

Disambiguating Verbs with the WordNet Category of the Direct Object

Eric V. Siegel

Department of Computer Science
Columbia University
New York, NY 10027
evs@cs.columbia.edu

Abstract

In this paper, I demonstrate that verbs can be disambiguated according to aspect by rules that examine the WordNet category of the direct object. First, when evaluated over a corpus of medical reports, I show that WordNet categories correlate with aspectual class. Then, I develop a rule for distinguishing between *stative* and *event* occurrences of *have* by the WordNet category of the direct object. This rule, which is motivated by both linguistic and statistical analysis, is evaluated over an unrestricted set of nouns. I also show that WordNet categories improve a system that performs aspectual classification with linguistically-based numerical indicators.

1 Introduction

The verb *have* is semantically ambiguous. It can denote a possessive relationship, as in, *I had a car*, or endow a quality, as in, *I had anxiety*. Further, *have* can describe an act of creation, as in, *I had a baby*, or an undertaking, as in, *I had lunch*. Broadly, all uses of *have* either denote a *state*, i.e., a situation that is not dynamic, or an *event*, i.e., a dynamic occurrence that entails change or activity. This semantic distinction, *stativity*, is fundamental to many domains, e.g., distinguishing symptoms and diagnoses from procedures in the medical domain.

Stativity is the first distinction for the semantic hierarchy of verb phrases known as *aspect*. This hierarchy is linguistically established to enable reasoning about time, i.e., *temporal reasoning*. Aspectual classification further distinguishes events according to *completedness* (i.e., telicity), which determines whether an event reaches a culmination point in time at which a new state is introduced. For example, *I made a fire* is culminated, whereas, *I gazed at the sunset* is non-culminated.

Aspectual classification is necessary for interpreting temporal modifiers and assessing temporal entailments (Moens and Steedman, 1988; Dorr, 1992; Klavans, 1994), and is therefore a necessary component for applications that perform certain language interpretation, summarization, information retrieval, and machine translation tasks. Aspectual classification is a difficult problem because many

verbs, like *have*, are aspectually ambiguous.


In this paper, I demonstrate that verbs can be disambiguated according to aspect by the semantic category of the direct object. To this end, WordNet, the largest publicly available on-line lexical database (Miller et al., 1993), is used to provide semantic categories for direct objects. When applied over a corpus of medical reports, I show that WordNet categories correlate with aspectual class. Furthermore, I develop a rule for aspectual classification by the WordNet category of the direct object. This rule is specialized for the verb *have*, which presents a more prevalent disambiguation problem in medical reports than any other verb. The design of this rule is guided by both linguistic and statistical analysis. The rule is evaluated over an unrestricted set of nouns. WordNet categories are also shown to improve a system that performs aspectual classification with linguistically-based numerical indicators.

The following section further discusses the semantic entailments of aspect and Section 3 discusses the problem of aspectual ambiguity. Section 4 describes the corpus used for this study, and Section 5 describes our approach to disambiguating *have*. Section 6 then evaluates this approach and Section 7 describes the use of WordNet for linguistic indicators. Finally, Section 8 provides conclusions and describes future work.

2 Aspect in Natural Language

Aspectual classification is a key component of models that assess temporal constraints between clauses (Moens and Steedman, 1988; Hwang and Schubert, 1991; Dorr, 1992; Hitzeman et al., 1994). For example, stativity must be identified to detect temporal constraints between clauses connected with *when*. For example, in interpreting (1),

(1) She *had* good *strength* when objectively *tested*.
the following temporal relationship can hold between the *have*-state and the *test*-event:




However, in interpreting (2),

Table 1: Four aspectual markers and their linguistic constraints on aspectual class.

If a clause can occur:	then it is:
with a temporal adverb (e.g., <i>then</i>)	event
in <i>progressive</i>	event
with a duration <i>in</i> -PP (e.g., <i>in an hour</i>)	culminated event
in the <i>perfect</i> tense	culm. event or state

(2) She *had* a seizure when objectively *tested*.
the temporal relationship is between two events, and can be different:



Certain temporal adjuncts and tenses are constrained by and contribute to the aspectual class of a clause (Vendler, 1967; Dowty, 1979). Tables 1 lists four example linguistic constraints. Each entry in this table describes an aspectual *marker* and the constraints on the aspectual category of any clause that appears with that marker. For example, if a clause appears in the progressive tense, it must be an event, e.g.,

(3) He is prospering. (event),
which contrasts with,

(4) *You are resembling your mother. (state).

As a second example, an event must be *culminated* to appear in the perfect tense, for example,

(5) She had made an attempt. (culminated),
which contrasts with,

(6) *He had stared at me. (non-culminated)

Such constraints linguistically validate the aspectual hierarchy of semantic classes, provide semantic constraints for natural language generation and understanding, and provide guidelines for aspectual corpus analysis.

3 Aspectually Ambiguous Verbs

While some verbs appear to connote only one aspectual class regardless of context, e.g., *stare* (non-culminated event), many verbs are aspectually ambiguous. For example, *show* denotes a state in, *His lumbar puncture showed evidence of white cells*, but denotes an event in, *He showed me the photographs*. This ambiguity presents a difficulty for automatically classifying a verb because the aspectual class of a clause is a function of several clausal constituents in addition to the main verb (Dowty, 1979; Moens and Steedman, 1988; Pustejovsky, 1991). However, previous work that numerically evaluates aspectual classification has looked at verbs in isolation (Klavans and Chodorow, 1992; Siegel, 1997).

The verb *have* is particularly problematic. In the medical domain, *have* occurs as the main verb of clauses frequently (8% of clauses) and is aspectually ambiguous, occurring 69.9% of the time as a state, and 30.1% of the time as an event. Most other ambiguous verbs are more highly dominated by one sense in this domain (Siegel, 1998).

In this section, I examine factors contributing to aspectual ambiguity. First, I exam the interaction between a verb and its arguments in determining aspectual class. The semantic category of open class words plays a large role in this process. And second, I describe a semantic hierarchy of statively ambiguous verb. This hierarchy groups together verbs that tend to interact with their arguments in similar ways.

3.1 How Clausal Constituents Contribute to Aspectual Class

The presence, syntactic categories, lexical heads, and plurality of a verb's arguments influence aspectual class. This is illustrated in Table 2, which shows example *clausal features* that influence aspectual class. The effect of each feature is illustrated by showing two similar sentences with distinct aspectual classes.

The number of ways in which clausal constituents interactively influence aspect is unknown. However, syntax alone is not sufficient, and the lexical head of multiple constituents (e.g., the verb phrase and the direct object) are often factors. Moreover, the semantic category of these features can also play a role. For example, *Sue played the piano* is non-culminated, while *Sue played the sonata* signifies a culminated event (this example comes from Moens and Steedman (1988)).

3.2 Classes of Ambiguous Verbs

Placing aspectually ambiguous verbs into semantic categories will help predict how these verbs combine with their arguments to determine aspectual class. This is because many verbs with related meanings combine with their arguments in similar ways. In general, there is a correlation between a verb's subcategorization frame and semantic class (Levin, 1993), and this applies to aspect in particular.

For example, *look* and *weigh* can each appear as events, e.g.,

I looked at the baby. (event)

I weighed the baby. (event)

and can also appear as states, as in,

The baby looked heavy. (state)

The baby weighed a lot. (state)

As this illustrates, these two verbs have similar subcategorization frames that determine their aspectual class. There is also a relationship between their meanings, since each describes a type of perception or measurement.

Table 2: Example clausal features and how they can influence aspectual class. “P” means process (i.e., non-culminated event), “C” means culminated event, and “S” means state.

Feature:	Example:	class:	Contrasting Example:	class:
Predicate adj	<i>John drove the car.</i>	P	<i>John drove the car ragged.</i>	C
Particle	<i>John drove the car.</i>	P	<i>John drove the car up.</i>	C
Dir obj cat	<i>John saw Sue.</i>	P	<i>John saw that Sue was happy.</i>	C
Dir obj head	<i>Judith played the piano.</i>	P	<i>Judith played the sonata.</i>	C
Dir obj det	<i>John ate fries.</i>	P	<i>John ate the fries.</i>	C
Ind obj det	<i>Kathy showed people her car.</i>	P	<i>Kathy showed the people her car.</i>	C
Ind obj head	<i>Kathy showed people her car.</i>	P	<i>Kathy showed Sal her car.</i>	C
Prep obj head	<i>Judith looked around the store.</i>	P	<i>Judith looked around the corner.</i>	C
Prep obj det	<i>Kathy shot at deer.</i>	P	<i>Kathy shot at the deer.</i>	C
Tense	<i>Sal said that it helps.</i>	C	<i>Sal says that it helps.</i>	S

Table 3: Groups of verbs that are statively ambiguous.

Group:	Example verbs:	Event sentence:	State sentence:
communication	<i>admit, confirm, indicate, say</i>	<i>I said, “Hello.”</i>	<i>I say it is correct.</i>
cognition	<i>judge, remember, think, wish</i>	<i>I thought about them.</i>	<i>I think they are nice.</i>
perception	<i>feel, see, smell, weigh</i>	<i>I felt the tablecloth.</i>	<i>I felt terrible.</i>
psych-movement	<i>astonish, dismay, please, surprise</i>	<i>You surprised me.</i>	<i>That surprises me.</i>
location	<i>hold, lie, sit, stand</i>	<i>I lay on the bed.</i>	<i>The book lies on the bed.</i>
metaphorical	<i>work, run</i>	<i>I worked hard.</i>	<i>The machine works.</i>
carrier	<i>continue, remain</i>	<i>I continued to talk about it.</i>	<i>I continued to feel good.</i>

Table 3 shows the top level of a hierarchy of statively ambiguous verbs. Seven semantic groups are shown, each with a set of example verbs, and two sentences illustrating contrasting uses of an example verb from that group. Each verb in the first group, **communication**, can appear as either an event or state. Intuitively, this is because each verb can convey a communicative *act*, e.g.,

She showed me the photos. (event)

or, alternatively, a non-dynamic situation, e.g.,

The xrays show no sign of growth. (state)

Verbs in the second group in Table 3, **cognitive**, can convey a mental event, e.g.,

When he mentioned bananas, she remembered Edward. (event)

or, alternatively, a mental state, e.g.,

I’ll always remember Disney World. (state)

The groups **perception** and **psych-movement** are subgroups of **cognition**. The **perception** and **communication** groups have previously been identified with respect to aspect in particular (Vendler, 1967; Dowty, 1979), and those and **psych-movement** for general purposes beyond aspectual ambiguity (Levin, 1993). The fifth group, **locative**, has previously been identified as “lay-verbs.” (Dowty, 1979)

The group **metaphorical** in Table 3 contains **event** verbs with idiomatic uses that are stative. These idiomatic uses correspond to a metaphorical interpretation of the event reading (Alexander D. Chaffee, personal communication). For example,

I ran down the street. (event)

It runs in the family. (state)

Finally, **carrier** verbs simply reflect the aspectual class of their clausal argument.

4 Corpus: Medical Reports

Our experiments are performed across a corpus of 3,224 medical discharge summaries comprised of 1,159,891 words. A medical discharge summary describes the symptoms, history, diagnosis, treatment and outcome of a patient’s visit to the hospital. Aspectual classification is necessary for several medical report processing tasks, since these reports describe events and states that progress over time (Friedman et al., 1995).

These reports were parsed with the English Slot Grammar (McCord, 1990), resulting in 97,973 clauses that were parsed fully with no self-diagnostic errors (error messages were produced on some of this corpus’ complex sentences). Parsing is needed to identify the main verb and direct object of each clause, as well as the presence of aspectual markers for related statistical work, described below in Section 7.

Be and *have* are the two most popular verbs, covering 31.9% of the clauses in this corpus. Clauses with *be* as their main verb, composing 23.9% of the corpus, always denote a state. Clauses with *have* as their main verb, composing 8.0% of the corpus, are statively ambiguous. In this domain, most clauses with main verbs other than *be* and *have* can be aspectually classified by the the main verb only, e.g., by

using numerical linguistic indicators (Siegel, 1998)

In order to produce *supervised* data with which to develop and evaluate our approach, a batch of 206 *have*-clauses from the parsed corpus were manually marked according to stativity. As a linguistic test for marking, each clause was tested for readability with, *What happened was...* In a separate study, a comparison between two human markers using this test to classify clauses over all verbs showed an agreement of approximately 91% (Siegel, 1998). The marked clauses, divided equally into training and testing sets of 103 clauses each, were used to develop and evaluate our approach, respectively.

5 Applying WordNet

I have manually designed a rule for classifying *have*-clauses according to stativity by the WordNet category of the direct object. To design this rule, the following were observed:

- Distributions of objects of *have* over the corpus.
- Linguistic intuition regarding WordNet categories and aspectual class.
- Correlations between the WordNet category of the direct object and stativity over the supervised training data.

To accumulate this information, WordNet was queried for each direct object of the parsed corpus. In particular, each noun was placed into one of the 25 categories at the top of WordNet’s semantic hierarchy, listed in Table 4. Many nouns have multiple entries corresponding to multiple senses. As an initial approach, we take the first WordNet category listed, i.e., the most frequent sense. Pronouns such as *him* and *it* were assigned their own category, **pronoun**.

As shown, in Table 5, the most frequent objects of *have* are primarily specific to the medical domain. This table shows the high level semantic category assigned by WordNet and the classification of *have*-clauses with each noun as a direct object. WordNet is able to handle this technical domain since 89.1% of *have*-clauses have direct objects that are widely-known medical terms and non-technical terms.

The rule shown in Table 6 classifies *have*-clauses based on the semantic category of their direct object. In particular, clauses with direct objects that belong to the categories **event**, **act**, **phenomenon**, **communication**, **possession** and **food** are classified as events, and all others are classified as states.

Linguistic insights guided the design of this rule. For example, if the direct object of *have* denotes an event, such as *seizure*, the clause describes an event. For this reason, it is clear why the WordNet categories **event**, **act**, **phenomenon** and **communication** each indicate an event clause. Note that nominalized event verbs, e.g., *resolution*, are placed in these four categories by WordNet. The category **possession**

WordNet class	as state	as event
location	0	1
event	2	5
act	6	6
artifact	5	3
phenomenon	2	1
entity	2	1
attribute	3	1
measure	3	1
N/A	5	1
cognition	11	1
state	19	0
time	9	0
substance	5	0
relation	3	0
person	2	0
communication	1	0
causalagent	1	0
possession	1	0
group	1	0
food	1	0
shape	0	0
natural object	0	0
feeling	0	0
animal	0	0
plant	0	0
motivation	0	0

Table 4: WordNet categories of direct objects of *have* in the supervised training data.

direct object	n	WordNet class	class of clause
<i>history</i>	624	time	state
<i>episode</i>	280	event	event
<i>pain</i>	192	cognition	state
<i>fever</i>	123	cognition	state
<i>temperature</i>	113	attribute	*state
<i>allergy</i>	109	state	state
<i>movement</i>	106	act	*event
<i>course</i>	96	act	*event
<none>	91	<none>	*state
<i>symptom</i>	81	cognition	*state
<i>complaint</i>	73	state	*state
<i>seizure</i>	72	event	event
<i>nausea</i>	67	cognition	*state

Table 5: Frequent objects of *have*, their WordNet category, and the aspectual class of *have*-clauses with the object. Asterisks (*) denote classifications that were intuitively derived, since these examples did not occur in the training cases.

was selected since, as shown in Table 6, most occurrences of **possession** as a direct object of *have* are instances of *loss*, e.g., *The patient had blood loss* describes an event. The category **food** was selected to cover idioms such as *The patient had lunch* (event).

Furthermore, this classification rule is quantitatively supported over the supervised training data.

If object is a(n):	then class is:	n	Frequent nouns
act	event	1,157	<i>movement</i> (106) <i>course</i> (96) <i>difficulty</i> (66) <i>scan</i> (61) <i>admission</i> (60)
event		655	<i>episode</i> (280) <i>seizure</i> (72) <i>pulse</i> (28) <i>recurrence</i> (25) <i>onset</i> (24)
phenomenon		242	<i>pressure</i> (52) <i>x-ray</i> (30) <i>flatus</i> (21) <i>response</i> (19) <i>intake</i> (15)
communication		194	<i>sign</i> (25) <i>resolution</i> (22) <i>effusion</i> (18) <i>section</i> (17) <i>electrocardiogram</i> (12)
possession		59	<i>loss</i> (27) <i>amount</i> (15) <i>residual</i> (5) <i>insurance</i> (4) <i>cut</i> (3)
food		17	<i>bun</i> (5) <i>coffee</i> (2) <i>vitamin</i> (1) <i>sugar</i> (1) <i>scotch</i> (1)
cognition	state	1,146	<i>pain</i> (192) <i>fever</i> (123) <i>symptom</i> (81) <i>nausea</i> (67) <i>test</i> (54)
state		875	<i>allergy</i> (109) <i>complaint</i> (73) <i>infection</i> (56) <i>disease</i> (56) <i>problem</i> (40)
N/A		860	<i>echocardiogram</i> (51) <i>hematocrit</i> (41) <i>ultrasound</i> (34) <i>stenosis</i> (29)
time		636	<i>history</i> (624) <i>rhythm</i> (8) <i>past</i> (3) <i>gestation</i> (1)
artifact		415	<i>catheter</i> (20) <i>stool</i> (19) <i>tube</i> (17) <i>output</i> (16) <i>PPD</i> (15)
attribute		349	<i>temperature</i> (113) <i>shortness</i> (46) <i>tenderness</i> (26) <i>level</i> (22) <i>sound</i> (16)
entity		209	<i>chest</i> (20) <i>head</i> (13) <i>abdomen</i> (13) <i>artery</i> (12) <i>shunt</i> (7)
measure		205	<i>count</i> (41) <i>increase</i> (18) <i>bout</i> (15) <i>fall</i> (12) <i>day</i> (9)
substance		182	<i>blood</i> (29) <i>thallium</i> (15) <i>sodium</i> (11) <i>urine</i> (10) <i>fluid</i> (9)
relation		116	<i>change</i> (40) <i>rate</i> (32) <i>function</i> (12) <i>aspirate</i> (5) <i>relationship</i> (3)
person		115	<i>child</i> (13) <i>aide</i> (13) <i>son</i> (8) <i>patient</i> (8) <i>temp</i> (6)
group		84	<i>culture</i> (41) <i>series</i> (7) <i>meeting</i> (6) <i>progression</i> (4) <i>panel</i> (4)
location		49	<i>area</i> (8) <i>post</i> (7) <i>left</i> (6) <i>state</i> (4) <i>lab</i> (4)
feeling		48	<i>appetite</i> (18) <i>relief</i> (7) <i>chill</i> (6) <i>preference</i> (3) <i>feeling</i> (3)
pronoun		39	<i>which</i> (18) <i>this</i> (8) <i>her</i> (4) <i>them</i> (3) <i>it</i> (3)
animal		12	<i>dog</i> (3) <i>paceer</i> (2) <i>pet</i> (1) <i>fetus</i> (1) <i>emu</i> (1)

Table 6: Aspectual classification rule for *have*-clauses. Counts are over all *have*-clauses in the medical reports corpus, from which the supervised training and testing data were extracted.

For each WordNet category, Table 4 shows the distribution of event and stative *have*-clauses with a direct object belonging to that category. As shown, each WordNet category linked to states with our rule occurs at least as frequently in stative clauses as they do in event clauses within the training set, with the exception of **communication**, **possession** and **food**. However, these categories occur only one time each in the training data, which is too sparse to counter linguistic intuition.

6 Results

There is a strong correlation between the WordNet category of a direct object, and the aspectual class of *have*-clauses it appears in. When using the classification rule established in the previous subsection, the WordNet categories that appear more than five times in the supervised test data correctly predict the class of *have*-clauses with an average precision of 82.7%. Specifically, **act** and **event** predict event *have*-clauses 85.7% and 66.7% correctly, respectively, and states are predicted with **artifact** (62.5% precision), **cognition** (88.2%), **state** (93.3%) and **time** (100.0%).

For evaluating the rule’s overall performance, there is a baseline of 69.9% and a ceiling of 84.5% accuracy. The baseline is achieved simply by classifying each clause as a state, since this is the dominant class over the supervised test cases.¹ However,

¹Similar baselines for comparison have been used for many classification problems (Duda and Hart, 1973), e.g., part-of-

	overall acc	States		Events	
		recall	prec	recall	prec
C	84.5%	93.1%	85.9%	64.5%	80.0%
R	79.6%	84.7%	85.9%	67.7%	65.6%
B	69.9%	100.0%	69.9%	0.0%	100.0%

Table 7: Performance of a rule (R) that uses the WordNet category of the direct object to aspectually classify *have*-clauses, versus ceiling (C) and baseline (B) approaches.

this approach classifies all event clauses incorrectly, achieving an *event recall* of 0.0%. The ceiling of 84.5% is the maximum achievable by a rule such as ours since the first WordNet category of the direct object is not always enough to resolve aspectual ambiguity; the same category appears in both stative and event test cases.

Overall classification performance using WordNet categories is greatly improved over the baseline method. As shown in Table 7, an accuracy of 79.6% was achieved. A binomial test showed that this improvement over the baseline is significant ($p < .04$).

An event greater improvement over the baseline is illustrated by the increase in the number of event clauses correctly classified, i.e. *event recall*. As shown in Table 7, an event recall of 67.7% was achieved by the classification rule, as compared to

speech tagging (Church, 1988; Allen, 1995).

the 0.0% event recall achieved by the baseline, while suffering no loss in overall accuracy. This difference in recall is more dramatic than the accuracy improvement because of the dominance of stative clauses in the test set. A favorable tradeoff in recall with no loss in accuracy presents an advantage for applications that weigh the identification of non-dominant instances more heavily (Cardie and Howe, 1997). For example, it is advantageous for a medical system that identifies medical procedures to identify event clauses, since procedures are a type of event.

There are several problematic cases that illustrate limitations to our approach. In particular, lexical ambiguity is misleading for the task of classifying *have*-clauses. For example, *The patient had Medicaid* denotes a state, but WordNet categorizes *Medicaid* as an **act**. Similarly, *PET*, *EMU* and *CAT* are categorized as **animal**, as shown in Table 6. This would be solved by recognizing these as proper nouns or acronyms due to capitalization. However, other ambiguous objects are more difficult to address. For example, *The patient had an enema* describes an **event**, but WordNet lists *enema* as **artifact** before **act**. As another example, *The patient had a urine culture* is an **event**, but WordNet’s first sense of *culture* is **group**. Furthermore, the direct object of 10.9% of *have*-clauses in the medical reports are unknown to WordNet (“N/A”). This includes medical terminology, e.g., *anticonvulsants* and *vitrectomy*, as well as certain expressions parsed by the English Slot Grammar that require further post-processing, such as *between 32 and 29*.

7 WordNet for Linguistic Indicators

Aspectual classification is a large-scale, domain-dependent problem. Although a complete aspectual lexicon of verbs may suffice to classify many clauses by their main verb only, a verb’s primary class is often domain-dependent. For example, while many domains primarily use *show* as an event, its appearances in medical discharge summaries primarily denote states. Therefore, it is necessary to produce a specialized lexicon for each domain.

One statistical approach is to measure *linguistic indicators* over a corpus (Siegel, 1998). These indicators measure how frequently each verb appears with markers such as those in Table 1. For example, a verb that appears more frequently in the progressive is more likely to describe an event than a state (Klavans and Chodorow, 1992). However, this approach attempts to classify verbs independent of their context.

Incorporating additional constituents of a clause could alleviate this problem. For example, indicators could be measured over verb-object pairs. However, since both the main verb and the head of the direct object are open-class categories, indicators would be sparsely measured (*enjoy-turnips* is rare).

To alleviate sparsity, but retain information about

	overall acc	Culm		Non-Culm	
		recall	prec	recall	prec
W	71.1%	81.5%	75.0%	53.1%	62.5%
V	68.5%	86.2%	70.6%	38.1%	61.4%
B	63.3%	100.0%	63.3%	0.0%	100.0%

Table 8: Comparison of indicators computed over the main verb (V), indicators over verb and object’s WordNet category pairs (W), and a baseline (B).

the direct object, we measured indicators over verb-object-category pairs, using WordNet to derive the semantic category of each object. I describe such experiments briefly here; Further details regarding these experiments is given by Siegel (1998).

Fourteen such indicators were evaluated for distinguishing clauses according to completedness over an unrestricted set of verbs and direct objects. A corpus of 75,289 parsed clauses from ten novels was used to measure indicator values. 307 training cases (196 culminated) and 308 test cases (195 culminated) were manually annotated using linguistic tests. Decision tree induction was performed over the training cases to combine the indicators.

Indicators measured over the main verb and direct object category achieved a more favorable recall tradeoff than those measured over the verb only, with comparable performance in accuracy. As shown in Table 8, indicators measured over the main verb and direct object category achieved a non-culminated recall of 53.1%, as compared to 38.1% achieved by the verb-only indicators. The baseline of 63.3% accuracy is achieved by simply classifying every clause as culminated.

8 Conclusions and Future Work

The semantic category of the direct object plays a major role in determining the aspectual class of a clause. To demonstrate this, a rule was developed that uses WordNet categories to classify *have*-clauses according to stativity. When evaluated over an unrestricted set of nouns, this rule achieved an accuracy of 79.6%, compared to the baseline performance of 69.9%. Moreover, a favorable tradeoff in recall was achieved, attaining 67.7% event recall, compared to the the baseline’s 0.0%. More specifically, frequent WordNet categories were shown to predict aspectual class with an average precision of 82.7%. These results are impressive, considering the unresolved semantic ambiguity of direct objects, and the technical terminology of the medical domain.

WordNet categories also improved the classification performance of linguistic indicators for completedness. Although more sparsely measured, the accuracy achieved by indicators measured over multiple constituents is comparable to that of indicators measured over the verb only, with a favorable trade-

off in recall. Therefore, the noise introduced by this more sparse measurement of indicators is more than compensated for by the ability to resolve aspectually ambiguous verbs.

Furthermore, I have derived a semantic hierarchy of statively ambiguous verbs in order to predict verbs' subcategorization frames. This in turn guides the disambiguation of such verbs. Future work will investigate whether rules such as that developed for *have* could apply over multiple verbs that share subcategorization behavior. Additionally, it is possible that WordNet's categorization of verbs could automatically place verbs into these semantic groups.

Finally, disambiguating the direct object according to WordNet categories, e.g., Resnik (1995), would improve the accuracy of using these categories to disambiguate verbs.

Acknowledgements

Kathleen R. McKeown was extremely helpful regarding the formulation of our work and Judith Klavans regarding linguistic techniques. Alexander D. Chaffee, Vasileios Hatzivassiloglou, Dragomir Radev and Dekai Wu provided many helpful insights regarding the evaluation and presentation of our results. James Shaw first brought to my attention that *have* is statively ambiguous, and, along with Eleazar Es- kin and Regina Barzilay, provided useful feedback on an earlier draft of this paper.

This research is supported in part by the Columbia University Center for Advanced Technology in High Performance Computing and Communications in Healthcare (funded by the New York State Science and Technology Foundation), the Office of Naval Research under contract N00014-95-1-0745 and by the National Science Foundation under contract GER-90-24069.

References

- J. Allen. 1995. *Natural Language Understanding*. Benjamin/Cummings, Redwood City, CA.
- C. Cardie and N. Howe. 1997. Improving minority class prediction using case-specific feature weights. In D. Fisher, editor, *Proceedings of the Fourteenth International Conference on Machine Learning*. Morgan Kaufmann.
- K. Church. 1988. A stochastic parts program and noun phrase parser for unrestricted text. In *Proceedings of the 2nd Conference for Applied Natural Language Processing*, pages 136–143.
- B.J. Dorr. 1992. A two-level knowledge representation for machine translation: lexical semantics and tense/aspect. In James Pustejovsky and Sabine Bergler, editors, *Lexical Semantics and Knowledge Representation*. Springer Verlag, Berlin.
- D.R. Dowty. 1979. *Word Meaning and Montague Grammar*. D. Reidel, Dordrecht, W. Germany.
- R. O. Duda and P.E. Hart. 1973. *Pattern Classification and Scene Analysis*. Wiley, New York.
- C. Friedman, G. Hripcsak, W. DuMouchel, S.B. Johnson, and P.D. Clayton. 1995. Natural language processing in an operational clinical information system. *Natural Language Engineering*, 2(1).
- J. Hitzeman, M. Moens, and C. Grover. 1994. Algorithms for analysing the temporal structure of discourse. Technical report, University of Edinburgh.
- C.H. Hwang and L.K. Schubert. 1991. Interpreting temporal adverbials. Technical report, University of Rochester.
- J.L. Klavans and M. Chodorow. 1992. Degrees of stativity: the lexical representation of verb aspect. In *Proceedings of the 14th International Conference on Computation Linguistics*.
- J.L. Klavans. 1994. Linguistic tests over large corpora: aspectual classes in the lexicon. Technical report, Columbia University Dept. of Computer Science. unpublished manuscript.
- B. Levin. 1993. *English Verb Classes and Alternations*. University of Chicago Press, Chicago, IL.
- M.C. McCord. 1990. SLOT GRAMMAR. In R. Studer, editor, *International Symposium on Natural Language and Logic*. Springer Verlag.
- G.A. Miller, R. Beckwith, C. Felbaum, D. Gross, and K. Miller. 1993. Introduction to wordnet: An online lexical database. Technical report.
- M. Moens and M. Steedman. 1988. Temporal ontology and temporal reference. *Computational Linguistics*, 14(2).
- J. Pustejovsky. 1991. The syntax of event structure. *Cognition*, 41(103):47–92.
- P. Resnik. 1995. Disambiguating noun groupings with respect to WordNet senses. In *Third Workshop on Very Large Corpora*, June.
- E.V. Siegel and K.R. McKeown. 1996. Gathering statistics to aspectually classify sentences with a genetic algorithm. In K. Oflazer and H. Somers, editors, *Proceedings of the Second International Conference on New Methods in Language Processing*, Ankara, Turkey, Sept. Bilkent University.
- E.V. Siegel. 1997. Learning methods for combining linguistic indicators to classify verbs. In *Proceedings of the Second Conference on Empirical Methods in Natural Language Processing*, Providence, RI, August.
- E.V. Siegel. 1998. *Linguistic Indicators for Language Understanding: Using machine learning methods to combine corpus-based indicators for aspectual classification of clauses*. Ph.D. thesis, Columbia University.
- Z. Vendler. 1967. Verbs and times. In *Linguistics in Philosophy*. Cornell University Press, Ithaca, NY.