A Framework for Representing Event Semantics of Verb Word Senses

Chris Irwin Davis
University of Texas at Dallas
cid@hlt.utdallas.edu

Abstract

In this paper we present a framework for representing event semantics as a set of semantic entities connected by binary relations. In contrast to frame semantics, our use of a fixed small number of entities and relations facilitates easy decomposition of event semantics into constituent parts, as well as allowing for integration into other systems and resources that rely on binary relations. We map each event representation onto a WordNet verb synset or cluster of related synsets. Thus, event semantics may be indexed by WordNet verb entries. Next, we describe a high-level taxonomy for the categorization of events based upon the semantic roles of the verb arguments. Finally, we briefly discuss acquisition techniques.

1 Introduction

Semantic representation of text is an important aspect of text understanding, reasoning, and identifying inferences. *Semantic relations* are a succinct and formal way to represent semantics in text and are the building blocks for creating the semantic structure of a sentence. In general, semantic relations are unidirectional connections between two concepts or *entities*. For example, the noun phrase "car engine" entails a PART-WHOLE relation: the *engine* is a part of the *car*.

Fillmore (1968) introduced the notion of thematic roles giving a semantic label to the connection between a verb and its arguments. He proposed a set of nine roles: AGENT, EXPERIENCER, INSTRUMENT, OBJECT, SOURCE, GOAL, LOCATION, TIME and PATH (Fillmore 1971).

Recently, Helbig (2005) proposed a classification of semantic entities in order to formally define the sorts of concepts that are valid for a given relation. This classification is performed by an ontology of entities, which defines a hierarchy of concept types following a semantic criteria. For example, in our framework AGENT holds between animate concrete objects (my wife, the president) and situations (arrive, decide) . It is therefore inapplicable to talk about inanimate objects (chair, rock) or abstract objects (yesterday, pain) being the Agent of a Situation.

Much work has been done in the development of resources for the representation of semantics. WordNet (Fellbaum, 1998)[1], FrameNet (Baker et al., 1998)[2], and PropBank (Palmer et al., 2005)[3] are three of the most widely used resources within the research community during recent years.

Our approach is like FrameNet in that we have predefined roles to be filled within a given archetypal semantic representation. In this respect, our semantic entities may be likened to frame elements. However, our use of binary relations allows for the decomposition of an event into constituent parts and the integration with other resources that utilize binary relations.

Like PropBank, our approach is verb focused, and centers on verbs and their arguments. Though unlike PropBank, our framework is not so tightly coupled with the syntactic domain.

Each semantic representation of an event we define is mapped to a WordNet verb synset (or cluster of related verb synsets). In this way we encode our representation within its associated synset as an extension to WordNet. This approach is inspired by the logic form transforms (LFTs) of WordNet glosses that are part of Extended WordNet (Moldovan and Novischi, 2002)[4]. In a similar fashion, we encode our extensions to synsets as a separate but parallel resource to WordNet.

[1] http://wordnet.princeton.edu

[2] http://framenet.icsi.berkeley.edu

[3] http://verbs.colorado.edu/~mpalmer/projects/ace.html

[4] http://xwn.hlt.utdallas.edu

2 Semantic Entities and Relations

The upper ontology for semantic entities and the set of 26 binary relations that we use are taken from Blanco et al. (2010). That work is concerned with combining two semantic relations that share a common argument and inferring a third relation between the remaining two unshared arguments. By contrast, we are representing the semantics of an entire discrete event archetype as the potentially complex combination of several semantic relations.

2.1 Upper Ontology of Entities

The root of the semantic entity ontology is simply Entity, and refers to *anything about which we can say something*. A diagram of the upper ontology of entities is shown in Figure 1.

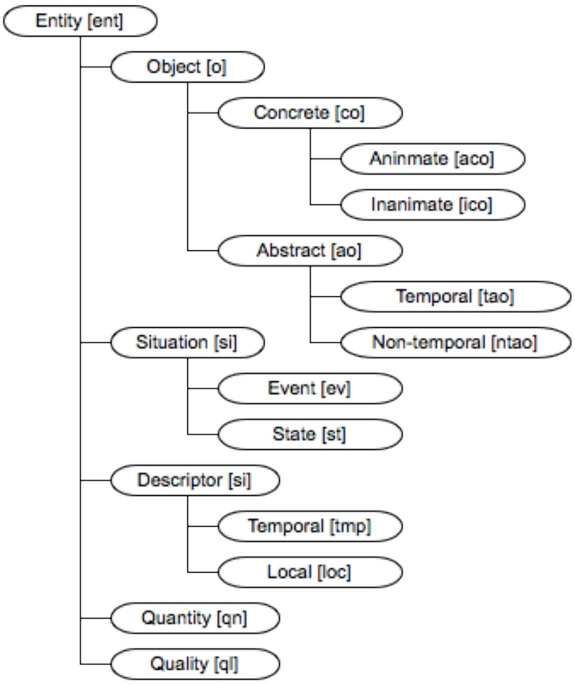

Figure 1. Entity upper ontology

We do not constrain semantic entities by part-of-speech or even to single lexical entries (an Entity may also be a syntactic phrase or clause). We do map WordNet noun synsets to entity types. This facilitates the identification of event arguments, either directly or through headword detection. The types of entity in the upper ontology are:

Objects can be either concrete or abstract. Concrete Objects are tangible things that exist in the physical universe. Abstract Objects are non-physical concepts that exist as a product of cognition. Concrete Objects may be either animate or inanimate. Animate Concrete Objects have agency, whereas Inanimate Concrete Objects do not. Abstract objects are either temporal or non-temporal. The former correlate with ideas regarding points or periods of time (e.g. *1984*, *tomorrow*), whereas the latter may be any other abstraction (e.g. *morality*, *illness*). Abstract objects can be perceived sensually (e.g., *pain*, *aroma*).

Situations are anything that occurs at a time and place. That is, if one can identify the time and location of an Entity, then it is a Situation. Events (e.g. *learn*, *dissolve*) imply a change in the status of other entities, States (e.g. *lying down*) do not. Although Situations can be expressed by either verbs or nouns, our framework focuses upon verbs.

Descriptors express temporal or spatial properties about entities. They may include an optional non-content word (e.g. a preposition) that indicates the temporal or spatial context in relation to another entity.

Qualities are qualitative properties than can be associated with entities. They can be either relative, (e.g. *wealthy*, *small*) or absolute, (e.g. *awake*, *invisible*).

Quantities are quantitative properties of entities (e.g., *750ml*, *a couple of dollars*).

2.2 Binary Semantic Relations.

A binary semantic relation is a relationship between two Entities that is expressed REL(x, y) and may be read "x is REL of y". We constrain the arguments of a binary semantic relation by entity type. In addition to having a more clear and concise definition for each semantic relation, defining the types of concepts that can be part of the DOMAIN and RANGE of a relation has several advantages: (a) helping to discard potential relations that do not hold. For example, inanimate objects cannot have INTENT. (b) aiding in the combining of semantic relations. By checking domain and range compatibilities, valid combinations of relations can be determined. The complete list of 26 relations is depicted in Table 1 along with DOMAIN and RANGE restrictions and examples of valid arguments.

3 Framework Description

In this section we describe the framework for combining the binary relations and semantic entities. We also expand the usage of three important entities from the upper ontology: State, Event, Non-temporal Abstract Object, and the ways they may interconnect in addition to binary relations.

Cluster	Relation Type	Abbreviation	Domain × Range	Examples
Reason	CAUSE	CAU	[si] × [si]	CAU (*virus*, *influenza*)
	JUSTIFICATION	JST	[si ∪ ntao] × [si]	JST (*it is illegal*, *not speeding*)
	INFLUENCE	IFL	[si] × [si]	IFL (*missing classes*, *poor grade*)
Goal	INTENT	INT	[si] × [aco]	INT (*teach*, *professor*)
	PURPOSE	PRP	[si ∪ ntao] × [si ∪ co ∪ ntao]	PRP (*storage*, *garage*)
Object Modifiers	VALUE	VAL	[ql] × [o ∪ si]	VAL (*smart*, *kids*)
	SOURCE	SRC	[loc ∪ ql ∪ ntao ∪ ico] × [o]	SRC (*Spanish*, *student*)
Syntactic Subjects	AGENT	AGT	[aco] × [si]	AGT (*John*, *bought*)
	EXPERIENCER	EXP	[o] × [si]	EXP (*John*, *heard*)
	INSTRUMENT	INS	[co ∪ ntao] × [si]	INS (*the hammer*, *broke*)
Direct Objects	THEME	THM	[o] × [si]	THM (*a car*, *bought*)
	TOPIC	TPC	[o ∪ si] × [si]	TPC (*agenda*, *discuss*)
	STIMULUS	STI	[o] × [si]	STI (*symphony*, *heard*)
Association	ASSOCIATION	ASO	[ent] × [ent]	ASO (*salt*, *pepper*)
	KINSHIP	KIN	[aco] × [aco]	KIN (*John*, *his wife*)
None	IS-A	ISA	[o] × [o]	ISA (*sedan*, *car*)
	PART-WHOLE	PW	[o] × [o] ∪ [l] × [l] ∪ [t] × [t]	PW (*handlebar*, *bicycle*)
	MAKE	MAK	[co ∪ ntao] × [co ∪ ntao]	MAK (*cars*, *BMW*)
	POSSESSION	POS	[co] × [co]	POS (*Ford F-150*, *John*)
	MANNER	MNR	[ql ∪ st ∪ ntao] × [si]	MNR (*quick*, *delivery*)
	RECIPIENT	RCP	[co] × [ev]	RCP (*Mary*, *gave*)
	SYNONYMY	SYN	[ent] × [ent]	SYN (*a dozen*, *twelve*)
	AT-LOCATION	AT-L	[o ∪ si] × [loc]	AT-L (*party*, *John's house*)
	AT-TIME	AT-T	[o ∪ si] × [tmp]	AT-L (*party*, *last Saturday*)
	PROPERTY	PRO	[ntao] × [o ∪ si]	PRO (*height*, *John*)
	QUANTIFICATION	QNT	[qn] × [o ∪ si]	QNT (*a dozen*, *eggs*)

Table 1. The 26 Binary Relations

3.1 State

We augment the usage of the State entity by allowing it to be associated with a semantic relation to indicate the State where a particular relation holds. For example, a state *s* may be associated with the binary relation POS(*x*, *y*) to indicate the State constituted by *x* being possessed by *y*.

3.2 Non-temporal Abstract Objects

Like the State entity, we also allow Non-temporal Abstract Objects to be associated with a semantic relation. This indicates a concept that is the object of a cognitive process (i.e. *thought*, *idea*, *belief*). That is, such a Non-temporal Abstract Object is the cognitive *concept* of the associated semantic relation.

We have noted that non-temporal abstract objects may also be entities that are sensually perceived (pain, odor, fear). We put a finer point on this by further identifying cognitive perceptions (*idea*, *belief*, *thought*). To represent these, we allow non-temporal abstract objects to be associated with a semantic relation to indicate the cognitive *concept* of that relation as an entity.

3.3 Events

We define an Event as either an ongoing, continuous change in status of an entity (e.g. *growing*, *rotating*), or the transition of one discrete State to another State. In the latter case, we note additional properties of a State transition related to lexical aspect: Durativity and Telicity. A State changing Event may be durative and take place over a time period (e.g. *drive*, *eat*) or non-durative (e.g. *sneeze*, *hit*). Telicity indicates whether an Event has a defined goal or completion. "Repaired cars for a week" is atelic, while "repaired *a* car last week" is telic. These event properties and the terms associated with their intersection are summarized in Table 2.

	Durative	***Non-durative***
Telic	Accomplishment	Achievement
Atelic	Activity / Process	Semelfactive

Table 2. Event lexical aspect properties

3.4 Semantic Widgets

To more easily illustrate the event semantics of our framework we introduce the notion of *semantic widgets*, an informal and convenient way of graphically viewing semantic relationships. Our use of the term *widget* is suggestive of how it may be combined with other compatible widgets in the representation of larger blocks of text. Figure 2 shows an illustration of the semantic relationships entailed by the verb "find", whose WordNet synset is "find, regain". This example demonstrates the major components of the framework.

Entities are represented by encircled letter variables. The negation character "!" is a shorthand way of indicating a distinct entity that is of the same type, but whose value is disjoint with its counterpart. In Figure 2, k and $!k$ indicate two different States, where if $k = s_1$ and $!k = s_2$, then $s_1 \neq s_2$.

Relations are symbolized by a directional shape that indicates the order of the arguments.

Events that illustrate the transition of one State to another have an equivalence double line connecting it to a triangle. In this example, the Event is z. The triangle is in turn connected to the two States, $!k$ and k, and shows the direction of the transition.

Finally, the non-temporal abstract object c, is the *concept* of y being at location p.

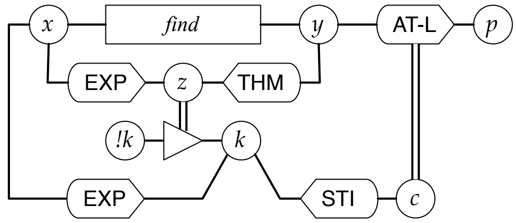

Figure 2. Semantic widget for the verb "find"

To get a better sense for the distinction between semantic frames and our semantic representation of archetypal events, we present a specific example. The FrameNet analog for the representation in Figure 2 is the verb entry `find.v`, a member of the LOCATING frame whose definition states "A [PERCEIVER] determines the [LOCATION] of a [SOUGHT_ENTITY] within a [GROUND]". An example sentence from the LOCATING frame is shown in Figure 3.

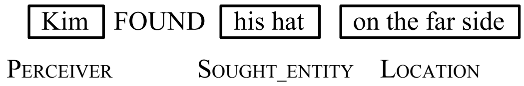

Figure 3. FrameNet example for LOCATING frame.

Figure 4 then shows the same sentence annotated with typical binary semantic relations.

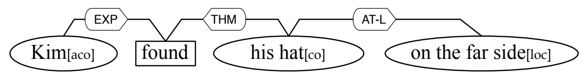

Figure 4. FrameNet LOCATING frame example sentence with typical semantic relations.

Figure 5 shows the same FrameNet example from the LOCATING frame represented using our framework. We do not account for the FrameNet concept of GROUND. However, we do account for more detail in the semantic nature of the event itself. The "found" event is an *achievement*, both telic (a realized goal) and non-durative (the transition from not found to found is instantaneous).

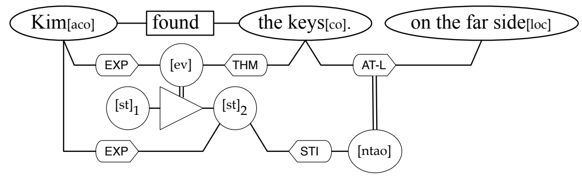

Figure 5. FrameNet LOCATING frame example sentence in our framework.

We show another widget in Figure 6 to illustrate how States may be complex, being associated with multiple relations simultaneously.

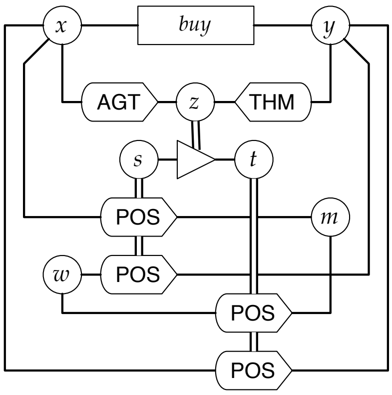

Figure 6. Example of an event with complex states.

Here we see that the possession of a good *y* and of monies *m* are exchanged by a buyer *x* and seller *w*.

3.5 Representation Encoding

The example widget diagram from Figure 2 is more formally encoded as below. Each element has been placed on a separate line for clarity.

```
find(x,y) ->
        EXP(x,z) &
        THM(y,z) &
        EXP(x,k) &
        STI(c,k) &
        x[aco] &
        y[co] &
        k[st] &
        z[ev]=(!k,k) &
        p[loc] &
        c[ntao]=AT-L(y,p)
```

It indicates the valid entity type for each variable, binary relations on those variables, and Event and Non-temporal Abstract Object associations. This is the form used in our resource that is mapped to WordNet verb synsets.

We indicate the entity typing of a particular variable with a bracketed suffix, e.g. *x*[aco]. This indicates that the entity variable *x* is an Abstract Concrete Object.

3.6 Taxonomy

For our taxonomy, we categorize by pairs of binary relations based upon verb arguments, one each from the Subject and Direct Object clusters of semantic relations. These relations share a common semantic Situation denoted by the variable *z* in Figures 2 and 6.

We write these categories in shorthand by the abbreviation for the Subject and Direct Object relations combined with a plus sign. Thus, EXP+THM and AGT+THM would symbolize the categories for the event structures in Figures 2 and 6 respectively.

Those relations that are compatible are shown in Table 3 with a check mark. Those combinations without a check are syntactic combinations that do not hold semantically. We recognize seven high-level categories.

		Direct Object		
		THM	***TPC***	***STI***
Subject	***AGT***	√	√	√
	EXP	√		√
	INS	√		√

Table 3. Event categories

4 Acquisition and Evaluation

Since event attributes may be inherited through associated synset hypernymy, acquisition of event semantics can be automatically propagated down the WordNet hierarchy. We therefore used semantic coverage created manually at the top of a synset tree to automatically seed constituent hyponyms.

Additionally, we note that some languages have properties that better allow for the automatic detection of particular semantic features than others. For instance, there are linguistic features of Modern Persian that allow for the automatic detection of lexical aspect (Folli, et al., 2003).

Farreres, et al. (2010) provide a theoretical foundation for mapping WordNet synsets between languages. Using standard bilingual alignment techniques (Och, 1999) and (Och and Ney, 2000), we mapped synsets between English and Persian, propagating (a) lexical aspect features back to English that were automatically acquired from Persian verbs, and (b) semantic entity information from Persian nouns.

	Precision	**Recall**	**F-measure**
Verb	69.23%	52.94%	60.0%
Noun	84.0%	53.85%	65.63%

Table 4. English-Persian synset mapping accuracy

Table 4 shows an evaluation of our English-Persian WordNet synset mapping of verbs associated with semantic event representations and nouns associated with semantic entities.

References

Baker, Collin F., Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. In *Proceedings of the 17th international conference on Computational Linguistics*, Montreal, Canada.

Blanco, Eduardo, Hakki C. Cankaya and Dan Moldovan. 2010. Composition of Semantic Relations: Model and Applications. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, Beijing, China.

Farreres, Javier, Karina Gibert, Horacio Rodríguez, and Charnyote Pluempitiwiriyawej. 2010. Inference of lexical ontologies. The LeOnI methodology. *Artificial Intelligence*, Vol. 174, No. 1. (25 January 2010), pages 1-19.

Fellbaum, Christiane. 1998. *WordNet: An Electronic Lexical Database*. MIT Press

Fillmore, Charles J. 1968. The case for case. *Universals in Linguistic Theory*, pages 1–88, April.

Fillmore, Charles J. 1971. Some problems for case grammar. *Monograph Series on Languages and Linguistics*, 24:35–36.

Folli, Raffaella, Heidi Harley, and Simin Karimi. 2003. Determinants of event type in persian complex predicates. In *Cambridge Working Papers in Linguistics*, pages 101–126.

Helbig, Hermann. 2005. *Knowledge Representation and the Semantics of Natural Language*. Springer, 1st edition.

Moldovan, Dan and Adrian Novischi. 2002. Lexical Chains for Question Answering. In *Proceedings of the 19th International Conference on Computational Linguistics (Coling 2002)*. Taipei, Taiwan.

Moldovan, Dan I. and Vasile Rus. 2001. Logic Form Transformation of WordNet and its Applicability to Question Answering. In *Proceedings of ACL 2001*. Pages 394-401

Och, Franz Josef. 1999. An Efficient Method for Determining Bilingual Word Classes. *Ninth Conference of the European Chapter of the Association for Computational Linguistics*, pages 71-76, EACL'99.

Och, Franz Josef and Hermann Ney. 2000. Improved Statistical Alignment Models. In *Proceedings of ACL 2000*, pages 440-447.

Palmer, Martha, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106