Annotating and Recognizing Event Modality in Text

Roser Saurí and **Marc Verhagen** and **James Pustejovsky**
Lab for Linguistics and Computation
Computer Science Department
Brandeis University
{roser,marc,jamesp}@cs.brandeis.edu

Abstract

Current results in basic Information Extraction tasks such as Named Entity Recognition or Event Extraction suggest that we are close to achieving a stage where the fundamental units for text understanding are put together; namely, predicates and their arguments. However, other layers of information, such as event modality, are essential for understanding, since the inferences derivable from factual events are obviously different from those judged as possible or non-existent. In this paper, we first map out the scope of modality in natural language; we propose a specification language for annotating this information in text; and finally we describe two tools that automatically recognizing modal information in natural language text.

Motivation

Basic tasks of Information Extraction such as named entity recognition, event extraction, or even semantic role labeling, are at the core of a variety of NLP applications, ranging from those strictly oriented to building lexical resources, to others which require some degree of text understanding, such as Question Answering or Summarization tasks.

Current performance levels in those three tasks are encouraging. The best system in the CoNLL-2003 shared tasks of language-independent named entity recognition reported an F1 measure of 88.76% and 72.41% on English and German data respectively (Sang & Meuler 2003). Along similar lines, the winning system in the CoNLL-2005 shared task on Semantic Role Labeling attained an F1 measure of 77.92% (Carreras & Màrquez 2005). On the other hand, current domain independent event extractors are reaching an F1 measure around 80% (e.g., Saurí *et al.* 2005).

Such results suggest that we are close to achieving a basic coverage of these tasks, arriving at a stage where the most elementary elements for text understanding are put together. Nevertheless, there are other layers of information that are fundamental for tasks involving text understanding. Temporal and causal information, for instance, play a significant role in the context of question answering systems aiming beyond factoid questions (Pustejovsky *et al.* 2005). Consider:

(1) a. What happened in Iraq after the Fallujah attack?
b. Why did the US take custody of Nick Berg?

Both questions require identifying a set of events related to the referred attack on Fallujah (for question (1a)) or taking custody of Nick Berg in (1b), and subsequently ordering them temporally. Question (1b), in addition, requires the identification of relations of a causal nature.

Another level of information necessary for text understanding is modality, which expresses the speaker's degree of commitment to the events being referred to in a text. Our use of the term modality here will encompass event factuality as well. Events in natural language discourse can be characterized along a veridicality axis that ranges from truly factual to counter-factual. Between these two extremes, there is a complex spectrum of modal types, expressing:

(2) a. Degrees of possibility: *These results indicate that Pb2+ **may** inhibit neurite initiation by inappropriately stimulating protein phosphorylation by CaM kinase.*
b. Belief: *Chinese analysts **believe** that the United States will continue to provoke North Korea.*
c. Evidentiality: *Subcomandante Marcos **said** that the Mexican government is not interested in putting an end to the conflict.*
d. Expectation: *Hans Blix **wants** the US to allow UN inspectors back into Iraq to verify any weapons found by coalition forces.*
e. Attempting: *George Mallory and Andrew Irvine first **attempted** to climb Everest in 1924.*
f. Command: *John Murtha **called** for the immediate withdrawal of U.S. troops from Iraq.*

Modality is a necessary component for representing events in discourse, together with other levels of information such as argument structure or temporal information. Inferences derived from events that have not happened or have only possibly happened, are different from those derivable from events that are judged as factual in nature.

The need for a more sophisticated approach sensitive to this additional level of information is just now becoming apparent in highly domain-oriented disciplines such as bioinformatics (Light, Qiu, & Srinivasan 2004). The modality marker *may* in (2a), for example, has effects on the pathway between bioentities that can be abstracted from the reported data.

Modality information is also starting to be taken into account in more genre-oriented applications, such as Question Answering. For example, several of the systems that participated in the pilot evaluation on Text Entailment, developed under the Question Answering initiative program AQUAINT, attempted to cope with modality information to some degree. Consider example (2e). Any QA system that would disregard the attempting context in which the event of Mallory and Irvine climbing Everest in 1924 is embedded, would be led to erroneously report (3b) as the answer to question (3a):

(3) a. When did George Mallory and Andrew Irvine first climb Everest?
b. #In 1924.

The present paper aims at identifying the scope of modality in natural language and proposes a solution for its automatic identification. To that end, the following section will review the most common grammatical resources underpinning modality information (focusing on English). Then, we will introduce TimeML, a specification language suitable for annotating it, and finally we present EvITA an SlinkET, two tools developed specifically for recognizing and annotating modality in text and discourse.

Modality in Natural Language

Event modality in natural language is marked by a variety of different strategies and constructions. In English, these include both lexical items and syntactic constructions.

Lexical modality markers:

At the lexical level, modality can be introduced by what we refer to as **Situation Selecting Predicates** (SSPs). These are predicates (either verbal, nominal, or adjectival) that select for an argument denoting and event (or situation) of some sort. Syntactically, they subcategorize for a *that*-, gerundive, or infinitival clause, but also an NP headed by an event denoting noun. Some examples are verbs like *claim, suggest, promise, offer, avoid, try, delay, think*, nouns like *promise, hope, love, request*, and adjectives such as *ready, eager, able*:

(4) a. The Human Rights Committee **regretted** that discrimination against women persisted in practice.
b. Uri Lubrani also **suggested** Israel was **willing** to withdraw from southern Lebanon.
c. Kidnappers **kept** their **promise** to kill a store owner they took hostage.

SSPs are interesting because part of their lexical semantics is projected as modality information onto the event denoted by its argument (underlined in examples (4)) by syntactic means. The event denoted by the argument is then marked as:

- **Not totally certain**: This is the case of the complements to the so-called weak assertive predicates (Hooper 1975), such as *think, believe*, and *suppose*;
- **Certain according to a source**: Such is the case of complements of reporting predicates (Bergler 1992);
- **Factual**: Complements of *regret* and *forget* (Kiparsky & Kiparsky 1970; Karttunen 1970, 1971);
- **Counterfactual**: Arguments of *avoid* and *prevent*;
- **Possibility**: Arguments of volition and commitment predicates, among others.

Also at the lexical level, there are **modal auxiliaries** of possibility (5a), obligation (5b), necessity (5c), and so on.

(5) a. *could, may*;
b. *must, have to*;
c. *need to*.

Clausal and sentential **adverbial modifiers** may express similar modal information:

(6) a. Possibility: *probably, perhaps*;
b. Frequency: *usually, always*.

Finally, **negative polarity** particles are important because they express the counterfactual nature of the event that is referred to by negated expressions:

(7) a. It became clear controllers could **not** contact the plane.
b. **No** one reached the site in time.

Syntactic modality contexts:

Syntactic structures introducing modality involve the presence of two clauses, generally one embedded within the other. The following list, although not exhaustive, gives an indication of how pervasive this phenomenon is.

- **Relative clauses:** The event denoted by the relative clause (underlined in the following example) is presupposed as true (e.g., *Rice, who became secretary of state two months ago today, took stock of a period of tumultuous change.*)
- **Cleft sentences:** Similarly, the event of the embedded clause is presupposed as true (e.g., *It was Mr. Bryant who, on July 19, 2001, asked Rep. Bartlett to pen and deliver a letter to him.*)
- **Subordinated temporal clauses:** Again, the event in the temporal clause is presupposed as true (e.g., *While Chomsky was revolutionizing linguistics, the rest of the social sciences was asleep.*
- **Purpose clauses:** In this case, the event denoted by the clause is intensional in nature. (e.g., *The environmental commission must adopt regulations to ensure people are not exposed to radioactive waste.*)
- **Conditional constructions:** The event denoted by the consequent clause (underlined) is intensional and dependent on the factuality of the event denoted in the antecedent clause (bold face), which is also intensional (e.g., *On Dec. 2 Marcos promised to return to the negotiating table if the conflict zone was* **demilitarized**.)

A Specification Language for Annotating Modality

In this section, we introduce TimeML as a specification language that is already adequate for annotating modality information in discourse.[1] TimeML aims at capturing the richness of temporal and event related information in language.

[1]TimeML has been developed under the ARDA-funded TARSQI research framework (Temporal Awareness and Reasoning Systems for Question Interpretation).

As such, it has not been specifically designed for representing modality information in text, but for annotating it in the wider context of event and temporal annotation. Given the focus of our work, we will concentrate only on the TimeML portion relevant for encoding event modality. For a complete view of the spec, refer to Pustejovsky et al. 2003a; 2005.

Event information in TimeML

TimeML identifies as events those event-denoting expressions that participate in the narrative of a given document and which can be temporally ordered. This includes all dynamic situations (punctual or durative) that happen or occur in the text, but also states in which something obtains or holds true, if they are temporally located in the text (see Saurí *et al.* (2004) for a more exhaustive definition of the criteria for event candidacy in TimeML).

Event-denoting expressions are found in a wide range of syntactic expressions, such as finite clauses (*that no-one from the White House* **was involved**), nonfinite clauses (*to* **climb** *Everest*), noun phrases headed by nominalizations (*the young industry's rapid* **growth**, *several anti-war* **demonstrations**) or event-referring nouns (*the controversial* **war**), and adjective phrases (*fully* **prepared**).

Event expressions in TimeML are annotated by means of the tag `EVENT`. In addition, TimeML distinguishes between event **types** and event **instances**, which are annotated using the non-consuming tag `MAKEINSTANCE`. As shown in (8a) below, where we focus only on the event denoting expression *cross*, two `MAKEINSTANCE` tags must be created: one referring to the expected event (*expected to cross*), the other one expressing a negative event (*couldn't cross*) :

(8) a. Jeremy Landesberg expected to **cross** the Charles river, but couldn't because of the unexpected rains.

b.
```
Jeremy Landesberg expected to
<EVENT eID="e1">cross</EVENT>
the Charles river, but couldn't
because of the unexpected rains.

<MAKEINSTANCE eiid="ei1" eID="e1"
tense="INFINITIVE"/>
<MAKEINSTANCE eiid="ei2" eID="e1"
modality="could" polarity="NEG"/>
```

Tense, aspect, polarity, and modality of each event instance will be represented in the `MAKEINSTANCE` tag as well.

TimeML and Modality

Modality information in TimeML is encoded within several different tags, depending on its origin. The general rule of thumb is that information originating at the lexical level should be annotated by text consuming tags that span over the lexical item; on the other hand, information triggered by syntactic means is represented by non-consuming tags with attributes pointing to the clauses involved (more specifically, the clausal head).

Modality at the Lexical Level: One of the most relevant lexical modality markers in natural language is the class of **Situation Selecting Predicates** (SSPs). TimeML distinguishes them by means of the attribute `class` within the `EVENT` tag. SSPs will be identified as events of class `I_ACTION` (9a) or `I_STATE` type (9b), depending on the dynamic or stative nature of the event they refer to. However, two subgroups of SSPs, perception and reporting predicates, are classified with the more specific values of `PERCEPTION` (9c) and `REPORTING` (9d), respectively, due to their role in providing evidentiality (either by introducing a witness or an informant source) to support the factuality nature of the subordinated event.

(9) a. Companies such as Microsoft or a combined worldcom MCI are **trying** to monopolize Internet access.

b. Analysts also **suspect** suppliers have fallen victim to their own success.

c. Some neighbors told Birmingham police they **saw** a man running.

d. No injuries were **reported** over the weekend.

Two other lexical-based modality markers in English are **modal auxiliaries** and **negative polarity** particles. As shown in example (8a), they may take scope over only one of the event instances referred to by an event-denoting expression (such as *cross* in (8a)). Thus, polarity and modal information will be captured by the attributes, `polarity` and `modality` in the `MAKEINSTANCE` tag.

Modality at the Syntactic Level: In order to capture modality information brought about at the syntactic level, TimeML provides a non-consuming tag, `SLINK` (for Subordination Link), which makes explicit the relation between the two clauses. In particular, reference to each clause is expressed by a pointer to their respective event heads (through the attributes `eventInstanceID` and `subordinatedEventInstance`), whereas the relation type is conveyed by means of the attribute `relType`.

SLINKs are primarily used for annotating the modality feature that SSPs project onto the events denoted by their subordinated arguments. Furthermore, SLINKs are also employed to represent purpose clauses and conditional constructions. The `relType` value of the `SLINK` tag captures the type of modality projected in each case onto the event denoted by the subordinated clause. It can be any of the following types:

1. `FACTIVE`: When the argument event is entailed or presupposed. Here is an annotated example:[2]

(10)
```
The Human Rights Committee
<EVENT eID="e1" class="I_ACTION">
regretted </EVENT>
that discrimination against women
<EVENT eID="e2" class="ASPECTUAL">
persisted</EVENT>
in practice.

<SLINK eventInstanceID="e1"
subordinatedEventInstance="e2"
relType="FACTIVE"/>
```

2. `COUNTERFACTIVE`: When the SSP presupposes the non-veracity of its argument:

[2] For the sake of simplicity, in this and the following examples we obviate the annotation of `MAKEINSTANCE` tags.

(11)
```
A Time magazine reporter
<EVENT eID="e1" class="I_ACTION">
avoided</EVENT>
<EVENT eID="e2" class="STATE">
jail</EVENT> at the last minute...

<SLINK eventInstanceID="e1"
subordinatedEventInstance="e2"
relType="COUNTERFACTIVE"/>
```

3. EVIDENTIAL: Typically introduced by REPORTING or PERCEPTION events.
4. NEGATIVE_EVIDENTIAL: Introduced by REPORTING and PERCEPTION events conveying negative polarity.
5. MODAL: For annotating events introducing a reference to possible world. This is also the value used for the relation between the event in a purpose clause and the one in the main clause that is being modified.

(12)
```
Uri Lubrani also
<EVENT eID="e1" class="I_ACTION">
suggested</EVENT>
Israel was
<EVENT eID="e2" class="I_STATE">
willing</EVENT>
to
<EVENT eID="e3" class="OCCURRENCE">
withdraw</EVENT>
from southern Lebanon.

<SLINK eventInstanceID="e1"
subordinatedEventInstance="e2"
relType="MODAL"/>
<SLINK eventInstanceID="e2"
subordinatedEventInstance="e3"
relType="MODAL"/>
```

6. CONDITIONAL: For annotating conditional constructions.

Recognizing Modality in Text

In what follows, we introduce EviTA and SlinkET, two tools developed under the TARSQI research framework for automatically identifying and tagging events in text, as well as characterizing them with the appropriate modality features which are triggered by a specific context.

EvITA: A tool for identifying events

EvITA ('Events In Text Analyzer') is an event recognition system aiming at a robust coverage of events in text (Saurí *et al.* 2005). In that sense, EvITA is not limited to any pre-established list of relation types (events), nor is it restricted to a specific domain, contrary to the more standard approach to the event extraction task. That will allow, at a later stage, the use of a subset of the extracted events for identifying certain modality contexts; precisely, those introduced by SSPs.

The functionality of EvITA breaks down into two parts: event identification and analysis of the event-based grammatical features that are relevant for temporal reasoning. Both tasks rely on a preprocessing step which performs part-of-speech tagging and chunking, and on a module for clustering together chunks that refer to the same event.

Event identification in EvITA is based on the notion of event as defined in the previous section. Only lexical items tagged by the preprocessing stage as either verbs, nouns, or adjectives are considered event candidates. Different strategies are used for identifying events in these three categories.

Event identification in verbal chunks is based on lexical look-up, accompanied by minimal contextual parsing in order to exclude weak stative predicates, such as 'be'. For every verbal chunk in the text, EvITA first applies a pattern-based selection step that distinguishes among different kinds of information: the chunk head; the modal auxiliary sequence, if any (e.g., *may have to*); the sequence of *do*, *have*, or *be* auxiliaries, marking for aspect, tense and voice; and finally, any item expressing the polarity of the event. The last three pieces of information will be used later, when identifying the event grammatical features.

The identification of nominal and adjectival events is also initiated by the step of information selection. For each noun and adjective chunk, their head and polarity markers, if any, are distinguished.

Identifying events expressed by nouns involves two parts. First, EvITA uses a Bayesian classifier that was trained on TimeBank1.2[3] and SemCor (Miller *et al.* 1994). The training corpus was built by using all nominal event contexts in TimeBank. Given its limited volume, we then enlarged it with: (a) the SemCor contexts of nominals tagged as events in TimeBank, and (b) the SemCor contexts of nominals tagged with the same WN synset as the nouns selected in the immediate previous step. For the classifier training, the features used were: (1) definiteness of the NP; and (2) number of the NP head.

For those nouns for which there was not enough data in the training corpus, EvITA then applies a lexical look-up to WN . The criterion used here is to take as events those nouns whose most common synset in WN is an event.

Finally, the task of recognizing events that are expressed as adjectives takes a conservative approach; namely , tagging as events only those adjectives that were annotated as such in TimeBank1.2, whenever they appear as the head of a predicative complement. Thus, in addition to the use of corpus-based data, the subtask relies again on minimal contextual parsing which is capable of identifying the complements of copular predicates.

Identifying the grammatical features of events follows different procedures, depending on the part of speech of the event-denoting expression, and whether the feature is explicitely realized by the morphology of such expressions. In event-denoting expressions that contain a verbal chunk, *tense*, *aspect*, and *non-finite morphology* values are directly derivable from the morphology of this constituent, which in English is quite straightforward. Thus, the identification of these features is done by first extracting the verbal constituents from the verbal chunk (disregarding adverbials,

[3]TimeBank1.2 is a gold standard corpus of around 200 news report documents from various sources, annotated with TimeML temporal and event information. A previous version, TimeBank1.1, can be downloaded from http://www.timeml.org/. For additional information see Pustejovsky *et al.* (2003b).

punctuation marks, etc.), and then applying a set of over 140 simple linguistic rules, which define different possible verbal phrases and map them to their corresponding tense, aspect, and non-finite morphology values. Figure 1 illustrates the rule for verbal phrases of future tense, progressive aspect, which bear the modal form *have to* (as in, e.g., *Participants* **will have to be working** *on the same topics*):

```
[form in futureForm],
[form=='have'],
[form=='to', pos=='TO'],
[form=='be'], [pos=='VBG'],
==>
[tense='FUTURE',
aspect='PROGRESSIVE',
nf_morph='NONE']
```

Figure 1: Grammatical Rule

For event-denoting expressions with no verbal chunk, tense and aspect are established as null (`'NONE'` value), and non-finite morphology is `'noun'` or `'adjective'`, depending on the part-of-speech of their head.

Modality and *polarity* are the two remaining morphology-based features identified here. EvITA extracts the values of these two attributes using basic pattern-matching techniques over the approapriate verbal, nominal, or adjectival chunk.

The identification of event *class*, however, cannot rely on linguistic cues such as the morphology of the expression. Instead, it requires a combination of lexical resource-based look-up and word sense disambiguation. At present, this task has been attempted only in a very preliminary way, by tagging events with the class that was most frequently assigned to them in TimeBank1.2. Despite the limitations of such a treatment, the accuracy is fairly promising.

The most recent evaluations of EvITA have been carried out by comparing its performance against TimeBank1.2.1. The current performance of EvITA is at 74.55% precision, 78.61% recall, for a resulting F1-measure of 76.53%. The Accuracy ratio (i.e., the percentage of values EvITA marked according to the gold standard) on the identification of event grammatical features is as shown in Table 1.

Feature	Accuracy
polarity	98.03%
aspect	97.72%
modality	97.04%
nf_morph	92.64%
tense	86.81%
class	79.01%

Table 1: Accuracy of Grammatical Features

SlinkET: An Algorithm for Identifying Modal Contexts

SlinkET (Slink Events in Text) (Saurí, Verhagen, & Pustejovsky 2006) is a tool developed under the TARSQI research framework with the specific goal of automatically identifying and tagging contexts of subordination that involve some type of modality; namely, those contexts that TimeML refers to as SLINKs. In addition to identifying these contexts, SlinkET assigns them the appropriate `relType` value which expresses the modal nature of the subordinated event: `factive`, `counterfactive`, `evidential`, `negative_evidential`, or `modal`, based on the modal force associated with the class of the subordinating event.

Lexically-based SLINKs are triggered by a well-delimited subgroup of verbal, nominal, and adjectival predicates (such as *regret*, *promise* and *be capable*), which select for a situation-denoting argument, expressed as either an embedded clause or an NP. First, lexical information is used for preselecting the candidates that introduce SLINKs. This information is based on corpus-induced knowledge from TimeBank, as well as standard linguistic classifications of such predicates (e.g., Kiparsky & Kiparsky 1970; Karttunen 1970, 1971; Hooper 1975, Bergler 1992, and subsequent elaborations of that work). Next, a finite-state syntactic module identifies the subordinated event in the clause, using subcategorization knowledge of the subordinating event, derived largely from corpus analytics, which is subsequently compiled into normalized dictionary entries. For each event, the dictionary specifies its possible subordinating contexts and its SLINK types. This is critical for disambiguating the modal force of such predicates. For example, *investigate* introduces an SLINK of type `modal` when subordinating an *if/whether*-clause (13a), but an SLINK of type `factive` when subcategorizing for an event-denoting NP (13b):

(13) a. Officials are investigating whether Rudolph **participated** in all three attacks.
b. Officials are investigating all three **attacks**.

Syntactically based SLINKs are introduced by purpose clauses and the head events modified by them. So far only a few of these are identified by SlinkET: those triggered by verbs with a strong tendency to be modified by such structures; e.g., *address*:

(14) The President addressed the nation to **announce** the election.

Both types of SLINKs are derived using lexical and basic structural features:

1. Lexical form.
2. POS tag.
3. Whether the item refers to an event.
4. Finite vs. non-finite morphology, in the case of event-denoting expressions.
5. Subordinating predicate class.
6. Chunking.
7. Sentence boundaries.

Event tags and finite/non-finite morphology are features obtained from EvITA, whereas predicate classes abstract from the mapping between predicate subcategorization patterns and their modal force in each of these patterns, as encoded in the above mentioned dictionary (*investigate*, for example, will belong to two different classes depending on the syntax of its object).

SlinkET uses that knowledge for identifying and wrapping the subordinated event with the appropriate modal information (as illustrated in Figure 2). It is currently rule-based, but we have already started working on modal parsing

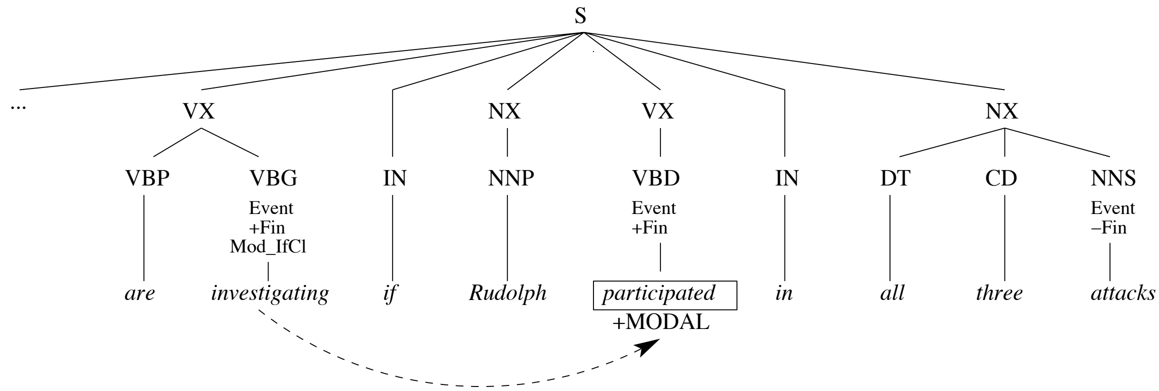

Figure 2: **SlinkET processing**

using machine learning algorithms, i.e., Maxent and Conditional Random Fields.

The current performance of SlinkET has been calculated over 10% of the TimeBank corpus containing a total of 218 SLINKs and 681 events. Precision is at 92%, Recall at 56%, with an F1-measure of 70%. Precision is high, however Recall leaves room for improvement, which will be achieved by enriching the dictionary and adding syntactic patterns to the FSA module.

Conclusion

Event modality is a fundamental information component for NLP tasks requiring some degree of text understanding. We identified the scope of modality in natural language text and proposed a solution for its automatic identification.

TimeML, a well-established specification language for annotating event and temporal information in text, has been shown as descriptively adequate for tagging event modality. Although it does not yet account for some of the syntactic constructions triggering modal information, it has the appropriate data structures for annotating the relevant textual pieces. Namely, text-consuming tags for expressions introducing modality at the lexical level, as well as non-consuming tags for information syntactically triggered.

We have also presented EvITA and SlinkET, two tools for automatically identifying and tagging events in text, as well as identifying any appropriate features of modality which are triggered by the specific context.

Acknowledgements

This work was supported by a grant from ARDA, a U.S. Government entity sponsoring research to the Intelligence Community which includes but is not limited to the CIA, DIA, NSA, NIMA, and NRO.

References

Bergler, S. 1992. *Evidential Analysis or Reported Speech*. Ph.D. Dissertation, Brandeis University.

Carreras, X., and Màrquez, L. 2005. Introduction to the conll-2005 shared task: Semantic role labelling. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*.

Hooper, J. B. 1975. On assertive predicates. In Kimball, J., ed., *Syntax and semantics, IV*. New York: Academic Press. 91–124.

Karttunen, L. 1970. Implicative verbs. In *Language*, 340–358.

Kiparsky, P., and Kiparsky, C. 1970. Fact. In Bierwisch, M., and Heidolph, K. E., eds., *Progress in Linguistics. A Collection of Papers*. The Hague, Paris: Mouton. 143–173.

Light, M.; Qiu, X. Y.; and Srinivasan, P. 2004. The language of bioscience: Facts, speculations, and statements in between. In *BioLINK 2004: Linking Biological Literature, Ontologies, and Databases*, 17–24.

Miller, G. A.; Chodorow, M.; Landes, S.; Leacock, C.; and Thomas, R. G. 1994. Using a semantic concordance for sense identification. In *Proceedings of the ARPA Human Language Technology Workshop*, 240–243.

Pustejovsky, J.; Castano, J.; Ingria, R.; Saurí, R.; Gaizauskas, R.; Setzer, A.; and Katz, G. 2003a. Timeml: Robust specification of event and temporal expressions in text. In *IWCS-5, Fifth International Workshop on Computational Semantics*.

Pustejovsky, J.; Hanks, P.; Saurí, R.; See, A.; Gaizauskas, R.; Setzer, A.; Radev, D.; Sundheim, B.; Day, D.; Ferro, L.; and Lazo, M. 2003b. The timebank corpus. In *Proceedings of Corpus Linguistics 2003*, 647–656.

Pustejovsky, J.; Knippen, B.; Littman, J.; and Saurí, R. 2005. Temporal and event information in natural language text. *Language Resources and Evaluation* 39(2).

Sang, E. F. T. K., and Meuler, F. D. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*.

Saurí, R.; Littman, J.; Knippen, R.; Gaizauskas, R.; Setzer, A.; and Pustejovsky, J. 2004. *TimeML Annotation Guidelines*. http://www.timeml.org.

Saurí, R.; Knippen, R.; Verhagen, M.; and Pustejovsky, J. 2005. Evita: A robust event recognizer for qa systems. In *Proceedings of the HLT/EMNLP 2005*.

Saurí, R.; Verhagen, M.; and Pustejovsky, J. 2006. Slinket: A partial modal parser for events. In *Proceedings of LREC 2006*.