grammars. In *Proceedings of the AAAI Fall Symposium on Probabilistic Approaches to Natural Language*, Cambridge, MA, pp. 33–38. AAAI.

Chodorow, M., R. Byrd, and G. Heidorn (1985). Extracting semantic hierarchies from a large on-line dictionary. In *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics*, University of Chicago, pp. 299–304.

Church, K., W. Gale, P. Hanks, and D. Hindle (1991). Using statistics in lexical analysis. In U. Zernik (Ed.), *Lexical Acquisition: Exploiting On-Line Resources to Build a Lexicon*, Chapter 6, pp. 115–164. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. Publishers.

Church, K., W. Gale, P. Hanks, D. Hindle, and R. Moon (1994). Lexical substitutability. In B. Atkins and A. Zampoli (Eds.), *Computational Approaches to the Lexicon*, Chapter 6, pp. 153–177. Oxford University Press.

Clark, E. and H. Clark (1979). When nouns surface as verbs. *Language 55*(4), 767–811.

Collins COBUILD English Dictionary (1995). New edition.

Copestake, A. (1992). *The representation of lexical semantic information*. Ph. D. thesis, Sussex University. Cognitive Science Research Paper CSRP 280.

Copestake, A. (1993). The compleat LKB. ACQUILEX-II Deliverable 3.1. Available from cide@cup.cam.ac.uk.

Copestake, A. (1995). Representing lexical polysemy. In J. Klavans (Ed.), *Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity*, Menlo Park, California, pp. 21–26. 1995 AAAI Symposium: AAAI Press. Technical Report SS-95-01.

Copestake, A. and T. Briscoe (1995). Semi-productive polysemy and sense extension. *Journal of Semantics 12*(1), 15–68.

Copestake, A. and A. Lascarides (1997). Integrating symbolic and statistical representations: The lexicon pragmatics interface. In *Proceedings of the Thirty Fifth Annual Meeting of the Association for Computational Linguistics*, Madrid. ACL.

Corley, S. (1996). CORSET: A corpus search toolkit. User manual. Unpublished ms, Department of Artificial Intelligence, University of Edinburgh.

Cottrell, G. W. (1989). *A Connectionist Approach to Word Sense Disambiguation*. Research Notes in Artificial Intelligence. London: Pitman.

Cruse, D. (1986). *Lexical Semantics*. Cambridge Textbooks in Linguistics. Cambridge University Press.

Dale, R. and M. Milosavljevic (1996). Authoring on demand: Natural language generation in hypertext documents. In *Proceedings of the First Australia Document Computing Symposium*, Melbourne, Australia, pp. 20–27.

Davidson, D. (1967). The logical form of action sentences. In *Essays on Actions and Events*, pp. 105–122. Oxford: Clarendon Press (1980).

Davidson, D. (1980). *Essays on Actions and Events*. Oxford: Clarendon Press.

Davis, A. (1995). *Lexical Semantics and Linking in the Hierarchical Lexicon*. Ph. D. thesis, Stanford University.

Di Tomaso, V. (1996). On the interaction between motion verbs and spatial prepositions. In *Quaderni del Laboratorio di Linguistica*. Scuola Normale Superiore, Pisa.

Dowty, D. R. (1979). *Word Meaning and Montague Grammar*. Dordrecht: Reidel.

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

Everaert, M. (1992). Auxiliary selection in idiomatic constructions. ms, Utrecht University.

Fillmore, C. J. (1968). The case for case. In E. Bach and R. Harms (Eds.), *Universals in Linguistic Theory*, pp. 1–90. New York: Holt, Rinehart and Winston.

Fillmore, C. J. (1988). The mechanisms of "Construction Grammar". In *Proceedings of the 14th annual meeting of the Berkeley Linguistics Society*, pp. 35–55.

Freed, A. F. (1979). *The Semantics of English Aspectual Complementation*. Dordrecht, Holland: D. Reidel Publishing Company.

Fujii, A., K. Inui, T. Tokunaga, and H. Tanaka (1996). Selective sampling of effective example sentence sets for word sense disambiguation. In *Proceedings of the Fourth Workshop on Very Large Corpora*, pp. 56–69. Available from the cmp-lg archive as cmp-lg/9702010.

Fukumoto, F. and J. Tsujii (1995). Representation and acquisition of verbal polysemy. In J. Klavans (Ed.), *Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity*, Menlo Park, California, pp. 39–44. 1995 AAAI Symposium: AAAI Press. Technical Report SS-95-01.

Gawron, J. M. (1986). Situations and prepositions. *Linguistics and Philosophy 9*(3), 327–382.

Gazdar, G. (1987). Linguistic applications of default inheritance mechanisms. In P. Whitelock, H. Somers, P. Bennett, R. Johnson, and M. McGee Wood (Eds.), *Linguistic Theory and Computer Applications*, pp. 37–68. London: Academic Press.

Godard, D. and J. Jayez (1993). Towards a proper treatment of coercion phenomena. In *Proceedings of the Sixth Conference of the European Chapter of the ACL*, Utrecht, pp. 168–177.

Goldberg, A. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*. The University of Chicago Press.

Green, G. (1974). *Semantics and Syntactic Regularity*. Bloomington: Indiana University Press.

Grice, H. (1975). Logic and conversation. In P. Cole and J. Morgan (Eds.), *Syntax and Semantics*, Volume 3: Speech Acts. New York: Academic Press.

Grimshaw, J. (1990). *Argument Structure*. Cambridge, MA: MIT Press.

Gruber, J. (1965). *Studies in Lexical Relations*. Ph. D. thesis, MIT, Cambridge, MA.

Guthrie, J. A., L. Guthrie, Y. Wilks, and H. Aidinejad (1991). Subject-dependent co-occurrence and word sense disambiguation. In *Proceedings of the 29th Meeting of the Association of Computational Linguistics*, Berkeley, pp. 146–152.

Haegeman, L. (1991). *Introduction to Government and Binding Theory*. Blackwell Publishers.

Hankamer, J. and I. Sag (1976). Deep and surface anaphora. *Linguistic Inquiry 7*(3), 391–428.

Hinrichs, E. (1986). Temporal anaphora in discourse of English. *Linguistics and Philosophy 9*, 63–82.

Hirst, G. (1987). *Semantic Interpretation and the Resolution of Ambiguity*. Studies in Natural Language Processing. Cambridge University Press.

Hobbs, J. (1985). On the coherence and structure of discourse. Technical Report CSLI-85-37, Centre for the Study of Language and Information.

Hobbs, J. R., M. E. Stickel, D. E. Appelt, and P. Martin (1993). Interpretation as abduction. *Artificial Intelligence 63*, 69–142.

Jackendoff, R. (1972). *Semantic Interpretation in Generative Grammar*. Cambridge, MA: The MIT Press.

Jackendoff, R. (1983). *Semantics and Cognition*. Cambridge, Massachusetts: The MIT Press.

Jackendoff, R. (1990). *Semantic Structures*. Current Studies in Linguistics series. Cambridge, Massachusetts: The MIT Press.

Jackendoff, R. (1996). Conceptual semantics and cognitive linguistics. *Cognitive Linguistics 7*(1), 93–129.

Jackendoff, R. (to appear). Twistin' the night away. *Language*.

Jacobson, P. (1990). Raising as function composition. *Linguistics and Philosophy 13*, 423–76.

Johnston, M. (1994). *The Syntax and Semantics of Adverbial Adjuncts*. Ph. D. thesis, University of California at Santa Cruz.

Johnston, M. (1995). Semantic underspecification in lexical types: Capturing polysemy without lexical rules. In *Proceedings of the ACQUILEX Workshop on Lexical Rules*, Cambridge, UK.

Johnston, M. and F. Busa (1996). Qualia structure and the compositional interpretation of compounds. In E. Viegas (Ed.), *Breadth and Depth of Semantic Lexicons*, Santa Cruz, California, pp. 77–88. Proceedings of a workshop of the Special Interest Group on the Lexicon (SIGLEX), Association for Computational Linguistics.

Kamp, H. and U. Reyle (1993). *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language*, Volume 42 of *Studies in Linguistics and Philosophy*. Dordrecht: Kluwer Academic Publishers.

Kasper, R. (1993). Adjuncts in the mittelfeld. In J. Nerbonne, K. Netter, and C. Pollard (Eds.), *German in Head-Driven Phrase Structure Grammar*, Lecture Note Series, Chapter 4, pp. 39–69. Stanford: CSLI.

Keil, F. (1989). *Concepts, Kinds, and Cognitive Development*. Learning, Development, and Conceptual Change series. Cambridge, MA: The MIT Press.

Kempson, R. M. (1977). *Semantic Theory*. Cambridge Textbooks in Linguistics. Cambridge: Cambridge University Press.

Kilgarriff, A. (1992). *Polysemy*. Ph. D. thesis, University of Sussex.

Kilgarriff, A. (forthcoming 1997b). Putting frequencies in the dictionary. *International Journal of Lexicography*.

Kilgarriff, A. (submitted 1997a). "I don't believe in word senses". *Computational Linguistics*. Special Issue on Word Sense Disambiguation.

Krifka, M. (1989). Nominal reference, temporal constitution and quantification in event semantics. In R. Bartsch, J. van Bentham, and P. van Emde Boas (Eds.), *Semantics and Contextual Expression*, Volume 11 of *Groningen-Amsterdam Studies in Semantics*, pp. 75–115. Dordrecht: Foris Publications.

Krovetz, R. (1991). Lexical acquisition and information retrieval. In U. Zernik (Ed.), *Lexical Acquisition: Exploiting On-Line Resources to Build a Lexicon*, Chapter 3, pp. 45–64. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. Publishers.

Langacker, R. W. (1987). *Foundations of Cognitive Grammar*, Volume 1: *Theoretical Prerequisites*. Stanford, CA: Stanford University Press.

Lascarides, A. and N. Asher (1991). Discourse relations and defeasible knowledge. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics (*ACL-91*)*, Berkeley, California, pp. 55–63.

Lascarides, A. and N. Asher (1993). Temporal interpretation, discourse relations and common sense entailment. *Linguistics and Philosophy 16*(5), 437–493.

Lascarides, A., E. Briscoe, N. Asher, and A. Copestake (1996). Persistent order independent typed default unification. *Linguistics and Philosophy 19*(1).

Lascarides, A. and A. Copestake (1995). The pragmatics of word meaning. In *Proceedings of Semantics and Linguistic Theory V*.

Lascarides, A. and A. Copestake (1997). Default representation in constraint-based frameworks. Submitted to *Computational Linguistics*.

Lascarides, A., A. Copestake, and E. Briscoe (1996). Ambiguity and coherence. *Journal of Semantics 13*(1), 41–65.

Leech, G. (1997). Recent work on the British National Corpus: Improved automatic tagging. Seminar given at the Centre for Cognitive Science, University of Edinburgh.

Leech, G., R. Garside, and M. Bryant (1994). CLAWS4: The tagging of the British National Corpus. In *Proceedings of the 15th International Conference on Computational Lingustics (COLING 94)*, Kyoto, Japan, pp. 622–628.

Lehrer, A. (1990). Polysemy, conventionality and the structure of the lexicon. *Cognitive Linguistics 12*, 207–246.

Lesk, M. (1986). Automated sense disambiguation using machine-readable dictionaries: How to tell a pine cone from an ice cream cone. In *Proceedings of the 1986 SIGDOC Conference*.

Levin, B. (1993). *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press.

Levin, B. and M. Rappaport Hovav (1995). *Unaccusativity: At the Syntax-Lexical Semantics Interface*. Number 26 in Linguistic Inquiry Monographs. MIT Press.

Light, M. (1996). Morphological cues for lexical semantics. In *Proceedings of the Thirty Fourth Annual Meeting of the Association for Computational Linguistics*, University of California Santa Cruz, pp. 25–31. ACL: Morgan Kaufmann Publishers.

Longman Dictionary of Contemporary English (1995). 3rd Edition.

Lyons, J. (1977). *Semantics*. Cambridge University Press.

Markantonatou, S. and L. Sadler (1995). Linking indirect arguments and verb alternations in English. In D. Godard and J. Jayet (Eds.), *Proceedings of CSSP '95*, Paris.

Markantonatou, S. and L. Sadler (1996). English resultatives revisited. Unpublished ms, University of Essex.

McRoy, S. W. (1992). Using multiple knowledge sources for word sense discrimination. *Computational Linguistics 18*(1), 1–30.

Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist 44*, 1469–1481.

Miller, G. A. (1978). Semantic relations among words. In M. H. et al. (Ed.), *Linguistic Theory and Psychological Reality*, pp. 60–117. Cambridge, MA: MIT Press.

Miller, G. A. (1990). Five papers on WordNet. *International Journal of Lexicography 3*(4). (Special Issue).

Moens, M. and M. Steedman (1988). Temporal ontology and temporal reference. *Computational Linguistics 14*(2), 15–28.

Mooney, R. J. (1996). Comparative experiments on disambiguating word senses: An illustration of the role of bias in machine learning. In *Proceedings of the 1996 Conference on Empirical Methods in Natural Language Processing*, Philadelphia, PA, pp. 82–91.

Nam, S. (1995). Argument orientations of locative PPs in english. In J. M. Fuller, H. Han, and D. Parkinson (Eds.), *ESCOL'94, Proceedings of the Eleventh Eastern States Conference on Linguistics*, Ithaca, pp. 234–245. Cornell University, Dept. of Modern Languages and Linguistics.

Nunberg, G. (1978). *The pragmatics of reference*. Bloomington, Indiana: Indiana University Linguistics Club.

Ostler, N. and B. Atkins (1991). Predictable meaning shift: Some linguistic properties of lexical implication rules. In *Lexical Semantics and Knowledge Representation*, pp. 76–87. Association of Computational Linguistics. SIGLEX workshop.

Partee, B. H. (1984). Nominal and temporal anaphora. *Linguistics and Philosophy 7*, 243–286.

Pedersen, T., R. Bruce, and J. Wiebe (1997). Sequential model selection for word sense disambiguation. In *Proceedings of the 5th Conference on Applied Natural Language Processing*, Washington, DC.

Perlmutter, D. (1970). On the two verbs *begin*. In R. Jacobs and P. Rosenbaum (Eds.), *Readings in English Transformational Grammar*. Waltham, Massachusetts: Ginn.

Pinkal, M. (1995). *Logic and Lexicon*. London: Oxford University Press.

Pinker, S. (1989). *Learnability and Cognition: The Acquisition of Argument Structure*. Learning, Development, and Conceptual Change series. Cambridge, Massachusetts: The MIT Press.

Polanyi, L. (1985). A theory of discourse structure and discourse coherence. In W. Eilfort, P. Kroeber, and K. Peterson (Eds.), *Papers from the General Session at the Twenty-First Regional Meeting of the Chicago Linguistics Society*.

Pollard, C. and I. Sag (1987). *Information-based syntax and semantics*, Volume 1: Fundamentals. Stanford, CA: Center for the Study of Language and Information.

Pollard, C. and I. Sag (1994). *Head-Driven Phrase Structure Grammar*. University of Chicago Press.

Pulman, S. G. (1983). *Word Meaning and Belief*. London: Croom Helm.

Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics 17*(4), 409–441.

Pustejovsky, J. (1995a). *The Generative Lexicon*. Cambridge, MA: MIT Press.

Pustejovsky, J. (1995b). Linguistic constraints on type coercion. In P. Saint-Dizier and E. Viegas (Eds.), *Computational Lexical Semantics*, pp. 71–97. Cambridge: Cambridge University Press.

Pustejovsky, J. and B. Boguraev (1994). A richer characterization of dictionary entries: The role of knowledge representation. In B. Atkins and A. Zampoli (Eds.), *Computational Approaches to the Lexicon*, pp. 295–312. Oxford University Press.

Pustejovsky, J. and P. Bouillon (1995). Aspectual coercion and logical polysemy. *Journal of Semantics 12*(2), 133–162.

Resnik, P. (1995). Disambiguating noun groupings with respect to WordNet senses. In *Proceedings of the Third Workshop on Very Large Corpora*, Cambridge.

Sablayrolles, P. (1995). *Formal semantics of expressions of motion. From lexical semantics to the calculus of the discourse structure in French.* Ph. D. thesis, University Paul Sabatier.

Sag, I. A. (in press). English relative clause constructions. *Journal of Linguistics*.

Salton, G. (1968). *Automatic Information Organization and Retrieval*. New York: McGraw-Hill.

Sanfilippo, A. (1995). Lexical polymorphism and word disambiguation. In *Representation and Acquisition of Lexical Knowledge: Polymsemy, Ambiguity and Generativity*, Stanford, CA. 1995 AAAI Symposium: AAAI Press. Judith Klavans, Chair. Technical Report SS-95-01.

Schilder, F. (1997). *Temporal Relations in English and German Narrative Discourse*. Ph. D. thesis, University of Edinburgh.

Simpson, J. (1983). Resultatives. In L. Levin, M. Rappaport, and A. Zaenen (Eds.), *Papers in Lexical Functional Grammar*, pp. 143–157. Bloomington, Indiana: Indiana University Linguistics Club.

Szpakowicz, S., S. Matwin, and K. Barker (1997). WordNet-based word sense disambiguation that works for short texts. Unpublished. Available from http://www.csi.uottawa.ca/~szpak/WSD_CL/WSD_CL.html.

Talmy, L. (1983). How language structures space. In H. Pick and L. Acredolo (Eds.), *Spatial Orientation: Theory, Research, and Application*. Plenum.

Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science 12*, 49–100.

van Noord, G. and G. Bouma (1994). The scope of adjuncts and the processing of lexical rules. In *Proceedings of the 15th International Conference on Computational Linguistics (COLING'94)*, Kyoto, Japan.

Verkuyl, H. (1989). Aspectual classes and aspectual composition. *Linguistics and Philosophy 12*, 39–94.

Véronis, J. and N. Ide (1995). Large neural networks for the resolution of lexical ambiguity. In P. Saint-Dizier and E. Viegas (Eds.), *Computational Lexical Semantics*, pp. 251–269. Cambridge: Cambridge University Press.

Verspoor, C. M. (1994). A cognitively relevant lexical semantics. Master's thesis, University of Edinburgh.

Vossen, P., W. Meijs, and M. den Broeder (1989). Meaning and structure in dictionary definitions. In B. Boguraev and T. Briscoe (Eds.), *Computational Lexicography for NLP*, Chapter 8, pp. 171–192. London: Longman.

Wechsler, S. (1991). *Argument Structure and Linking*. Ph. D. thesis, Stanford University.

Wechsler, S. (1995). Preposition selection outside the lexicon. In R. Aranovich, W. Byrne, S. Preuss, and M. Senturia (Eds.), *Proceedings of the Thirteenth West Coast Conference on Formal Linguistics*, pp. 416–431. Stanford Linguistics Association.

Wechsler, S. (1996). Explaining resultatives without unaccusativity. Unpublished manuscript.

Wilks, Y. (1975). A preferential pattern-seeking semantics for natural language inference. *Artificial Intelligence 6*, 53–74.

Wilks, Y., D. Fass, C.-M. Guo, and J. McDonald (1989). A tractable machine dictionary as a resource for computational semantics. In B. Boguraev and T. Briscoe (Eds.), *Computational Lexicography for NLP*, Chapter 9, pp. 193–228. London: Longman.

Wilks, Y. and M. Stevenson (1997). Sense tagging: Semantic tagging with a lexicon. In *Proceedings of the SIGLEX workshop "Tagging Text with Lexical Semantics"*. Association for Computational Linguistics. Also available in the cmp-lg archive as cmp-lg/9705016.

Yarowski, D. (1995). Unsupervised word sense disambiguation rivaling supervised methods. In *Proceedings of the 33rd ACL Meeting*, Cambridge, pp. 189–196. Association of Computational Linguistics.

Zaenen, A. (1988). Unaccusatives in Dutch and the syntax-semantics interface. CSLI report 123, Stanford University.

Zaenen, A. (1993). Unaccusativity in Dutch: Integrating syntax and lexical semantics. In J. Pustejovsky (Ed.), *Semantics and the Lexicon*, Chapter 9, pp. 129–161. Kluwer Academic Publishers.

Zernik, U. (Ed.) (1991). *Lexical Acquisition: Exploiting On-Line Resources to Build a Lexicon*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. Publishers.

Zhai, C. (1997). Exploiting context to identify lexical atoms – A statistical view of linguistic context. In *Proceedings of the International and Interdisciplinary Conference on Modelling and Using Context (CONTEXT-97)*, Rio de Janeiro, pp. 119–129.