Class-Based Construction of a Verb Lexicon

Karin Kipper, Hoa Trang Dang, and Martha Palmer

Department of Computer and Information Sciences
University of Pennsylvania
200 South 33rd Street
Philadelphia, PA 19104, USA
{kipper,htd,mpalmer}@linc.cis.upenn.edu

Abstract

We present an approach to building a verb lexicon compatible with WordNet but with explicitly stated syntactic and semantic information, using Levin verb classes to systematically construct lexical entries. By using verb classes we capture generalizations about verb behavior and reduce the effort needed to construct the lexicon. The syntactic frames for the verb classes are represented by a Lexicalized Tree Adjoining Grammar augmented with semantic predicates, which allows a compositional interpretation.

Introduction

Despite many different approaches to lexicon development (Pustejovsky 1991), (Copestake & Sanfilippo 1993), (Lowe, Baker, & Fillmore 1997), (Dorr 1997), the field of Natural Language Processing (NLP) has yet to develop a clear consensus on guidelines for computational verb lexicons, which has severely limited their utility in NLP applications. Many approaches make no attempt to associate the semantics of a verb with its possible syntactic frames. Others list too many fine-grained sense distinctions due to a lack of a systematic account of verb polysemy. Even WordNet (Miller 1985), one of the most widely used online lexical databases in NLP applications, has elicited much criticism about its representation of verbs.

We address these problems by creating VerbNet, a verb lexicon compatible with WordNet but with explicitly stated syntactic and semantic information, using Levin verb classes (Levin 1993) to systematically construct lexical entries. While previous research on tying semantics to Levin classes (Dorr 1997) has not made explicit the close relation between syntax and semantics hypothesized by Levin, our lexical resource combines traditional lexical semantic information such as thematic roles and semantic predicates, with syntactic frames and selectional restrictions. We have used Lexicalized Tree Adjoining Grammar (LTAG) (Joshi 1985; Schabes 1990) to capture the syntax associated with each verb class. We also show how regular extensions of verb meaning can be achieved through the adjunction of particular syntactic phrases. We base these regular extensions on intersective Levin classes, a fine-grained variation on Levin classes, as a source of semantic components associated with specific adjuncts (Dang *et al.* 1998). Unlike work by Saint-Dizier (1999), we do not attempt to automatically construct classes from syntactic criteria. We concentrate on building lexical entries only for the word senses already classified by Levin; future work may add verbs to the classes, as well as construct additional classes for senses not covered by Levin.

Levin Classes and WordNet

Two current approaches to English verb classifications are WordNet and Levin classes. WordNet is an on-line lexical database of English that currently contains approximately 120,000 sets of noun, verb, adjective, and adverb synonyms, each representing a lexicalized concept. A synset (synonym set) contains, besides all the word forms that can refer to a given concept, a definitional gloss and - in most cases - an example sentence. Words and synsets are interrelated by means of lexical and semantic-conceptual links, respectively. Antonymy or semantic opposition links individual words, while the super-/subordinate relation links entire synsets. WordNet was designed principally as a semantic network, and contains little syntactic information. Even as a semantic resource, however, it is missing some of the information that has traditionally been required by NLP applications, including explicit predicate-argument structures. In addition, WordNet senses are often too fine-grained, lacking an underlying notion of semantic components and a systematic extension of basic senses to produce these more fine-grained senses.

On the other hand, the Levin verb classification explicitly states the syntax for each class, but also falls short of assigning semantic components to each class. 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 (diathesis alternations) (Levin 1993). The fundamental assumption is that the syntactic frames are a direct reflection of the underlying semantics. However, Levin classes exhibit inconsistencies that have hampered researchers' ability to reference them directly in applications. Many verbs are listed in multiple classes, some of which have conflicting sets of syntactic frames. Dang et al. (1998) showed that multiple listings could in some cases be interpreted as regular sense extensions, and defined intersective Levin classes, which are a more fine-grained, syntactically and semantically coherent refinement of basic Levin classes. We implement these verb classes and their regular sense extensions in the Lexicalized Tree Adjoining Grammar formalism.

Lexicalized Tree-Adjoining Grammars

Lexicalized Tree-Adjoining Grammars (LTAGs) consist of a finite set of initial and auxiliary elementary trees, and two operations to combine them. The minimal, non-recursive linguistic structures of a language, such as a verb and its complements, are captured by initial trees. Recursive structures of a language, such as prepositional modifiers which result in syntactically embedded verb phrases (VP), are represented by auxiliary trees.

Elementary trees are combined by the operations of substitution and adjunction. Substitution is a simple operation that replaces a leaf of a tree with a new tree. Adjunction is a splicing operation that replaces an internal node of an elementary tree with an auxiliary tree. Every tree is associated with a lexical item of the language, called the *anchor* of the tree. The tree represents the domain over which the lexical item can directly specify syntactic constraints, such as subject-verb number agreement, or semantic constraints, such as selectional restrictions, all of which are implemented as features in Feature-Based LTAG (FB-LTAG) (Vijay-Shanker & Joshi 1991). An example of a simple transitive tree is shown in Figure 1. Alternative syntactic realizations of a lexical item are grouped together into *tree families*. Each family represents a basic argument structure (e.g., intransitive, transitive) and has a set of trees corresponding to transformations (e.g., passive, wh-movement) of the basic structure. The semantic constraints automatically apply to the same arguments in the alternative trees.

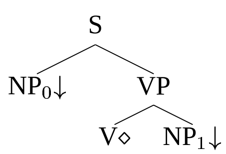

Figure 1: Transitive Tree

Previous work in incorporating semantics into TAG trees has been done. Stone and Doran (1997) described a system used for generation that simultaneously constructs the semantics and syntax of a sentence using LTAGs. Each lexical item is associated with a syntactic tree or family of trees, and each tree has logical forms representing the semantic and pragmatic information selected for that particular lexical item and tree. The meaning of a sentence is the conjunction of the meanings of the elementary trees used to derive it, once the arguments are recovered.

Palmer et al. (1999) and Bleam et al. (1998) defined compositional semantics in FB-LTAG for classes of verbs, representing general semantic components (e.g., motion, manner) as features on the nodes of the trees.

Joshi and Vijay-Shanker (1999), and Kallmeyer and Joshi (1999), describe the semantics of a derivation tree as a set of attachments of trees. The semantics of these attachments is given as a conjunction of formulae in a flat semantic representation. The order of the attachments is not relevant and does not need to be reflected in the semantics. They provide a specific methodology for composing semantic representations much like Candito and Kahane (1998), where the directionality of dominance in the derivation tree should be interpreted according to the operations used to build it.

Description of the Verb Lexicon

VerbNet has two aspects: static and dynamic. The static aspect refers to the verb entries and how they are organized. This aspect provides the characteristic descriptions of a verb sense or a verb class. The dynamic aspect of the lexicon constrains the entries to allow a compositional interpretation in LTAG derivation trees. By incorporating adjuncts we also capture extended verb meanings.

Each verb entry refers to a set of classes, corresponding to the different senses of the verb. For example, the manner of motion sense of "run" is a member of the *Manner of Motion* class, whereas "run" as in "the street runs through the district" is a member of the *Meander* class. For each verb sense there is a verb class as well as specific selectional restrictions (e.g., an instrument of "kick" must be of type $foot$) and semantic characteristics (e.g., a particular manner of directed motion) that may not be captured by the class membership. In order to provide a mapping to other dictionaries, we are also including links to WordNet synsets. Because WordNet has more fine-grained sense distinctions than Levin, each verb sense in VerbNet references the set of WordNet synsets (if any) that captures the meaning appropriate to the class.

Verb classes allow us to capture generalizations about verb behavior. This reduces not only the effort needed to construct the lexicon, but also the likelihood that errors are introduced when adding a new verb entry. Each verb class lists the thematic roles that the predicate-argument structure of its members allows, and provides descriptions of the syntactic frames corresponding to licensed constructions, with selectional restrictions defined for each argument in each frame. Each frame also includes semantic predicates describing the participants during various stages of the event described by the 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, which is facilitated by using intersective Levin classes. In addition, a particular verb may add more semantic information to the basic semantics of its class.

We decompose each event E into a tripartite structure according to Moens and Steedman (1988). We introduce a time function for each predicate specifying whether the predicate is true in the preparatory ($during(E)$), culmination ($end(E)$), or consequent ($result(E)$) stage of an event. The tripartite event structure allows us to express the semantics 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)***).

HIT class

⟨⟨MEMBERS ⟩⟩	$[\langle hit, 1\rangle, \langle kick, 1\rangle, \langle slap, 1\rangle, \langle tap, 1\rangle, \ldots]$
⟨⟨THEMATIC ROLES ⟩⟩	Agent(A), Patient(P), Instrument(I)
⟨⟨SELECT RESTRICTIONS ⟩⟩	Agent[+animate], Patient[+concrete], Instrument[+concrete,-animate]

⟨⟨FRAMES and PREDICATES ⟩⟩

Basic Transitive	A V P	manner(during(E),directedmotion,A) ∧ manner(end(E),forceful,A) ∧ contact(end(E),A,P)
Transitive with Instrument	A V P with I	manner(during(E),directedmotion,I) ∧ manner(end(E),forceful,I) ∧ contact(end(E),I,P)
Together reciprocal	A V P[+plural] together	manner(during(E),directedmotion,P_i) ∧ manner(during(E),directedmotion,P_j) ∧ manner(end(E),forceful,P_i) ∧ manner(end(E),forceful,P_j) ∧ contact(end(E),P_i,P_j)
Resultative	A V P Adj	manner(during(E),directedmotion,A) ∧ manner(end(E),forceful,A) ∧ contact(end(E),A,P) ∧ Pred(result(E),P)
Resultative	A V P Adj with I	manner(during(E),directedmotion,I) ∧ manner(end(E),forceful,I) ∧ contact(end(E),I,P) ∧ Pred(result(E),P)
Resultative	A V P PP	manner(during(E),directedmotion,A) ∧ manner(end(E),forceful,A) ∧ contact(end(E),A,P) ∧ Pred(result(E),P)
Resultative	A V P PP with I	manner(during(E),directedmotion,I) ∧ manner(end(E),forceful,I) ∧ contact(end(E),I,P) ∧ Pred(result(E),P)
Conative	A V at P	manner(during(E),directedmotion,A)
Conative	A V at P with I	manner(during(E),directedmotion,I)
With/against alternation	A V I against/on P	manner(during(E),directedmotion,I) ∧ manner(end(E),forceful,I) ∧ contact(end(E),I,P)
Body-part object or reflexive object	A V I[+body-part/+refl]	manner(during(E),directedmotion,I) ∧ manner(end(E),forceful,I) ∧ contact(end(E),I,?)
Body-part object or reflexive object	A V I[+body-part/+refl] against/on P	manner(during(E),directedmotion,I) ∧ manner(end(E),forceful,I) ∧ contact(end(E),I,P)
Transitive	I V P	manner(during(E),directedmotion,I) ∧ manner(end(E),forceful,I) ∧ contact(end(E),I,P)
Resultative	I V P Adj	manner(during(E),directedmotion,I) ∧ manner(end(E),forceful,I) ∧ contact(end(E),I,P) ∧ Pred(result(E),P)
Resultative	I V P PP	manner(during(E),directedmotion,I) ∧ manner(end(E),forceful,I) ∧ contact(end(E),I,P) ∧ Pred(result(E),P)

Figure 2: Example entry for the *Hit* class

Transfer of a Message - level 1 class

⟨⟨MEMBERS⟩⟩	$[\langle ask, 1\rangle, \langle cite, 1\rangle, \langle demonstrate, 1\rangle, \langle dictate, 1\rangle, \ldots]$
⟨⟨THEMATIC ROLES⟩⟩	Agent(A), Recipient(R), Theme(T)
⟨⟨SELECT RESTRICTIONS⟩⟩	Agent[+animate], Recipient[+animate], Theme[+message]

⟨⟨FRAMES and PREDICATES⟩⟩

Transitive with Theme	A V T	transfer_info(during(E),A,?,T)
Theme and Recipient	A V T to R	transfer_info(during(E),A,R,T)

Transfer of a Message - level 2 class

⟨⟨PARENT⟩⟩	Transfer of a Message - level 1
⟨⟨MEMBERS⟩⟩	$[\langle ask, 1\rangle, \langle dictate, 1\rangle, \langle quote, 1\rangle, \langle read, 1\rangle, \langle show, 1\rangle, \langle teach, 1\rangle, \langle tell, 1\rangle, \langle write, 1\rangle]$

⟨⟨FRAMES and PREDICATES⟩⟩

Ditransitive	A V R T	transfer_info(during(E),A,R,T)

Figure 3: Example entries for the *Transfer of a Message - levels 1* and *2* classes

Figure 2 shows the entry for the *Hit* class. This class allows for three thematic roles: Agent, Patient and Instrument, with constraints on these roles.[1] The Agent is generally animate; the Patient, concrete; and the Instrument, concrete but inanimate. These selectional restrictions refer to a feature hierarchy where animate subsumes animal and human, concrete subsumes both animate and inanimate, and so forth. The representation does not suffer from some drawbacks of theta role analysis because our roles are not global primitives, but are only meaningful within a class.

The strength of our description comes from the explicit relation between syntax and semantics captured in each entry. The example shows the syntactic frames allowed for the class, with thematic roles as descriptors, which are mapped into predicate-arguments as well as the arguments in a TAG representation. It may be the case that a construction is only possible when a particular lexical item is present; for example, the conative construction[2] needs the lexical item "at", and the *together* reciprocal alternation requires the lexical item "together" to be present. In LTAGs these lexical items would be co-anchors in the initial tree.

Using Levin's description of the allowed frames for the verbs in the *Hit* class, each frame is represented by an ordered sequence of thematic roles. There are two possible simple transitive frames, defined as A V P (Agent Verb Patient) and I V P (Instrument Verb Patient). The *together* reciprocal alternation, as in "John hit the sticks together", requires a plural direct object and the presence of the lexical item "together". The resultative constructions incorporate an adjective(Adj), as in "John kicked the door open", or a prepositional phrase (PP), as in "John kicked the door into an open position". In the *with/against* alternation, in order to have the instrument reading for the direct object, we require a prepositional phrase "against/on Patient", as in "John hit the stick against the fence", where "the stick" is the instrument and "the fence" is the patient.

We use a flat semantic representation in which the semantics of each frame is captured by a conjunction of predicates. Many of the frames allow for a "with Instrument" prepositional phrase, which modifies the arguments of some of the predicates. The basic predicates for the *Hit* class are *manner(during(E),directedmotion,X)*, which specifies that during event E, either the agent or the instrument is in motion, and *manner(end(E),forceful,X)* ∧ *contact(end(E),X,P)*, which captures the idea that at the end of event E the agent or the instrument establishes contact with the patient in a forceful way. The conative construction does not have predicates for contact and manner at the end of the event, since the intended goal of contact is never satisfied. The resultative adds another predicate ($Pred$, instantiated by the particular adjectival or prepositional resultative phrase), which indicates the new resulting state achieved by the patient at the end of the event. Reciprocity for multiple patients selected by a plural direct object is captured by having as many manner predicates as there are patients, since each one is supposedly in motion and forceful. The body-part object and reflexive object alternations are the same as the basic predicates, except that the patient may be unspecified ("John hit his elbow" does not contain any information about what the agent hit his elbow against.)

The hierarchical organization of VerbNet is illustrated in Figure 3. The *Transfer of a Message* verb class 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 Theme (message) as the direct object (Agent Verb Theme), as in "John explained trigonometry", and a frame

[1] These constraints are more like preferences that generate a preferred reading of a sentence. They may be relaxed depending on the domain of a particular application.

[2] In the conative construction there is an intention of a goal during the event, that is not achieved at the end of the event.

for Theme and Recipient (Agent Verb Theme 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 Theme) in addition to the frames and predicates inherited from the parent class. A subset of these verbs (*ask*, *show*, *teach*, *tell*, *write*) can take yet another frame, transitive with Recipient as direct object (Agent Verb Recipient), defining a third level class.

Compositional Semantics

The static description of frames given in the previous section is mapped onto TAG elementary trees, and the semantic predicates are associated with each tree, as was done by Stone and Doran. By using TAGs we get the additional benefit of an existing parser that yields derivations and derived trees from which we can construct the compositional semantics of a given sentence. Kipper et al. (2000), describes the dynamic aspect of VerbNet in greater detail.

Initial trees capture the semantics of the basic senses of verbs in each class. For example, many verbs in the Levin *Run* 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 4 shows the initial trees for the causative and inchoative variants for the *Run* class, along with their semantic predicates.

Predicates are associated with not only the verb trees, but also the auxiliary trees. For example, the predicates for a path prepositional phrase headed by “across” specify that an object is in $motion$ with a path via some other object (the object of the preposition) during the event. The semantics of a sentence is the conjunction of the semantic predicates of the trees used to derive the sentence, in a manner similar to that used by Kallmeyer and Joshi. The ability of verbs to take on extended senses in sentences based on their adjuncts is captured in a natural way by the TAG operation of adjunction and our conjunction of semantic predicates. In cases where the path PP adjoins onto a tree representing a basic motion event, such as for the *Run* class, the $motion$ predicate contributed by the PP is redundant but doesn’t conflict with existing predicates. However, this $motion$ predicate is useful for verb classes such as the *Hit* class, which does not include movement of the direct object as part of the meaning of hit (only sudden contact has to be established); by adjoining a path PP such as “across the room”, we get an extended meaning and change in Levin class membership.

The trees are organized into tree families, as is done in the Xtag grammar (XTAG Research Group 1995). Each thematic role is mapped to an indexed node in the basic syntactic tree. For example, *A V P* maps to a structure such as the transitive tree shown in Figure 1, with Agent mapped to NP_0 and Patient to NP_1. The selectional restrictions of each role are expressed as semantic features in the nodes. The correspondence between thematic roles and indexing of syntactic arguments in TAG trees is preserved within the tree family, 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).

Conclusion

We have presented a class-based approach to building a verb lexicon that makes explicit the close association between syntax and semantics, as postulated by Levin. By using verb classes we capture generalizations about verb behavior and reduce not only the effort needed to construct the lexicon, but also the likelihood that errors are introduced when adding new verbs.

Our use of logical forms gives a detailed semantics for each syntactic frame, so that for an event involving motion, it is possible to know not only that the event has a *motion* semantic component, but also which entity is actually in motion. This level of detail is necessary for applications such as animation of natural language instructions (Bindiganavale *et al.* 2000). Another important contribution is that by dividing each event into a tripartite structure, we permit a more precise definition of the associated semantics. Finally, the operation of adjunction in TAGs provides a principled approach to representing the type of regular polysemy that has been a major obstacle in building verb lexicons.

Acknowledgments

This paper is based on research conducted under CAPES grant 0914-95 and NSF grants IIS-9800658 and IIS-9900297.

References

Bindiganavale, R.; Schuler, W.; Allbeck, J. M.; Badler, N. I.; Joshi, A. K.; and Palmer, M. 2000. Dynamically Altering Agent Behaviors Using Natural Language Instructions. *Fourth International Conference on Autonomous Agents*.

Bleam, T.; Palmer, M.; and Vijay-Shanker, K. 1998. Motion Verbs and Semantic Features in TAG. In *Proceedings of the Fourth TAG+ Workshop*.

Candito, M.-H., and Kahane, S. 1998. Can the TAG Derivation Tree Represent a Semantic Graph? An Answer in the Light of Meaning-Text Theory. In *Proceedings of the Fourth TAG+ Workshop*, 21–24.

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

Dang, H. T.; Kipper, K.; Palmer, M.; and Rosenzweig, J. 1998. Investigating Regular Sense Extensions Based on Intersective Levin classes. In *Proceedings of COLING-ACL98*.

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

Joshi, A. K., and Vijay-Shanker, K. 1999. Compositional Semantics with Lexicalized Tree-Adjoining Grammar: How Much Under-Specification Is Necessary? . In

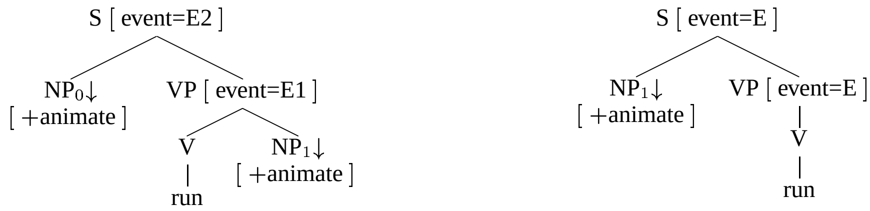

$cause(during(E2), Y, E1) \wedge motion(during(E1), X)$ $\qquad$ $motion(during(E), X)$

Figure 4: Causative/Inchoative alternation for the *Run* verbs

Proceedings of the Third International Workshop on Computational Semantics (IWCS-3), 131–145.

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., eds., *Natural language parsing: Psychological, computational and theoretical perspectives*. Cambridge, U.K.: Cambridge University Press. 206–250.

Kallmeyer, L., and Joshi, A. 1999. Underspecified Semantics with LTAG . In *Proceedings of Amsterdam Colloquium on Semantics*.

Kipper, K.; Dang, H. T.; Schuler, W.; and Palmer, M. 2000. Building a Class-Based Verb Lexicon Using TAGs . In *Proceedings of Fifth TAG+ Workshop*.

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

Lowe, J.; Baker, C.; and Fillmore, C. 1997. A Frame-Semantic Approach to Semantic Annotation. In *Proceedings 1997 Siglex Workshop/ANLP97*.

Miller, G. 1985. WORDNET: A Dictionary Browser. In *Proceedings of the First International Conference on Information in Data*.

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

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

Pustejovsky, J. 1991. The Generative Lexicon. *Computational Linguistics* 17(4).

Saint-Dizier, P. 1999. Alternations and Verb Semantic Classes for French: Analysis and Class Formation. In Saint-Dizier, P., ed., *Predicative Forms in Natural Language and in Lexical Knowledge Bases*. Kluwer Press.

Schabes, Y. 1990. *Mathematical and Computational Aspects of Lexicalized Grammars.* Ph.D. Dissertation, Computer Science Department, University of Pennsylvania.

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

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

XTAG Research Group. 1995. A Lexicalized Tree Adjoining Grammar for English. Technical report, University of Pennsylvania.