VERBNET:
A BROAD-COVERAGE, COMPREHENSIVE VERB LEXICON

Karin Kipper Schuler

A DISSERTATION

in

Computer and Information Science

Presented to the Faculties of the University of Pennsylvania

in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2005

Martha S. Palmer
Supervisor of Dissertation

Benjamin C. Pierce
Graduate Group Chair

Acknowledgments

First and foremost I would like to thank my advisor, Martha Palmer, whose knowledge and friendship were invaluable for the completion of this thesis. I also would like to express my gratitude and admiration to the members of my committee, Aravind Joshi, Christiane Fellbaum, Norm Badler, and John Trueswell, for their guidance and encouragement throughout this process. I am especially grateful to the many people who contributed to this work in a variety of ways: Hoa Trang Dang, Benjamin Snyder, Neville Ryant, Edward Loper, Dan Gildea, Scott Cotton, Paul Kingsbury, Anna Korhonen, and Olga Babko-Malaya. The years at Penn were filled with friendship and I would like to thank William Schuler, Hoa Trang Dang, Jan Allbeck, Susan Converse, Alexandra Kinyon, David Chiang and the people from the lab for that. Finally, I am forever indebted to my family: to my husband William for his understanding, endless patience, and encouragement when it was most required; and to my daughters Claire and Anna for being here at the right time, for their unconditional love, and for allowing to have the time that I would have otherwise spent with them. This thesis is dedicated to my mother who did not see its completion. Her courage and confidence will always inspire me.

ABSTRACT

VERBNET:

A BROAD-COVERAGE, COMPREHENSIVE VERB LEXICON

Karin Kipper Schuler

Supervisor: Martha S. Palmer

Despite the proliferation of approaches to lexicon development, the field of natural language processing has yet to develop a clear consensus on guidelines for computational verb lexicons, which has severely limited their utility in information processing applications. James Pustejovsky's Generative Lexicon has concentrated on nouns rather than verbs. WordNet does not provide a comprehensive account of possible syntactic frames and predicate argument structures associated with individual verb senses and ComLex provides syntactic frames but ignores sense distinctions. Dorr's LCS lexicon attempts to address these limitations, but does not provide broad coverage of syntactic frames or different senses or links to actual instances in corpora.

In order to address this gap, we created VerbNet, a verb lexicon compatible with WordNet but with explicitly stated syntactic and semantic information, using Levin verb classes to systematically construct lexical entries. Classes are hierarchically organized to ensure that all their members have common semantic and syntactic properties. Each class in the hierarchy is characterized extensionally by its set of verbs, and intensionally by syntactic frames and semantic predicates and a list of typical verb arguments.

One of VerbNet's primary applications has been as a basis for Parameterized Action Representations (PARs), which are used to animate the actions of virtual human agents in a simulated 3D environment. In order to support the animation of the actions, PARs have to make explicit many details that are often underspecified in the language. This detailed level of representation also provides a suitable pivot representation for generation in other

natural languages, i.e., a form of interlingua.

To evaluate VerbNet's syntactic coverage it has been mapped to the Proposition Bank. VerbNet syntactic frames account for over 84% exact matches to the frames found in PropBank.

VerbNet provides mappings between its verbs and WordNet senses and between its verbs and FrameNet II frames, and mappings between the syntactic frames and Xtag tree families. All these resources are complementary and can be used as extensions of each other.

The original set of classes described by Levin has been refined and extended in many ways through systematic efforts: the coverage experiment against PropBank corpus instances proposed a large set of new syntactic frames and a better treatment of prepositions; new classes from Korhonen and Briscoe's resource were integrated into the lexicon; and new members from the LCS database were added.

Taking advantage of VerbNet's class-based approach automatic acquisition methods were investigated. Additional verbs derived from Kingsbury's clustering experiments and from Loper's VerbNet-WordNet correlation experiment were integrated into the lexicon. These experiments show that it is possible to semi-automatically supplement and tune VerbNet with novel information from corpus data. These approaches reduce the manual classification and enable easy adaptation of the lexicon to specific tasks and applications.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The difficulty of achieving adequate semantic representations has limited natural language processing to applications that can be contained within well-defined domains. Such applications have limited needs which are not difficult to fulfill; the real problem is to scale up to a broad-coverage resource which can be used automatically in computational applications.

In particular, since verbs often convey the main idea of a sentence, such a resource must represent verb meanings. These require a particularly precise and well defined representation that captures both their predicate-argument structure as well as their semantic content. Resources such as verb lexicons are frequently language and domain specific, not always available to the whole community, and are expensive and time-consuming to build.

Many approaches to verb lexicon development make no attempt to associate the semantics of a verb with its possible syntactic frames. Others list an overwhelming number of sense distinctions. Even WordNet (Miller, 1985; Fellbaum, 1998), one of the most widely used on-line lexical databases in natural language applications does not provide a comprehensive account of possible syntactic frames and predicate argument structures associated with individual verb senses. Other resources such as Comlex (COMLEX, 1994) provide syntactic frames but ignore sense distinctions, and CoreLex (Buitelaar, 1998) based on the Generative Lexicon (Pustejovsky, 1991; Pustejovsky, 1995) has concentrated on nouns rather than verbs. The LCS lexicon (Dorr, 2001) attempts to revise these limitations, but does not provide broad coverage of syntactic frames or different senses or links to actual instances in corpora.

Verb classes such as the ones proposed by the Levin classification (Levin, 1993) are important for their ability to capture generalizations that go beyond a single verb instance, thus reducing not only the effort needed to construct a lexicon, but also the likelihood that errors are introduced when adding a new verb entry. But more importantly, these generalizations are based on similar meaning constituents and similar syntactic behavior of words (Pinker, 1989; Jackendoff, 1990a; Levin, 1993) and are credited with incorporating a wider range of linguistic properties than if they were based on semantics alone. Verb classes can be identified throughout language and are asserted to exist across languages since their basic meaning components can be applied cross-linguistically (Jackendoff, 1990a).

From a practical point of view, these types of classifications can be particularly helpful. Verb classes can be used to reduce the ambiguity in the lexicon and to fill gaps in lexical knowledge. Although the association of syntax and semantics is not always perfect, these classes provide close correlations between syntactic and semantic behavior therefore enabling inferences about them.

A variety of natural language tasks including machine translation, language generation (Dorr, 1997), document classification (Klavans and Kan, 1998), lexicography (Sanfillppo, 1994), semantic role labeling (Gildea and Jurafsky, 2002), word sense disambiguation (Dang, 2004), and subcategorization acquisition (Korhonen, 2002) have benefited from the use of verb classes. Some of this work however has been done in a small scale and/or has dealt with restricted domains. A large-scale lexical resource that exploits the notion of verb classes is needed for real-world tasks.

We address problems of existing resources by creating VerbNet, a hierarchical domain-independent, broad-coverage verb lexicon that is compatible with other existing resources but with explicitly stated syntactic and semantic information, using Levin's verb classes to systematically construct lexical entries. VerbNet is freely available on-line and currently has descriptions for over 5,200 verbs and 237 verb classes.

The lexicon is the place where all the information about the representation of words is stored. In our case, we are describing verbs: a class of words for which syntactic and semantic argument structures are particularly diverse. To limit the problem of this representation to a manageable task, we opted to use verb classes to capture generalizations about verb behavior. Verb classes have been shown to provide enough descriptive information for use in natural language applications.

Naturally, although the lexicon is an important repository of semantic information, it does not contain all the *contextual* information necessary to understand language. Nonetheless, we aim to provide a representation that captures key components of both the syntax and the semantics of these classes, and that makes explicit the links between the two, since we believe that both sources of knowledge are necessary for large scale natural language processing applications such as information extraction and machine translation.

Because we are using a class-based approach to construct the lexicon, we provided key elements of information that reflects the syntax and semantics of all members of a (sub)class and not that of specific verbs. Additionally, due to the very nature of a domain independent resource, we chose to provide concise representations which can be extended by the systems using it. We never lost sight however of the minimum level of detail necessary to instruct a virtual character in a simulated environment, since this was our first application.

Thematic roles are used to describe the verb arguments as opposed to more generic labels such as numbered arguments. These roles describe lexical and semantic patterns in the behavior of verbal predicates and provide part of the semantics for the classes. In addition they also help disambiguate between classes with similar frames.[1] Semantic predicates which denote the relations between participants and events are used to convey

[1] A more detailed argumentation is provided in Chapter 3.

the meaning of each class. As the classes may be distinguished by their temporal (aspectual) characteristics (e.g., *Verbs of Assuming a Position* vs. *Verbs of Spatial Configuration*), it is also necessary to convey information about when each of the predicates apply. In order to capture this information, we included a time function in the semantic predicate adapted from the event decomposition from Moens and Steedman (1988).

VerbNet combines traditional lexical semantic information such as thematic roles and semantic predicates, with syntactic frames and selectional restrictions.

Each verb class is completely described by:

- a set of members;
- thematic roles for the predicate-argument structure of the verbs in the class;
- selectional restrictions on the roles; and
- a set of frames consisting of:
  - a brief description;
  - an example;
  - a syntactic description (syntactic frame);
  - a set of semantic predicates including a temporal function specifying whether the predicate is true in the preparatory *(during(E))*, culmination *(end(E))*, or consequent *(result(E))* stage of an event.

The verb classes are hierarchically organized, with 194 new subclasses added to the original Levin classes. These new subclasses were added during various stages of the lexicon extension. This refinement ensures that all the members of each (sub)class have a common semantics and share a common set of thematic roles and syntactic frames. The information presented in the class is strictly monotonic, a child subclass inherits all the information from its parent class, and adds information to it, which can impose further restrictions on the roles, or add syntactic frames or semantic predicates to the subclass. The monotonicity and hierarchical design of the lexicon greatly facilitates its integration with other lexical resources.

One of VerbNet's primary applications has been as a basis for Parameterized Action Representations (PARs), which are used to animate the actions of virtual human agents in a simulated 3D environment. Currently, the use of verb classes is being attested by the numerous researchers using VerbNet in a variety of applications such as automatic verb acquisition (Swift, 2005), semantic role labeling (Rambow et al., 2003; Hensman and Dunnion, 2003; Swier and Stevenson, 2004), robust semantic parsing (Shi and Mihalcea, 2005), building conceptual graphs (Hensman and Dunnion, 2004), building a corpus of instructions annotated with lexical semantic information (Terenzi and Di Eugenio, 2003), word sense disambiguation (Dang, 2004; Girju et al., 2005), and creating a unified lexical resource for knowledge extraction (Croch and King, 2005).

VerbNet's coverage has been evaluated through a systematic mapping to the Proposition Bank's Framesets and corpus instances. Both VerbNet and PropBank have explicit syntactic frames associated with each verb, which allowed an automatic mapping between the two resources for over 78,000 instances. VerbNet syntactic frames account for over 84% exact matches to the frames found in PropBank.

VerbNet has been mapped to other lexical resources. It currently has mappings between its verbs and WordNet senses and between its verbs and FrameNet II frames, and mappings between the syntactic frames and Xtag tree families. All these resources can be viewed as complementary and the mappings between them allow them to be combined. The association of VerbNet verbs to WordNet synsets provides our lexicon with rich semantic information which can be derived from WordNet's relations. The mappings to FrameNet give us a more fine-grained set of semantic roles for VerbNet's verbs and allows for different perspectives on the events described by the verbs. The correspondence to Xtag indirectly provides a much larger syntactic coverage to our lexicon by incorporating the syntactic transformations of the basic frames into our syntactic frames.

The initial set of classes and members has been extended in many ways in an effort to make VerbNet a more robust lexical resource: we supplied a much more detailed account of subclasses than that originally proposed by the Levin classification; we integrated over 40 new classes derived from Korhonen and Briscoe's resource, most of these new classes deal with verbs that take sentential complements which were still largely excluded originally; we also investigated the addition of new members to existing classes from a resource based on the same verb classification, the LCS database. Based on these efforts, we added over 700 verbs and 41 classes to VerbNet. The careful design of the lexicon, with a strong hierarchical structure is to be credited for such a successful integration.

Taking advantage of VerbNet's class-based approach we have also analyzed some of the automatic methods that were proposed to extend VerbNet's coverage. In particular, we investigated Kingsbury's results suggesting additional members from clustering verbs based on frequencies of subcategorization frames observed in PropBank and Loper's results of correlating VerbNet and WordNet resources. This two experiments added another 300 new members to the lexicon. Unsupervised approaches such as the ones presented still require some refinements and face a number of obstacles such as gaps in syntactic context, lack of sense disambiguation in the corpus data, and lack of semantic class annotations for the verbs' arguments.

These recent experiments however show that it should be possible to supplement and adapt VerbNet semi-automatically with novel information derived from corpora in the future. The use of automatic and semi-automatic methods to increase the lexicon's coverage will enable reducing the expensive overhead of manual classification and will allow tuning the resource for specific tasks and applications. The increased coverage of Verbnet will address many of the issues with respect to coverage limitations.

The remainder of this thesis is organized as follows. Chapter 2 describes background on lexical semantic information such as thematic roles and verb classes, and discusses related work on lexical resources such as WordNet and the LCS lexicon. In Chapter 3 we present VerbNet, our domain-independent verb lexicon with associated syntax and semantics based on Levin verb classes. Chapter 4 shows how VerbNet can be used to derive representations of actions for a virtual agent in a 3D-simulated environment. In Chapter 5 we present VerbNet's evaluation against PropBank. In Chapter 6 we introduce the mappings between our lexicon and other resources, more specifically, we show the mappings between VerbNet verbs and WordNet's synsets; the mappings between VerbNet and FrameNet; and the mappings between our syntactic frames and Xtag tree families. Chapter 7 describes the integration of Korhonen and Briscoe's recently proposed verb classes into VerbNet and discusses the challenges of such task; this chapter also shows the results from integrating LCS verbs into our lexicon. Chapter 8 discusses results of two experiments which attempt to automatically augment VerbNet's coverage.

Chapter 2

Background

In the first part of this chapter we present background on lexical semantic information such as roles and verb classes, and the second part presents related work which will be contrasted with VerbNet in Chapter 3.

2.1 Thematic Roles

Thematic roles refer to the underlying semantic relationship between a predicate and its arguments. They were introduced in the mid-60s (Gruber, 1965),(Fillmore, 1968),(Jackendoff, 1972) in order to create a closed set of participant types for a predicate's arguments. These roles are used to describe lexical and semantic patterns in the behavior of verbs.

Thematic roles have been criticized because there is no consensus about which set of roles is necessary to exhaustively characterize argument types of verbs, there are no clear criteria for determining which role should be applied to a particular argument of a predicate, and because definitions for these roles are often vague. A list of common roles and the properties associated with them is given below:

- **Agent**: an active instigator of an action or event. Some authors propose tests such as the 'voluntarily' test (*Tom voluntarily broke the cup* vs. **Tom voluntarily feels* *sick*) and the 'promise'-construction test (*Tom promised to break the cup* vs. **Tom promised to feel sick*) for agentive roles.

- **Patient**: a participant undergoing a process or affected by an action. The emphasis is on change of state. Patients can occur either in subject (*'the ice' melted*) or object position (*they melted 'the ice'*) and a common test for a patient role is "what happened to X?".

- **Theme**: a participant that is located in a place or that is seen as moving from one place to another. The emphasis is on locational (or possessional) properties. Objects of verbs such as *give*, and subjects of *walk* and *die* are examples of themes.

- **Experiencer**: a participant characterized as aware of or experiencing something. Many psychological and emotion verbs have Experiencer as an argument role either in subject (e.g., *love, admire*) or object (e.g., *amuse, perturb*) position.

- **Stimulus**: the event or object that brings about a psychological response in an Experiencer. Subjects or objects of psych-verbs may have a Stimulus role as in *'The heavy rain' scared the kids.* This role is sometimes characterized as Causer or inanimate Agent.

- **Instrument**: a physical force or object that makes a change in something usually by coming into direct contact with it. Often a tool used by an animate agent to bring about a change. *With*-prepositional phrases such as in *We hit the ball with 'the bat'* are associated with this role.

- **Location**: participant that expresses a location, usually introduced by a prepositional phrase.

- **Source**: participant that is the starting point of the motion expressed by the predicate. Examples include *The Martians left 'home'* and *She came from 'another planet'*.

Source is also used in change of possession predicates.

- **Goal**: participant to which the motion proceeds. For example, subject of *receive*, *buy*, dative objects of *tell*, *give* as in *The Martians reached 'home'* and *She went to 'another planet'.* This role is also frequently used in predicates of change pf possession.

- **Recipient**: participant that is target of the transfer of some concrete or abstract entity.

- **Benefactive**: entity benefiting from some action, introduced by a *for*-prepositional phrase.

These roles can be independent of the syntactic encoding of situations. As an example, in both sentences in (1):

(1) John hit Bill
Bill was hit by John

Bill has the thematic role *patient* of the hitting action and *John* has the thematic role *agent*. Moreover, the use of thematic roles as descriptors may entail restrictions which could be used to further constrain arguments of verbs, such as *animacy* for agentive roles.

2.1.1 Approaches

Fillmore (Fillmore, 1968) in his deep case theory, argues in favor of a set of universal semantic roles (cases), such as agentive, instrumental, dative, factive, locative, and objective, they are used to describe linguistic generalizations in terms of meaning, not syntax. Verbs are classified according to the relationships they have with their arguments, in other words, according to the case-frames they take.

Jackendoff (Jackendoff, 1983) bases his framework of Lexical Conceptual Structure on the previous work of Gruber (Gruber, 1965), which investigated similarities and extensions of thematic relations across domains. For Jackendoff, thematic roles are defined in terms of primitive semantic properties of predicates, and thematic relations are not to be considered part of syntax, but as an integral part of a level of semantic and conceptual structure (Jackendoff, 1987). Jackendoff suggests that thematic relations should be defined in terms of the three conceptual primitives CAUSE, CHANGE, and BE which constitute some of the primitive building blocks of lexical meanings. According to Jackendoff (1972) "*Agent* is the argument of CAUSE, *Theme* is an argument of CHANGE, *Source* and *Goal* are the initial arguments of CHANGE. *Location* is defined by a further semantic function BE that takes an individual (the *Theme*) and a state (the *Location*)." Subsection 2.4.3 contains a more detailed description of Jackendoff's Lexical Conceptual Structures.

Sanfilippo (1990) points out that a rigorous definition of thematic roles in terms of predicate decomposition remains a goal to be achieved, and that assumptions regarding the semantic content of roles and their representation in grammar tend to have a very speculative character. As Dowty (1989) notes, when roles such as Jackendoff's are subject to closer observation they tend to fragment into a number of independent role types (for example, *Theme* can be either dynamic or static) in such a way that the hypothesis that there is a small and discrete set of such roles is undermined.

In contrast with this view, Dowty (Dowty, 1991) proposes a weaker definition of thematic roles, where the use of discrete role types is abandoned and replaced by the relation between role types and sets of possible entailments. He defines thematic roles for a verb as the set of all properties that this verb allows for its argument positions. Here, only two broad classes of thematic role types for verb-arguments are characterized: Proto-Agent and Proto-Patient. The choice of which lexical items are assigned to which roles is made through the Principle of Verbal Argument Selection that states that if a predicate has both a subject and an object, the argument with the greatest number of Proto-Agent properties will be lexicalized as the subject and the one with the greatest number of Proto-Patient properties will be the object. Proto-Agent properties include sentience or perception, volitional involvement in the event or state described, the ability to cause an event or change of state in another participant, or movement relative to the position of another participant. The Proto-Patient role entails that the NP undergoes change of state, is causally affected by another participant, or is stationary relative to the movement of another participant.

2.2 Levin verb classification

Levin classes are a current approach to English verb classification. This classification explicitly states the syntax for each class, but falls short of assigning semantic components. The classes are based on the ability or inability of a verb to occur in pairs of syntactic frames that are in some sense meaning preserving (referred to as diathesis alternations) (Levin, 1993). The sets of syntactic frames associated with a particular Levin class are supposed to reflect underlying semantic components that constrain allowable arguments and adjuncts. For example, as shown in the alternations in (1) and (2), *break* verbs and *cut* verbs are similar in that they can all participate in the transitive and middle constructions. However, as illustrated in alternations (3) and (4), only *break* verbs can also occur in the simple intransitive, and *cut* verbs can occur in the conative construction (attempted action, with no result achieved), where *break* verbs cannot. The explanation given is that *cut* describes a series of actions directed at achieving the goal of separating some object into pieces. It is possible for these actions to be performed without the end result being achieved, but where the *cutting* manner can still be recognized (i.e., "John cut at the loaf"). For *break*, the only thing specified is the resulting change of state where the object becomes separated into pieces. If the result is not achieved, no attempted *breaking* action can be recognized.

(1) Transitive construction

(a) John broke the window.

(b) John cut the bread.

(2) Middle construction

(a) Glass breaks easily.

(b) This loaf cuts easily.

(3) Intransitive construction

(a) The window broke.

(b) *The bread cut.

(4) Conative construction

(a) *John broke at the window.

(b) John valiantly cut/hacked at the frozen loaf, but his knife was too dull to make a dent in it.

Although Levin classes group together verbs with similar argument structures, the meanings of the verbs are not necessarily synonymous. Some classes such as *Break (break, chip, crack, crash, crush, fracture, rip, shatter, smash, snap, splinter, tear)* and *Cut (chip, clip, cut, hack, hew, saw, scrape, scratch, slash, snip)* contain verbs that are quite synonymous, but other classes, such as *Braid* (*bob, braid, brush, clip, coldcream, comb, condition, crimp, crop, curl*, etc.) are clearly not intended to be synonymous.

The fundamental assumption is that the syntactic frames are a direct reflection of the underlying semantics, but the association of sets of syntactic frames with individual verbs in each class is not as straightforward as one might suppose. Many verbs are listed in multiple classes, some of which have conflicting sets of syntactic frames. For instance, *Carry* verbs are described as not taking the conative (*"The mother carried at the baby"), and yet many of the verbs in the *Carry* class *(push, pull, tug, shove, kick)* are also listed in the *Push/Pull* class, which does take the conative. In subsection 2.2.1 we describe an extension of the basic Levin verb classes, intersective Levin classes (Dang et al., 1998), that clarifies some of the issues around multiple listings and competing sets of alternations, as well as more precisely highlighting and isolating the meaning components of a verb class.

2.2.1 Intersective Levin classes

Because some of the original Levin classes contain members that exhibit a diverse range of possible semantic components, and due to fact that many verbs are present in multiple classes, as discussed above in Section 2.2, the original Levin classes were augmented with a set of *intersective* classes, which were created by grouping together subsets of existing classes with overlapping members. All subsets were included which shared a minimum of three members. If only one or two verbs were shared between two classes, we assumed this might be due to homophony, an idiosyncrasy involving individual verbs rather than a systematic relationship involving coherent sets of verbs.

Figure 2.1 shows an example of the intersective class created from classes *Split*, *Carry* and *Push/Pull*, verbs shown in parentheses although not listed by Levin in all classes, participate in all the alternations for these classes.

2.2.2 Cross-linguistic verb classes

As discussed previously, the ability of a verb to participate in an alternation is related to its semantic content. We examined a small set of Portuguese verbs which behaved more similarly to their English counterparts than we expected. Many of the verbs participate in alternations that are direct translations of the English alternations. However, there are some interesting differences as to which sense extensions are allowed.

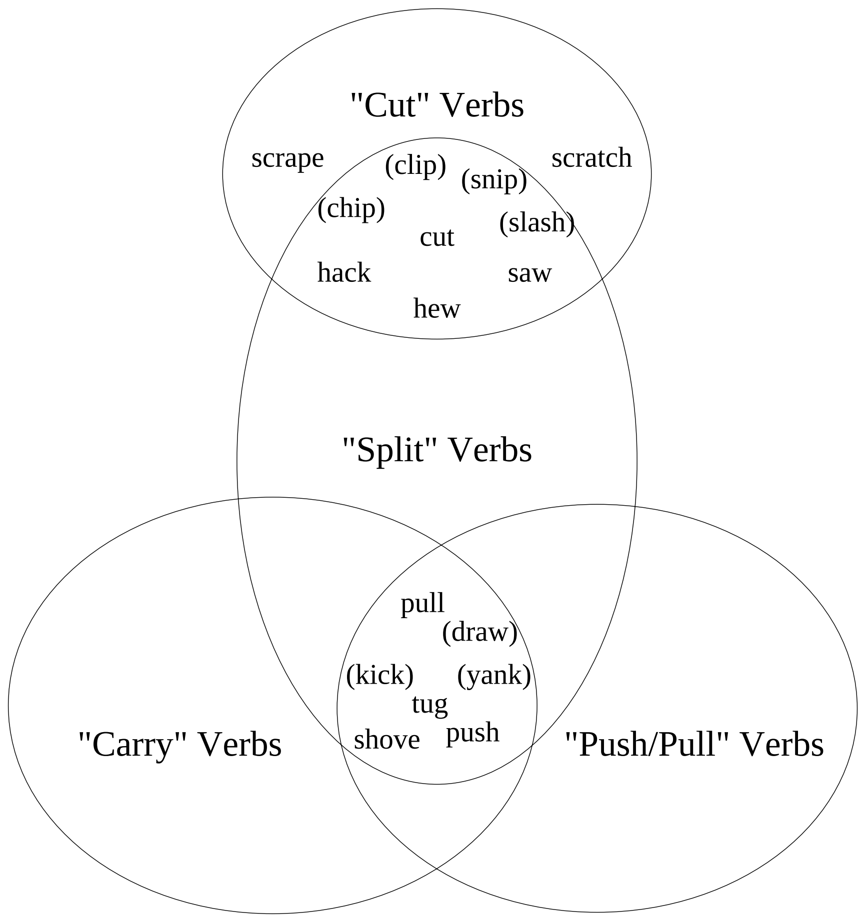

Figure 2.1: Intersective class formed from Levin *Carry*, *Push/Pull* and *Split* verbs

Similar sense extensions

A preliminary study of the Portuguese translation of the *Carry* verb class indicates that as in English, these verbs seem to take different alternations. Table 2.1 shows how these Portuguese verbs naturally cluster into two different subclasses, based on their ability to take the conative and apart alternations as well as path prepositions. These subclasses correspond very well to the English subclasses grouped by the intersective class.

The conative alternation in Portuguese is mainly *contra* (against), and the apart alternation is mainly *separando* (separating). For example, *Eu puxei o ramo e o galho separando-os* (*I pulled the twig and the branch apart*), and *Ele empurrou contra a parede* (*He pushed*

English	Portuguese	Conat.	Apart	Path
carry	levar	no	no	yes
drag	arrastar	no	yes	yes
haul	fretar	no	no	yes
heft	levantar com			
	dificuldade	no	no	yes
hoist	içar	no	no	yes
lug	levar com			
	dificuldade	no	no	yes
tote	levar			
	facilmente	no	no	yes
tow	rebocar	no	no	yes
shove	empurrar com			
	violência	yes	yes	yes
push	empurrar	yes	yes	yes
draw	puxar	yes	yes	yes
pull	puxar	yes	yes	yes
kick	chutar	yes	yes	yes
tug	puxar com			
	forca	yes	yes	yes
yank	arrancar	yes	yes	yes

Table 2.1: Portuguese *carry* verbs with their alternations

against the wall).

Changing class membership

We also investigated the Portuguese translation of some intersective classes of motion verbs. We selected the *slide/roll/run*, *meander/roll* and *roll/run* intersective classes. Most verbs have more than one translation into Portuguese, so we chose the translation that best described the meaning or that had the same types of arguments as described in Levin's verb classes.

The elements of the *slide/roll/run* class are *rebater* (*bounce*), *flutuar* (*float*), *rolar* (*roll*) and *deslizar* (*slide*). The resultative in Portuguese cannot be expressed in the same way as in English. It takes a gerund plus a reflexive, as in *A porta deslizou abrindo-se* (*The door slid opening itself*). Transitivity is also not always preserved in the translations. For

	rebater (bounce)	flutuar (float)	rolar (roll)	deslizar (slide)	derivar (drift)	planar (glide)
dative	yes		yes	yes		
∗conative	no		no	no		
caus./inch.	yes		yes	yes		
middle	yes		yes	yes		
accept. coref.	yes		yes	yes		
caus./inch.	yes		yes	yes		
resultative	yes	yes	yes	yes	yes	yes
adject. part.	yes	yes	yes	yes	yes	yes
ind. action	yes	yes	yes	yes	no	yes
locat. invers.	yes	yes	yes	yes	yes	yes
measure	yes	yes	yes	yes	yes	yes
∗adj. perf.	no	no	no	no	no	no
∗cogn. object	no	no	no	no	no	no
zero nom.	yes	yes	no	yes	yes	yes

Table 2.2: Portuguese *slide/roll/run* and *roll/run* verbs with their alternations

example, *flutuar* does not take a direct object, so some of the alternations that are related to its transitive meaning are not present. For these verbs, we have the induced action alternation by using the light verb *fazer* (*make*) before the verb, as in *Maria fez o barco flutuar* (*Mary floated the boat*).

As can be seen from Table 2.2 the alternations for the Portuguese translations of the verbs in this English intersective class indicate that they share similar properties with the English verbs, including the causative/inchoative. The exception to this, as previously noted, is *flutuar* (*float*). The result of this is that *flutuar* should move out of the *slide* class, which puts it with *derivar* (*drift*) and *planar* (*glide*) in the closely related *roll/run* class. As in English, *derivar* and *planar* are not externally controllable actions and thus don't take the causative/inchoative alternation common to other verbs in the *roll* class. *Planar* doesn't take a direct object in Portuguese, and it shows the induced action alternation the same way as *flutuar* (by using the light verb *fazer*). *Derivar* is usually said as "estar a deriva" ("to be adrift"), showing its non-controllable action more explicitly.

There are still many questions that require further investigation. First, since our experiment was based on a translation from English to Portuguese, we can expect that other verbs in Portuguese would share the same alternations, so the classes in Portuguese should by no means be considered complete. Second, since the translation mappings may often be many-to-many, the alternations may depend on which translation is chosen, potentially giving us different clusters, but it is uncertain to what extent this is a factor, and it would require further investigation. For this preliminary study, we chose the Portuguese verb that is most closely related to the description of the English verb in the Levin class.

We expect these cross-linguistic features to be useful for capturing translation generalizations between languages as discussed in the literature (Palmer and Rosenzweig, 1996), (Copestake and Sanfilippo, 1993), (Dorr, 1997).

2.3 Event Structure

Moens and Steedman (1988) elaborate Vendler's work (Vendler, 1967) on temporal and aspectual categories, and discuss the following types of events:

- **culmination**: a punctual or instantaneous event, which is accompanied by a transition to a consequent state

  Ex: Harry (has) reached the top.

- **point**: an indivisible event (not necessarily instantaneous) with no consequent.

  Ex: John hiccupped.

- **process**: an event which is extended in time but that does not lead to any particular conclusion or culmination. This type of event can take a *for*-adverbial phrase (e.g., *He climbed for an hour*), but will not allow for an *in*-adverbial phrase (e.g., **He climbed in an hour*).

  Ex: Harry climbed.

<table>
<tr><td rowspan="2"></td><td colspan="2">EVENTS</td><td>STATES</td></tr>
<tr><td>atomic</td><td>extended</td><td rowspan="3">(understand
love, know,
resemble)</td></tr>
<tr><td>+conseq</td><td>CULMINATION

(recognize, spot,
win the race)</td><td>CULMINATED
PROCESS
(build a house,
eat a sandwich)</td></tr>
<tr><td>-conseq</td><td>POINT
(hiccup,
tap, wink)</td><td>PROCESS
(run, swim, walk,
play the piano)</td></tr>
</table>

Table 2.3: Events and States (Moens and Steedman (1988))

- **culminated process**: state of affairs that extends in time but does have a particular culmination associated with it at which a change of state takes place. This type of event can take an *in*-adverbial phrase (e.g., *He climbed to the top in an hour*), but will not allow a *for*-adverbial phrase (e.g., **He climbed to the top for an hour*).
  Ex: Harry climbed to the top.

Table 2.3 (from that work), shows schematically how events can be characterized by atomicity or extension over time, and by whether they have a consequent state or not. Events can be inter-related through culmination, preparatory processes, iteration, and progressive links between them.

The relations between subparts of events lead to the notion of a single event structure, called a nucleus. The nucleus can be seen as a tri-partite event structure comprised of a preparatory process, a culmination and a consequent stage, as shown in Figure 2.2.

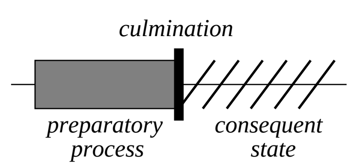

Figure 2.2: Moens and Steedman's tripartite structure of events

Verbs can be classified according to which stages of the event structure they participate in. Verbs of motion such as *run* and *bounce*, have a preparatory stage, but no culmination or consequent; verbs of contact such as *hit* and *kick* have both a preparatory stage and a culmination (when the contact between objects is established); and verbs of change of state, such as *break* participate in all three stages. These groupings can be easily related to the verb classes from Levin. Moreover, this decomposition also meets well the needs of animation, which requires a detailed characterization of each stage of an action in order to instruct a virtual human, as discussed in Chapter 4.

2.4 Related Work

The following subsections present work related to VerbNet, either in terms of building blocks for VerbNet or as different approaches from our work that will be contrasted in Chapter 3. Subsection 2.4.1 examines WordNet, an on-line lexical database to which verbs in VerbNet have mappings to. Subsection 2.4.2 discusses FrameNet, a frame-based database with a similar notion of verb groupings and a more detailed notion of thematic roles than VerbNet. Subsection 2.4.3 presents Jackendoff's work on Lexical Conceptual Structures (LCS), as well as Dorr's database with LCS descriptions for English verbs. Subsections 2.4.4 and 2.4.5 briefly discuss the generative lexicon view of Pustejovsky and the Acquilex LKB. Finally, the last two subsections 2.4.6 and 2.4.7 are devoted to English syntactic lexicons, Comlex and the Xtag grammar.

2.4.1 WordNet

WordNet (Miller, 1985; Fellbaum, 1998) is an on-line hierarchical lexical database for English nouns, verbs, adjectives, and adverbs. This database is organized into synonym sets (synsets), where each synset contains word forms referring to a given concept, as well as a definition gloss and an example sentence. The hierarchy is defined through hypernym and hyponym links (super-ordinate and subordinate relations) between synsets. It is a widely used resource for a variety of natural language processing tasks, it provides a domain-independent, large-coverage (over 120,000 lexical items are included) lexicon, and it is freely available.

WordNet's **verb** hierarchy is comprised of:

- verbs that denote actions and events, divided into 14 more specific semantic domains *(motion, perception, contact, communication, competition, change, cognition, consumption, creation, emotion, perception, possession, bodily care and functions, social behavior and interactions)*;

- verbs that denote states (elaborations of "be", such as *resemble, belong, suffice*; auxiliaries, control verbs such as *want, fail, prevent, succeed*; and aspectual verbs such as *begin*).

Verbs are grouped together as sets of synonyms, in which the concept of synonym is a little more relaxed than mere substitution of word pairs. Lexical items that can be used interchangeably within a context are considered synonyms. Idioms are included in a synset with an associated meaning (e.g., *kick the bucket* is grouped with *die*). The set of relations described in the verb hierarchy is derived from psychological experiments and is quite comprehensive, including hypernyms and hyponyms, as well as relations between individual lexical items such as synonyms, antonyms, troponyms (manner elaboration), entailment, and causation.

WordNet was designed mainly as a semantic network, and contains little syntactic information. The network contains explicit information about verbs, in the form of the relations amongst them, and implicit information about semantic and syntactic properties of these verbs in the examples and glosses, which can be related to approaches of meaning decomposition (such as Lexical Conceptual Structures), and verb classes. It does not provide a comprehensive account of possible syntactic frames and predicate argument structures associated with individual verb senses.

There is a vast amount of work done in developing new networks and extending WordNet for other languages. For example, the creation of EuroWordNet (Vossen, 2003), Chinese WordNet (Chen et al., 2002), MultiWordNet (Pianta et al., 2002), BalkaNet (Tufts, 2000), and CoreNet (Choi, 2003), among many others. In India, wordnets are being built for Hindi, Marathi, Tamil, Gujarathi, Oriya and Sanskrit, and Bengali.

2.4.2 FrameNet

The Berkeley FrameNet project (Baker et al., 1998) contains information about nouns, adjectives, and verbs based on semantically hand-annotated corpora. Word senses are grouped into conceptual structures, called "frames", which share certain semantic properties, and are representations of cognitive concepts.

FrameNet's database contains descriptions of the semantic frames and lexical items associated with their syntactic and semantic representations in contexts ("valency representations"). The project also includes example sentences for the frames described. The semantic frames are defined as schematic representations of situations involving various participants, propositions, and other conceptual roles (Fillmore, 1976). Examples of these high-level domain-independent frames include the JUDGMENT frame, with roles like JUDGE, EVALUEE, and REASON, and the STATEMENT frame with roles like SPEAKER, ADDRESSEE, and MESSAGE, as shown in sentences (2) and (3):

(2) [$_{Judge}$ She] **blames** [$_{Evaluee}$ the Government] [$_{Reason}$ for failing to do enough to help].

(3) [$_{Message}$ "I'll knock on your door at quarter to six"] [$_{Speaker}$ Susan] **said**.

Much like approaches that use semantic roles, FrameNet describes the argument structures of a lexical item in terms of roles called "frame elements". But in FrameNet these elements are specific to the concepts described by the frames, and can be obligatory or optional depending on the lexical item. Because of the specificity of frame elements, there is a large number of them, creating a vast sparse data problem for machine learning applications.[1] Several types of relations between frames are underway, including USE (a kind of partial inheritance), and efforts to assign links that establish relations between different frames.

[1]It is exactly for this reason that Gildea and Jurafsky (2002) mapped these roles to a more manageable set of 18 roles in their automatic semantic role labeling experiments.

Figure 2.3 shows a commercial transaction frame which involves frame elements BUYER, SELLER, MONEY, and GOODS. Verbs like *buy*, *sell*, *charge*, *cost*, *lease*, *pay*, etc. and nouns like *buyer*, *cost*, *goods*, *seller*, etc. are part of this frame.

BUYER buys GOODS from SELLER for MONEY SELLER sells GOODS to BUYER for MONEY	
Frame elements	Example sentences
BUYER	[$_{buyer}$ Kim] bought the sweater from Pat for $50 Pat sold the sweater to [$_{buyer}$ Kim] for $50
GOODS	Kim bought [$_{goods}$ the sweater] from Pat for $50 Pat sold [$_{goods}$ the sweater] to Kim for $50
MONEY	Kim bought the sweater from Pat for [$_{money}$ $50] Pat sold the sweater to Kim for [$_{money}$ $50]
SELLER	Kim bought the sweater from [$_{seller}$ Pat] for $50 [$_{seller}$ Pat] sold the sweater to Kim for $50

Figure 2.3: Commercial Transaction frame

Both WordNet and FrameNet contain a dictionary and a thesaurus. FrameNet aims to have a more comprehensive set of examples than WordNet and to recognize relationships between words with different parts of speech in a single frame. FrameNet does not provide a mapping between its entries and WordNet due to differences in what constitutes a word sense.

2.4.3 Lexical Conceptual Structures

Jackendoff (Jackendoff, 1983; Jackendoff, 1990b) argues for a semantic decomposition approach to verb semantics, in terms of their Lexical Conceptual Structures (LCS). The main components of the LCS are conceptual constituents, conceptual primitives, and semantic fields:

- conceptual constituents belong to a set of categories, such as thing, event, state, place, path, property, purpose, manner, amount, time.

- semantic fields are features such as +loc, +temp, +poss, that behave as selectional restrictions.
- conceptual primitives are elementary concepts such as BE (for states), GO (for events), STAY (BE with duration), CAUSE (causality), INCH (inchoative), EXT (spatial extension), etc.
  - conceptual primitives also describe prepositions, such as AT, IN, ON, FROM, etc.
  - conceptual primitives can be further restricted by adding semantic fields to them, such asGO_{+loc} to describe change of location, or GO_{+poss} to describe change in possession.

A semantic decomposition of verbs in terms of their lexical conceptual structures explains their syntactic properties. The meaning components specified for a verb make a strong case for the types of arguments it prefers (Talmy, 1985; Levin, 1993). Levin and Rapport (1988) argue that verbs with similar LCS components share syntactic behavior such as diathesis alternations, and as an example, show that only verbs with CAUSE can have a middle construction.

Applications of the LCS formalism include the PUNDIT/KERNEL text processing system from Palmer et al. (1993), which uses LCS as a way of encoding the predicate-argument structure and selectional preferences of verbs.

One of the best known uses of LCS is Bonnie Dorr's machine translation system which used this representation as an interligua (Dorr, 1993). Dorr also used semantic decomposition to create a database (the LCS database) for a set of verbs based on Levin's groupings (Dorr, 2001). Verb semantics are described by lexical conceptual structures, and their argument structure is described by a small set of thematic roles, with arguments being either obligatory or optional in certain positions. Currently this database contains 492 classes and over 4,400 verbs.

2.4.4 Generative Lexicon

Pustejovsky (Pustejovsky, 1991; Pustejovsky, 1995) argues for a generative view of lexical semantics, as opposed to decomposition of a lexical item into a fixed set of primitives. In this view, the lexical item carries a minimal set of information (as fields in a template) and operations are defined on how to combine this information.

The meaning of a sentence is given by the composition of the lexical items involved in the sentence. Pustejovsky argues for four levels of lexical semantic representation: *argument structure*, *event structure*, *qualia structure* and *inheritance structure*.

Verbs refer to a characterization of event structure in which three types of elements are considered: states, events and transitions. These elements can be interrelated and may overlap, precede or follow each other in the event structure. The argument structure reflects the possible arguments allowed for a verb. Four types of arguments are considered: true arguments which are required for the verb; optional arguments which may or not be present in the syntactic realization; shadow arguments that are usually already realized by the verb itself; and adjunct arguments. Nouns are described by a qualia structure, which is formed from four different roles: *constitutive*, *formal*, *telic* and *agentive*. The CoreLex (Buitelaar, 1998) lexicon based on the Generative Lexicon from Pustejovsky lexicon has concentrated on nouns rather than verbs, although descriptions for both are presented.

2.4.5 Acquilex

The Acquilex Lexical Knowledge Base (LKB) (Copestake, 1992) contains syntactic and semantic multilingual information extracted from machine readable dictionaries. It uses a typed-based structure feature mechanism, and makes use of default inheritance, default unification and lexical rule application as operations. It does not support more general forms of inference which would be needed for reasoning. The lexical semantic information for nouns is encoded through a set of features based on *qualia structure* (Pustejovsky, 1991).

2.4.6 Comlex

Comlex (COMLEX, 1994) is an English database, developed at New York University. It is a very rich syntactic resource but it is not in the public domain. Each lexical item has syntactic features and complements. There are 92 sub-categorization features described for the verbs. These features describe syntactic sub-categorization as well as several types of control such as *object control, subject control, variable control, arbitrary control.* The complement types reflect the syntactic pattern in which the item can appear, such as "to-inf-sc" (infinitival with subject control) as in *I wanted to come*; or "to-inf-rs" (infinitival with raising) as in *They seemed to win.* Comlex contains syntactic descriptions for about 38,000 head words (21,000 nouns, 6,000 verbs) but no explicit semantic information for these items is provided.

2.4.7 Xtag Grammar

The Xtag project (XTAG Research Group, 2001) developed a comprehensive English grammar based on the Tree Adjoining Grammar formalism (TAGs) (Joshi, 1985). The Xtag grammar and tools are public domain resources, with very rich descriptions for English verbs. Each lexical item is assigned a set of trees describing its basic syntactic properties. The trees reflect the predicate-argument structure of these items. Trees can be combined by two operations: substitution and adjunction. The grammar includes syntactic descriptions for 33,000 lexical items (9,000 verbs, 14,500 nouns), using 1,300 trees described and organized into 70 tree families. Because of the detailed syntactic characterization of its lexical items, the Xtag grammar is being used primarily as a large-scale parsing resource in natural language applications.

Chapter 3

VerbNet

VerbNet is a hierarchical verb lexicon with syntactic and semantic information for English verbs, using Levin verb classes (Levin, 1993) to systematically construct lexical entries. By using verb classes we capture generalizations about verb behavior, thus reducing not only the effort needed to construct the lexicon, but also the likelihood of introducing errors when adding a new verb entry. The first level in the hierarchy consists of the original Levin classes, with each class subsequently refined to account for further semantic and syntactic differences within a class. Each node in the hierarchy is characterized extensionally by its set of verbs, and intensionally by a list of the arguments of those verbs and syntactic and semantic information about the verbs. The argument list consists of thematic roles and possible selectional restrictions on the arguments expressed using binary predicates. The syntactic information maps the list of thematic arguments to deep-syntactic arguments (i.e., normalized for voice alternations, and transformations). The semantic predicates describe the participants during various stages of the event described by the syntactic frame.

Verb classes are hierarchically organized, ensuring that each class is coherent enough so that all its members have a common semantics and share a common set of thematic roles and basic syntactic frames. This requires some manual restructuring of the original Levin classes, and further refinements, which is facilitated by using intersective Levin classes (Dang et al., 1998). Furthermore, a particular verb may add more semantic information to the basic semantic components of its class. This hierarchical structure of the lexicon, inspired by the strong hierarchical organization of the Acquilex Lexical Knowledge Base (LKB) (Copestake, 1992), has facilitated the addition of new classes and new members as discussed in Chapter 7 and Chapter 8.

3.1 Components

Each verb class is completely described by its set of members, thematic roles for the predicate-argument structure of its members, selectional restrictions on these arguments, and frames consisting of a brief description, an example, a syntactic description, and a set of semantic predicates with a temporal function specifying whether the predicate is true in the initial ($start(E)$), preparatory ($during(E)$), culmination ($end(E)$), or consequent ($result(E)$) stage of an event. In deciding how to describe the semantic components of the classes, we were influenced by the requirements of the Parameterized Action Representation (PARs) (discussed in Chapter 4), which requires a very detailed level of representation. The PARs are used for planning applications and require that parts of the event such as the start and the end of actions, and the results of that action be explicitly stated. The Moens and Steedman (1988) event decomposition formalism matches exactly the requirements of the representation.

A class may be subdivided according to specific syntactic frames or semantic predicates which are true only for a subset of the class members. The information presented in the class is strictly monotonic, with each subclass always adding more information to its parent, or in the case of selectional restrictions, imposing further restrictions to the ones already present in the thematic role.

3.1.1 Thematic Roles

VerbNet's argument list consists of a set of 23 thematic roles, which have been created to map verb arguments for all classes. It includes commonly used roles, as well as roles which are more class-specific in order to better convey key semantic components for some classes.

Although it is difficult to determine a well-motivated set of thematic roles, our goal is to provide as much information as possible for classes, without details of each specific verb. As with other choices we made, we aim to capture generalizations about verb behavior within classes. The choice of using thematic roles instead of generic labels such as numbered arguments is based on the capacity of these roles to provide richer semantic information for the class members than it would be afforded by numbers alone. In other words, the specification of roles supplies part of the semantic description for the class which would otherwise not be made available by numbered arguments. Moreover, we try to use these roles consistently across classes and certain roles need to have specific characteristics to be present. For example, only *Agent* can be the volitional initiator of an event, and *Patient* is someone or something that is necessarily undergoing a process or being affected in some way because of the event. The additional information provided by the use of roles as opposed to numbers is made clear when distinguishing between classes which take the same syntactic frame.[1] For example, class *Admire-31.2* as well as other classes of *Psych verbs* use roles such as *Experiencer* and *Theme* for their arguments shown in Example (4):

(4) *The tourists admired the paintings*
Experiencer V Theme

whereas *Contact verbs* such as the members of the *Hit-18.1* class use roles such as *Agent* and *Patient* to describe the same transitive frame, this can be seen in Example (5):

[1]This extra information is also useful when dealing with causative verbs such as members of the *Run-51.3.2*, where the subject of the intransitive frame *John runs* would have to be described as Arg1.

(5) *Paula hit the ball*

Agent V Patient

There is no evidence that any small set of thematic roles will cover all the possible arguments for all kinds of verbs, and we also do not claim our set of roles to be exhaustive. We have found however, that for the 5,200 verb senses included in our lexicon this set of 23 roles offered enough descriptive information. These roles are not atomic nor do they fall naturally into hierarchical groupings, which would allow for subsumption relations by other roles. [2]

Most of the roles used for VerbNet are generally recognized. A list of the roles and example classes where they appear is given:

- **Actor:** used for some communication classes (e.g., *Chitchat-37.6*, *Marry-36.2*, *Meet-36.2*) when both arguments can be considered symmetrical (pseudo-agents).[3]
- **Agent:** generally a human or an animate subject. Used mostly as a volitional agent, but also used in VerbNet for internally controlled subjects such as forces and machines.
- **Asset:** used for the *Sum of money alternation*, present in classes such as *Build-26.1*, *Get-13.5.1*, and *Obtain-13.5.2* with 'currency' as a selectional restriction.[4]
- **Attribute:** attribute of Patient/Theme refers to a quality of something that is being changed, as in $(\textit{The price})_{att}$ *of oil soared.* At the moment, we have only one class using this role *Calibratable_cos-45.6* to capture the *Possessor subject possessor-attribute*

[2]However, there are, of course, a few exceptions such as Material/Source, Product/Goal, Topic/Theme, but we decided to keep those roles distinct for better characterization of the classes.

[3]Talmy (1985) argues that the symmetric verbs are not really symmetric (at least psychologically), but that the subject is a "figure" compared to a "ground". Dowty has a different point of view on symmetrical verbs and argues that changes in the order of the arguments does not affect the meaning of the sentence. We follow the latter.

[4]Currency is not restricted to money but anything that could be used for a trade, as in *Three bags of sheep's wool will buy you a computer.*

factoring alternation. The selectional restriction 'scalar' (defined as a quantity, such as mass, length, time, or temperature which is completely specified by a number on an appropriate scale) ensures the nature of Attribute.

- **Beneficiary:** the entity that benefits from some action. Used by classes such as *Build-26.1*, *Get-13.5.1*, *Performance-26.7*, *Preparing-26.3*, and *Steal-10.5*. Generally introduced by the preposition 'for', or double object variant in the *benefactive alternation*.

- **Cause:** used mostly by classes involving *Psychological verbs* and *Verbs involving the Body*.

- **Location, Destination, Source:** used for spatial locations

  - **Destination:** end point of the motion, or direction towards which the motion is directed. Used with a 'to' prepositional phrase by classes of change of location, such as *Banish-10.2*, and *Verbs of Sending and Carrying*. Also used as location direct object in classes where the concept of destination is implicit (and location could not be Source), such as *Butter-9.9*, or *Image_impression-25.1*.

  - **Source:** start point of the motion. Usually introduced by a source prepositional phrase (mostly headed by 'from' or 'out of'). It is also used as a direct object in classes such as *Clear-10.3*, *Leave-51.2*, and *Wipe_instr-10.4.2*.

  - **Location:** underspecified destination, source, or place, in general introduced by a locative or path prepositional phrase.

- **Experiencer:** used for a participant that is aware or experiencing something. In VerbNet it is used by classes involving *Psychological verbs*, *Verbs of Perception*, *Touch*, and *Verbs involving the Body*.

- **Extent:** used only in *Calibratable-45.6* class, to specify the range or degree of change, as in *The price of oil soared* $(10\%)_{ext}$. This role may be added to other classes.

- **Instrument:** used for objects (or forces) that come in contact with an object and cause some change in them. Generally introduced by a 'with' prepositional phrase. Also used as a subject in *Instrument Subject Alternation* and as a direct object for the *Poke-19* class in the *Through/With Alternation*, and in the *Hit-18.1* class for the *With/against Alternation*.

- **Material and Product:** used in the *Build* and *Grow* classes to capture the key semantic components of the arguments. Used by some classes from *Verbs of Creation and Transformation* which allow for the *Material/Product Alternation*.

  - **Material:** start point of transformation.

  - **Product:** end result of transformation.

- **Patient:** used for participants that is undergoing a process or that have been affected in some way. Verbs that explicitly (or implicitly) express changes of state have Patient as their usual direct object. We also use Patient1 and Patient2 for some classes of *Verbs of Combining and Attaching* and *Verbs of Separating and Disassembling*, where there are two roles which undergo some change with no clear distinction between them.

- **Predicate:** used for classes with a predicative complement.

- **Recipient:** target of the transfer. Used by some classes of *Verbs of Change of Possession*, *Verbs of Communication*, and *Verbs involving the Body*. The selection restrictions on this role always allow for animate and sometimes for organization recipients.

- **Stimulus:** used by *Verbs of Perception* for events or objects that elicits some response from an Experiencer. This role usually imposes no restrictions.

- **Theme:** used for participants in a location or undergoing a change of location. Also, Theme1 and Theme2 used for a few classes where there seems to be no distinction between the arguments, such as *Differ-23.4*, and *Exchange-13.6* classes.

- **Time:** class-specific role, used in *Begin-55.1* class to express time.

- **Topic:** topic of communication verbs to handle theme/topic of the conversation, or transfer of message. In some cases like the verbs in the *Say-37.7* class, it would seem better to have 'Message' instead of 'Topic', but we decided not to proliferate the number of roles.[5]

Each verb argument is assigned one (usually unique) thematic role within the class. A few exceptions to this uniqueness are classes which contain verbs with symmetrical arguments, such as *Chitchat-37.6* class, or the *ContiguousLocation-47.8* class. These classes have indexed roles such as *Actor1* and *Actor2*, as explained above.

Because the roles are assigned to all verbs in a class, verbs present in more than one class may have different roles. Sometimes this happens because they are different senses of the verb, and sometimes because that particular class is better expressed with a different set of roles. The verb *grow*, for example, belongs to several classes, taking different roles depending on its sense. In class *Appear-48.1.1* it has the sense of *develop*, *take shape*, with thematic roles Location and Theme; in class *Calibratable_cos-45.6*, it has the sense of *changing value*, and thematic roles Attribute, Extent, and Patient; in class *Grow-26.2* and *Build-26.1* this verb has the sense of *grow*, *farm*, *develop* (as in "grow vegetables"), and the roles are class-specific with Agent, Material, and Product to best capture the key semantic

[5]FrameNet has both TOPIC and MESSAGE as frame elements in the communication classes (COMMUNICATION, COMMUNICATION MANNER, COMMUNICATION NOISE, COMMUNICATION RESPONSE, STATEMENT, REQUEST, etc.)

components of these two classes.

3.1.2 Selectional Restrictions

VerbNet's selectional restrictions on the thematic role arguments are based on EuroWordNet (Vossen, 2003) top level entries. Our decision to use EuroWordNet was a practical one, we wanted to associate our resource with an ontology that was publicly available and that was widely used. Our selectional restrictions ontology is a hierarchy (*is-a*, mostly), with multiple inheritance, and no cycles in the inheritance ordering. The lack of cycles ensures that the transitive reflexive closure on the is-a relation is a partial order. The hierarchy is shown in Fig. 3.1.

Location is subdivided into three fields:

- *region* in order to express prepositional phrases such as 'from under the rug'
- *place* such as 'in Paris', 'under the rug'
- *object* such as 'on the table'

The location restriction on many roles is described as *[+location -region]* in order to avoid expressions such as 'from under the rug'. Some *Source* thematic roles allow for both expressions as in "She swept the crumbs from under the rug" and "She swept under the rug", thus consisting of only the *[+ location]* restriction.

3.1.3 Syntactic Frames

The syntactic frames in VerbNet act as a short-hand description for the surface realizations allowed for the members of the class. They describe constructions such as transitive, intransitive, prepositional phrases, resultatives, and a large set of Levin's alternations.

A syntactic frame consists of the thematic roles in their preferred argument position around the verb, the verb itself, and other lexical items which may be required for a

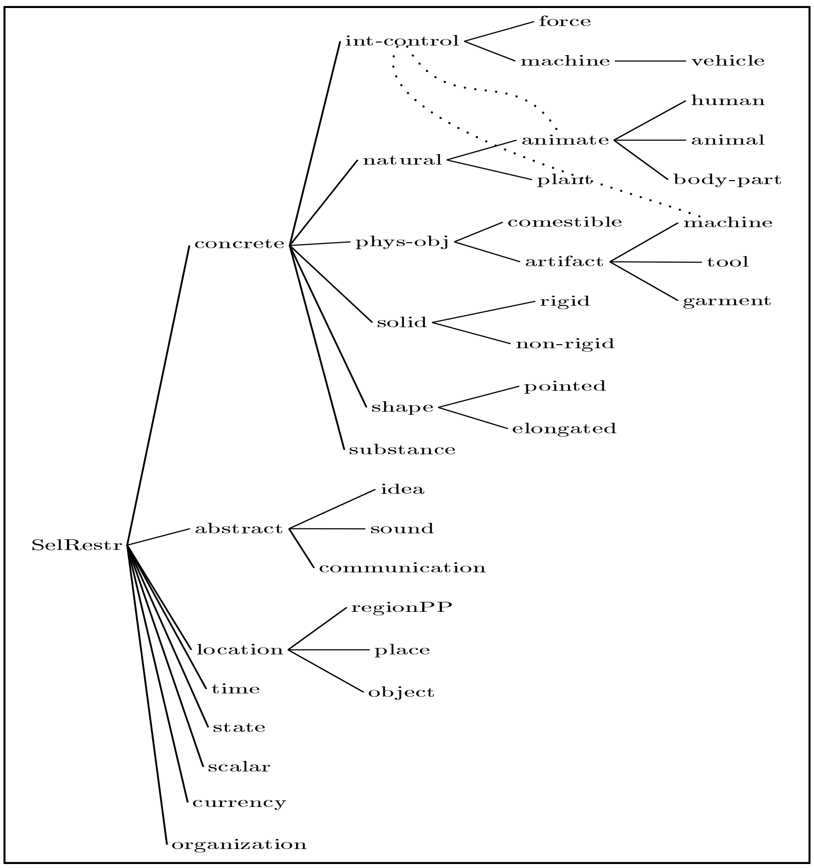

Figure 3.1: Selectional Restrictions used in VerbNet

particular construction or alternation. Additional restrictions may be further imposed on the thematic roles (quotation, plural, infinitival, etc.). Illustrations of syntactic frames are shown in examples 6, 7, and 8.

(6) *Agent V Patient*

(John hit the ball)

(7) *Agent V at Patient*

(John hit at the window)

(8) *Agent V Patient[+plural] together*

(John hit the sticks together)

There is also a hierarchy of prepositions, in order to specify which prepositions are possible in a particular frame. This is expected since many of Levin's alternations require specific prepositions such as 'as' or 'with/against'. A simplified version of the hierarchy is shown in 3.2. This hierarchy is derived from an extended version of work described in Jones and Boguraev (1987).

3.1.4 Semantic Predicates

Semantic predicates which denote the relations between participants and events are used to convey the key components of meaning for each class. The semantic information for the verbs in VerbNet is expressed as a conjunction of semantic predicates, such as *motion, contact, transfer_info*, which can be negated (expressed using a '!' in front of the predicate) or not. These can take arguments over the verb arguments, as well as over implicitly existentially quantified event variables. The semantic predicates in VerbNet are not overloaded (i.e., they always have the same number of arguments), which causes difficulties when the same predicate is applied to cases with plural subjects *property(Patient1+Patient2,Prop)*

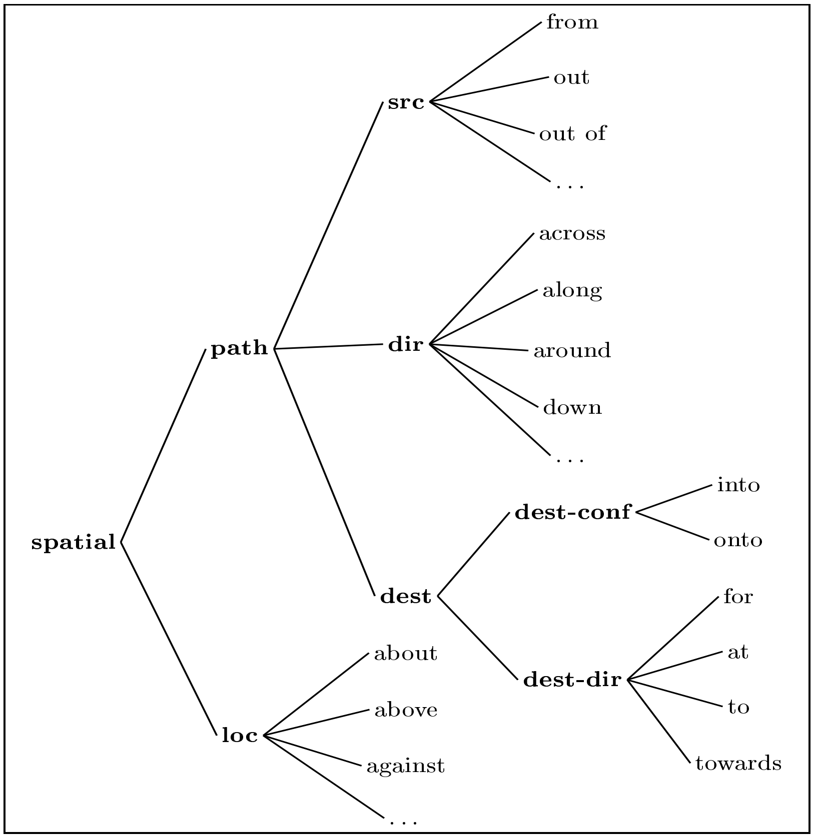

Figure 3.2: Prepositions in VerbNet

versus *property(Patient,Prop)*. A possible solution to this problem would be to have arguments defined as lists, such as in *property([Patient1, Patient2],Prop)*, in a similar way to the PUNDIT/KERNEL system (Palmer et al., 1993).

As the classes may be distinguished by their temporal (aspectual) characteristics (e.g., *Verbs of Assuming a Position* vs. *Verbs of Spatial Configuration*), it is also necessary to convey information about when each of the predicates apply. In order to capture this information, we included a time function in the semantic predicate adapted from the event decomposition from Moens and Steedman (1988). Our choice was influenced by the detailed level of representation of actions required by the PARs. The PAR representation needs to make explicit the pre-conditions, post-conditions, and results of actions for planning purposes. In VerbNet the time function specifies whether the predicate is true at all times in the event (E), at the start ($start(E)$), in the preparatory ($during(E)$), culmination ($end(E)$), or consequent ($result(E)$) stage of the event. This division of the event variable structure permits us to express the key semantics components of classes of verbs like change of state verbs whose adequate description requires reference to a complex event structure. In the case of a verb such as "break", it is important to make a distinction between the state of the object before the end of the action ($during(E)$), and the new state that results afterwards ($result(E)$).

VerbNet's semantic predicates fall into four classes:

- **General predicates** include predicates such as *motion* and *cause* and are assumed to be generic across a large class of verbs and across languages.

- **Variable predicates** are predicates whose meaning is assumed to be in a one-to-one relation with a set of words in the language. We are currently using *Prep, Adv, and Pred* as variable predicates.

- **Specific predicates** are predicates that carry specific verbal meaning, such as *suffocate*, which are common to the set of verbs in the *Suffocate-40.7* class.

- **Predicates for multiple events** are predicates used to express relations between events such as *throw a ball across the room*, where the *throwing* event is followed immediately by a change of location of the ball. A few of these types of predicates are used in VerbNet and they are derived from Allen and Ferguson's work on temporal logic (Allen and Fergunson, 1994). While the Moens and Steedman theory is used to express the decomposition of a single event, the Allen and Ferguson theory deals with the relation between two events. Figure 3.3 (taken from that work) shows the possible relations between time periods. We also include *equals(E0, E1)* for events that happen simultaneously.

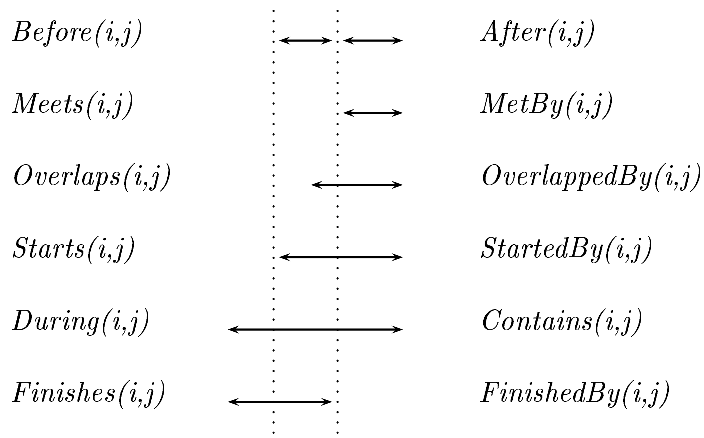

Figure 3.3: The possible relations between time periods (equality not shown)

The arguments in the semantic predicates can be of the following types:

- **Event:** the event variable E or a subpart of it (start(E), during(E), end(E), result(E)), describing for which part in the event the semantic predicate holds true.

- **Constant:** an argument in a predicate that specifies a property of that class but is not one of the roles. This type of arguments are used in order to allow the predicates to the used in different classes. Examples include *forceful* and *directedmotion* used in the *manner* predicate.

- **ThemRole:** which include thematic roles present in the syntactic frame, and thematic roles not present in this particular syntactic frame but available for the class (preceded by '?').

- **Verb Specific:** used for arguments which are instantiated differently by verbs in the class. Examples of these are *Direction* (direction of the exertion of force during the event for Put_direction-9.4 class), and *Form* (physical form of the result for *Break-45.1* and *Bend-45.2* classes)

Relations between verbs (or between verb classes) such as antonomy and entailment present in WordNet can be predicted upon verification of the predicates used. For example, classes with change of location of the object, *Pocket* and *Remove*, have the same predicates *cause* and *location* used differently (negated in different places). Moreover, relations between verbs (and verb classes) such as the ones predicted in FrameNet, can also be verified by the semantic predicates, for instance all of the *Communication* classes have the same predicates of *cause* and *transfer_info*.

Aspect in VerbNet is captured by the time function argument present in the predicates, according to Moens and Steedman tri-partite event structure. Verbs that denote activities or processes, such as motion verbs, only have predicates referring to the *during(E)* stage of the event (except for resultative frames). Verbs that denote activities with an end point, such as verbs of contact, have predicates that refer to both the *during(E)* and the *end(E)* stage of the event. Verbs that denote accomplishments, like the creation verbs, have in addition to the other stages, predicates referring to the *result(E)* stage. States are handled by having predicates that refer to the whole event *(E)*. Verbs which are aspectually ambiguous in English, such as *read*, are described as being activities, and would need a result added to them if an "in"-adverbial is adjoined to the basic frame (e.g., *John read the book in an hour.*)

3.1.5 Example

A simplified example entry for the class *Hit-18.1* is given in Fig. 3.4. Because the semantic predicates are valid for all members of the class, they lack verb-specific information; for example, in the case of *kick*, we do not have an explicit representation of the body-part or instrument used to establish the kicking contact (in this case foot). A natural extension to the class predicates are verb-specific predicates added as daughters of the class predicates, similarly to the construction of subclasses. In the case of the verb *kick*, for example, a new predicate such as *instrument(foot)* could be added, with all other predicates inherited from

Hit-18.1.

3.2 Refinement of classes and the hierarchical nature of VerbNet

In order for the members of each class to be coherent with respect to the thematic roles, selectional restrictions, syntactic frames, and semantics they allow, an extensive refinement of the original classes was necessary with 194 new subclasses with a depth never greater than 3 added. Some of the original verbs from Levin are not present, because they presented conflicting information. More details of this refinement can be found in (Kipper et al., 2004a; Kipper et al., 2004b).

This refinement and addition of new subclasses provides a hierarchical organization for VerbNet which is illustrated in Figure 3.5. A child subclass inherits all the information from its parent class, and adds information to it, which can be in terms of imposing further restrictions on the roles, or adding syntactic frames or semantic predicates to the subclass.

The *Transfer of a Message* verb class in the example is subdivided into three levels. At the top level are thematic roles, syntactic frames and semantic predicates shared by all members of the class. In this particular case, there is a transitive frame with the Topic (message) as the direct object (Agent Verb Topic), as in "John explained trigonometry", and a frame for Topic and Recipient (Agent Verb Topic to Recipient), as in "John taught math to Mary". Both syntactic frames have semantic predicates expressing the transfer of information event, but in the first case the Recipient is underspecified. Some of the verbs in this class are able to participate in other syntactic frames as well. Verbs at the second level can take the ditransitive frame (Agent Verb Recipient Topic) in addition to the frames and predicates inherited from the parent class.

Class	*Hit-18.1*			
Parent	—			
Themroles	Agent Patient Instrument			
Selrestr	Agent[+int_control] Patient[+concrete] Instrument[+concrete]			
Frames	Name	Example	Syntax	Semantics
	Basic Transitive	Paula hit the ball	Agent V Patient	cause(Agent, E) manner(during(E), directedmotion, Agent) !contact(during(E), Agent, Patient) manner(end(E),forceful, Agent) contact(end(E), Agent, Patient)
	Resultative	Paul kicked the door open	Agent V Patient Adj	cause(Agent, E) manner(during(E), directedmotion, Agent) !contact(during(E), Agent, Patient) manner(end(E), forceful, Agent) contact(end(E), Agent, Patient) Pred(result(E), Patient)
	Resultative	Paul hit the window to pieces	Agent V Patient Prep[to/into] Oblique[+state]	cause(Agent, E) manner(during(E), directedmotion, Agent) !contact(during(E), Agent, Patient) manner(end(E), forceful, Agent) contact(end(E), Agent, Patient) Pred(result(E), Patient)
	Conative	Paul hit at the window	Agent V at Patient	cause(Agent, E) manner(during(E), directedmotion, Agent) !contact(during(E), Agent, Patient)

Figure 3.4: Simplified VerbNet entry for *Hit-18.1* class

Class	Transfer_mesg-37.1		
Parent	—		
Members	cite(1,3,4), demonstrate(1), ...		
Themroles	Agent Topic Recipient		
Selrestr	Agent[+animate] Topic[+message] Recipient[+animate]		
Frames	Name	Syntax	Semantic Predicates
	Transitive	Agent V Topic	transfer_info(during(E),Ag,?,Top)∧ cause(Ag,E)
	Dative (to-PP variant)	Agent V Topic Prep(to) Recipient	transfer_info(during(E),Ag,Rec,Top)∧ cause(Ag,E)

Class	Transfer_mesg-37.1-1		
Parent	Transfer_mesg-37.1		
Members	dictate(2), quote(1), read(3)		
Themroles			
Selrestr			
Frames	Name	Syntax	Semantic Predicates
	Dative (ditrans variant)	Agent V Recipient Topic	transfer_info(during(E),Ag,Rec,Top)∧ cause(Ag,E)

Figure 3.5: Example entries for the *Transfer of a Message - levels 1* and *2* classes

3.3 Status of VerbNet

VerbNet currently has descriptions for over 5,200 verbs (over 3,800 lemmas) distributed in 237 top-level classes, and 194 new subclasses. For these descriptions, we use 23 thematic roles which can be constrained through a set of 36 selectional restrictions. For the syntax and semantic descriptions, we use 357 syntactic frames, and 94 semantic predicates. There is also a shallow hierarchy of prepositions with 57 entries.

Table 3.1 shows the various stages of the lexicon and the contribution of each of the extensions. V1.0 is the initial lexicon with refinements. V1.5 refers to the lexicon after the extensive revisions resultant from the experiment with PropBank as presented in Chapter 5. V2.0 and V2.2 are versions of the lexicon after integration with Korhonen and Briscoe's resource as discussed in Section 7.1 and after the incorporation of verbs from the LCS database as described in Section 7.2, respectively. V-clu and V-wn reflect versions of the lexicon extended by the latest experiments, clustering and correlations with WordNet,

	senses	lemmas	classes
v1.0	4173	3007	191
v1.5	4227	3007	191
v2.0	4526	3175	237
v2.2	4955	3601	237
v-clu	5002	3611	237
v-wn	5257	3819	237

Table 3.1: Status of VerbNet

presented in Chapter 8.

3.4 Related work

VerbNet's broad-coverage with explicit syntax and semantics, attempts to address several gaps present in other resources. This subsection presents comparisons between VerbNet and other lexicons.

WordNet was designed mainly as a semantic network, and contains little syntactic information. The senses are fine-grained, without a mapping to an underlying notion of semantic components and a systematic extension of basic senses. WordNet also does not provide a comprehensive account of possible syntactic frames and predicate argument structures associated with individual verb senses which may required by many natural language applications. VerbNet, in contrast, includes explicit predicate argument structures for verbs in their classes, as well as a way to systematically extend those senses based on the key semantic components of each class. Furthermore, these two resources present relations (such as antonomy and entailment) among lexical items in different ways, they are made explicit in WordNet and are implicit in our lexicon. VerbNet verbs have been mapped to their corresponding WordNet synsets as presented in Section 6.1.

FrameNet and VerbNet both contain the notion of verb groupings. FrameNet is based on shared semantics and makes no claim about common syntactic behavior being caused by similar semantic components. Verbs belonging to the same frame do not necessarily participate in the same diathesis alternations, and syntactic behavior of these verbs do not determine their class membership. The members do not have a more explicit semantics other than what is provided by the semantic labels of the frame. Unlike VerbNet which uses a small set of thematic roles, FrameNet uses frame elements which are particular to a lexical item or to small groups of frames, making the annotation and addition of new data an expensive and time-consuming task. This large set of roles also make automatic machine learning tasks difficult. In addition, FrameNet provides relationship between different parts of speech, while VerbNet concentrates on verbs. VerbNet's verbs have been mapped to FrameNet frames as discussed in Section 6.2.

Like VerbNet, Dorr's LCS database (Dorr, 2001) includes a large number of verbs organized into semantic classes derived from Levin's classification. This database was augmented from the original classes and includes 492 classes, mostly refinements of the original classes covering over 4,400 verbs. Although this lexicon attempts to address the limitations of previous work by including a simplified notion of syntactic frames, it does not provide links to actual instances in corpora and the relation between syntax and semantics is not made explicit. We analyzed the new additions to this database, and incorporated a large set of them into VerbNet as discussed in Subsection 7.2. Much of that data however had been automatically acquired and presented inconsistencies that barred their inclusion in our lexicon.

The Generative Lexicon (Pustejovsky, 1991; Pustejovsky, 1995) includes both a syntactic and a semantic representation of verbs, with an explicit characterization of event structure. The elements which distinguish among Pustejovsky's states, events and transitions are captured in VerbNet by the temporal function in the semantic predicates specifying the stages of the event in which that predicate holds. The majority of the work in the CoreLex (Buitelaar, 1998) however, has concentrated on nouns, rather than verbs.

Xtag (XTAG Research Group, 2001) and Comlex (COMLEX, 1994) are large-scale resources with explicit and well-characterized syntactic descriptions for lexical items. These two resources make no distinction between verb senses and do not at this moment provide explicit semantics. We mapped VerbNet's syntactic frames to Xtag tree families as discussed in Section 6.3 thus providing VerbNet with a greater syntactic coverage by allowing transformations to the basic frames.

Chapter 4

Describing actions for a simulated environment

One of VerbNet's primary applications has been as a basis for Parameterized Action Representations (PAR) (Badler et al., 1999) which are used to animate the actions of virtual human agents in a simulated 3D environment. PARs provide a conceptual representation of different types of actions that involve changes of state, changes of location (kinematic) and exertion of force (dynamic). PARs are hierarchical, parameterized structures that facilitate both visual and verbal expressions (Badler et al., 2000). In order to support the animation of the actions, PARs have to make explicit many details that are often underspecified in the language. This detailed level of representation also provides a suitable pivot representation for generation in other natural languages, i.e., a form of interlingua. This will be illustrated by examples of how certain divergences in machine translation can be handled, focusing specifically on *verb-framed* and *satellite-framed* languages as described in Talmy (1991).

4.1 PAR representation

We use *parameterized action representations* to animate the actions of virtual human agents. The PAR for an action includes the action's *participants* (its agent and objects), [1] as well as kinematic properties such as its *path*, *manner* and *duration*, and dynamic properties, such as its *speed* and *force* (see Fig. 4.1). The representation also allows for traditional state-space properties of actions, such as *applicability conditions* and *preparatory actions* that have to be satisfied before the action can be executed, and *termination conditions* and *post assertions* which determine when an action is concluded and what changes it makes to the environment state.

CONTACT PAR

$$\begin{bmatrix} activity : \begin{bmatrix} \texttt{ACTION} \end{bmatrix} \\[1em] participants : \begin{bmatrix} agent : & \texttt{AGENT} & \\ objects : & \texttt{OBJ1}, & \texttt{OBJ2} \end{bmatrix} \\[1em] applic\_cond : \begin{bmatrix} \texttt{reachable(OBJ1)} \\ \texttt{have(AGENT,OBJ2)} \end{bmatrix} \\[1em] preparatory\_spec : [\texttt{get(AGENT,OBJ2)}] \\[1em] termination\_cond : [\texttt{contact(OBJ1,OBJ2)}] \\[1em] post\_assertions : [\texttt{contact(OBJ1,OBJ2)}] \\[1em] path, duration, motion, force \\[1em] manner : \begin{bmatrix} \texttt{MANNER} \end{bmatrix} \end{bmatrix}$$

Figure 4.1: A PAR schema for actions of contact

It is possible to create a hierarchy of actions, based on VerbNet, exploiting the idea

[1] Objects and agents are stored in a hierarchy and have a number of properties associated with them. Properties of the objects may include their location and status. Agents have capabilities, such as the ability to walk or swim, and properties such as their strength and height.

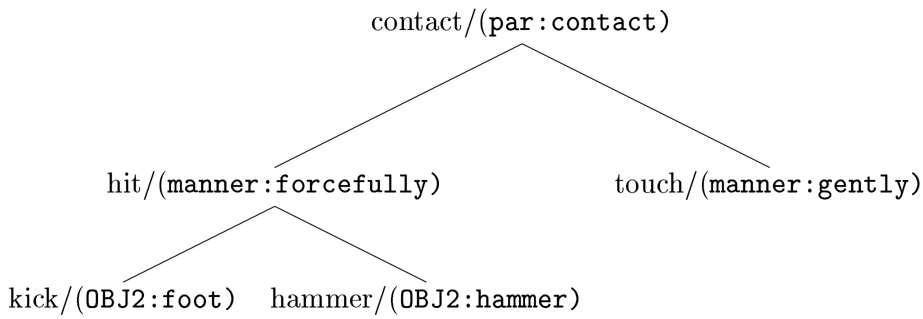

Figure 4.2: A lexical/semantic hierarchy for actions of contact

that verbs can be represented in a lattice that allows semantically similar verbs, such as motion verbs or verbs of contact, to be closely associated with each other under a common parent that captures the properties these verbs all share (Dang et al., 1998). The highest nodes in the hierarchy are occupied by generalized PAR *schemas* which represent the basic predicate-argument structures for entire groups of subordinate actions. The lower nodes are occupied by progressively more specific schemas that inherit information from the generalized PARs, and can be instantiated with arguments from natural language to represent a specific action such as *John hit the ball with his bat.* The example in Figure 4.1 is a generalized PAR schema for contact actions between two objects. This schema specifies that the 'contact' action has an agent and two objects, and that the action is concluded when the two objects come together.[2] The preparatory specification of getting the second object is tested and carried out if the object is not already in the agent's possession. In order to describe a specific action, say *hammer*, we would combine all of its ancestor representations in the action hierarchy, as shown in Figure 4.2, and add the information specific to that action. Since *hammer* inherits from the PAR for *hit*, and ultimately from the PAR for *contact*, its representation would use the generalized 'contact' PAR, with a forceful manner, and a hammer as the instrument. The action *hit* does not specify

[2] In this example, the second object is the instrument with which the action is performed.

any instrument, but inherits the forceful manner and generalized contact PAR from its ancestors, and the action *contact* leaves both the instrument and the manner unspecified, and is associated only with the generalized contact PAR.

The PAR is intended to provide slots for information that is typically conveyed in modifiers or adjuncts in addition to internal verb arguments. As such, it is often the case that several different syntactic realizations can all map to the same PAR schema. For example, *John hit the ball*, *John hit the ball with a bat* and *John swung mightily and his bat hit the ball with a resounding crack* would all map to the same schema.[3]

4.2 PAR as an IL

The PAR representation for an action can be seen as a general template. PAR schemas include, as part of the basic sub-categorization frame, properties of the action that can occur linguistically either as the main verb or as adjuncts to the main verb phrase. This captures problems of divergences, such as the ones described by Talmy (Talmy, 1991), for verb-framed versus satellite-framed languages.

Additional information may be provided by a sentence that can modify the action's inherent properties, such as in *John hit the ball slowly*, where 'slowly' is not part of the initial representation of the action 'hit'. This new information needs to be added to the PAR schema.

Verb- versus Satellite-framed languages

Verb-Framed Languages (VFL) map the motion (path or path + ground location) onto the verb, and the manner either onto a satellite or an adjunct, while Satellite-Framed Languages (SFL) map the motion into the satellite, and the manner onto the main verb.

[3]The relationship between PARs and alternations may become much more complicated when we consider other verb classes such as change of state verbs.

English and other Germanic languages are considered satellite-framed languages, expressing the path in the satellite; Spanish, among other Romance languages, is a verb-framed language and expresses the path in the main verb. The pairs of sentences (9) and (10) from Talmy (1991) show examples of these divergences. In (9), in English, the exit of the bottle is expressed by the preposition *out*, in Spanish the same concept is incorporated in the main verb *salir* (to exit). In (10), the concept of *blowing out* the candle is represented differently in English and Spanish.

(9) *The bottle floated out*
La botella salió flotando
(the bottle exited floating)

(10) *I blew out the candle*
Apagué la vela soplándola
(I extinguish the candle blowing)

4.2.1 Motion

In order to capture generalizations about motion actions, we have a generalized PAR schema for motion, and our hierarchy includes different types of motion actions such as inherently directed motion and manner of motion actions that inherit from the more general schema, as shown in Figure 4.3. Directed motion actions, such as *enter* and *exit*, do not bring with them the manner by which the action is carried out but they have a inherent termination condition. For example, the instruction 'enter a room' may be accomplished by walking, crawling or flying depending on the agents' capabilities, but it should end when the agent is in the room. In contrast, manner of motion verbs express the action explicitly and do not have an intrinsic termination condition.

Motion is a type of framing event where the path is in the main verb for VFLs and in the satellite for SFLs. In (11), we see the English sentence expressing the 'enter' idea in

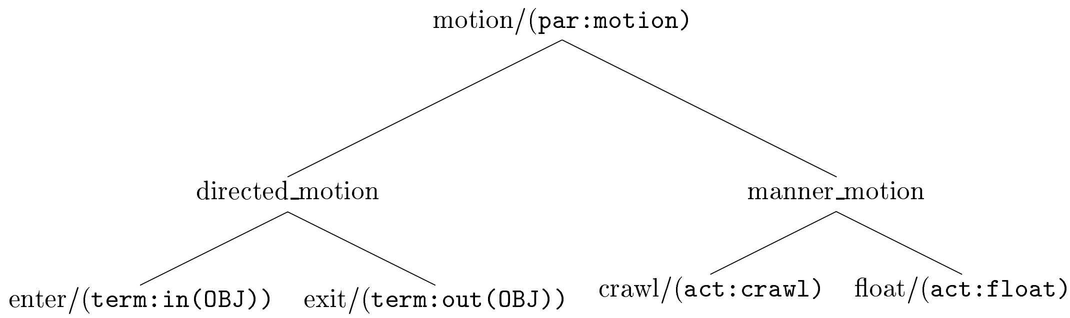

Figure 4.3: PAR schema hierarchy for motion actions

the preposition *into* whereas the Spanish sentence expresses it in the main verb *entrar* (to enter).

(11) *The bottle floated into the cave*
La botella entró flotando a la cueva
(the bottle entered floating the cave)

The PAR schemas do not distinguish the representation for these sentences because there is a single schema which includes both the manner and the path without specifying how they are realized linguistically. Mappings from the lexical items to the schemas or to constraints in the schemas can be seen in Figure 4.4.[4] Independently of which is the source language, the PAR schema selected is `motion`, the *activity* field, which determines how the action is performed (in this case, by floating), is filled by *float* (the main verb in English, or the adjunct in Spanish). The termination condition, which says that action ends when the agent is in the object, is added from the preposition in English and is part of the semantics of the main verb *to enter* in Spanish.

Because all of the necessary elements for a translation are specified in this representation, it is up to the language specific component to transform it into a surface structure

[4] A lexical item may have several mappings to reflect its semantics. For instance, *float* in English can be used also in the non-motion sense, in which case there will be two entries to capture that distinction.

EN float/[par:motion,activity:float]
into/[term:in(AG,OBJ)]

SP entrar/[par:motion,term:in(AG,OBJ)]
flotar/[activity:float]

Figure 4.4: Entries for the example sentences in (11)

MOTION PAR

$$\begin{bmatrix} \textit{activity} : \texttt{float} \\ \\ \textit{participants} : \begin{bmatrix} \textit{agent} : & \texttt{bottle} \\ \textit{object} : & \texttt{cave} \end{bmatrix} \\ \\ \textit{termination\_cond} : \texttt{in}(\texttt{bottle}, \texttt{cave}) \end{bmatrix}$$

Figure 4.5: A (simplified) PAR schema for the sentences in (11)

that satisfies the grammatical principles of the destination language.

Comparison with other work

This approach diverges considerably from the approach outlined in Palmer et al. (1999) which discusses the use of Feature-Based Tree Adjoining Grammars (Joshi, 1985),(Vijay-Shanker and Joshi, 1991) to capture generalizations about manner-of-motion verbs. They do not propose an interlingua but use a transfer-based mechanism expressed in Synchronous Tree Adjoining Grammars to capture divergences of VFL and SFL through the use of semantic features and links between the grammars. The problem of whether or not a prepositional phrase constitutes an argument to a verb or an adjunct (described by Palmer et al.) does not constitute a problem in our representation, since all the information is recovered in the same template for the action to be animated.

The PAR approach is much more similar to the LCS approach, used as an interlingua representation by Dorr (Dorr, 1993). Based on the assumption that motion and manner of motion are conflated in a matrix verb like *swim*, the use of LCS allows separation of the concepts of motion, direction, and manner of motion in the sentence *John swam across the lake*. Each one of these concepts is represented separately in the interlingua representation, as GO, PATH and MANNER, respectively. Our approach allows for a similar representation and the end result is the same, namely that the event of *swimming across the lake* is characterized by separate semantic components, which can be expressed by the main schema and by the *activity* field. In addition, our representation also incorporates details about the action such as applicability conditions, preparatory specifications, termination conditions, and adverbial modifiers. It is not clear to us how the LCS approach could be used to effect the same commonality of representation.

4.2.2 Instrument

The importance of the additional information such as the termination conditions can be more clearly illustrated with a different set of examples. Another class of actions that presents interesting divergences involves instruments where the instrument is used as the main verb or as an adjunct depending on the language. The sentence pair in (12) shows this divergence for English and Portuguese. Because Portuguese does not have a verb for *to spoon*, it uses a more general verb *colocar* (to put) as the main verb and expresses the instrument in a prepositional phrase. Unlike directed motion actions, a *put with hand-held instrument* action (e.g., spoon, scoop, ladle, etc.) leaves the *activity* field unspecified in both languages. The specific action is generated by taking the instrument into account. A simplified schema is shown in Figure 4.6.

(12) *Mary spoons chocolate over the ice cream*
Mary coloca chocolate sobre o sorvete com a colher
(Mary puts chocolate over the ice cream with a spoon)

Notice that the connection between *to spoon* and its Portuguese translation in the termination condition where the object of the verb, *chocolate*, has a new location which is

PUT3 PAR

$$\begin{bmatrix} \mathit{activity}: - \\ \\ \mathit{participants}: \begin{bmatrix} \mathit{agent}: & \texttt{Mary} \\ \mathit{objects}: & \texttt{chocolate}, \\ & \texttt{icecream}, \\ & \texttt{spoon} \end{bmatrix} \\ \\ \mathit{preparatory\_spec}: \texttt{get}(\texttt{Mary}, \texttt{spoon}) \\ \\ \mathit{termination\_cond}: \texttt{over}(\texttt{chocolate}, \texttt{icecream}) \\ \end{bmatrix}$$

Figure 4.6: Representation of the sentences in (12)

over the ice cream, in addition to the instrument.

4.3 Aggregate Movements

VerbNet has also been used as the natural language basis to the *ACUMEN* project which has as a goal the building of efficient computational models for creating and recognizing aggregate activities within a collection of entities (Allbeck et al., 2002). This system is designed to provide control of aggregate entities for simulations of military exercises, crowds, and urban environments in a manageable way, through the use of an aggregate movement recognizer, a feature-based recognition procedure, and a verb classification scheme.

In order to avoid code-specific action recognizers, which would be a time-consuming task and difficult to scale, we use natural language to communicate aggregate activities to the lower-level attractors and repulsors responsible for such animations. We decompose higher-level concepts, such as *assembling* and *dispersing*, into a small set of features that can be quickly and robustly detected in large aggregate populations.

The choice of features used to decompose the aggregate actions is based on the EMOTE system (Chi et al., 2000; Zhao, 2001), which incorporates Laban Movement Analysis (LMA)

EMOTE	VERB CLASSES				
	Gathering	Dispersing	Obj. Referential	Formation	Milling
Shape Advancing					
Retreating					
Spreading		x			
Enclosing	x				
Effort Slow					
Fast					
Sudden					
Sustained					x
Direct	x	x	x	x	
Indirect					x
Free					
Bound					
Other Obj Referent			x		
Structured				x	

Table 4.1: Verb classes for aggregate actions defined using EMOTE features

(Laban, 1971) to parameterize and to modulate action performance by changing the qualities of a given behavior. The appeal of such an approach is the small number of parameters needed to control a much larger set of actions. The EMOTE work focuses on two components of LMA, Effort and Shape. Effort comprises four motion factors: Space, Weight, Time, and Flow. Each motion factor is a continuum between two extremes: (1) *indulging* in the quality and (2) *fighting* against the quality. Shape changes in movement can be described in terms of three dimensions: Horizontal, Vertical, and Sagittal.

After extensively analyzing the individual behavior of verbs that describe activities of aggregates to determine features that span a broad and expressive action space, and deciding to adopt the idea of verb classes used in VerbNet, we grouped sets of these aggregate verbs into classes, extending the EMOTE features to group movement. Since EMOTE is designed for human arm gestures, the features had to be revised for aggregate entity movements. By using the EMOTE features as primitives, we are able to capture both generalizations and distinctions among sets of verbs as shown in Table 4.1. Verb classes created include verbs of *Gathering*, *Dispersing*, and *Milling*; and examples can be seen in Figure 4.7.

Gathering movements have an enclosing shape and direct effort, which means that the density of the aggregate is increasing and the movement has a focus.[5] *Assembling*, *congregating*, and *getting together* are quasi-synonyms of *gathering*. Our hierarchical classification scheme provides different levels of granularity which allow us to treat these terms as synonyms but, if required, also a more detailed classification. *Dispersing* movements are similar to gathering movements except that the shape is spreading instead of enclosing. *Dissipate*, *scatter*, and *spread out* have meanings similar to *dispersing*. Object Referential actions such as *surrounding* and *encircling* syntactically and semantically require an object which is the focus of the action. Unlike *dispersing* and *gathering*, which can have an implicit focus, object referential actions require an explicit focus. *Milling* actions are sustained actions lacking focus. These actions progress over a period of time in a wandering or meandering fashion. Some similar concepts in the milling class can be further distinguished by other Effort parameters. The action of *bustling*, for example, implies a faster movement than *milling*. *Formations* are aggregate actions with structure, such as lines or columns (e.g., a group of soldiers may be organized in a two-by-four formation, or a triangular formation). We have determined ways to recognize this structure in aggregate movements based on the position and orientation of the individuals. It is, however, difficult to generate these formations in a general way, many of them require precise positions and can not be easily characterized by a simple verb.

We have found that the Effort dimensions (slow-fast; sudden-sustained; direct-indirect; free-bound) have a meaningful correlation to aggregate behavior and are useful for distinguishing verb classes semantically. We have also found that two of the Shape dimensions

[5]This focus can be explicit, as in *gather around the door* or implicit, as in *gather together*.

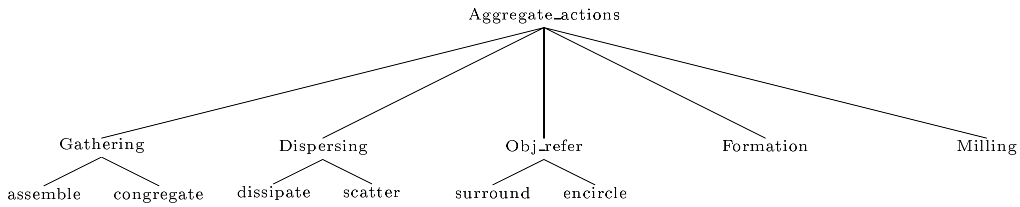

Figure 4.7: Examples of Aggregate actions

```
assemble    / ARG0-v
            / is_concrete(ARG0)
              is_plural(ARG0)
             !together_group(start(e),ARG0)
              transl_motion(during(e),ARG0)
              shape_enclosing(during(e),ARG0)
              effort_direct(during(e),ARG0)
              together_group(end(e),ARG0)
```

Figure 4.8: PAR schema for 'assemble'

(advancing-retreating; spreading-enclosing;) correlate to aggregate shapes. Other dimensions seem less appropriate (rising-sinking; left-right), most aggregate entities, unlike individual humans, do not have an inherent top and bottom or left and right. Other factors, such as the focus of attraction, are geometric features that have been found crucial to the proper characterization of group actions.

Figure 4.8 shows a PAR schema for the intransitive frame of *assemble* as in *The kids assemble (in a location)*, which has only one participant (a plural concrete entity). As can be seen from the example entry, the description of the action is closely related to a VerbNet entry, with syntactic frames and semantic predicates specified for a verb. Many of the semantic predicates for these aggregate actions are derived directly from the EMOTE features. The semantics of *assemble* establish that at the start of the event the participants are not together as a group, during the event the participants move, the shape of the group

Class	*Herd-47.5.2*			
Parent	—			
Members	accumulate aggregate amass assemble cluster collect congregate convene flock gather group herd huddle mass			
Themroles	Agent[+animate] Theme[+concrete +plural]			
Frames	Name	Example	Syntax	Semantics
	Intransitive	The kids are assembling	Theme V	!together(start(E),physical,Theme) together(end(E),physical,Theme)
	Causative	The teacher gathered the kids	Agent V Theme	cause(Agent,E) !together(start(E),physical,Theme) together(end(E),physical,Theme)

Figure 4.9: VerbNet entry for *Herd-47.5.2* class

is enclosing and the movement is direct, at the end of the event the participants should be together. The VerbNet entry for the verb *assemble* and other members of *Herd-47.5.2* class is given in Figure 4.9 for comparison. As can be seen, the argument structure and the predicates for the PAR schema are derived from the VerbNet class, with more specific predicates using EMOTE features added to better describe the group movement.

These verb definitions for describing movements of oriented entities are used in a scenario to program the activities of aggregate entities in a simulation. The scenario takes place in a schoolyard with eight children and a supervising teacher (see Figure 4.10).

Instructions which had to be performed by the children (e.g., *When the teacher opens the door, assemble*), and instructions given to the teacher (e.g., *If the bell rings, assemble the children*) are present in the simulation. The types of instructions given fall into three different categories: actions of an individual, individual actions given to an aggregate, and actions of an aggregate. Actions such as *blowing a whistle* or *opening a door* are included in instructions for an individual and performed only by that individual. Actions such as *panicking* are included in instructions for an entire group, but are performed separately by each member of the group. Finally, actions such as *gathering*, *assembling*, *dispersing*, *surrounding*, and *milling* can only be performed by aggregate entities and are described through our set of features. The complete script for this scenario is shown in Figure 4.11.

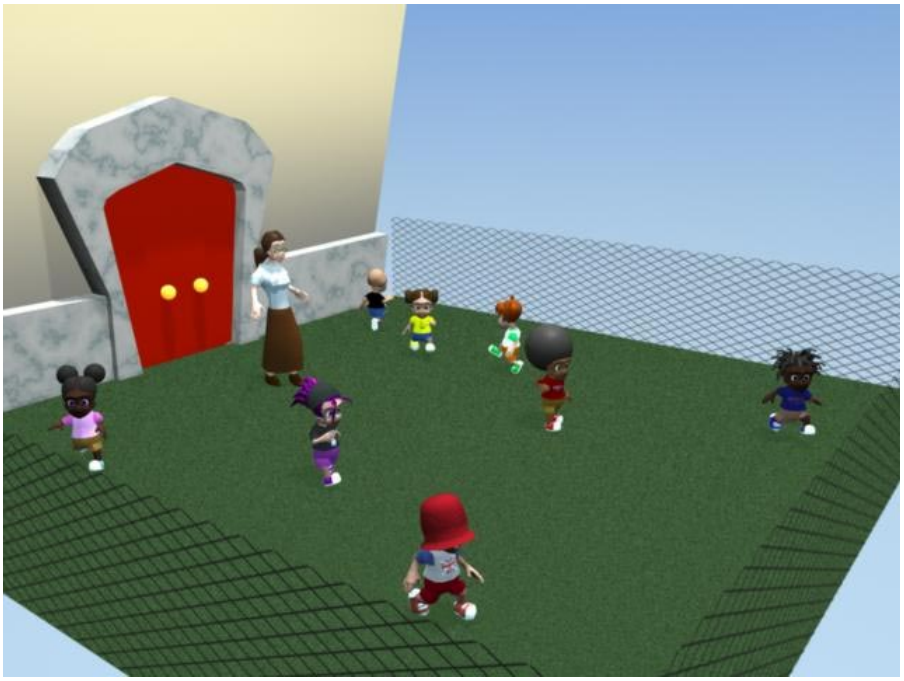

Figure 4.10: Playground scenario

Instructions given to the children:

1. When the teacher blows the whistle, disperse.
2. When the teacher opens the door, assemble.
3. When the teacher waves, gather in front of her.
4. If a skunk enters the playground, panic.

Instructions given to the teacher:

1. If the children are surrounding Ralph, blow a whistle.
2. If the children gather around the door, open it.
3. If the bell rings, assemble the children.

Figure 4.11: Script used in the Playground scenario

4.4 Conclusion

This chapter showed how VerbNet is used to provide semantics of actions of virtual human agents in a simulated 3D environment. We described VerbNet as a basis for Parameterized Action Representations and its implications as a form of interlingual representation to account for divergences in machine translation such as the ones described in Talmy (1991) between *verb-framed* and *satellite-framed* languages. It also showed how VerbNet, extended with the EMOTE features, is used in the *ACUMEN* project to account for descriptions of aggregate actions.

Chapter 5

Evaluating syntactic coverage against PropBank

5.1 Overview of PropBank

PropBank (Kingsbury and Palmer, 2002) is an annotation of the Wall Street Journal portion of the Penn Treebank II with dependency structures (or 'predicate-argument' structures), using sense tags for each word and argument labels for each dependency. An important goal is to provide consistent predicate-argument structures across different syntactic realizations of the same verb, as in the *window* in *[$_{ARG0}$ John]* **broke** *[$_{ARG1}$ the window]* and *[$_{ARG1}$ The window]* **broke**. PropBank can provide frequency counts for (statistical) analysis or generation components in a machine translation system, but provides only a shallow semantic analysis in that the annotation is close to the syntactic structure and each verb is its own predicate. In addition to the annotated corpus, PropBank provides a lexicon which lists, for each broad meaning of each annotated verb, its *frameset*, i.e., the possible arguments in the predicate and their labels and all possible syntactic realizations.

In PropBank, semantic roles are defined on a verb by verb basis. An individual verb's semantic arguments are simply numbered, beginning with 0. Polysemous verbs have several framesets, corresponding to a relatively coarse notion of word senses, with a separate set of numbered roles, a roleset, defined for each frameset. A frameset also includes a "descriptor" field for each role which is intended for use during annotation and as documentation, but which does not have any theoretical standing. In addition, each frameset entry is complemented by a set of examples, which attempt to cover the range of syntactic alternations afforded by that roleset. The collection of frameset entries for a verb is referred to as the verb's *frame*. As examples of a PropBank entries, we give the frame for the verbs *accept* in Figure 5.1, and *kick* in Figure 5.2.

<table>
<tr><td>ID</td><td>accept.01</td></tr>
<tr><td>Name</td><td>take willingly</td></tr>
<tr><td>VerbNet classes</td><td>13.5.2 29.2</td></tr>
<tr><td>Roles</td><td><table><tr><th>Number</th><th>Description</th></tr><tr><td>0</td><td>Acceptor</td></tr><tr><td>1</td><td>Thing Accepted</td></tr><tr><td>2</td><td>Accepted-from</td></tr><tr><td>3</td><td>Attribute</td></tr></table></td></tr>
<tr><td>Example</td><td>[ARG0 He] wouldn't accept [ARG1 anything of value] from [ARG2 those he was writing about].</td></tr>
</table>

Figure 5.1: A PropBank roleset for *accept*

While most framesets have three or four numbered roles, as many as six can appear, in particular for certain verbs of motion:

edge (sense: MOVE SLIGHTLY)

Arg0:causer of motion	Arg3:start point
Arg1:thing in motion	Arg4:end point
Arg2:distance moved	Arg5:direction

In addition to the arguments described in the framesets, verbs can take any of a set of general, adjunct-like arguments (ARGMs):

<table>
<tr><td>ID</td><td>kick.01</td></tr>
<tr><td>Name</td><td>drive or impel with the foot</td></tr>
<tr><td>VerbNet classes</td><td>11.4-2, 17.1, 18.1, 23.2, 40.3.2, 49</td></tr>
<tr><td>Roles</td><td><table><tr><th>Number</th><th>Description</th></tr><tr><td>0</td><td>Kicker</td></tr><tr><td>1</td><td>Thing kicked</td></tr><tr><td>2</td><td>Instrument (defaults to foot)</td></tr></table></td></tr>
<tr><td>Example 1</td><td>But [two big New York banks]$_{\text{ARG0}}$ seem *trace* to have kicked [those chances]$_{\text{ARG1}}$ [away]$_{\text{ARGM}-\text{DIR}}$, [for the moment]$_{\text{ARGM}-\text{TMP}}$, with [the embarrassing failure of Citicorp and Chase Manhattan Corp. to deliver $7.2 billion in bank financing for a leveraged buy-out of United Airlines parent UAL Corp]$_{\text{ARG2}}$</td></tr>
<tr><td>Example 2</td><td>[John]$_i$ tried [*trace*$_i$]$_{\text{ARG0}}$ to kick [the football]$_{\text{ARG1}}$, but Mary pulled it away at the last moment</td></tr>
</table>

Figure 5.2: The unique PropBank roleset for *kick*

LOC: location

EXT: extent, numer. role

DIS: discourse connectives

ADV: general-purpose

NEG: negation marker

MOD: modal verb

CAU: cause

TMP: time

PRP: purpose

MNR: manner

DIR: direction

Arg0 is roughly equivalent to Agent and Arg1 is usually similar to Theme or Patient. However, argument labels are not necessarily significant across different verb meanings of the same verb, or across different verbs, as thematic roles are usually taken to be.

The criteria to distinguish framesets is based on both semantics and syntax. Two verb meanings are distinguished as different framesets if they have distinct subcategorization frames, for example the verb *leave* has two framesets, as shown in examples (13), and (14).

(13) Roleset **leave.01** "move away from":

Arg0: entity leaving

Arg1: place left

Arg3: attribute

Ex: [$_{ARG0}$ The move] [$_{rel}$ left] [$_{ARG1}$ the companies] [$_{ARG3-as}$ as outside bidders.]

(14) Roleset **leave.02** "give":

Arg0: giver

Arg1: thing given

Arg2: beneficiary

Ex: [$_{ARG0}$ John] [$_{rel}$ left] [$_{ARG1}$ cookies] [$_{ARG2-for}$ for Mary]

However, alternations which preserve verb meanings, such as causative/inchoative or object deletion are considered to be one frameset only, as shown for in example (15) for *open.01*. Both, the transitive and intransitive uses of the verb *open*, correspond to the same frameset, with some of the arguments left unspecified.

(15) Roleset **open.01** "cause to open":

Arg0: agent

Arg1: thing opened

Arg2: instruments

Ex1: [$_{ARG0}$ John] [$_{rel}$ opened] [$_{ARG1}$ the door]

Ex2: [$_{ARG1}$ The door] [$_{rel}$ opened]

Ex3: [$_{ARG2}$ The wind] [$_{rel}$ opened] [$_{ARG1}$ the door]

Moreover, differences in the syntactic type of the arguments do not constitute criteria for distinguishing between framesets, for example, *see.01* allows for both an NP object or a clause object, as illustrated in example (16).

(16) Roleset **see.01** "view":

Arg0: viewer

Arg1: thing viewed

Ex1: [$_{ARG0}$ John] [$_{rel}$ saw] [$_{ARG1}$ the president]

Ex2: [$_{ARG0}$ John] [$_{rel}$ saw] [$_{ARG1}$ the President collapsed]

5.2 Mapping VerbNet to PropBank

PropBank consists of shallow, theory-neutral, semantic annotation for naturally occurring text that is intended to provide training data for statistical semantic parsers and taggers. Its effectiveness in this guise is attested to by the successful training of Gildea's semantic argument tagger (Gildea and Jurafsky, 2002). However, any 1M word corpus, no matter how diverse, will have only limited representation of many lexical phenomena. The true benefit of the semantic annotation will not be realized by statistical systems until they can be appropriately generalized from individual instances to groups of instances, and from individual verbs to classes of verbs. This is where the hierarchical structure of VerbNet can be of assistance, by suggesting classes for back-off purposes. VerbNet also suggests additional semantic predicates which, if applicable, can enrich the shallow PropBank semantics and might prove useful for applications such as Information Extraction, Question Answering and Machine Translation.

In turn, by comparing VerbNet's theoretically motivated sets of syntactic frames for an individual verb with the actual data, we can evaluate both the coverage of VerbNet and its theoretical underpinnings. There are many questions to be addressed with respect to coverage: *Do the predicted syntactic frames occur? Do other, unpredicted frames occur as well? Do the predicted prepositions occur? Do other, unpredicted prepositions occur as well? Are the relative frequencies of occurrence of the syntactic frames in line with expectations?* Depending on the answers to these questions, verbs may be inserted into or deleted from specific classes and entire classes may be restructured.

Roleset id="install.01" Roleset name="put in place"		
Lexicon class *Put-9.1*		
Roles	Description	Lexicon mapping
0	putter	Agent
1	thing put	Theme
2	LOC, where put	Destination

Figure 5.3: Mapping between thematic roles of class *Put-9.1* class and roleset install.01

Since VerbNet contains many of the same verbs that have been framed and annotated for PropBank, we have a unique opportunity to measure how well the linguistic intuitions motivating VerbNet are attested to in the actual data. The first step is to put the thematic roles of VerbNet into correspondence with the individual PropBank framesets. To this end, two things were done: a mapping between the framesets and VerbNet's verb classes; and a mapping between the argument labels in each roleset of a frameset to the thematic roles in VerbNet classes.

The process of assigning a verb class to a frameset was performed manually during the creation of new PropBank frames. The thematic role assignment, on the other hand, is a semi-automatic process which finds the best match for the argument labels, based on their descriptors, to the set of thematic role labels of VerbNet. This process required a certain amount of human intervention due to the variety of descriptors provided by the PropBank labels, the fact that the argument label numbers are not consistent across verbs, and the gaps in frameset to verb class mappings. Some of the mappings were very straightforward, whereas others reveal significant differences in approach.

Figure 5.3 shows an example of the mapping of roleset *install.01* with VerbNet class *Put-9.1*. In this case, all of PropBank's arguments have a mapping to the lexicon's thematic roles. Figure 5.4, gives an incomplete mapping, where arguments 2 and 3 are not mapped to any of the VerbNet's roles since they are not characteristic of other members of the *Pocket-9.10* class.

Roleset id="jail.01" Roleset name="put in jail"		
Lexicon class *Pocket-9.10*		
Roles	Description	Lexicon mapping
0	court, judge, jury	Agent
1	criminal	Theme
2	term	
3	cause, misdeed	

Figure 5.4: Mapping between thematic roles of *Pocket-9.10* class and roleset jail.01

5.3 Syntactic Coverage

We performed an experiment to evaluate the syntactic coverage of VerbNet against the syntactic frames encountered in a corpus (Kipper et al., 2004a). We used the Penn Treebank data, with PropBank annotation. For this evaluation we used 78,968 instances of PropBank, corresponding to 2,164 verb entries in VerbNet (1,595 lemmas). These entries comprised a large number of VerbNet classes (185 classes).

5.3.1 Finding the Syntactic Coverage of PropBank

The goal of this evaluation is to verify the syntactic coverage of VerbNet against a resource created empirically and independently. We aim for both a quantitative and qualitative analysis of how the syntactic frames in VerbNet reflect the syntactic frames in PropBank, per verb and per class. These results highlight gaps and inconsistencies in both resources. We expected to detect the basic frames (e.g., transitives, intransitives, ditransitives) for the classes, but the syntactic coverage depends on the actual distribution of frames in the corpus. As the results show VerbNet provided exact matches of syntactic frames in 84.67% of the instances, and this result goes up to 86.30% if we allow for more relaxed criteria of matching, such as ignoring preposition mismatches, or allowing ARGMs to match against specific roles.

For each class, we gather all the sentences in which verbs of that class appear (for example, for class *Put-9.1* we would look for all sentences of the verbs *put*, *arrange*, *immerse*, *install*, etc. which are members of this class). For each sentence, we analyze the sentence structure and follow the traces (a depth-first search in the treebank tree), we also undo transformations such as passive and create a simpler structure retaining the argument numbers from the original sentence. Example (17) shows a sentence found in the annotated corpus for the verb *accept*, and example (18) shows the revised annotation for that sentence, after this treatment.

(17) wsj/00/wsj_0051.mrg 15 18:
But the growing controversy comes as many practices historically accepted * as normal here – such as politicians accepting substantial gifts from businessmen or having extramarital affairs – are coming under close ethical scrutiny.

(18) [$_{ARG0}$ politicians] [$_{rel}$ accept] [$_{ARG1}$ substantial gifts] [$_{ARG2\_from}$ businessmen]

We then retrieve the mappings between the argument labels and the thematic roles for each class and generate a syntactic frame in a format similar to VerbNet from the corpus sentence. Because a verb may be mapped to different classes, different frames are generated (see example (19)). In this example, the verb *accept* is mapped to two different classes (*Obtain-13.5.2* and *Characterize-29.2*). For the second class (29.2), there is no thematic role mapping for ARG2.

(19) 13.5.2: Agent V Theme Prep(from) Source
29.2: Agent V Theme Prep(from) NP

Once the set of syntactic frames associated with all the instances of verbs in a given class has been retrieved (the "instance frames"), we assess their level of match to the syntactic frames of that class present in VerbNet. To do this, we allow an individual corpus instance (and its associated frame) to match the given verb class in three different ways:[1]

[1]Initially, we also had a fourth criteria (D) match to a frame in the verb class after Subject (Agent,

(A) exact match to a frame in the verb class (ignoring unmatchable fragments such as ARGM's, which do not have mappings to syntactic elements in VerbNet);

(B) match to any value for prepositions;

(C) retaining ARGMs for LOC, DIR, and EXT in the instance frame, with the following assumed mappings to thematic roles:
LOC = Location, Destination, or Source
DIR = Destination, Location
EXT = Extent

Table 5.1 shows results for class *Price-54.4* and its members *appraise*, *value*, *fix*, *rate*, *price*, *estimate*, and *assess*. The last row indicates total results for the class. There are 864 instances present in the corpus with PropBank annotation and mappings to VerbNet, 454 of which do not match under any of criteria (A)-(C). The next four columns indicate the number of instances which match under the respective match criteria discussed above, and the final column, "revised miss", indicates the number of misses which cannot be accounted for by a match with one of the other verb classes to which the instance is mapped. The percentages at the top of the table are computed as follows:
matches (using misses) = 1 - (misses/total)
matches (using revised misses) = 1 - (revised miss/(total - (miss - revised miss)) [2]

The remaining rows of the table indicate the breakdown of numbers by individual verbs. In this example, verbs *appraise*, *fix*, and *assess*, have few misses; the most frequent frames found in the corpus are *(Agent V Theme)* and *(Agent V Theme Prep(at) Value)* which are present in class *Price-54.4*. Many of the matches for the verbs in this class are exact matches (criterion A), for the verb *rate*. An analysis of the initial class revealed tha often

Theme, etc.) has been added to the instance frame in order to account for passive instances which often have a dropped subject in the corpus. The experiment has been modified to include a more detailed account of transformations, especially passive transformations and therefore this criteria is no longer necessary.

[2] We subtract from the total number of instances the number of such instances which are being discounted as matching another class.

Class 54.4 *Price*						
Matches: 47.45%						
Revised Matches: 59.33%						
verb	total	miss	(A)	(B)	(C)	revised miss
appraise	2	0	2	2	2	0
value	114	102	12	12	9	1
fix	65	10	55	55	54	2
rate	69	64	5	5	5	2
price	278	167	111	111	109	167
estimate	309	108	201	201	201	108
assess	27	3	24	24	23	1
Totals	864	454	410	410	403	281

Table 5.1: Example of counts per class

instance frames in the corpus are found with Value and no preposition *(Agent V Theme Value)*, this frame was not initially present in VerbNet and was added.

5.3.2 Evaluation of VerbNet

In the previous subsection, we described a method for calculating the match results of particular verb classes. Now we describe how we calculated the match results of VerbNet as a whole, and discuss these results.

Since particular instances in the PropBank corpus are often mapped to more than one verb class, there are essentially two methods for calculating total results: consider a particular instance a "match" if it matches at least one of the verb classes to which it is mapped; and consider a particular instance a "match" if it matches all of the verb classes to which it is mapped.

Of course, whether or not a particular instance matches a particular class again depends on which of the three criteria (A-C) discussed in Subsection 5.3.1 is used. Table 5.3.2 summarizes these results, with (A-C) representing a match to any one of the three criteria.

It is expected that VerbNet will contain frames never found in the PropBank corpus; of more concern, however, are those frames found in the corpus which VerbNet does not

	Matching any mapped class	Matching all mapped classes
A	66,866 (84.67%)	51,908 (65.73%)
B	68,086 (86.22%)	53,733 (68.04%)
C	62,378 (78.99%)	48,623 (61.57%)
(A-C)	68,149 (86.30%)	53,898 (68.25%)

Table 5.2: Results for the lexicon including exact and partial matches

predict. Ideally, we would like each corpus instance to match at least one of the verb classes to which it is mapped.

As can be seen from Table 5.3.2, 84.67% of the instances bear an exact match (criterion A) to at least one of the frames in one of their mapped verb classes, and 86.30%[3] match at least one verb class by at least one of the three match criteria (A-C). A few comments should be added regarding the different numbers resulting from the match criteria: we consider criterion (A) the baseline. Criterion (B) increases the match rate by allowing prepositions not anticipated in VerbNet to match. It is not clear, however, if adding a new preposition to the frames of a verb class is always the correct thing to do, to this end we investigated in detail the role of prepositions, as described in Subsection 5.3.3. Regarding criterion (C), we had no way of knowing a priori whether this criterion would increase or decrease the match rate. In fact, adding the assumed thematic role mappings for ARGM's significantly decreased the match rate.[4]

Several classes whose members are very frequent in the Wall Street Journal perform well, classes of verbs of Communication (classes 37.x) have average results over 82%; verbs of Change of State (classes 45.x) have syntactic frames predicted accurately in 84.8% of the cases; verbs of Hitting (classes 18.x, except class 18.4), average 83.58%; all classes of verbs of Appearance (classes 48.x), have accuracy above 92%. Oddly enough, some classes with few members and infrequent occurrences also performed very well, such as *Cooking-45.3*,

[3]This includes an increase of 0.5% which comes from allowing resultatives.

[4]We also added "ADV ARGM-MNR" to see if the resulting frame matches an expected frame with middle construction with no difference in the score, no middle construction was found in the corpus which reflected the middle construction frames expected.

Illustrate-25.3, *Exhale-40.1.3*, and *Animal_sounds-38*, all with results above 94%.

Previously, we had performed an experiment which provided us with a baseline of 78% exact matches, and about 80% relaxed matches. This initial experiment provided us with a large set of syntactic frames that occurred in the corpus and which were missing in VerbNet. The classes have been extended substantially by adding syntactic frames and revising prepositions based on this investigation. The frames that were lacking initially were in part a reflection of how we interpreted Levin's predictions of the syntactic frames allowed for the classes. For example, members of class 14 (*Learn* verbs) which include *acquire*, *read* and *learn* have a simple transitive frame *(Agent V Topic)* in over half of the instances, and VerbNet did not include this frame since it was not present explicitly in Levin.

The proper mapping between arguments and roles is not always possible. In class 45.6 (*Calibratable Change of State* verbs), all syntactic frames include both Patient and Attribute in order to express the divergences between *Oil's price soared* and *The price of oil soared* (Possessor Subject Possessor - Attribute Factoring Alternation) that distinguishes this class, but Patient (*oil*) and Attribute (*price*) are conflated as one argument in PropBank allowing for no possible syntactic frame match. In class 55.1 (*Begin* verbs), all VerbNet's frames require a Time role, which has no correspondence in PropBank. Verbs in class 12 (*Push/Pull* verbs), often appear in PropBank with ARG2 (direction), but the semantics of this particular class expresses only "exertion of force" and not motion. Some verbs in this class can be used with direction, and are cross-listed in class 11.4 *Carry* verbs, where direction is captured by the thematic role *Destination*. This last is also an example of why the second column in Table 5.2 shows a lower result than the first column.

These mismatches are a consequence of the fact that these two resources have been built independently, stem from different sources, and currently have disjoint coverage. This experiment provides important insights into what kinds of revisions will increase the robustness of both resources.

5.3.3 Using prepositions from the corpus to refine verb classes

In this subsection we compare and discuss preposition mismatches found while doing the initial syntactic coverage evaluation described in the previous Section. The discussion here is related to the initial experiment using the partially annotated version of PropBank, with almost 50,000 instances. The criteria used for matches in that experiment included both a notion of exact match, as well as a more relaxed notion of frame match except for preposition. VerbNet matches over 78% of all the syntactic frames found in PropBank. However, when restricting the frames found in PropBank to those without prepositions, the resulting match rate is almost 81%. This difference hints at the difficulty of accounting for semantically significant prepositions in sentences, and a proper account of this preposition-semantic relationship seems essential to us in order to build a more robust lexical resource. This detailed analysis of preposition mismatches is reported in (Kipper et al., 2004b).

Prepositions in the Corpus

Verb occurrences are partitioned according to whether a preposition occurs or not in the instance frame, and according to how well the constructed frame matches a VerbNet frame. Almost 4/5 of the verb instances studied do not contain a significant preposition in their PropBank annotation (and consequently their constructed frames do not include any prepositions).[5] On these instances, we obtained a 81% match rate under the strict criterion.

Of the 49,073 verb instances we are looking at, 9,304 instances had a significant preposition, with constructed frames including one or more prepositional items. For those we obtain match rates of 65% and 76% (depending on whether preposition mismatches were allowed or not).

The difference between the 81% match rate of the frames without prepositions and the 65%-76% match rate in the frames with prepositions is substantial enough to lead us

[5] We consider a preposition "significant" if the preposition object is a PropBank argument with a mapping to a thematic role, excluding preposition "by".

to believe that a close examination of the sentences containing a preposition and their comparison to VerbNet frames would allow us to improve the coherence of our verb classes.

Prepositional Mismatch

For the instances with significant prepositional items, 65% (6,033 instances) have constructed frames with an exact match to VerbNet. Of the remaining 3,271 instances, 1,015 are relaxed matches, and 2,256 do not bear any matches to VerbNet frames.

We focused on those verb instances which would have matched a VerbNet frame if only a different preposition had been used in the sentence or if the VerbNet frame had included a wider range of prepositions. In addition to the 1,015 instances, we looked at 652 verb instances, all of which share the following two properties: (i) that the verb in question is contained in multiple VerbNet classes, and (ii) that although the constructed frame matches one of those VerbNet classes exactly, there is at least one other class where it matches only under the relaxed criterion (when the value of the preposition is ignored). These instances are important because the value of the preposition in these cases can help decide which is the most appropriate VerbNet class for that instance. This information could then be used for coarse-grained automatic sense tagging – either to establish a PropBank Frameset or a set of WordNet senses for those instances, since verbs instances in our verb lexicon are mapped to that resource.

These 1,667 verb instances (1,015 preposition mismatches + 652 exact matches) comprise 285 unique verbs and are mapped to a total of 97 verb classes.

5.3.4 Explanation of Mismatch

After a close examination of these 1,667 instances, we verified that the mismatches can be explained and divided into the following cases:

1. cases where a preposition should be added to a VerbNet class (in some of these cases, a refinement of the class into more specific subclasses is needed, since not all members take the included preposition);

2. cases where the particular usage of the verb is not captured by any VerbNet entry (this is the case with metaphorical uses of certain verbs);

3. incorrect mappings between PropBank and VerbNet;[6]

4. cases where the PropBank annotation is inconsistent;

5. cases where the particular instance belongs to another VerbNet class (which are expected since the PropBank data used does not yet provide sense tags).

As an example, in the PropBank annotated corpus we find the sentence:

"Lotus Development Corp. feeds its evaluations into a computer...",

The verb *to feed* is present in four VerbNet classes. The frame resulting from translating the PropBank annotation to a VerbNet-style frame *Agent V Theme Prep(into) Recipient* bears a resemblance to a frame present in one of the classes (*Give-13.1*, syntactic frame *Agent V Theme Prep(to) Recipient*). This is a case where a VerbNet class requires refinements (with addition of new subclasses) to account for prepositions unique to a subset of the verbs in the class. It is an open question whether such refinements, taken to completion, would result in subclasses that are so fine-grained they have a membership of one. If so, it may be more appropriate to add verb-specific preposition preferences to existing classes.

Another example is the following use of *build* in the PropBank corpus:

[6]We asserted an error of 6.7% for the automatic mappings in a random sample of the data.

"...to build their resumes through good grades and leadership roles ..."

This sentence yields the frame *Agent V Product Prep(through) Material* after translating the PropBank annotation to a VerbNet-style frame. This frame bears a relaxed match to the *Agent V Product Prep(from, out of) Material* syntactic frame found in the *Build-26.1* class. In VerbNet, the phrase *"..through good grades ..."* is considered an adjunct and therefore not relevant for the syntactic frame. In PropBank, however, this phrase is annotated as an argument (Arg2), which maps to the *Material* thematic role in VerbNet. This example shows, as expected, mismatches between argument and adjuncts in the two resources.

As a final example, consider the following use of the verb *lease*:

"The company said it was leasing the site of the refinery from Aruba."

Two frames are constructed for this verb instance, one for each of the VerbNet classes to which the PropBank *lease* Frameset is mapped. Its membership in class *Get-13.5.1*, and class *Give-13.1* respectively yield the following two VerbNet-style frames:

(a) 13.1: Agent V Theme Prep(from) Recipient
(b) 13.5.1: Agent V Theme Prep(from) Source.

The first frame bears a relaxed match to a frame in its class (*Agent V Theme Prep(to) Recipient*) whereas the second is an exact match to a frame in the second class. In this instance, the preposition 'selects' the appropriate VerbNet class.[7] In fact, we expect this to happen in all the 652 instances with exact matches, since in those instances, the constructed frame bears an exact match to one VerbNet class, but a relaxed match to another. The

[7]It was pointed out that a possible interpretation is that "from Aruba" is linked to the "refinery" argument, in which case this instance would be translated as *Agent V Theme* and therefore have a perfect match to the *Give-13.1 class.*

different Framesets of a verb are typically mapped to distinct sets of VerbNet classes. If the preposition present in the sentence matches frames in only a subset of those VerbNet classes, then we are able to rule out certain Framesets as putative senses of the instance in a sense tagging task.

We presented a detailed account of how prepositions taken from a semantically annotated corpus can be used to extend and refine a hand-crafted resource with syntactic and semantic information for English verbs. That the role of prepositions should not be neglected can be clearly seen from the differential in match rates between those sentences with prepositions and those without. The significance of prepositions and their relation with verbs is of the utmost importance for a robust verb lexicon, not only as a syntactic restrictor, but also as a predictor of semantic content.

5.3.5 Conclusion

VerbNet has been considerably extended and currently matches frames in over 86% of the instances in the over 78,000 instances of the corpus annotated by PropBank. We believe that the two resources are complementary in many ways, and this result offers an insightful view on how both resources can be made more robust.

In addition, by having the mappings between the verb senses (framesets to verb classes) and the argument labels in the PropBank to VerbNet's thematic roles, it is also easy to experiment with other resources such as WordNet. All of VerbNet entries are mapped to WordNet synsets, so it is possible to verify which WordNet senses are really present in the corpus.

This experiment presented an evaluation of VerbNet and showed us its strengths and weaknesses. It allowed us to considerably extend and refine VerbNet, adding a large number of subclasses and syntactic frames to the lexicon, going far beyond basic Levin Classes.

Chapter 6

Mapping to other resources

There is currently much interest in training supervised systems to perform shallow semantic annotation tasks such as word sense tagging and semantic role labeling. These systems are typically trained on annotated corpora such as the Penn Treebank (Marcus, 1994), and perform best when they are tested on data from the same genre. A more long-term goal is to develop systems that will perform equally well on diverse genres, and that will also be able to perform additional, more complex, semantic annotation tasks. These systems will need to rely on robust resources not tied to specific corpora. With this end in mind, we have added mappings from our verbs to WordNet senses (Miller, 1985; Fellbaum, 1998) and between our verbs and FrameNet II frames (Baker et al., 1998), and mappings between our syntactic frames and Xtag (XTAG Research Group, 2001) trees. All these resources can be considered complementary and the mappings between them allow users and researchers to combine VerbNet information with the information available in these other well-known broad-coverage resources.

6.1 Mappings to WordNet

Each verb in VerbNet is mapped to its corresponding synset(s) in WordNet, if available. In addition to allowing VerbNet to be used as a cross-reference resource with implicit links between e.g., Propbank and WordNet, these mappings uncovered a set of synsets which could be used as groupings generating more coarse-grained sense distinctions. These more coarse-grained sense distinctions can significantly help improve word sense disambiguation for both humans and machines as attested by the experiments described in Palmer, Dang, and Fellbaum (Palmer et al., 2005).

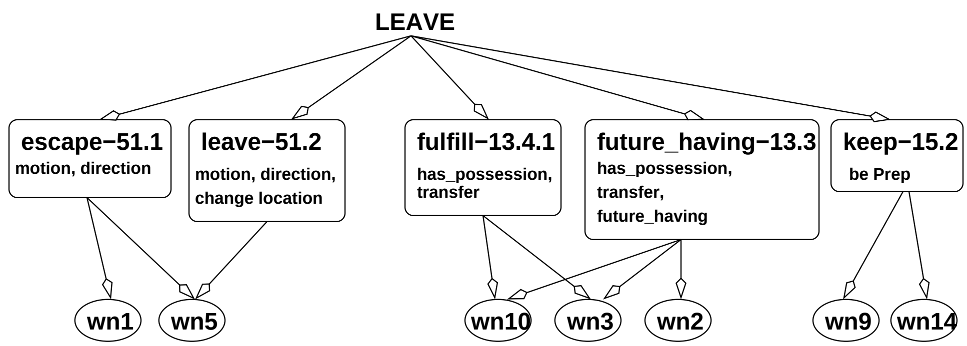

Figure 6.1: Mappings between VerbNet *leave* and WordNet synsets

Figure 6.1 shows mappings for the verb *leave* present in 5 distinct VerbNet classes (*Escape-51.1*, *Leave-51.2*, *Fulfill-13.4.1*, *Future_having-13.3*, and *Keep-15.2*). *Leave* has different semantics illustrated by the predicates describing these classes. It is a **motion** verb in the first two classes (e.g., *John left*, *John left the country*), a **transfer of possession** verb in the third and fourth classes (e.g., *John left a book for Mary*) and describes a **presence** in the last class (e.g., *John left the papers on the table*). When mapping to WordNet these same distinctions are retained. Moreover classes which exhibit similar semantics have

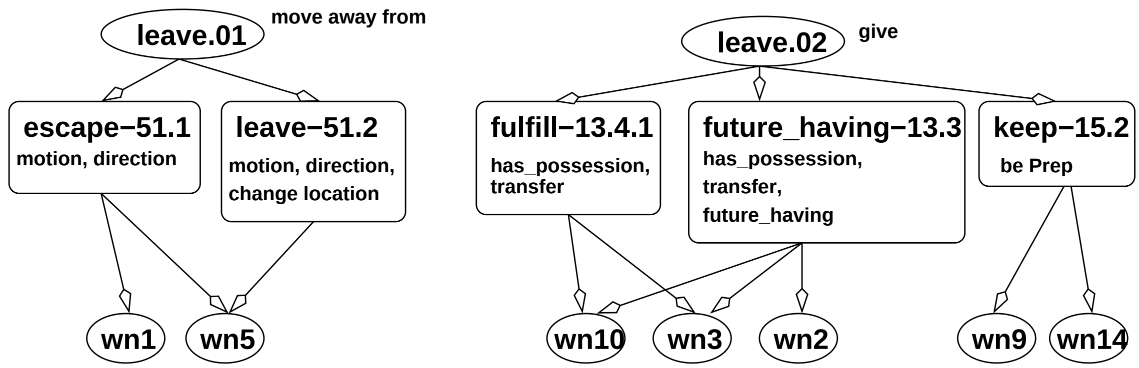

Figure 6.2: Implicit mappings between PropBank and WordNet verb senses

mappings to the same synsets. These findings are ideal for aggregating synsets with similar meaning components into more course-grained groupings which may be useful for word sense disambiguation tasks such as the ones needed for machine translation and question and answering systems, for example.

Figure 6.2 shows how mappings between PropBank and WordNet can be derived from mappings from VerbNet to PropBank, and from VerbNet to WordNet. PropBank makes very coarse-grained sense distinctions, based mostly on the verb-arguments, in contrast WordNet makes fairly fine-grained sense distinctions based on detailed semantic differences, VerbNet classes tend to exhibit sense distinctions at a level of granularity that falls in between these other two resources.

WordNet is probably the most widely used on-line lexicon currently available. The mapping between VerbNet verbs and WordNet senses provides existing applications that make use of WordNet with explicit semantics and clear syntactic descriptions. It also allows us to investigate new ways to expand VerbNet's coverage as discussed in Section 8.2.

6.2 Mappings to FrameNet II

Pursuing the goal of creating a broad-coverage lexical resource compatible to others in the field, we also mapped VerbNet to FrameNet. We used the data from FrameNet II available on-line as of January 2004 (data last updated 2004.01.06). Although FrameNet contains information about several parts of speech, only verbs were considered for these mappings. Moreover, verbs with particles, prepositions, and other multiword expressions present in FrameNet were discarded to keep compatibility with VerbNet.

The mapping between these two independently built resources consisted of two tasks:

1. mapping VerbNet verb senses to FrameNet;
2. mapping VerbNet thematic roles to the equivalent (if present) FrameNet frame elements for the corresponding classes/frames uncovered during task 1.

A total of 1952 verb senses from VerbNet representing 172 classes were successfully mapped to FrameNet frames. This resulted in 458 VerbNet-FrameNet mappings and 233 unique FrameNet frames being assigned to VerbNet verbs. VerbNet and FrameNet are complementary resources and these mappings give us different levels of representation for the events these verbs represent.

6.2.1 Mappings between VerbNet verb senses and FrameNet frames

In the first task, VerbNet verb senses were mapped to corresponding FrameNet senses, if available. Each verb member of a VerbNet class was assigned to a (set of) FrameNet frames according to semantic meaning and to the roles this verb instance takes. The syntactic descriptions present in the FrameNet frame assigned may not be exactly the same as the ones described in the VerbNet class of which the verb is a member.[1]

[1] This is expected given the inherent difference of the two resources, FrameNet does not give an extensive account of syntactic possibilities and is not intended as a database of alternations as VerbNet is.

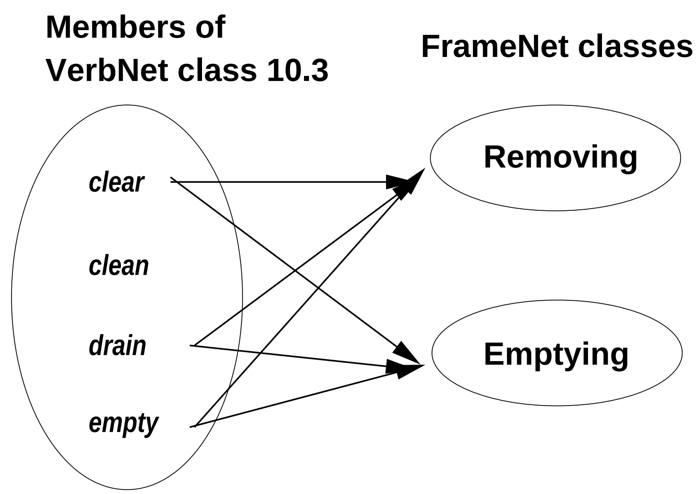

Figure 6.3: Example of mappings between VerbNet class *Clear-10.3* and FrameNet frames REMAINING and EMPTYING

Figure 6.3 shows how members of VerbNet class *Clear-10.3* map to FrameNet frames. It can be noticed that some verbs, such as *clean* do not map to any FrameNet frame, but that other verbs such as *clear*, *drain*, and *empty* map to several FrameNet frames. For example, the verb *clear* with the meaning expressed by class *Clear-10.3* has senses of emptying (e.g., *The strong winds cleared the sky*), and removing (e.g., *Doug cleared the dishes from the table*) in the FrameNet frames EMPTYING and REMOVING, as do verbs *empty* and *drain.*

These mappings are not one-to-one since VerbNet and FrameNet were built with distinctly different design philosophies. VerbNet verb classes are constructed by grouping verbs based mostly on the basis of their participation in diathesis alternations (though semantic and morphological properties also play a role in some classes) following Levin (1993). In contrast, Baker and Ruppenhofer (2002) point out that FrameNet is designed to group verbs, nouns, and adjectives based entirely on frame semantics. While Levin's investigation focused on grouping verbs with identical behavior with respect to valence alternations, a single FrameNet frame may contain sets of verbs with closely related senses but somewhat or even vastly different subcategorization properties and sets of verbs with similar syntactic behavior may be split up among multiple frames.

Baker and Ruppenhofer (2002) noted 4 general types of discrepancies between the Levin classification and FrameNet:

1. Levin class roughly equivalent to FrameNet frame

   (for example, Levin's *Cooking* class and FrameNet's APPLY_HEAT frame)

2. Levin class narrower than FrameNet frame

   (for example, Levin's classes *Pocket* and *Put* are combined by FrameNet's PLACING frame)

3. Levin class broader than FrameNet frame

   (for example, Levin's class *Other_cos*, which is exceptionally broad)

4. Overlapping Groupings

   (for example, the various Levin classes dealing with communicative acts)

These generalizations were reflected in the mappings, as we frequently observed cases where a single VerbNet class was spread out among multiple FrameNet frames or where FrameNet frames subsumed multiple VerbNet classes.

Verbs *load* and *fill* are considered canonical examples of an alternator versus a non-alternator and serve to illustrate one of these discrepancies. *Load* participates in the Locative Alternation, which describes how the *Theme* and *Location* arguments are expressed. In each variant of this alternation one of the arguments is recorded as a prepositional phrase while the other is not. On the other hand, *fill*, while similar in meaning, does not participate in this alternation. Examples (20) and (21) show the verbs participating in the Locative Alternation.

(20) Alternating

Jessica loaded the boxes on the wagon. (locative variant)

Jessica loaded the wagon with boxes. (with variant)

(21) Non-Alternating

**Leslie filled boxes in/on the room. (locative variant)*

Leslie filled the room with boxes. (with variant)

As *fill* and *load* exhibit different behavior with respect to the Locative Alternation and alternators and non-alternators cannot generally appear in the same class, *fill* appears as a member of class *Fill-9.8* while *load* appears as a member of *Spray-9.7* class in VerbNet whereas in FrameNet both appear as members of the FILLING frame.

VerbNet class *Poison-42.2* is another example of such design differences. VerbNet makes no distinction between harming, killing, and executing, however FrameNet does. This is a subtle semantic difference and it is not reflected syntactically. Verbs *asphyxiate, crucify, drown, garrote, smother,* and *strangle* map to FrameNet frame KILLING, verbs *electrocute, knife,* and *stab* map to FrameNet frame CAUSE_HARM, and verb *hang* maps to EXECUTION.

Additionally, two frames in FrameNet may be distinguished by whether the verbs are expressing a causative or an inchoative variant. By contrast in VerbNet this ability is considered an alternation and therefore this distinction is expressed by having both syntactic frames within the same class.

Naturally, there were also cases where a VerbNet class is mapped to a single FrameNet frame. VerbNet class *Put-9.1* (in Table 21) is an example of that. Three verbs from that class *(immerse, lodge, mount)* map to the FrameNet PLACING. The verb *arrange* is present in FrameNet with a different meaning (ARRANGING) and therefore does not have a mapping here, and the verb *sling* does not appear in FrameNet. An analysis of these discrepancies can be used to expand FrameNet's coverage.

6.2.2 Mapping of Thematic Roles to Frame Elements

The second task consisted of mapping VerbNet thematic roles to FrameNet frame elements for the pairs of classes found in the first task. Again, the mapping was not always one-to-one

VN class	VN member	FN frame
9.1	arrange	(diff. sense)
9.1	immerse	Placing
9.1	lodge	Placing
9.1	mount	Placing
9.1	sling	

Table 6.1: Mapping members of Class *Put-9.1* to FrameNet frames

as FrameNet tends to record much more fine-grained distinctions than VerbNet regarding argument structure.

In some cases VerbNet conflates as one thematic role what FrameNet considers something to be two distinct frame elements. A common example occurs in frames that involve some causative aspect such as the FrameNet frame PLACING, which share members with VerbNet class *Put_spatial-9.2*. FrameNet distinguishes the frame elements *Agent*, *Cause*, *Theme*, and *Goal* for this event whereas VerbNet describes the thematic roles *Agent*, *Theme*, and *Destination*. In VerbNet *Agent* and *Cause* are subsumed under the single thematic role *Agent* while in FrameNet a distinction is made between *Agents*, which imply volitional properties, and *Cause* (for instance, FrameNet documentation for this frame indicates that some verbs, i.e. *thump*, explicitly disallow *Agent* and others, i.e. *thud*, explicitly disallow non-agentive *Causes*). In most cases our VerbNet *Agent* is a generalization of FrameNet's two frame elements. The complete correspondence of VerbNet thematic roles and FrameNet frame elements in this case is as follows:

Similar situations arise in mappings between VerbNet class *Pour-9.5* and FrameNet frame FLUIDIC_MOTION (VerbNet *Location* corresponds to FrameNet frame elements *Area* and *Goal*), VerbNet *Steal-10.5* and FrameNet frame THEFT (VerbNet *Source* corresponds to FrameNet frame elements *Source* and *Victim*), VerbNet *Spank-18.3* and FrameNet frame CAUSE_HARM (VerbNet *Patient* corresponds to FrameNet frame elements *Victim* and *Body_part*), numerous other cases show similar generalizations.

VNclass 9.2	FNclass "Placing"
FN role	VN role
Agent	Agent
Cause	Agent
Goal	Destination
Theme	Theme

Table 6.2: Relations between FrameNet frame elements and Roles for class 9.2

The opposite situation also occurs, where VerbNet expresses a single FrameNet frame element by multiple thematic roles as in mapping between VerbNet class *Future_having-13.3* and FrameNet frame GIVING. In this instance the FrameNet frame element *Donor* corresponds to both the VerbNet thematic roles *Agent* and *Cause*.

In general, FrameNet has a large set of frame specific conceptual roles whereas VerbNet has gone the opposite direction and used a limited set of more abstract thematic roles. It is no surprise, then, that this situation would lead to mismatches between the descriptions of arguments in the resources, most frequently realized by FrameNet providing multiple frame elements corresponding to a single VerbNet thematic role, but occasionally occurring in the opposite direction as well.

6.2.3 Results

A total of 4,170 VerbNet verb senses distributed in 191 first level classes were used for the mappings. Of those 46.8% (1,952 verb senses) from 172 classes (out of 191) were successfully mapped to FrameNet II frames. A large number of VerbNet members (1741) were not present in FrameNet, and a smaller but still significant number (754) were present with a different sense than the one intended in the VerbNet class they belonged to. Table 6.3 shows these results: members of 4 VerbNet classes are completely absent from FrameNet, members of 3 VerbNet classes all have different senses in FrameNet, and 12 VerbNet classes had both absent and different sense members.

	Total	Mapped	Not in FN	FN different sense	Both
Verbs	4170	1952	1741	754	
Classes	191	172	4	3	12

Table 6.3: Number of VerbNet verbs and classes in the mappings

Number of VN classes	Number of FN frames mapped
19	0
72	1
38	2
28	3
7	4
11	5
3	6
6	7
3	8
3	10
1	29

Table 6.4: Distribution of VerbNet to FrameNet mappings

On average a VerbNet class was mapped to 2.5 FrameNet frames.[2] Table 6.4 shows in detail the distribution of the 458 unique mappings between VerbNet classes and FrameNet frames. As can be seen, almost half of the classes (72) map to a single FrameNet frame, and 19 classes had no mappings to FrameNet frames (either because the members were not present in FrameNet or because they were present with a different sense).

Of those 19 classes not mapped:

- 4 classes had 0 members in FrameNet

  Mine-10.9 (2 members)

  Dressing_well-41.3.2 (4 members)

[2]The real average is closer to 2.66 but because one class, *Other_COS-45.4* which serves as a pool for "miscellaneous change of state" verbs, was assigned 29 mappings, we considered this number misleading.

Being_dressed-41.3.3 (4 members)

Weekend-56 (9 members)

The first three classes are very small, and the last class *Weekend-56* has members which are zero-related to nominals and that tend to used most frequently as nouns (e.g., *summer, holiday, weekend*).

- 3 classes where all members had different senses from the FrameNet senses

  Stalk-35.3 (4 members)

  Bulge-47.5.3 (3 members)

  Rush-53.2 (3 members)

  For example, some members of the *Stalk-35.3* class *smell, taste* appear in FrameNet with their more common PERCEPTION-EXPERIENCE meaning; and members of *Rush-53.2* (*hasten, hurry, rush*) all appear in FrameNet in the SELF-MOTION frame.

These 458 VerbNet-FrameNet mappings between classes resulted in 233 unique FrameNet frames being assigned to VerbNet. The frames that appeared most frequently are naturally a factor of the classes that had the most semantically coherent members. This was the case with verbs that mapped to frames MAKE_NOISE, SELF_MOTION, and EXPERIENCER_OBJ, for example. MAKE_NOISE mapped to a single (large) class *Sound_Emission-43.2*; SELF_MOTION mapped across several motion classes including *Meander-47.7*, *Swarm-47.5.1*, *Run-51.3.2* and *Waltz-51.5*. EXPERIENCER_OBJ mapped to the psych-verb classes *Amuse-31.1* and *Marvel-31.3*.

Over 43% of the classes (83) had at least half of their members assigned to FrameNet frames. 13 classes of those 83 were completely mapped, they were small classes where all their members were monosemous. These include some classes of *Verbs of Perception*, a few classes of *Verbs of Communication*, *Verbs of Ingesting*, and several classes of *Verbs of*

Motion, among others.

This experiment of mapping VerbNet verbs to FrameNet frames, two lexical resources with a similar notion of classes yet with disjoint coverage, open several possibilities to increase VerbNet's robustness:

- Augment VerbNet's coverage: many VerbNet classes map to a single (or a few) FrameNet frame(s) (e.g., some classes of *Putting* and *Psychological verbs*). Verbs members of these frames which are not yet present in VerbNet are potential candidates for membership in our lexicon and merit a syntactic and semantic evaluation.
- Evaluation of subclasses: classes which map to different frames can be evaluated with respect to their subclasses to verify if they make similar distinctions.
- Clean up class members: in some classes just a few members map to some FrameNet frame, whereas other members are either not present or map to different frames. Investigate whether these verbs are spurious members of these classes or not.

However, FrameNet can also profit from these mappings to add new members to their frames (e.g., from classes such as *Put-9.1* and *Pocket-9.10*) and to evaluate class membership.

6.2.4 Experiment in Semantic Role Labeling

Dan Gildea used FrameNet's semantic roles (frame elements) to automatically assign semantic labels to a corpus (Gildea and Jurafsky, 2002). Because the set of frame elements were shown to be too fine-grained for machine learning applications, his work also included a mapping between FrameNet frame elements (for 67 frames) and a set of 18 commonly used roles. This approach assigned labels correctly to 87.4% instances (the recall was 50.1%).

The University of Colorado performed a similar experiment using Gildea's automatic semantic role labeler on the task of assigning semantic labels to the combined corpus of

PropBank and FrameNet, with results around 70% (precision and recall). For this experiment, our roleset and the University of Colorado rolesets (using the 18 roles) were mapped and some adjustments were required. These included replacing *Product* and *Material* from our roleset to *Theme* and *Result* for verbs in the *Creation and Transformation* class, and adding *Actor1* and *Actor2* for symmetrical verbs on their roleset. Although beyond the scope of this thesis, we hypothesize that if this experiment is redone using the currently available mappings, the results may improve significantly.

6.3 Mappings to Xtag

This section presents the mappings between the syntactic information of VerbNet to Xtag trees.[3] A natural extension of VerbNet's syntactic frames is to incorporate the possible transformations of each frame, and the Xtag grammar (XTAG Research Group, 2001) presents a large existing grammar for English verbs that accounts for just such richness of constructions. This mapping between complementary resources allowed us to increase the syntactic coverage of our verb lexicon by capturing transformations of the basic syntactic description of the verbs present in VerbNet. In addition, having a mapping between these two resources could allow the semantic predicates present in our lexicon to be used to disambiguate Xtag verb senses.

VerbNet, while providing an explicitly constructed verb lexicon with syntax and semantics, offers limited syntactic coverage since it describes only the declarative frame for each syntactic construction or alternation. The Xtag grammar, on the other hand, is a lexical resource with well-characterized syntactic descriptions for lexical items but makes no distinctions between verb senses and currently contains no explicit semantics. An obvious way to extend VerbNet's syntactic coverage is to incorporate the coverage of Xtag, accounting for the possible transformations of each declarative frame. Presumably, transformations of

[3] An earlier version of this work has been presented as a poster in the TAG+7 Workshop (Ryant and Kipper, 2004).

VerbNet's syntactic frames are recoverable by mapping onto elementary trees of TAG tree families. Then, for any verb in VerbNet each thematic role can be mapped to an indexed node in the basic syntactic tree and the selectional restrictions on VerbNet thematic roles to features on the nodes. In addition to increasing the coverage of VerbNet, this provides us with a pre-existing parser for computing derived and derivation trees to which our semantic predicates can be added and therefore sense distinctions can be made more explicit.

6.3.1 Extending VerbNet with Xtag

Each frame in VerbNet is described by 4 components: 1) a brief text description (such as *Transitive, Resultative*), 2) an example sentence, 3) a syntactic frame, 4) a semantic description using a set of semantic predicates. Text descriptions and syntactic frames are very much interrelated, but the text description is independent of the roles assigned to the verb's arguments as can be seen in Examples (22) and (23) where both frames have *Basic Transitive* as their description

(22) *Basic Transitive*

"Jackie accompanied Rosie."

Agent V Theme

(23) *Basic Transitive*

"The clown amused the children."

Cause V Experiencer

These text descriptions consist of both primary and secondary descriptions which were made completely consistent for the whole VerbNet lexicon prior to these mappings. Examples of primary descriptions include *Transitive, Material/Produce Alternation, and NP-PP*. Many of these descriptions are based on Briscoe's subcategorization frames as described in Briscoe (2000) . These frames were added at the time of the integration of Korhonen and Briscoe's (2004) new classes into VerbNet as discussed in Section 7.1. Secondary descriptions provide additional information about the semantics and/or syntax. These might specify the types of prepositional phrases that a verb may take or the existence of restrictions on a complement (often secondary descriptions are used to distinguish between different types of sentential complements or different thematic roles in object or subject position). Examples (24) and (25) show the primary (*Material/Product Alternation Intransitive*) and secondary descriptions (*Material Subject* and *Product Subject*) for members of class *Grow-26.2*.

(24) *Material/Product Alternation Intransitive (Material Subject)*
"That acorn will grow into an oak tree."
Material V Prep(into) Product

(25) *Material/Product Alternation Intransitive (Product Subject)*
"An oak tree will grow from that acorn."
Product V Prep(from out of) Material

Generally, the VerbNet syntactic frame specified by the primary description corresponds to the surface syntactic realization of an Xtag elementary tree. Each VerbNet syntactic frame was mapped to a corresponding Xtag tree family, with the index of the tree family recorded in the VerbNet entry. In theory we should be able to annotate each unique VerbNet syntactic frame with a mapping to an Xtag elementary tree.

6.3.2 Discussion

A number of VerbNet syntactic frames do not correspond to an Xtag elementary tree. Some of VerbNet classes contain syntactic frames that specify multiple adjuncts. As a case in point consider VerbNet class *Turn-26.6.1*, with members *alter, deform, metamorphose, mutate, transform, transmute, change, convert, and turn.* Each of these can appear in the two frames presented in (26) and (27):

(26) *NP-PP-PP (Causative variant, Material-PP Product-PP)*

"The witch turned him from a prince into a frog."

Agent V Patient Prep(from) Material Prep(into) Product

(27) *PP-PP (Inchoative variant, Material-PP Product-PP)*

"He turned from a prince into a frog."

Patient V Prep(from) Material Prep(into) Product

In the Xtag grammar, the frame presented in (26) corresponds to no elementary tree of any tree family. One might disagree over what elementary tree it is derived from. For instance, (26) can be seen as a transitive sentence with PP adjuncts (and thus belonging to tree family Tnx0Vnx1), as a ditransitive taking a PP complement with another PP adjunct (tree family Tnx0Vnx1pnx2), or as a resultative with a PP anchor and an additional PP adjunct (and thus belonging to tree family TRnx0Vnx1Pnx2). Similarly, the frame presented in (27) can be seen as either an intransitive sentence with optional PP adjuncts (Tnx0V), or as a resultative with ergative verb and PP anchor (TREnx1VPnx2). In the current version of VerbNet we have 27 syntactic frames that fall into this category and are not treated in our mappings, these include Examples (28) – (32).

(28) *ADVP (here/there)*

The books lean there.

(29) *NP-P-ING (from-PP)*

The rules forbid us from smoking.

(30) *PP-NP (Goal-PP)*

Nora brought to lunch the book.

(31) *PP-QUOT (Recipient-PP)*

Ellen complained to Sara, 'The mail didn't come today.'

(32) *PP-TO-INF-OC (Recipient-PP)*

Susan whispered to Rachel to come.

6.3.3 Extending our mappings

We provided mappings between the syntactic frame descriptions of VerbNet to the Xtag tree families. In order to further explore these mappings one needs to map each thematic role to an indexed node in the basic syntactic tree. For example, *Agent V Patient* would map to a structure such as the transitive tree shown in Figure 6.4, with Agent mapped to NP_0 and Patient to NP_1. The selectional restrictions of each role could be expressed as semantic features in the nodes. The correspondence between thematic roles and indexing of syntactic arguments in TAG trees is preserved within each tree family (which contains possible transformations of a basic syntactic structure), so by specifying the mapping to the basic tree, we also get the mapping to all the transformations applicable to the tree family (e.g., passivization, wh-movement).

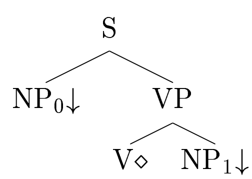

Figure 6.4: Transitive Tree

Initial trees capture the semantics of the basic senses of verbs in each class. For example, many verbs in the *Run-51.3.2* class can occur in the causative/inchoative alternation, in which the subject of the intransitive sentence has the same thematic role as the direct object in the transitive sentence. Figure 6.5 shows the initial trees for the causative and inchoative variants for this class, along with VerbNet's semantic predicates.

Semantic predicates could also be associated with each tree, as was done by Stone and Doran (1997) , Kallmeyer and Joshi (2003), and Kallmeyer and Romero (2004). The

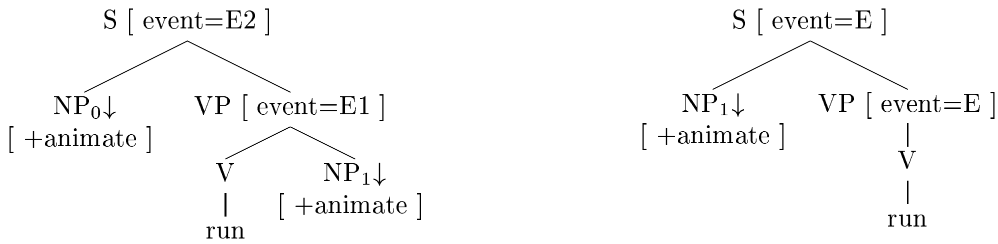

$cause(during(E2), Y, E1) \wedge motion(during(E1), X)$ $\quad$ $motion(during(E), X)$

Figure 6.5: Causative/Inchoative alternation for the *Run* verbs

details of mapping this approach to the current proposed compositional semantics for TAGs is beyond the scope of this thesis, however.

6.3.4 Coverage

In the latest version of VerbNet v2.2 there are 357 unique frames (as distinguished by primary and secondary description), and 131 distinct primary descriptions. Of these 131 frames used for the mappings, all but 27 correspond exactly to some Xtag tree family. For these 104 VerbNet syntactic frames that map exactly to an Xtag tree family, only 19 of the 57 Xtag tree families were used. A detailed inspection on the 41 Xtag tree families with no corresponding VerbNet frame, revealed that 22 of them deal with small clauses, 8 with idiomatic expressions, and 6 with other various classes.

6.4 Conclusion

This chapter presented a detailed account of the mappings between VerbNet and other large-scale resources. This effort greatly increases the robustness of VerbNet while permitting each resource to be used as an extension of the others. Because VerbNet is a general lexicon, one of its payoffs is that it allows the extension of coverage of resources tied to specific corpora (e.g., PropBank, FrameNet).

The association of our verbs to WordNet synsets provides our lexicon with rich semantic information while uncovering different relations between verbs which are not accounted for in VerbNet. It also provides WordNet senses with explicit semantics and syntax. The mappings to FrameNet add a new way of looking at the classes; in addition to syntactic patterns verbs can be grouped by their semantic frame elements. It also gives us a more fine-grained set of semantic roles for VerbNet's verbs. The correspondence to Xtag allows us to indirectly provide a much larger syntactic coverage by incorporating the transformations of the basic frames into our syntactic frames. In addition to increasing the coverage of VerbNet, the mappings to Xtag supply us with a pre-existing parser for computing derived and derivation trees to which our semantic predicates can be associated thus helping the task of producing semantic representation of sentences.

Chapter 7

Extending VerbNet's coverage

Levin-style lexical classes are useful for their ability to capture generalizations about a range of linguistic properties. Classes are defined in terms of shared meaning components and similar (morpho-)syntactic behavior of words (Pinker, 1989; Jackendoff, 1990a; Levin, 1993) and, thus, generally incorporate a wider range of linguistic properties than classes defined solely on semantic grounds (e.g., (Miller, 1990; Baker et al., 1998)).

Two additional available resources, the LCS database (Dorr, 2001) and Korhonen and Briscoe's (2004) new classes, are worth investigating as a means of extending VerbNet's coverage. These two resources are also based on Levin's classification and can potentially provide additional members and classes to VerbNet.

Korhonen and Briscoe (2004) proposed a substantial extension to Levin's original classification which incorporates 57 novel classes for verb types not covered in detail by Levin. In VerbNet, Levin's original taxonomy has gained considerably in depth, but not to the same extent in breadth. Verbs taking adjectival (ADJP) and adverbial phrases (ADVP), particles, predicative, control and sentential complements were still largely excluded, except where they showed interesting behavior with respect to NP and PP complementation. The new classes proposed by Korhonen and Briscoe include these verb types, many of which have high frequency.

The LCS database is a rich lexical resource which organizes verbs into classes according to Levin's classification. Verb semantics are described by lexical conceptual structures, and their argument structure is described by a small set of thematic roles in a manner similar to VerbNet.

The first section of this chapter describes the integration of the novel classes of Korhonen and Briscoe (2004) into the extensive and detailed VerbNet verb lexicon. This integration required resolving a number of conflicts between the two classification systems and required assigning the novel classes explicit semantic descriptions. The second section describes the new verbs that have been automatically assigned to LCS' Levin classes and their integration into VerbNet.

We believe that these extensions resulted in a much improved lexical resource which now provides the most comprehensive and versatile classification of Levin-style lexical classes for English.

7.1 Korhonen and Briscoe's new classes

The resource of Korhonen and Briscoe (2004) includes a substantial extension to Levin's classification proposing 57 novel classes for verbs and verb types not covered (comprehensively) by Levin as well as 106 new diathesis alternations created as a side product of constructing the new classes.

Korhonen and Briscoe have demonstrated the utility of these novel classes by using them to support automatic subcategorization acquisition and have shown that the resulting extended Levin classification has a fairly extensive coverage of the English verb lexicon. However, they have not refined and organized their novel classes into taxonomies which would incorporate different degrees of granularity or integrated them into any existing taxonomy of Levin-style classes. They have also not provided any specific semantic descriptions for the classes. This additional information would greatly enhance the utility of these classes, and contribute to a single coherent lexical-semantic resource for the research community.

The classes were created according to the following steps: [1]:

1. Verbs not extensively addressed by Levin were investigated against the subcategorization classification described in Briscoe (2000). Possible alternations between pairs of subcategorization frames (SCFs) in Briscoe's classification were considered thus uncovering a large set of new diathesis alternations not in Levin. Briscoe's classification incorporates 163 different subcategorization frames (SCFs) and is an extension of ANLT (Boguraev et al., 1987) and COMLEX Syntax dictionaries(Grishman et al., 1994).[2] The subcategorization frames used describe mappings from surface arguments to predicate-argument structures but leave specific particles and prepositions underspecified.

   Korhonen and Briscoe (2004) suggested 106 new alternations based on that classification using criteria similar to Levin's, i.e., the alternating subcategorization frames either preserve or systematically modify meaning. Table 7.1 shows examples of these alternations.

   Levin incorporates around 80 alternations in her classification, so these 106 alternations are not a surprising number considering that Korhonen and Briscoe cover a wider range of complement types than Levin does. Korhonen points out that although Levin classification may cover less alternation types than the ones suggested in Korhonen and Briscoe's resource, the alternations proposed by Levin do not necessarily cover less data since they involve the most frequent subcat frames in language:

[1] See Korhonen and Briscoe (2004) for the details of this approach and http://www.cl.cam.ac.uk/users/alk23/classes/ for the latest version of the classification.

[2] Briscoe's classification schema has also been mapped to the Xtag grammar.

Category	**Example Alternations**
Equi	*I advised Mary to go* ↔ *I advised Mary* *He helped her bake the cake* ↔ *He helped bake the cake*
Raising	*Julie strikes me as foolish* ↔ *Julie strikes me as a fool* *He appeared to her to be ill* ↔ *It appeared to her that he was ill*
Category switches	*He failed in attempting to climb* ↔ *He failed in the climb* *I promised Mary to go* ↔ *I promised Mary that I will go*
PP deletion	*Phil explained to him how to do it* ↔ *Phil explained how to do it* *He contracted with him for the man to go* ↔ *He contracted for the man to go*
P/C deletion	*I prefer for her to do it* ↔ *I prefer her to do it* *They asked about what to do* ↔ *They asked what to do*

Table 7.1: Example of alternating frames from Korhonen and Briscoe (2004)

intransitive and simple NPs and PPs.

2. A large group of candidate lexical-semantic classes from existing resources were investigated ((Rudanko, 1996; Rudanko, 2000), (Sager, 1981), (Levin, 1993) and (Dorr, 2001)) with 102 classes selected.

3. Each candidate class was evaluated by examining sets of subcategorization frames taken by its member verbs in syntax dictionaries (e.g. COMLEX) and whether these frames could be related in terms of diathesis alternations (106 novel ones or Levin's original ones). Where one or several alternations were found which preserve the original meaning, a new verb class was created.

The process of identifying the relevant alternations for establishing new classes helped to determine new subcategorization frames and additional alternations. For those candidate classes which had an insufficient number of member verbs, new members were searched for in WordNet.

The subcategorization frames and alternations discovered during the identification process were used to create the syntactic-semantic description of each novel class. For example, a new class was created for verbs such as *order* and *require*, which share the approximate

Class	Example Verbs
URGE	*ask, persuade*
FORCE	*manipulate, pressure*
WISH	*hope, expect*
ALLOW	*allow, permit*
FORBID	*prohibit, ban*
HELP	*aid, assist*
DEDICATE	*devote, commit*
LECTURE	*comment, remark*

Table 7.2: Examples of K&B's Verb Classes

meaning of "direct somebody to do something". This class was assigned the description shown below.[3]

3. ORDER VERBS

SCF 57:	NP-TOBE	*John ordered him to be nice*
SCF 104:	S	*John ordered that he should be nice*
SCF 106:	S-SUBJUNCT	*John ordered that he be nice*

Alternating SCFs: 57 $\leftrightarrow$ 104, 104 $\leftrightarrow$ 106

Korhonen and Briscoe's investigation resulted in 57 verb classes, including 5 entirely novel classes which were not taken from any of the resources initially used as described in Step 2. Table 7.2 shows a sample of these classes along with example verbs.

The evaluation of the novel classes showed that they can be used to support an NLP task and that the extended classification has a fairly good coverage of the English verb lexicon. However, as Korhonen and Briscoe pointed out, the classes still need further

[3] The numbers refer to the original Briscoe (2000) subcategorization frames classification.

refinement and integration into an existing taxonomy to yield a maximally useful resource for the research community.

7.1.1 Incorporating the proposed classes into VerbNet

Although the new classes are similar in style to the Levin classes already included in VerbNet, the integration of the two resources proved challenging. VerbNet provides very specific syntactic-semantic information about members of verb classes. This information is used to organize classes into sets of taxonomies which incorporate different degrees of granularity. This is an important quality given that the desired level of granularity varies from one natural language application to another. However, the new classes of Korhonen and Briscoe (K&B, herein) lack explicit semantic descriptions, have syntactic descriptions not directly compatible with VerbNet's descriptions, and are not consistent in terms of granularity (some of the classes are broad while others are fine-grained).

The proposed novel classes fall into three possible categories regarding Levin's original classification: 1) classes that could be subclasses of existing Levin classes; 2) classes that require a reorganization of the original Levin classes[4]; 3) entirely new classes.

The different sets of subcategorization frames used by the two systems required creating new semantic roles, re-evaluating the existing syntactic descriptions, and adding new restrictions to VerbNet. The set of subcategorization frames used in K&B is particularly broad in coverage and relies, in many cases, on more fine-grained treatment of sentential complements than present in VerbNet. Therefore, VerbNet's syntactic descriptions had to be enriched with a more detailed treatment of sentential complements. On the other hand, prepositional subcategorization frames in K&B do not provide VerbNet with explicit lists of allowed prepositions as required. In addition, unlike VerbNet frames, they do not include information about semantic (thematic) roles or restrictions on the arguments. Also, no

[4] Levin focused mainly on NP and PP complements, but many verbs classify more naturally in terms of sentential complements

syntactic description of the surface realization of the frames is included in K&B, although such information would be useful to show differences between sentential complements, for example.

Challenges imposed by the integration

Thematic Roles: In integrating the proposed new classes, it was found that none of the 21 existing VerbNet thematic roles seemed to appropriately convey the semantics of the arguments for some classes.

As an example, the members of the new URGE class describe events in which one entity exerts psychological pressure on another to perform some action (*John urged Maria to go home*). While the urger (*John*) is assigned the role *Agent* as the volitional agent of the action and the urged entity (*Maria*) is assigned *Patient* as the affected participant, it is unclear what thematic role best suits the urged action (of going home). A new *Proposition* role was included which seemed to more appropriately describe the semantics of the urging event. Similar situations arose in the attempted integration of 8 other classes. In the end, two new thematic roles were added to the VerbNet lexicon, *Content* and *Proposition*.

Syntactic Descriptions: Only 44 of VerbNet's syntactic frames had a counterpart in Briscoe's classification. This discrepancy is the by-product of differences in the design philosophy of the two resources. Briscoe abstracts over prepositions and particles whereas VerbNet differentiates between otherwise identical frames based on the precise types of prepositions that a given class of verbs subcategorizes for. Additionally, VerbNet may distinguish two syntactic frames depending on thematic roles (for instance, there are two variants of the Material/Product Alternation Transitive frame differing on whether the object is the *Material* or *Product*) as shown in Examples (24) and (25) in Subsection 6.3.1.

But regarding sentential complements the opposite occurs, with VerbNet conflating subcategorization frames that Briscoe's classification considers distinct. In integrating the proposed classes into VerbNet it was necessary to greatly enrich the set of possible syntactic restrictions VerbNet allows on clauses. The original hierarchy contained only the valences $\pm$ *sentential*, $\pm$ *infinitival*', and $\pm$ *wh_inf*.

The new set of possible syntactic restrictions consists of 55 such features accounting for object control, subject control, and different types of complements.

Examples (33), (34), (35), and (36) show the VerbNet realizations and the set of constraints for the proposed FORCE class which includes two frames with object control complements (OC).

(33) Basic Transitive

"I forced him."

Agent V Patient

(34) NP-P-ING-OC (into-PP)

"I forced him Prep(into) coming."

Agent V Patient into Proposition[+oc_ing]

(35) NP-PP (into-PP)

I forced John into the chairmanship."

Agent V Patient into Proposition[-sentential]

(36) NP-TO-INF-OC

"I forced him to come."

Agent V Patient Proposition[+oc_to_inf]

Semantic descriptions: Integrating the proposed classes also required enriching VerbNet's set of semantic predicates. Whenever possible, existing VerbNet predicates were reused. However, due to the nature of the incoming classes, many of them representing concepts entirely novel to VerbNet, it was necessary to introduce new predicates. In total 30 new predicates were required to adequately provide descriptions of the semantics of these integrated classes.

Discussion of the classes included

Each of K&B's classes was thoroughly investigated to determine the feasibility of it being added to VerbNet. Of the 57 classes proposed, two were rejected as being either insufficiently semantically homogeneous [5] or too small to be added to our lexicon, with the remaining 55 selected for incorporation. A total of 42 classes were added immediately as novel classes or subclasses. But 13 proposed classes overlapped significantly in some way with existing VerbNet classes (either too close semantically or syntactically) and required restructuring of the verb lexicon. The extent of our integration can be seen in Figure 7.1.

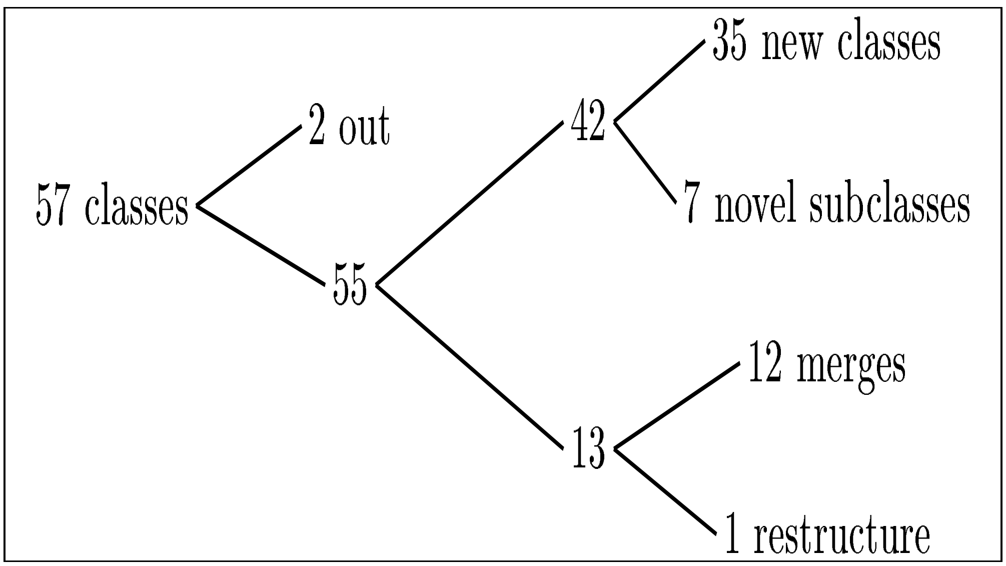

Figure 7.1: Integration of K&B's classes to VerbNet

I. No restructuring needed

As mentioned above, 42 could be added to the lexicon as novel classes or subclasses without any restructuring. Some of these classes overlapped to an extent with existing VerbNet classes semantically but syntactic behavior of the members was sufficiently distinctive to allow them to be added as new classes without restructuring of VerbNet. A

[5] Suggested class ENFORCE, for example, purports to contains verbs of enforcing and coercing. However, the class seems to contain 2 distinct subsets: verbs of enforcing (e.g. enforce, exert) and verbs of gambling (e.g., *risk, gamble*).

more specific breakdown shows that 35 novel classes were actually added while 7 others were added as novel subclasses (an additional novel subclass, *Continue-55.3*, was discovered in the process of subdividing *Begin-55.1*).

IIa. Novel Classes

The 35 classes all share the quality of not overlapping to any appreciable extent with a pre-existing VerbNet class from the standpoint of semantics. For instance, K&B's proposed classes of FORCE, TRY, FORBID, and SUCCEED express entirely new concepts as compared to VerbNet.

IIb. Novel Sub-Classes

Some of the proposed classes, such as CONVERT, SHIFT, INQUIRE, and CONFESS were considered sufficiently similar in meaning to current classes and were added as new subclasses to existing VerbNet classes. For example, both the proposed classes CONVERT and SHIFT are similar syntactically to the VerbNet class *Turn-26.6*. However, whereas the members of *Turn-26.6* exclusively involve total physical transformations, the members of the proposed class CONVERT invariably exclude physical transformation, instead having a meaning that involves non-physical changes such as changes in the viewpoint of the Theme (*I converted the man to Judaism.*) Similarly, the verbs of the proposed class SHIFT might be characterized as only taking the intransitive frames from CONVERT. Consequently, as both classes SHIFT and CONVERT are semantically similar, yet still distinct, from the existing VerbNet class *Turn-26.6*, they were added as subclasses to 26.6, yielding the new classification *Turn-26.6.1*, *Convert-26.6.2*, and *Shift-26.6.3*.

Both proposed classes INQUIRE and CONFESS relate to communication between entities as do other VerbNet classes of *Verbs of Communication*. INQUIRE members all relate to questioning, with an *inquirer* asking of an (optionally expressed) *inquiree* some (optionally expressed) *question*. This new class seems especially close to existing VerbNet class *Transf_msg-37.1*, while still being distinct, so this resulted in creating two new subclasses *Transf_msg-37.1.1* and *Inquire-37.1.2*. Similarly, proposed class CONFESS, while similar to other classes of communication in that it involves communication and transfer of ideas, also differs from them substantially both semantically and syntactically. The closest other subclasses in meaning would be either the *Tell-37.2* subclass or the *Say-37.7* subclass. However, for the members of the *Tell-37.2* class, the recipient of the message cannot be expressed as the object of a to-PP phrase when the message is a sentential complement while this is perfectly felicitous for the members of CONFESS:

(37) Tell

∗Ellen told to Helen how it happened.

(38) Confess

Ellen confessed to Helen how it happened.

While most members of CONFESS can occur in subcategorization frames where the verb is followed by a gerund whose logical subject is that of the matrix verb, the members of VerbNet class *Say-37.7* universally sound bad in the same frame.

(39) Say

∗Ellen said stealing money.

(40) Confess

Ellen confessed stealing money.

II. Classes where restructuring was necessary

In 13 of the the proposed classes, the overlap in sense was substantial enough to require some degree of restructuring of VerbNet. Classes such as WANT, PAY, and SEE obviously overlapped with existing VerbNet classes *Want-32.1*, *Give-13.1*, and *See-30.1* in terms of meaning. Nor could the proposed classes be distinguished from the existing classes by recourse to syntactic behavior. Adding such classes required restructuring of VerbNet to produce classes whose verb membership was the union of the overlapping proposed and existing classes and whose subcategorization frames, similarly, were the union of those for each of the overlapping classes.

Broadly, the process of integrating these classes can be divided into two categories: 1) merging proposed classes with the related VerbNet class; 2) adding the proposed class as a novel class but making modifications to existing VerbNet classes.

IIa. Cases involving merger of a proposed class and an existing class

In considering these classes for addition to VerbNet, it was observed that semantically their members patterned after a pre-existing class almost exactly.[6] In the cases where the frames from the new classes were a superset of the frames recorded in VerbNet, then the existing VerbNet class was restructured by adding the new members and by enriching its syntactic description with the novel frames.

For example, both K&B's proposed WANT class and the VerbNet class *Want-32.1* relate to the act of an *experiencer* desiring *something*. VerbNet class *Want-32.1* differs from the proposed WANT class in its membership and in that the class considers only alternations in NP and PP complements whereas the proposed class WANT also considered alternations in sentential complements, particularly control cases. Example frames from the original and the extended classes *Want-32.1* are shown in (41) and (42).

(41) **Original** *Want-32.1* **class:**

Dorothy needs new shoes.

Dorothy needs her for her skills.

(42) **Extended** *Want-32.1* **class:**

I need for her to be happy.

I need exercising.

[6] This is the case for Korhonen and Briscoe's classes WANT, PAY, SEE, LOVE, DISCUSS, SHOW, EXPLAIN, SUGGEST, OCCUR, AVOID, BEGIN, and COMPLETE.

She wanted the meat red.

I needed him here.

I need him cooking.

I need the children found.

I needed his cooking.

I needed to come.

I needed him to go.

I need him to be nice.

IIb. Added as new class and requiring restructuring of existing classes

K&B's work is of particular importance when considered in the context of classes of *Verbs With Predicative Complements*, whose members are frequent in language. These verbs classify more naturally in terms of sentential complements rather than NP or PP complements. The proposed class CONSIDER overlaps with four of VerbNet's classes (*Appoint-29.1*, *Characterize-29.5*, *Declare-29.4*, and *Conjecture-29.6*), none of which were originally very semantically homogeneous. The process of adding CONSIDER as another class of verbs with predicative complements gave us the opportunity to revise these four problematic classes making them more semantically homogeneous by using the more detailed coverage of complements presented in K&B. Examples and detailed frames for the added verb class *Consider-29.9* are shown in Table 7.3.

At the end 25 verbs were removed from these existing classes. However, the classes have become more semantically and syntactically homogeneous as a result.

7.1.2 Summary of VerbNet Extensions from Korhonen and Briscoe's classes

A limitation in Levin's classification is that it deals mainly with noun and prepositional

Example sentence	Syntactic description
I considered how he could become professor.	Agent V Theme[+how_extract]
I considered how to be a professor.	Agent V Theme[+wh_inf]
They considered him as being smart.	Agent V Theme[+np_p_ing]
They considered him to be the professor.	Agent V Theme[+to_be]
I considered the matter closed.	Agent V Theme[+np_ppart]
They considered that he was the professor.	Agent V Theme[+that_comp]
They considered him as smart.	Agent V Theme as ADJ
He considered whether he should come.	Agent V Theme[+wh_comp]
He considered what he should do.	Agent V Theme[+what_extract]
He considered what to do.	Agent V Theme[+what_inf]
He considered smoking.	Agent V Theme[+be_sc_ing]
They considered him smart.	Agent V Theme ADJ
They considered him the professor.	Agent V Theme Predicate[-sent]
They considered him for professor.	Agent V Theme for Predicate[-sent]
He considered whether to clean the house.	Agent V Topic[+whether_inf]
They considered the children found.	Agent V Theme[+np_ppart]

Table 7.3: Verb Class *Consider-29.9* added to VerbNet from K&B class `CONSIDER`

phrase complements with the result that verbs taking ADJP, ADVP, particle, predicative, control and sentential complements are not addressed in depth. Thus, K&B's work, which focused particularly on these very frequent verbs, is of particular interest.

While most of the 57 new Levin classes proposed by K&B could be added without changes to the current VerbNet, 13 had some degree of semantic overlap with current classes. These cases were considered on a class-by-class basis, with the typical result being restructuring of the existing class to account for any novel verbs or alternations appearing in the proposed classes. In certain cases the restructuring was drastic, such as in VerbNet classes 29.x, whose verbs are most properly classified by their patterns of sentential complements. A summary of how this integration affected VerbNet and the result of the extended lexicon is shown in Table 7.4.

	VerbNet 1.5	VerbNet Extended by K&B
Verbs	4227	4526
First-level classes	191	237
Thematic roles	21	23
Semantic Predicates	64	94
Select. Restr. (semantic)	36	36
Syntactic Restr. (on sent. compl.)	3	55
Syntactic Frames	314 (55 primary)	357 (131 primary)

Table 7.4: Summary of VerbNet's extensions from Korhonen and Briscoe's novel classes

7.2 Additions from the LCS database

The LCS Database is a lexical resource developed at the University of Maryland (Dorr, 2001). This database has verbs organized in verb classes according to a modified version of Levin's classification and contains syntactic descriptions described by a set of thematic roles in a manner similar to VerbNet. The verbs' semantics are described through lexical conceptual structures as explained in Subsection 2.4.3.

An example entry of this database can be seen in Figure 7.2 for the verb *abandon* from Levin class 51.2 (*Leave* verbs). The corresponding synsets in WordNet 1.5 and 1.6 are specified, and there are mappings to PropBank's arguments. Two thematic roles are used to describe the predicate-argument structure of this class *theme* and *source*, and these roles are also used for the verb's LCS semantic description.

Dorr and Jones (1995) devised an experiment which classified a large number of unknown verbs suggested by English glosses provided in bilingual dictionaries. Each existing Levin class was described by a *syntactic signature* specified by a set of positive and negative examples, and the prepositions available for the class, with a mapping created between these syntactic signatures and LDOCE (Procter et al., 1978) canonical codes. A new verb *V* was classified by looking up its synonym sets in WordNet, and assigning to *V* the Levin

```
;; Grid: 51.2#1#_th,src#abandon#abandon#abandon#abandon+ingly
#(1.5,01269572,01188040,01269413,00345378)(1.6,01524319,01421290,
01524047,00415625)##
#AD

(
 :DEF_WORD "abandon"
 :CLASS "51.2"
 :WN_SENSE (("1.5" 01269572 01188040 01269413 00345378)
            ("1.6" 01524319 01421290 01524047 00415625))
 :PROPBANK ("arg1 arg2")
 :THETA_ROLES ((1 "_th,src"))
 :LCS (go loc (* thing 2)
          (away_from loc (thing 2) (at loc (thing 2) (* thing 4)))
          (abandon+ingly 26))
 :VAR_SPEC ((4 :optional) (2 (animate +)))
)
```

Figure 7.2: An example entry from the LCS database

class with the closest match between its synonyms and the LDOCE canonical codes. If no synonyms in Levin are present a new class is hypothesized. The Dorr and Jones experiment tested 84 LDOCE verbs and 82% of the instances had the correct Levin class assigned. Dorr and Jones (2000) expanded the previous experiment to classify all of the verbs from LDOCE using a syntactic-filter paired with a semantic filter which uses synonym relations from WordNet. These techniques were used to further extend the LCS database.

This work is of special interest to VerbNet since the two resources are based on the same Levin classification. Since the LCS database is a large broad-coverage lexicon very much like VerbNet, we can benefit from integrating verbs already classified in the LCS Levin classes into our lexicon. In addition, these two resources can be seen as complements of each other, they provide different semantics for similarly grouped classes of verbs.

The LCS database version we used has 10,003 entries (with 4,324 lemmas) distributed in 231 classes. The 40 new classes are for the most part small and have a total of 250

Number of new classes	Number of members
9	1
14	2
4	3
2	4
2	5
2	6
1	8
1	10
1	11
2	14
1	37
1	42

Table 7.5: Distribution of members in LCS new classes

members distributed as shown in Table 7.5. 9 of the new classes contain only one member, 14 classes contain 2 members, 4 classes contain 3 members, 2 classes contain 4 or 5 or 6 members, 1 class contains 8 members, and only 6 classes contain 10 or more members.

Although the new classes seem promising, a more thorough investigation would be necessary to identify which classes are relevant to VerbNet. Of particular interest to us are the members of existing Levin classes that were automatically added by Dorr and Jones' experiments. The members of existing classes are good candidates for inclusion in Verbnet and could potentially be incorporated without major restructuring of our classes. A total of 1,266 verbs are in existing Levin classes and not present in VerbNet after the integration with Korhonen and Briscoe's (2004) resource. These verbs satisfied our criteria and and therefore were the focus of our attention.

Of these 1,266 verbs, 429 (426 lemmas) were initially integrated into our lexicon. The verbs that extend LCS had been automatically acquired and had not been hand checked against the existing classes. Because of this noise, many of them did not match all the criteria for our classes and would require further restructuring of VerbNet.

Chapter 8

Automatic techniques for extending coverage

Lexical resources for practical applications are limited by their coverage and by the difficulty of increasing this coverage in a systematic way. Linguists have believed for decades that a statistical analysis of naturally occurring text is the key to unlocking many of the mysteries of how words are classified and interpreted (Harris, 1962; Nevin, 2003). In spite of vast quantities of digital forms of text, such an automated analysis is still tantalizingly out of reach.

Although a number of supervised approaches to lexical acquisition have extracted a variety of information about verbs from corpora, such as semantic role labeling (Gildea and Jurafsky, 2002), verb classification (Dorr and Jones, 1996; Lapata and Brew, 1999; Merlo and Stevenson, 2001), and extraction of selectional preferences (Resnik, 1996), unsupervised machine learning techniques have made hardly a dent in the challenge of semantic characterization of natural languages (Hearst, 1999; Riloff, 1993). As pointed out by Stevenson and Joanis (2003) the development of unsupervised or minimally supervised methods is of special importance to automatically classify verbs in languages other than English, where substantial amounts of labeled data are not available to train classifiers.

This chapter discusses some of the automatic techniques used to augment VerbNet and the obstacles these automatic approaches face. In particular, it describes results of clustering verbs based on frequencies of subcategorization frames observed in a corpus and results of correlating VerbNet and WordNet resources.

8.1 Clustering

VerbNet allows semantic information to be specified for several verbs at a time by grouping verbs into classes that have the same subcategorization alternations. Because these classes capture generalizations about verbs, new candidates can be automatically added to a class by testing them against pre-defined class specifications in order to quickly increase the coverage of the lexicon.

8.1.1 Previous Work

Kingsbury and Kipper (2003), performed an initial experiment using a *k-means* clustering algorithm that measured similarity between VerbNet classes and clusters derived from PropBank. The clustering was tested on 921 verbs (not all present in VerbNet) and 200 syntactic patterns (such as Arg0 V Arg1). To avoid sparse data problems, only those senses which occurred 10 or more times in the annotated corpus were considered.

The algorithm proceeds by dividing the verbs into clusters with different coarseness (i.e., first it partitions the verbs into one cluster, then into 2 clusters, ..., up to 150 clusters). The partitions with closest similarity to VerbNet classes happen when the data is divided into 14 clusters and around 89 clusters. [1] The analysis with 14 clusters however uses very

[1] The similarity metric used to compare the clusters with VerbNet classes in this experiment is computed by $similarity(A, B) = \frac{|A \cap B|}{|A \cup B|}$. In other words, the similarity of these two sets of classes, A and B, can be measured as the number of elements shared by the two sets divided by the total number of unique elements in the two sets. For example, two identical sets {a, b}, {b, a} would have a similarity score of 1, because the intersection contains two elements, as does the union. Two sets {a, b} and {a, b, ..., z} would have a very low similarity score, since the intersection contains two elements, but that is divided by the union

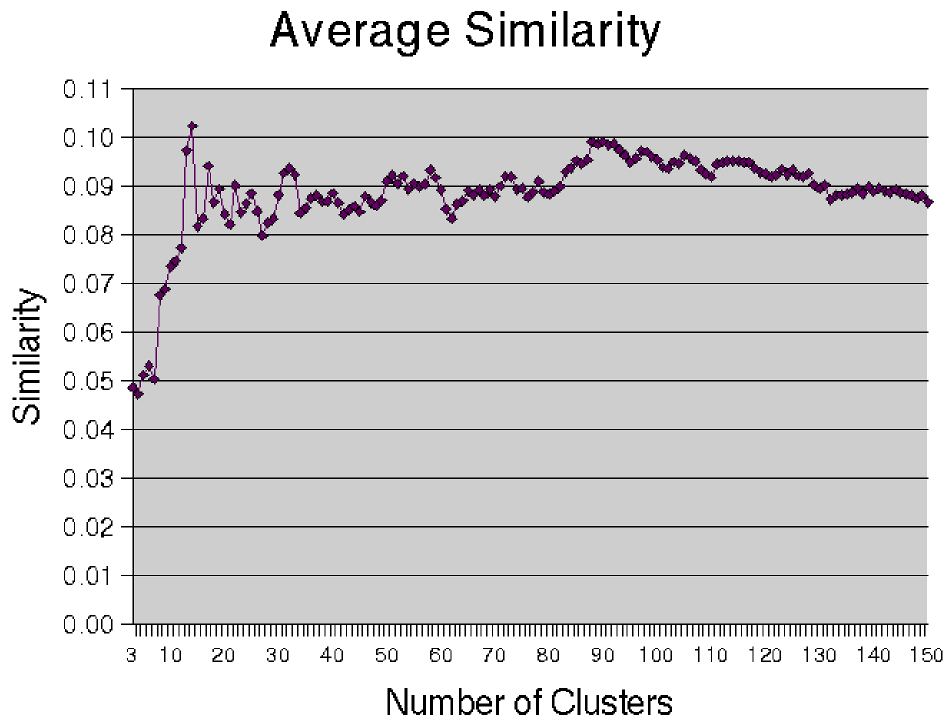

Figure 8.1: Outcome of clustering procedure

large clusters (average of 65 verbs/cluster) many of which are not good representatives of any VerbNet class. Figure 8.1 shows the outcome of this procedure.

When the partition of 90 clusters is used the derived classes tend to be very small, with 34 clusters containing two or fewer verbs. In contrast, a few clusters continue to have a very large membership, with 145 and 146 verbs, corresponding roughly to VerbNet classes 29.5 (*Conjecture* class) and 29.2 (*Characterize* class). Another interesting case is when the clustering process establishes multiple clusters which are all identified with the same VerbNet class, suggesting a subdivision of the original class. This happened with class 45.4 (*Other_change_state*), containing the 'miscellaneous causatives,' which started subdiving when we had only 9 clusters.[2] It is hardly surprising that a collection labeled 'miscellaneous'

with 26 elements.

[2]This shows that there is more variation in the syntactic patterns within this single class than between many of the other pairs of classes.

should contain such a wide variety of syntactic patterns. What is encouraging however, is the fact that this methodology identifies this miscellany.

An example of inclusion suggested by the algorithm is the cluster identified with VerbNet class 36.3 (*Meet* class), verbs of combative meeting. This cluster includes verbs present in the VerbNet class such as *fight* and *consult*, and suggests verbs such as *pull.02* (phrasal: *pull out*) and *withdraw*. *Pull out* is not in VerbNet (there were few phrasal verbs in the lexicon at the time of the experiment), and *withdraw* is a member of the *Remove* classes (10.1 and 10.5). Yet it might be just as natural to think of *withdraw* in the sense of exiting a meeting or an engagement rather than removing something from somewhere (e.g., *It was just another one of the risk factors that led to the company's decision to withdraw from the bidding*). These are two distinct senses of *withdraw* with the one suggested by the experiment appearing with higher frequency in the corpus due to the financial domain of the Wall Street Journal.

Based on the results of this experiment, we discussed the need for different approaches and different features to improve the partitions and to suggest better candidates for VerbNet. In particular, we decided that a more detailed analysis of different algorithms was necessary.

8.1.2 Additional clustering experiments

Following previous work, Kingsbury (2004) performed a comprehensive analysis of several clustering algorithms to group English verbs based on the subcategorization frames exhibited by these verbs in the PropBank annotated corpus. The version of PropBank used in this experiment contains annotations for over 3,200 lexical items, divided into 4,238 senses. As in the preliminary experiment (*k-means*) only those senses which occurred 10 or more times in the annotated corpus were considered.

The syntactic frames selected for the experiment are the exact surface realization of the arguments as found in the corpus, and may include Examples such as (43),(44),(45), and

verb sense	Arg0.rel.Arg1	rel.Arg1	Arg1.rel	Arg0.rel.Arg1.Arg2
soften.01	0.3684	0.1578	0.4736	0.0000
put.01	0.0592	0.0069	0.0034	0.0000
feel.02	0.9610	0.0129	0.0000	0.0000
feel.03	0.5000	0.3571	0.0000	0.0000
break.08	0.6250	0.0625	0.2500	0.0000
crowd.01	0.0000	0.2500	0.2500	0.0000

Table 8.1: Simplified input matrix for the clustering experiment

(46). In the end 2,507 different subcategorization frames (of many more available) were chosen for participation, thus avoiding the infrequent frames.

(43) *Arg0.rel.Arg1*

"John felt a bump."

(44) *rel.Arg1*

"a seed was planted."

(45) *Arg1.rel*

"the butter softened."

(46) *Arg0.rel.Arg1.Arg2-from*

"John obtained a secret from Mary."

The input consists of verb senses and the frequency with which they appear in each of the selected syntactic frames. A simplified example of the input matrix is given in Table 8.1.

This analysis compared six clustering algorithms from the R Developmental Team packages. Those techniques consisted of three variations (*Single*, *Complete*, and *Average*) from the hierarchical cluster algorithm package, two variations (*Clara* and *Pam*) from the k-means cluster package, and an EM algorithm.

In order to account for the accuracy of the clusters, comparisons with other known resources are provided. PropBank data clusters are compared to verb groupings by Stevenson and Merlo (2001), Stevenson and Joanis (2003), and the Levin and VerbNet verb classes as a test of replicability.

Stevenson and Merlo (2001) grouped 59 verbs into three clusters (unergative, unaccusative, and object-drop), and reported around 66% accuracy. The PropBank experiment used 18 verbs (there were not enough instances of the others in the corpus to replicate the experiment) divided roughly equally among the three clusters. Kingsbury (2004) used a similar measure as Stevenson and Merlo (2001) reporting 61% accuracy.

Stevenson and Joanis (2003) selected verbs from 28 Levin classes and merged them into 14 classes with a wider variety of semantic types. Each of the 14 classes contains 29 or more lexical items and each verb occurs at least a 100 times in the corpus. A total of 841 verbs were used, and the verbs were clustered under a variety of feature spaces and different cluster coarseness. The score measure used was the Adjusted Rand Index (Hubert and Arabie, 1985), and they reported a score of 0.11 in the 14-way cluster. The same methodology is applied to PropBank data: 14 most frequent classes are selected and monosemous verbs which appeared at least 10 times in the corpus are used. The adjusted scores for the 6 clustering algorithms varied between 0.015 (for the *HClust Single*) to 0.159 (for *Clust Pam*). Kingsbury (2004) points out that these numbers suggest that the PropBank annotation is a more informative set of features than the ones used by Stevenson and Joanis, since the results were better even though his clusters were smaller (which the metric should have penalized).

More interesting to us, however, is the comparison of the resulting clusters with Levin and VerbNet classes. Two main factors complicate a comparison with Levin classes directly: 1) the lack of sense disambiguation of the verbs present in these classes; 2) PropBank and Levin/VerbNet contain roughly the same number of verbs but only about half of the verbs are simultaneously present in both resources. Although there is not much to do about the coverage for the experiment, the first factor is resolved by the existing mappings between VerbNet and PropBank verbs.

The methodology used for deriving the clusters to be compared to VerbNet include choosing: a) verbs currently present in VerbNet to replicate existing classes, and b) additional members not present in VerbNet as candidates. A total of 484 verbs that occur in the corpus at least 10 times and belong to a unique VerbNet class were chosen for a). The 167 classes represented by these 484 verbs were further grouped to a more manageable number of 121 classes.

The six clustering algorithms were tested on the data given by the frequency of the syntactic frames for each verb, and two measures of similarity were used to score these clustering algorithms: the adjusted Rand Index, and a dissimilarity metric following Melia (2002), which in this case measured the inaccuracy of the test partition.

The algorithms created different numbers of clusters, from a single cluster (all members present in one cluster), to a maximum of 167 clusters. Using the dissimilarity metric *HClust Single* performed worse than chance, probably because the scoring metric penalizes clusters with a single member and this algorithm tends to create such clusters. All the algorithms perform steadily worse as the number of clusters grows, since with more clusters the number of clusters with a single member increases. The best results are achieved with the EM algorithm when it partitions the data into around 90 clusters and with the *Clust Pam* algorithm which has consistently high scores across all numbers of clusters.

Certain clusters received a low score because of the differences in the expected usage of verbs. For example, a cluster containing (*leap, advance, jump*, and *gain*) received a low adjusted score since these verbs are members of distinct VerbNet classes but given that it is constructed using sentences from the Wall Stree Journal these verbs are used in the corpus to describe performance of traded stocks and are not actually that dissimilar. Similar cases happen with certain verbs of communication, which are distributed in various classes in VerbNet but that have a similar range of syntactic realization and therefore are justifiably clustered together. Some of the groupings however such as cluster 3 (*transform, inject*, and *convert*) are odd and dominated by syntactic patterns (*verb X into Y*).

8.1.3 Incorporating the clusters into VerbNet

After analyzing these six algorithms and comparing them to VerbNet classes, the clusters derived from the EM algorithm are chosen to suggest new verbs for VerbNet.

A total of 1,278 verbs which occurred at least 10 times in the PropBank annotation was used as data. Of those, 484 were already associated with a unique VerbNet class, as previously stated. This left 824 potential candidates for inclusion in VerbNet classes. Each verb that is in VerbNet has the class information in addition to the features given by the subcategorization frames from the corpus. This is done so that the algorithm can try to reproduce the existing classes while adding verbs to them, but all of this information, including the class information, may be overridden in case other features are more expressive.

The algorithm partitions the data into 121 clusters. These clusters varied in size from 0 to 45 elements each, with an average membership of 10.56 verbs in each cluster. Table 8.2 shown the distribution of the members into the clusters. As can be observed, one of the clusters has 0 members and should have been discarded. The 10 clusters with a single member are not useful in suggesting new members either: all of these clusters contain verbs which are already in VerbNet and therefore these clusters do not add any suggestions to existing classes. In addition, 4 clusters have all their members assigned to the same VerbNet class and again do not add any new members. After this initial and very general cleanup, we are left with 105 clusters worth investigating.

There are 24 clusters (out of the 105 left) in which at least half of their members is already associated with unique VerbNet classes. These 24 clusters are good candidates to contribute additional members to these classes.

Cluster 53 is a good example of one of the 'well-behaved' clusters. It contains 5 members *serve.01, qualify.02, rank.01, act.01,* and *emerge.02*, and with the exception of the latter all verbs belong to VerbNet class *Masquerade-29.6*. Since *emerge.02 (come to be seen as)*

number of clusters	number of members
1	0
10	1
7	2
7	3
12	4
13	5
4	6
4	7
10	8
6	9
4	10
3	11
5	12
5	13
2	14
2	15
3	16
3	17
2	18
1	19
3	20
2	23
2	26
1	30
1	31
1	32
1	33
1	34
1	37
1	40
1	41
1	43
1	45

Table 8.2: Distribution of the Clusters

has the same syntactic frames as the others, it can be added to VerbNet.

But not all 'well-behaved' clusters immediately add new members to the VerbNet class suggested. Cluster 17 which represents class *Give-13.1* has 10 members, 8 of which already belong to that class, and suggests two new members for this class: *hedge.01* and *protect.01*. This class is characterized by the dative alternation and the verbs suggested do not take that alternation.[3]

Other examples of poorly performing clusters include Cluster 7. While 8 additions are suggested (*cap.02*, *spur.01*, *hurt.01*, *enhance.01*, *fuel.01*, *offset.01*, *hit.01*, *fund.01*) by the clustering experiment for this cluster (representing VerbNet class *Butter-9.9*), one (*fuel.01*) is already a member. Of the others, *spur.01* could be added as a new member. For instance, *spur.01* could be interpreted in this class as spurring an animal which involves the application of spurs to the animal. The other verbs suggested take the syntactic frames specified for the class, but differ significantly in semantics, moreover these verbs do not pattern in any way to suggest that they represent a separate cohesive class.

Some clusters were very unreliable and did not significantly overlap any existing VerbNet class. For example, cluster 98 and 111 are very difficult to understand. The first has 8 members, 4 of which appear in three distinct classes and do not give us any hint on what it is clustering for. Cluster 111 is the largest cluster with 45 members. Members of this cluster appear in 16 distinct classes, with very few of them in the same class. It seems that the common thread between these verbs is that they all merely take the transitive frame.

Almost all of the large clusters, the ones with 30 or more members have a high degree of entropy and are not predictive of any existing classes. Cluster 94 is one of the few exceptions, it represents *Miscellaneous Change of State* verbs in class *COS-45.4* and it is a very good cluster. This cluster suggests 4 new verbs for that class.

It is also interesting to note that certain classes are not suggested by the algorithm.

[3]These two verbs *protect* and *hedge* have been added to one of the new classes suggested by the Korhonen and Briscoe (2004) resource, class *Defend-83*.

Whether this is because these are very closed classes or because they are extracted from the Wall Street Journal in their development remains a question, although the most probable explanation is the data where these clusters are derived from. Example classes: *Touch-20*, *Coloring-24*, *Animal_sounds-38*. Several classes of *Verbs Involving the Body*, *Verbs of Grooming*, and *Manner of Motion Verbs* do not have clusters that represent them.

8.1.4 Comparison with extended VerbNet

We also examined in detail these clusters to see if the less cohesive clusters could be found to correspond to any of the other VerbNet classes, or to the classes suggested by the resource of Korhonen and Briscoe, or even to a possible new class. With the exception of a few isolated cases, the poorly performing clusters failed to cohere in any reasonable way. For instance, a few of the members of Cluster 13 (*commit.01* and *devote.01*) would fit into Korhonen and Briscoe's proposed DEDICATE class. Also, a sizable subset of the verbs in Cluster 29 appear to belong to a class not present in VerbNet but similar to Levin's 45.6 (these verbs differ from the calibratable change of state verbs in that they participate in the Causative Alternation). A large number of verbs in Cluster 79 fail to be appropriate for addition to class 37.9, but seem to correspond to Korhonen and Briscoe's URGE class.

8.1.5 Integration

Of the 824 possible candidates suggested by the algorithm, only about 5.6% of the verbs (47 verbs) were actually integrated into VerbNet. The algorithm did not explicitly provide any new classes, so all these new verbs were integrated into existing VerbNet classes. For the clusters where the algorithm performed well and predicted a VerbNet class with some degree of accuracy, very few new members were suggested. But for the most part, especially for the larger clusters, the algorithm did not return a coherent group of verbs which could be assigned to a single class (or even to a related group of classes).

As we come up with ways to filter the clusters which are clearly unhelpful and focus on the clusters for which the algorithm presents a more coherent classification, we are left with 374 verbs to consider. For these clusters we get about 12.6% verbs that can be incorporated into our lexicon. At this point there are still open questions about which features are the most predictive and about ways to properly filter the clusters. The current set of features is impoverished, making little or no use of sematic information. In addition, a number of gaps were uncovered by the experiment. These gaps reflect the current state of the resources used and once filled will result in more accurate automatic verb acquisition.

- Many verb senses given by the PropBank are not the same as the ones predicted in VerbNet. PropBank data is heavily skewed toward financial usages of verbs whereas VerbNet describes more basic meanings.
- There are gaps in syntactic context and syntactic frames. Much of the syntactic context which could disambiguate senses is missing.
- The experiment does not take advantage of semantic classes of arguments, since no sense tag is available for the nouns.

We are just beginning to scratch the surface of the problems of automatic verb classification. The small number of verbs assigned automatically to a VerbNet class hints at the difficulty of extracting alternations automatically from corpus data with a high degree of accuracy. It also shows how hard it is to find the right set of features needed to create a model that extracts the desired class distinctions directly from the corpus data.

8.2 Using WordNet synsets to extend coverage

The experiments with clustering techniques based on the PropBank annotation broadened our interest in experimenting with other resources for automatic acquisition of new members. The existing mappings between VerbNet, WordNet, and PropBank provided us with an unparalleled opportunity to further explore automatic acquisition of members for our lexicon.

WordNet is a vast repository with rich information for over 11,000 verbs organized into 13,000 synsets. Since VerbNet is mapped to WordNet synsets, and VerbNet is also mapped to PropBank framesets, Loper in a recent experiment examined WordNet as a source of potential candidates for inclusion in our lexicon based on syntactic contexts of these verbs in Propbank. In order to reduce the number of incorrect extensions, new words are filtered based on the grammatical patterns they occur in, and the relationship between those patterns and known members of each class.

All verbs that appear in the same WordNet synsets as members of a VerbNet class are initially candidates for inclusion in the lexicon.[4] Syntactic contexts of these verbs in corpora (in this case PropBank and other smaller corpora) are retrieved. Each candidate verb is assigned a cumulative score according to the instances found in the corpus. The score of each candidate for inclusion is adjusted depending on whether the syntactic frames observed appear in a subset of the frames attested by the VerbNet members. The algorithm also accounts for the frequency of examples and whether the verb appears with words that are the same as those observed with the VerbNet members. All verbs from WordNet whose cumulative score is above a certain threshold are chosen as candidates for inclusion in VerbNet classes.

Loper's initial experiment proposed a total of 707 new VerbNet verbs from WordNet synsets based on the methodology discussed above. Of those, about a third (255), could directly be integrated into our lexicon exactly in the subclasses for which they were suggested. Since the experiment was done using VerbNet v1.0, some of the classes had already significantly changed and some of the suggested verbs were already present in the extended VerbNet. Other factors that contributed to a low inclusion rate:

[4] If there is a mapping between VerbNet and WordNet for a particular member, only that synset is used.

- Although some verbs could not be integrated in the suggested subclass they may be members of other verb classes (or subclasses).

- WordNet as a comprehensive semantic network has different goals than VerbNet. Even though some entries as exemplified below are appropriately included in WordNet, they do not have a place in our lexicon and therefore should be filtered out as candidates for inclusion:

  - Many verbs were suggested with a preposition but in VerbNet the preposition is part of the syntactic frame. For example, the verb *chase* was suggested as *chase after* (as a synonym of *chase*). This verb however is already a member of class *Chase-51.6* with syntactic frame *Agent V Theme Prep[+path OR +loc] Location*.
  - Several instances were idiomatic expressions (e.g., *shoot a line*).
  - A large number of the proposed instances were simple morphological variants of the same verb (e.g., *criminalize vs. criminalise*, *canvas vs. canvass*, and *bourgeon vs. burgeon*).
  - Several of the proposed verbs are of archaic usages (e.g., *aggroup, coggle*) for which there is little intuition about syntactic preferences.

The algorithm suggested members that fit very well for some of the classes proposed, such as 23 new verbs included in class *Create-26.4*, 17 new verbs included in class *Destroy-44*, and 15 new verbs for class *Contribute-13.2*. Table 8.2 shows for each VerbNet class the number of new members included. As can be seen from this table, the genre of the Wall Street Journal plays an important role in which classes have additional members suggested.

For some classes however, the suggested verbs did not seem to fit. Examples of these classes and a few reasons why this happens include: VerbNet class *Send-11.1* where the

VerbNet class	Number of new additions
Put_spatial-9.2	3
Funnel-9.3	2
Put_direction-9.4	5
Spray-9.7	6
Banish-10.2	4
Wipe_manner-10.4.1	5
Carry-11.4	2
Push-12	2
Contribute-13.2	15
Equip-13.4.2	3
Poke-19	2
Mix-22.1	2
Amalgamate-22.2	12
Tape-22.4	2
Cling-22.5	2
Scribble-25.2	7
Illustrate-25.3	4
Build-26.1	7
Create-26.4	23
Dub-29.3	3
Conjecture-29.5	9
Orphan-29.7	3
See-30.1	13
Peer-30.3	5
Stimulus_subject-30.4	2
Long-32.2	6
Judgement-33	24
Investigate-35.4	14
Rummage-35.5	6
Transfer_mesg-37.1	4
Instr_communication-37.4	4
Complain-37.8	7
Devour-39.4	6
Flinch-40.5	4
Body_internal_states-40.6	5
Hurt-40.8.3	11
Murder-42.1	10
Poison-42.2	6
Destroy-44	17
Bend-45.2	7
Entity_specific_cos-45.5	12
Herd-47.5.2	7
Meander-47.7	5
Appear-48.1.1	17
Occurrence-48.3	4
Body_internal_motion-49	3
Price-54.4	4

Table 8.3: Number of verbs added to VerbNet classes from WordNet

verbs suggested are more appropriate to be added to a different subclass; the verbs suggested for class *Spray-9.7* are very archaic; class *Begin-55.1* has changed considerably between versions; and for other classes such as *Stalk-35.3* and *Correspond-36.1* the algorithm simply did not suggest any good inclusions.

This semi-automatic extension of VerbNet's coverage demonstrates one of the immediate benefits of having these three complementary resources, VerbNet, WordNet, and PropBank, mapped to each other. The experiment shows promising results as attested by the 255 new verbs (208 lemmas) added to our lexicon. However, the number of inclusions is still small given the number of potential candidates, which again shows the difficulties of automatically extracting discriminating contexts from corpora. In addition, for this particular experiment, a more comprehensive analysis of the filters used and possibly more annotated corpora is likely to produce candidate verbs that more closely resemble our classes.

This experiment was redone on VerbNet v2.2, and 9,302 verb senses were suggested (4,992 lemmas). Of those, we inspected only candidates which appear in the same subcategorization frames (same syntactic context) as the VerbNet members (413 senses). We incorporated 179 senses to VerbNet as a result (43.34%).

8.3 Conclusion

The development of unsupervised techniques for lexical acquisition is crucial for extending coverage of broad-coverage lexical resources. The experiments discussed in this chapter show the many obstacles faced by researchers tackling this difficult problem. Although it is in its early stages, we demonstrate that it is possible to semi-automatically supplement and tune VerbNet with novel information from corpus data. Such semi-automatic methods reduce the need for manual classification and enable easy adaptation of the lexicon to specific tasks and applications. The only reason these approaches to lexical acquisition are even possible is because of the class-based organization of VerbNet.

Chapter 9

Conclusion

This thesis describes a broad-coverage natural language resource that provides syntactic and semantic argument structures for English verbs: a class of words for which these structures are particularly diverse.

This resource differs from existing lexicons in that it:

1. focuses both on syntax and semantics and provides a clear association between the two;
2. is not domain specific;
3. is not tied to specific corpora;
4. is available to the whole community.

Our preliminary investigation into other languages such as Portuguese indicates that this methodology could also be applied cross-linguistically.

The verb entries were constructed using a class-based approach adapted from the Levin classification. According to Levin, verbs fall into classes which are described by the diathesis alternations they take. These classes, although mostly characterized by the verbs' syntax are believed to have some underlying semantic cohesiveness.

Verb classes are important for their ability to capture generalizations about language by grouping verbs with similar meaning constituents, and similar syntactic behavior. These classes are present throughout the lexicon and are claimed to exist across languages since their basic meaning components can be applicable cross-linguistically. Verb classes can also be used to reduce ambiguity in the lexicon and to fill gaps in lexical knowledge by highlighting the correspondence between syntax and semantics thus enabling inferences of syntactic and semantic behavior.

Our main contribution is the development of VerbNet, one of the most extensive verb lexicons currently available for English. VerbNet describes in detail the syntactic behavior and semantic components of classes of English verbs based on the Levin classification. It is a hierarchically organized lexicon that extends and refines the original classes substantially through the addition of subclasses thus providing more semantically and syntactically homogenous classes. Validation of the hierarchical structure and detailed computational design of the lexicon is attested by its easy integration with other resources. VerbNet uses roles with semantic restrictions to characterize verbs in a class; it describes the syntactic frames for the class members through explicit descriptions of allowed surface realizations using the roles as arguments; and it describes the key components of the semantics of each syntactic frame using semantic predicates with a time function for the event variable.

Since VerbNet is related to and can extend other known resources which also provide broad-coverage, another contribution of this work is the mappings between VerbNet and existing resources. These mappings present researchers with different views of these lexicons which complement each other. The mapping between VerbNet verbs and WordNet synsets allows WordNet verbs to have an explicit predicate-argument structure and to make an explicit use of semantic predicates, and it provides our lexicon with the rich semantic information which can be derived from WordNet's relations. The mapping from our verbs to FrameNet frames allows for a different view of events while providing FrameNet with additional semantic information than supplied by the semantic roles, and it gives VerbNet a more fine-grained set of roles for the verbs with more detailed semantic components. The mapping of VerbNet frames to Xtag trees provides the Xtag grammar a way of differentiate among verb senses while extending significatively our syntactic coverage.

We provided an evaluation of VerbNet's coverage through mappings to the Proposition Bank on over 78,000 instances and VerbNet syntactic frames account for more than 84% exact matches to the frames found in that rich lexical resource.

In addition, we investigated several ways to extend VerbNet:

1. New classes from Korhonen and Briscoe (2004) that had been semi-automatically acquired were incorporated into the lexicon, extensively increasing our coverage, especially that of verbs with sentential complements.

2. A large set of verbs from the LCS database was used to extend the coverage of existing classes.

3. Kingsbury's (2004) clusters derived from the subcategorization frames from PropBank were systematically examined and some verbs were incorporated into VerbNet classes.

4. Many new verbs suggested by Loper's experiment which correlates VerbNet and WordNet were added to the lexicon.

These recent experiments show that it should be possible to semi-automatically supplement and tune VerbNet with novel information from corpus data. Such an approach reduces the expensive overhead of manual classification and enables adapting the lexicon for specific tasks and applications. The computationally-based design and hierarchical structure of VerbNet greatly facilitate these extensions. In addition, it is possible to include in the lexicon statistical information concerning the relative likelihood of different classes, subcategorization frames, and alternations for verbs in the corpus. Such information can be highly useful for statistical natural language systems utilizing lexical-semantic classes.

References

Allbeck, J., Kipper, K., Adams, C., Schuler, W., Zoubanova, E., Badler, N., Palmer, M., and Joshi, A. (2002). ACUMEN: Amplifying control and understanding of multiple entities. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2002)*, Bologna, Italy.

Allen, J. and Fergunson, G. (1994). Actions and events in interval temporal logic. *Logic and Computation*, 4(5).

Badler, N. I., Bindiganavale, R., Allbeck, J., Schuler, W., Zhao, L., and Palmer, M. (2000). *Embodied Conversational Agents*, chapter Parameterized Action Representation for Virtual Human Agents. MIT Press.

Badler, N. I., Palmer, M., and Bindiganavale, R. (1999). Animation control for real-time virtual humans. *Communications of the ACM*, 42(7):65–73.

Baker, C. F., Fillmore, C. J., and Lowe, J. B. (1998). The Berkeley FrameNet project. In *Proceedings of the 17th International Conference on Computational Linguistics (COLING/ACL-98)*, pages 86–90, Montreal.

Baker, C. F. and Ruppenhofer, J. (2002). FrameNet's Frames vs. Levin's Verb Classes. In *In J. Larson and M. Paster (eds.) Proceedings of the 28th Annual Meeting of the Berkeley Linguistics Society*, pages 27–38.

Boguraev, B., Briscoe, T., Carroll, J., Carter, D., and Grover, C. (1987). The derivation of a grammatically-indexed lexicon from the Longman Dictionary of Contemporary English. In *Proceedings of the* 25$^{\text{th}}$ *annual meeting of the Association for Computational Linguistics*, pages 193–200, Stanford, CA.

Briscoe, T. (2000). *Dictionary and System Subcategorisation Code Mappings.* Unpublished manuscript, http://www.cl.cam.ac.uk/users/alk23/subcat/subcat.html, University of Cambridge Computer Laboratory.

Buitelaar, P. (1998). *CoreLex: Systematic Polysemy and Underspecification.* PhD thesis, Computer Science, Brandeis University.

Chen, H.-H., Lin, C.-C., and Lin, W.-C. (2002). Building a Chinese-English wordnet for translingual applications. *ACM Transactions on Asian Language Information Processing (TALIP)*, 1(2):103–122.

Chi, D., Costa, M., Zhao, L., and Badler, N. (2000). The EMOTE model for Effort and Shape. In *Proc. ACM SIGGRAPH*, pages 173–182, New Orleans, LA.

Choi, K.-S. (2003). Corenet: Chinese-japanese-korean wordnet with shared semantic hierarchy. In *Proceedings of the International Conference on Natural Language Processing and Knowledge Engineering*, pages 767–770.

COMLEX (1994). The COMLEX Lexical Database. Copyright Trustees of the University of Pennsylvania.

Copestake, A. (1992). *The Representation of Lexical Semantic Information.* PhD thesis, University of Sussex.

Copestake, A. and Sanfilippo, A. (1993). Multilingual Lexical Representation. In *Proceedings of the AAAI Spring Symposium: Building Lexicons for Machine Translation*, Stanford, California.

Croch, D. and King, T. H. (2005). Unifying lexical resources. In *In Proceedings of Interdisciplinary Workshop on the Identification and Representation of Verb Features and Verb Classes, edited by Katrin Erk, Alissa Melinger and Sabine Schulte im Walde*, Saarbruecken, Germany.

Dang, H. T. (2004). *Investigations into the Role of Lexical Semantics in Word Sense Disambiguation*. PhD thesis, CIS, University of Pennsylvania.

Dang, H. T., Kipper, K., Palmer, M., and Rosenzweig, J. (1998). Investigating regular sense extensions based on intersective Levin classes. In *Proceedings of ACL98*, Montreal, Canada.

Dorr, B. J. (1993). *Machine Translation: A View from the Lexicon*. MIT Press, Cambridge, Massachussets.

Dorr, B. J. (1997). Large-Scale Dictionary Construction for Foreign Language Tutoring and Interlingual Machine Translation. *Machine Translation*, 12:1–55.

Dorr, B. J. (2001). LCS Verb Database. In *Online Software Database of Lexical Conceptual Structures and Documentation*, University of Maryland.

Dorr, B. J. and Jones, D. (1995). Automatic extraction of semantic classes from syntactic information in online resources. Technical report, UMIACS-TR-95-65, CS-TR-3481, University of Maryland, College Park, MD.

Dorr, B. J. and Jones, D. A. (1996). Acquisition of semantic lexicons: Using word sense disambiguation to improve precision. In *Proceedings of SIGLEX*, Santa Cruz, California.

Dorr, B. J. and Jones, D. A. (2000). *Breadth and Depth of Semantic Lexicons*, chapter Acquisition of Semantic Lexicons: Using Word Sense Disambiguation to Improve Precision in Evelyn Viegas, pages 79–98. Kluwer Academic Publishers, Norwell, MA.

Dowty, D. (1989). On the semantic content of the notion thematic role. In Chierchia, G., Partee, B., and Turner, R., editors, *Properties, Types, and Meaning, vol. 2*. Kluwer Academic Publishers, Dordrecht.

Dowty, D. (1991). Thematic proto-roles and argument selection. *Language*, 67.3:547–619.

Fellbaum, C., editor (1998). *WordNet: An Eletronic Lexical Database*. Language, Speech and Communications. MIT Press, Cambridge, Massachusetts.

Fillmore, C. J. (1968). The Case for Case. In Bach, E. and Harms, R., editors, *Universals in Linguistic Theory*. Holt, Rinehart, and Winston, New York.

Fillmore, C. J. (1976). Frame semantics and the nature of language. In *Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech*, volume 280, pages 20–32.

Gildea, D. and Jurafsky, D. (2002). Automatic labeling of semantic roles. *Computational Linguistics*, 28(3):245–288.

Girju, R., Roth, D., and Sammons, M. (2005). Token-level disambiguation of VerbNet classes. In *In Proceedings of Interdisciplinary Workshop on the Identification and Representation of Verb Features and Verb Classes, edited by Katrin Erk, Alissa Melinger and Sabine Schulte im Walde.*, Saarbruecken, Germany.

Grishman, R., Macleod, C., and Meyers, A. (1994). COMLEX syntax: building a computational lexicon. In *Proceedings of the International Conference on Computational Linguistics*, Kyoto, Japan.

Gruber, J. (1965). *Studies in Lexical Relations*. PhD thesis, MIT (reprinted in Lexical Strutures in Syntax and Semantics. Amsterdam, North-Holland, 1976).

Harris, Z. S. (1962). *String Analysis of Sentence Structure*. Mouton, The Hague.

Hearst, M. (1999). Untangling text data mining. In *Proceedings of the 37th Annual Meeting of the ACL*, pages 3–10, College Park, Maryland.

Hensman, S. and Dunnion, J. (2003). Applying VerbNet for automatic semantic role identification. In *In Proceedings of the 14th Irish Conference of Artificial Intelligence and Cognitive Science (AICS 2003)*, pages 72–76, Dublin, Ireland.

Hensman, S. and Dunnion, J. (2004). Automatically building conceptual graphs using VerbNet and WordNet. In *In Proceedings of the 3rd International Symposium on Information and Communication Technologies (ISICT)*, pages 115–120, Las Vegas, NV.

Hubert, L. and Arabie, P. (1985). Comparing partitions. *Journal of Classification*, 2:193–218.

Jackendoff, R. (1972). *Semantic Interpretation in Generative Grammar.* MIT Press, Cambridge, Massachusetts.

Jackendoff, R. (1983). *Semantics and Cognition.* MIT Press, Cambridge, Massachusetts.

Jackendoff, R. (1987). The status of thematic relations in linguistic theory. *Linguistic Inquiry*, 18:369–411.

Jackendoff, R. (1990a). *Semantic Structures.* MIT Press, Cambridge, Massachusetts.

Jackendoff, R. (1990b). *Semantic Structures.* MIT Press, Cambridge, Massachusetts.

Jones, K. S. and Boguraev, B. (1987). A note on a study of cases. *American Journal of Computational Linguistics*, 13((1-2)):65–68.

Joshi, A. K. (1985). How much context sensitivity is necessary for characterizing structural descriptions: Tree Adjoining Grammars. In D. Dowty, L. K. and Zwicky, A., editors, *Natural language parsing: Psychological, computational and theoretical perspectives*, pages 206–250. Cambridge University Press, Cambridge, U.K.

Kallmeyer, L. and Joshi, A. K. (2003). Factoring predicate argument and scope semantics: Underspecified semantics with LTAG. *Research on Language and Computation*, 1(1–2):3–58.

Kallmeyer, L. and Romero, M. (2004). LTAG semantics with semantic unification. In *Seventh International Workshop on Tree Adjoining Grammar and Related Formalisms (TAG+7)*, Vancouver, Canada.

Kingsbury, P. (2004). Verb clusters from PropBank annotation. Technical report, University of Pennsylvania, Philadelphia, PA.

Kingsbury, P. and Kipper, K. (2003). Deriving verb-meaning clusters from syntactic structure. In *Workshop on Text Meaning, held in conjunction with HLT–NAACL 2003 Conference*, Edmonton, Canada.

Kingsbury, P. and Palmer, M. (2002). From Treebank to PropBank. In *Proceedings of the 3rd International Conference on Language Resources and Evaluation (LREC-2002)*, Las Palmas, Canary Islands, Spain.

Kipper, K., Snyder, B., and Palmer, M. (2004a). Extending a verb-lexicon using a semantically annotated corpus. In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC-04)*, Lisbon, Portugal.

Kipper, K., Snyder, B., and Palmer, M. (2004b). Using prepositions to extend a verb lexicon. In *Proceedings of the HLT/NAACL Workshop on Computational Lexical Semantics*, Boston, MA.

Klavans, J. and Kan, M.-Y. (1998). Role of verbs in document analysis. In *Proceedings of the 17th International Conference on Computational Linguistics (COLING/ACL-98)*, pages 680–686.

Korhonen, A. (2002). *Subcategorization Acquisition.* PhD thesis, University of Cambridge, UK.

Korhonen, A. and Briscoe, T. (2004). Extended Lexical-Semantic Classification of English Verbs . In *Proceedings of the HLT/NAACL Workshop on Computational Lexical Semantics*, Boston, MA.

Laban, R. (1971). *The Mastery of Movement.* Plays, Inc., Boston.

Lapata, M. and Brew, C. (1999). Using Subcategorization to Resolve Verb Class Ambiguity. In *Proceedings of the Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora (EMNLP'99)*, pages 266–274, College Park, MD.

Levin, B. (1993). *English Verb Classes and Alternation, A Preliminary Investigation.* The University of Chicago Press.

Levin, B. and Rapoport, T. (1988). Lexical subordination. In *Papers from the 24th Regional Meeting of the Chicago Linguistic Society*, pages 275–289, Chicago, IL.

Marcus, M. (1994). The Penn TreeBank: A revised corpus design for extracting predicate-argument structure. In *Proceedings of the ARPA Human Language Technology Workshop*, Princeton, NJ.

Melia, M. (2002). Comparing clusterings. Technical report, Technical Report 418, University of Washington, Department of Statistics, Seattle, WA.

Merlo, P. and Stevenson, S. (2001). Automatic verb classification based on statistical distribution of argum ent structure. *Computational Linguistics*, 27(3).

Miller, G. (1985). Wordnet: A dictionary browser. In *Proceedings of the First International Conference on Information in Data*, Waterloo, Ontario.

Miller, G. A. (1990). WordNet: An on-line lexical database. *International Journal of Lexicography*, 3(4):235–312.

Moens, M. and Steedman, M. (1988). Temporal Ontology and Temporal Reference. *Computational Linguistics*, 14:15–38.

Nevin, B., editor (2003). *The Legacy of Zellig Harris. Language and Information in the 21st century.* John Benjamins, Amsterdam, NL.

Palmer, M., Dang, H., and Fellbaum, C. (2005). Making fine-grained and coarse-grained sense distinctions, both manually and automatically. *Journal of Natural Language Engineering.* To appear.

Palmer, M. and Rosenzweig, J. (1996). Capturing motion verb generalizations with synchronous TAGs. In *Proceedings of the 2nd Conference of the Association for Machine Translation of the Americas (AMTA 1996)*, Montreal, Quebec.

Palmer, M., Rosenzweig, J., and Schuler, W. (1999). Capturing Motion Verb Generalizations in Synchronous Tree-Adjoining Grammar. In Saint-Dizier, P., editor, *Predicative Forms in Natural Language and in Lexical Knowledge Bases* . Kluwer Press.

Palmer, M. S., Passonneau, R. J., Weir, C., and Finin, T. W. (1993). The KERNEL text understanding system. *Artificial Intelligence*, 63(1-2):17–68.

Pianta, E., Bentivogli, L., and Girardi, C. (2002). MultiWordNet: developing an aligned multilingual database. In *In Proceedings of the First International Conference on Global WordNet*, Mysore, India.

Pinker, S. (1989). *Learnability and Cognition: The Acquisition of Argument Structure.* MIT Press, Cambridge, Massachusetts.

Procter, P., Illson, R., Ayto, J., and et al. (1978). *Longman Dictionay of Contemporary English.* Longman Group Ltd., Harlow, UK.

Pustejovsky, J. (1991). The Generative Lexicon. *Computational Linguistics*, 17(4):409–441.

Pustejovsky, J. (1995). *The Generative Lexicon.* MIT Press, Cambridge.

Rambow, O., Dorr, B. J., Kipper, K., Kuvcerova, I., and Palmer, M. (2003). Automatically Deriving Tectogrammatical Labels from Other Resources: A Comparison of Semantic Labels Across Frameworks. *The Prague Bulletin of Mathematical Linguistics*, 80.

Resnik, P. (1996). Selectional constraints: an information-theoretic model and its computational realization. *Cognition*, 61:127–159.

Riloff, E. (1993). Automatically constructing a dictionary for information extraction tasks. In *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI)*, pages 811–816, Washington, D.C.

Rudanko, J. (1996). *Prepositions and Complement Clauses.* State University of New York Press, Albany.

Rudanko, J. (2000). *Corpora and Complementation.* University Press of America.

Ryant, N. and Kipper, K. (2004). Assigning xtag trees to verbnet. In *Seventh International Workshop on Tree Adjoining Grammar and Related Formalisms (TAG+7)*, Vancouver, Canada.

Sager, N. (1981). *Natural Language Information Processing: A Computer Grammar of English and Its Applications.* Addison-Wesley Publising Company, MA.

Sanfilippo, A. (1990). *Grammatical relations, Thematic Roles and Verb Semantics.* PhD thesis, University of Edinburgh.

Sanfillppo, A. (1994). Word knowledge acquisition, lexicon construction and dictionary compilation. In *International Conference on Computational Linguistics*, pages 273–277, Kyoto, Japan.

Shi, L. and Mihalcea, R. (2005). Putting pieces together: Combining FrameNet, VerbNet and WordNet for robust semantic parsing. In *Proceedings of the Sixth International Conference on Intelligent Text Processing and Computational Linguistics*, Mexico City, Mexico.

Stevenson, S. and Joanis, E. (2003). Semi-supervised verb class discovery using noisy features. In *Proceedings of the 7th Conference on Natural Language Learning (CoNLL-03)*, Edmonton, Canada.

Stone, M. and Doran, C. (1997). Sentence Planning as Description Using Tree Adjoining Grammar. In *Proceedings of ACL-EACL '97*, Madrid, Spain.

Swier, R. and Stevenson, S. (2004). Unsupervised semantic role labelling. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing (EMNLP-04)*, pages 95–102, Barcelona, Spain.

Swift, M. (2005). Towards automatic verb acquisition from VerbNet for spoken dialog processing. In *In Proceedings of Interdisciplinary Workshop on the Identification and Representation of Verb Features and Verb Classes, edited by Katrin Erk, Alissa Melinger and Sabine Schulte im Walde.*, Saarbruecken, Germany.

Talmy, L. (1985). Lexicalization patterns: Semantic structure in lexical forms. In Shopen, T., editor, *Language Typology and Syntactic Descriptions*, volume 3, pages 57–149. Cambridge University Press, Cambridge, England.

Talmy, L. (1991). Path to realization–via aspect and result. In *Proceedings of the 17th Annual Meeting of the Berkeley Linguistic Society*, pages 480–519.

Terenzi, E. and Di Eugenio, B. (2003). Building lexical semantic representations for natural language instructions. In *HLT–NAACL03, 2003 Human Language Technology Conference*, Edmonton, Canada. (short paper).

Tufts, D., editor (2000). *BalkaNet - Design and Development of a Multilingual Balkan WordNet*, volume 7, number (1-2). Publishing House of the Romanian Academy.

Vendler, Z. (1967). Verbs and times. In *Linguistics in Philosophy*, pages 97–121. Cornell University Press, Ithaca, NY.

Vijay-Shanker, K. and Joshi, A. (1991). Unification Based Tree Adjoining Grammars. In Wedekind, J., editor, *Unification-based Grammars*. MIT Press, Cambridge, Massachusetts.

Vossen, P. (2003). EuroWordNet: Building a multilingual database with wordnets for several european languages. In *http://www.illc.uva.nl/EuroWordNet/*.

XTAG Research Group (2001). A lexicalized tree adjoining grammar for English. Technical Report IRCS-01-03, IRCS, University of Pennsylvania.

Zhao, L. (2001). *Synthesis and Acquisition of Laban Movement Analysis Qualitative Parameters for Communicative Gestures*. PhD thesis, CIS, University of Pennsylvania.