# Putting Semantics into Semantic Roles

**James F. Allen**[1,2] **and Choh Man Teng**[1]

[1]Institute for Human and Machine Cognition, 40 S. Alcaniz St., Pensacola FL 32502, USA

[2]Department of Computer Science, University of Rochester, Rochester NY 14627, USA

jallen@ihmc.us, cmteng@ihmc.us

## Abstract

While there have been many proposals for theories of semantic roles over the years, these models are mostly justified by intuition and the only evaluation methods have been inter-annotator agreement. We explore three different ideas for providing more rigorous theories of semantic roles. These ideas give rise to more objective criteria for designing role sets, and lend themselves to some experimental evaluation. We illustrate the discussion by examining the semantic roles in TRIPS.

## 1 Introduction

Semantic roles play a foundational role in most computational approaches to encoding meaning, yet they remain surprisingly ill-defined. For the most part, a role taxonomy is defined by an informal gloss and some examples. In other cases, semantic roles act purely as convenient names for arguments of a predicate, frame or event, but are otherwise uninterpreted. This paper starts from the belief that if we are going to base a representation on semantic roles, they should have consequences independent of the predicate or event they are used in. In particular, we will explore three different aspects that identify criteria one might want in a theory of semantic roles:

1. *Entailment*: We should be able to identify entailments from a role independent of the type that has such roles

2. *Integration with ontology*: Roles should obey the typical entailments in an ontology (e.g., inheritance of properties from parents in the ontology)

3. *Derivability*: The roles that a type has should be derivable from its semantic properties, as revealed by definitions of the type in sources such as dictionaries.

The first property helps ensure that roles are used consistently in a semantic lexicon, the second ensures consistency with an ontology used for reasoning, and the third evaluates the semantic interdependence of roles and the types they occur in.

This paper will examine a particular role set, the TRIPS roles sketched in Allen & Teng, (2017)[1]. The first two properties are more formal in nature, but the third allows empirical evaluation, namely, whether the roles of an unknown word sense can be derived from its definition.

## 2 Preliminaries

Semantic roles have a long history, originating in linguistics as thematic roles (e.g., Fillmore, 1968; Dowty, 1991) and widely adopted in computational linguistics for semantic representations because of their compatibility with frame-based and graph-based (i.e., semantic networks) representations of meaning.

Very roughly, computational approaches can be divided into two classes based on whether one believes there is a single universal set of roles (e.g., LiRICS (Bunt & Rosemary, 2002; Petukhova & Bunt, 2008), VerbNet (Kipper et al., 2008; Bonial et al., 2011)), or whether one believes each type may identify its own unique roles (e.g., FrameNet (Baker et al., 1998)). Straddling a middle ground is PropBank (Palmer et al., 2005), which uses a universal set of role names, but allows each type to define what their roles mean. Our interest is in defining a universal set of roles across all types.

A key distinction that most frameworks make is between the inner (or core or argument) roles and the outer (or relational or adjunct) roles. The core roles identify objects that are typically required to fully specify the content of the type, while relational roles are typically optional but add additional information. For instance, in

The snow melted into a puddle.

The subject of this sentence is clearly a critical argument to the melting event. In fact, we cannot

---

[1] More detail can be found in the TRIPS LF documentation: trips.ihmc.us/parser/LF%20Documentation.pdf

describe a melting event in a well-formed sentence without mentioning what melted (even if it is vague, as in *something melted*). On the other hand, the result construction realized by *into a puddle* is optional. Typically, roles of time and location are relational roles as they can be applied to almost any event and are usually optional.

But there are differences that go beyond optionality. The core roles semantically are like arguments to a predicate, although they might be optional. The relational roles semantically denote a relationship between an event and some other condition, not an argument. For instance, the phrase *into a puddle* above denotes first a semantic relationship called RESULTING-OBJECT (one of the senses of the word *into*) that takes as one argument the melting event and the other argument the object that the snow becomes, i.e., a puddle. We will elaborate on this later in the paper.

This paper describes three possible ways to establish a more rigorous semantics for a set of semantic roles, and explores each method in terms of the TRIPS semantic roles. In Section 4 we describe the start of an axiomatization of the roles. In Section 5 we describe the integration of the roles with an ontology. And in Section 6 we describe experiments involving derivability, i.e., can the roles that a given verb sense has be derived from its dictionary definition.

## 3   Overview of TRIPS Semantic Roles

The TRIPS core roles are shown in Table 1. These roles are defined to reflect the causal relationship between the role values and the events they are in. Informally, an AGENT of an event necessarily exerts causal influence to bring about the event. An AFFECTED of an event is necessarily affected or changed in some way by the event but does not cause the event. Objects filling the NEUTRAL and FORMAL roles are acausal, i.e., not necessarily causing or changed by the event. NEUTRAL objects are existent and can be created or destroyed. These include both physical (e.g., a box) and abstract (e.g., an idea) objects. In contrast, FORMAL objects have no temporal existence (e.g., a proposition). Finally, EXPERIENCER is a special class for sentient objects in stative events involving perceptual or cognitive states.

Table 1 summarizes the distinguishing features of each of the roles, and shows correspondences with VerbNet and LiRICS. Note the mappings are nowhere near one-to-one, reflecting differing criteria that are used to define each of the role sets. In some cases, the LiRICS role is not completely clear and is marked with a question mark.

Some differences are differences of granularity (cf: the hierarchical roles in Bonial et al. (2001) and Allen (1995)). For instance, the TRIPS AGENT role only requires a causal relationship to the event, and does not require intentionality of the agent (as in LiRICS which distinguishes between its AGENT and CAUSE roles). TRIPS takes this stance as it seems most verbs with agents would also allow non-intentional causes as well. We believe the intentionality distinction is not signaled in linguistic usage, and should be derived based on commonsense knowledge and reasoning. Thus, we do not make the distinction in the role set.

As another example, VerbNet identifies classes of STIMULUS/EXPERIENCER verbs, as in *The clown amused the children.* Roughly, the STIMULUS role plays a causal role similar to AGENT and the EXPERIENCER plays an AFFECTED role where the object is affected in some cognitive fashion. The TRIPS roleset stays at an abstract level of representation, assigning the clown to AGENT and children to AFFECTED. For similar reasons, TRIPS does not have INSTRUMENT as a core role (although there is a sense of the preposition *with* that captures accomplishing actions with a tool – e.g., *He opened the door with a key*).

Note TRIPS and VerbNet agree on the EXPERIENCER role in some cases, namely with stative verbs. In TRIPS, EXPERIENCER only occurs with stative verbs of perception and cognition.

There is a surprising variety in the VerbNet roles corresponding to the TRIPS roles, partly due to the principle in VerbNet that objects should fill the same role across all the different constructions supported by the verb. For instance, in *The horse jumped over the fence*, the horse is assigned a THEME role because *jump* supports another construction where the horse plays that role, as in *He jumped the horse over the fence*. Using the TRIPS criteria, however, *the horse* is clearly an AGENT in the first, and an AFFECTED in the second.

| TRIPS Roles | Properties | Example | VerbNet Roles | Example from VerbNet | LiRICS |
|---|---|---|---|---|---|
| **AGENT** | +causal | **He** pushed the box | AGENT | **Amanda** carried the package | AGENT |
| | | | CAUSER | **The bug** causes the chip to give wrong answers … | CAUSE |
| | | | INSTRUMENT | **The hammer** broke the window | INSTRUMENT |
| | | | STIMULUS | **The clown** amused the children | AGENT? |
| | | | THEME | **The horse** jumped over the fence | AGENT |
| **AFFECTED** | -causal, +affected | He pushed **the box** | PATIENT | Tony bent **the rod** | PATIENT |
| | | | THEME | Carla slid **the books** | |
| | | | DESTINATION | Lora buttered **the toast** | |
| | | | SOURCE | The doctor cured **Pat** of pneumonia | |
| | | | EXPERIENCER | Carrie touched **the cat** | |
| | | | PRODUCT | The contractor will build you **a house** | RESULT |
| | | | RESULT | David dug **a hole** | |
| **NEUTRAL** | -causal, -affected, +existent | I saw **the box** | PIVOT | **Dorothy** needs new shoes | PIVOT |
| | | | THEME | We avoided **the ball** | THEME |
| | | | STIMULUS | I saw **the play** | THEME |
| | | | EXPERIENCER | **I** loved to write. | PIVOT |
| | | | LOCATION | We avoided **the area** | THEME? |
| | | | TOPIC | Ellen said **a few words** | THEME? |
| **FORMAL** | -causal, -affected -existent | I want **to cry** | ATTRIBUTE | He appeared **crazy** | ATTRIBUTE |
| | | | STIMULUS | I loved **to write** | THEME |
| | | | THEME | I needed **to come** | THEME |
| | | | PREDICATE | The bug causes **the chip to give wrong answers …** | THEME |
| | | | RESULT | I forced him **to come** | THEME? |
| **EXPERIENCER** | -causal, -affected, +cognitive | **He** saw the explosion | EXPERIENCER | **I** saw the play | PIVOT |
| | | | AGENT | **The populace** feel that the RIAA has too much power | |

Table 1: Key Core Roles and Correlates in Other Rolesets

Note that based on the definitions of the roles, it is common that the same role appears more than once in a sentence. We distinguish these arguments by attaching numbers to them. Thus, in *The box is touching the table* we have two roles NEUTRAL and NEUTRAL1 for *the box* and *the table* respectively. VerbNet uses a similar scheme and labels these arguments THEME and CO-THEME.

Most of the remaining TRIPS roles are relational roles, which as discussed above relate an event to some other property. Linguistically, relational roles are realized by prepositional phrases and other adverbial constructions. Semantically, a relational role identifies causal temporal relationships between the event and the property denoted by the prepositional phrase. As an example, for the sentence *He pushed the box into the corner,* there is an event (*He pushed the box*) that results in a culmination state (*The box is in the corner*). The key characteristic of a RESULT relation is that the state is caused by the event and starts immediately at the end of the event. Table 2 shows a number of result related roles based on their temporal properties. These three roles differ only in the temporal properties of the caused state.

Note that many cases one might think are SOURCE roles are actually RESULT roles according to their temporal criterion. For instance, in *He lifted the bottle **out of the box,*** the state of being out of the box is true at the end of the event! Using our definitions the SOURCE role seems mostly limited to cases using the preposition *from*.

Note also that the prepositions in these constructions have fully independent word senses, so our representation does not conflate *He put the cup on the box* and *He put the cup in the box*. In contrast, VerbNet assigns *the box* to a DESTINATION role in both and ignoring the preposition.

Other relational roles, which we will not have the space to discuss here, correspond relatively

| TRIPS Role | Distinctive Properties | Example | VerbNet Role | Example from VerbNet | LiRICS Role |
|---|---|---|---|---|---|
| **RESULT** | E causes R to become true at end of E | He pushed the box **inside the closet** | DESTINATION | Amanda carried the package **to New York** | FINAL LOCATION |
| | | | PRODUCT | I kneaded the dough **into a loaf** | RESULT |
| | | | RESULT | Tony bent the rod **into a U** | RESULT |
| **SOURCE** | E causes R to become not true at start of E | He pushed the box **from the shelf** | INITIAL LOCATION | The book slid **from the table** | INITIAL LOCATION |
| | | | SOURCE | The thief stole the paint **from the museum** | SOURCE |
| **TRANSIENT-RESULT** | E causes R to be transiently true some time during E | He walked **by the school** | TRAJECTORY | Carla slid the books **across the table** | ATTRIBUTE |
| | | | LOCATION | He jumped the horse **over the fence** | THEME |

Table 2: Some Result Related Relational Roles and Correlates in Other Rolesets. E denotes the event and R denotes the role under discussion.

well to similar roles in VerbNet, LiRICS and PropBank (e.g., LOCATION, TIME, MANNER, EXTENT, FREQUENCY, …).

## 4 Axiomatizing Roles

Given the space constraints we cannot present a full axiomatization of the role set. To give a flavor of the axiomatization we look at one core role in particular, namely the AFFECTED role. This introduces most of the formal framework that is used to define all the roles. In addition, we will show the axiomatization of one of the key relational roles, namely the RESULT role.

### 4.1 The Framework

We start from the formalism developed in Allen & Teng (2013), extended from the interval temporal logic based framework in Allen & Ferguson (1994) and Allen (1984). In this framework, both events and property predicates are reified (cf. Davidson, 1967) with functional relations capturing semantic roles and arguments. Terms, rather than predicates, are temporally qualified. For example, $x@t$ represents "object $x$ over time $t$".

Objects filling roles of events are temporally situated. For example, *Jack lifted the ball* (over interval $t_1$) is represented as

$\exists e.(LIFT(e), time(e)=t_1,$
$AGENT(e)=jack_1@t_1,$
$AFFECTED(e)=ball_1@t_1)$

Key to this framework is a theory of scales. For example, $height(o@t)$ maps a temporally situated object to the set of values on the *height* scale that

this object takes over period $t$. Note that the object $o$ may take different values on a scale $sc$ over a given time interval $t$. Thus, $sc(o@t)$ is a *set* of values. If $sc$ is not a scale applicable to $o@t$, then $sc(o@t)$ is empty. For example, for all time intervals $t$, we have $mood(rock_1@t) = \emptyset$.

Adjectives in natural language are typically represented as *Scale Predicates*, which denote (often but not necessarily convex) subsets of values on a corresponding scale. Two examples are *ScalePred(temperature, Cold)* and *ScalePred(mood, Happy)*. Thus, it is true that *John is Happy today* is written as

$TrueOf(john_1@today, Happy)$.

In this paper, to describe relations between time intervals, we will make use of the *meets* relation, written $t_1:t_2$, and "*during or equal*", written $t_1 \subseteq t_2$, from Allen's temporal relations (Allen, 1983).

For more details, see Allen & Teng (2013).

### 4.2 The AFFECTED Role

In Allen & Teng (2013), existence is taken as a primitive in the formal framework. We will define this as a scale with dichotomous values: *existent(o@t) = true* if $o$ exists over the time interval $t$. An object can go in and out of existence at different times. This includes both physical objects such as tables and chairs as well as some abstract objects such as thoughts and orderings.

For some objects, it does not make sense to talk about their existence. These include propositions, properties, scales and scale values. (For example, hungry, five pounds.) For these objects, $\forall t.$ *exist-*

238

*ent(o@t) = ∅*. Note that the value is ∅, not *false*, meaning the existent scale is not applicable.

Now consider the AFFECTED role of an event. It denotes an object that is casually acted upon in the event by the AGENT of the event and is changed in some way by the event. In "I lifted the box", the box changed location. In "The snow melted", the snow changed state of matter. In "I molded the clay", the clay changed shape.

Each event type (e.g., LIFT, MELT, MOLD) is associated with a specific dimension of change such that every occurrence of a particular event type entails a change in the AFFECTED along that particular dimension. For example, for each MOLD event, the AFFECTED of MOLD changes shape, even though the resulting shape in each instance might differ. This is captured by specifying that all events of the same event type share a scale on which change occurs for the AFFECTED.

We can thus formulate the conditions for AFFECTED as follows. Let *E(e)* denote that *e* is an event occurrence of type *E*. Let *scale(sc)* denote that *sc* is a scale.

$$\forall E \; \exists sc \; \forall e, o, t.$$
$$E(e), \; scale(sc),$$
$$time(e) = t, \; \textbf{AFFECTED}(e) = o@t$$
$$\implies \exists t1, t2.$$
$$t1 \subseteq t, \; t1{:}t2,$$
$$sc(o@t1) \neq sc(o@t2)$$

This says the AFFECTED of *e* undergoes at least one change in a scale dimension *sc* characteristic of the event type *E*. There can be more changes, both along the same scale (e.g., the clay changes shape continuously while being molded), or along other dimensions (e.g., the snow changes both shape and volume while being melted).

In the above, $t_1$ is during or equal $t$, the time interval of the event, whereas $t_2$ is met by $t_1$ but could be during, overlaps, or at the end of *t*. This allows for changes that are intermittent and may not persist to the end of the event (e.g., flicker, wiggle) and also changes that occur only at the end of the event (e.g., the stick snapped). In the latter case, for AFFECTED objects that only come into existence at the end of the event (e.g., I drew a circle), it is possible that $sc(o@t_1) = \emptyset$.

Note that the semantics of the roles are defined with respect to completed events. This is captured by the predicate *E(e)*. For example, for *I was drawing a circle but did not finish*, even though the circle never came into existence, the roles are derived not from the progressive formulation but from the corresponding case in which the event has occurred.

## 4.3 Characteristic Properties of Core Roles

Table 1 lists several properties: *causal, affected, existent*. Each of the core roles can be characterized by a combination of the presence or absence of these properties. The axiomatization of the AFFECTED role indicates how the *affected* property can be defined. The *existent* property is captured by the *existent* scale discussed in the previous section.

The *causal* property is meant to indicate that an object exerts a causal influence. In the transitive case the AGENT causes the event to happen to the AFFECTED (I caused the lifting of the box). The *causal* property is often taken as primitive in other rolesets, but we will outline here how it might be formalized.

We take advantage of the intuition that if *X* causes an event *e*, then if we change *X* in some way we can change the event that occurs. More formally, let $sc_E$ be the scale characteristic of event type *E*, that is, when an event of type *E* occurs, its AFFECTED is changed along the scale dimension $sc_E$. Loosely speaking, when +causal obtains (for an AGENT $o_{ag}$), there exists some scale $sc^*$ such that a change of $o_{ag}$ on the $sc^*$ scale would entail a change of the AFFECTED $o_{aff}$ on the $sc_E$ scale. In other words, there is some property of the AGENT such that when this property is changed, regardless by what means, some property of the AFFECTED will vary accordingly.

## 4.4 Wrinkles

A few further considerations complicate the formulation above.

First, existent objects are constantly undergoing changes. For example, an object typically gets older (on the AGE scale) as time passes, even without being involved with any explicit AGENT or EVENT. Such changes (with their associated scales) would trivially satisfy the conditions in the formulation above, such that almost any (existent) object could be a candidate for being an AFFECTED of any EVENT. We call such innocuous changes *Background Changes* and exclude them from consideration.

Second, for some events although the occurrence of the event typically induces a change along the specified scale dimension, this is not always the case. For instance, for the PUSH event, typically the AFFECTED object changes location but in some cases the object might not move. For example, I pushed the door but it was locked. We need to define a notion of canonical or perhaps

counterfactual change. This is especially pertinent for events such as PREVENT (I prevented the accident) or MAINTAIN (The pump maintained the air pressure), in which the events in question are meant to induce no change in the object being acted upon.

Third, the NEUTRAL role mostly entails the absence of properties: *not* causal and *not* affected. It would seem then that any existent object could fill this role, even if this object bears no relationship to the event in question. Similarly, in LiRICS for example, the THEME role is defined as

> *"Participant in a state or event that is essential to the event taking place or the state being in effect. In an event, a theme does not have control over the way the event occurs and is not structurally changed by the event."* (Schiffrin & Bunt, 2007).

We need to define the notion of being *relevant* or *essential* to an event to select for the proper objects. The same consideration applies to FORMAL and EXPERIENCER roles as well.

The axiomatization above thus represents our first attempts to formalize the semantic role entailments. We expect further refinements as we explore these issues in depth.

### 4.5   The RESULT Relational Role

Intuitively a RESULT relates an EVENT and the eventual state. The RESULT only becomes true at the end of the EVENT. For example, in "The cat slid under the table", the cat was not under the table before or during the sliding event, but at the end of the event the cat is under the table.

In contrast, in many approaches the relationship between the event and the resulting spatial predicate is unclear. For example, VerbNet would treat "the table" as the DESTINATION without explicit representation of the spatial relation (under).

RESULTs are often spatial, but they can also be other general states. For example, in "I wiped the desk clean", the result is that the desk is clean.

Resultative constructions apply to intransitive events (e.g., slide) as well as transitive events (e.g., wipe). For intransitive events, the resulting state pertains to the subject of the event (e.g., In *The cat slid under the* table, the result is that *the cat* is under the table), whereas for transitive events, the resulting state pertains to the object of the event (e.g., from *I wiped the desk clean*, the RESULT is that *the desk* is clean).

It is possible to transform intransitive constructions into transitive constructions. For example, in "The dog barked the cat up the tree", the BARK event is normally intransitive, but in the above re-

sultative construction, the dog is the AGENT of the BARK event, the cat is the AFFECTED and the RESULT is that the cat is up the tree. (For a treatment of this, see Allen & Teng (2017).)

Thus, the RESULT role can be formalized as follows for a transitive event type E.

$$\exists sc \; \forall e, o, t, P.$$
$$\quad E(e), \; scale(sc), \; ScalePred(sc, P),$$
$$\quad time(e) = t,$$
$$\quad \mathbf{AFFECTED}(e) = o@t,$$
$$\quad \mathbf{RESULT}(e) = P$$
$$\Longrightarrow \exists t2. \; t{:}t2,$$
$$\quad\quad TrueOf(o@t2, P),$$
$$\quad\quad \forall t1 \subseteq t. \; \sim\!TrueOf(o@t1, P)$$

For example, for the WIPE event *e* in "I wiped the desk clean", we have *ScalePred(cleanliness, Clean)*, that is, the predicate *Clean* is defined on the *cleanliness* scale. In addition, AFFECTED*(e)* = $desk_1@t$ and RESULT*(e) = Clean*; that is, the desk is clean immediately after the WIPE event but not before.

Similarly, for the intransitive case, the AFFECTED above is replaced by the role filled by the subject, that is, the RESULT is a change in the subject at the end of the EVENT.

Note that an EVENT can have multiple RESULTs. For example, "The cat climbed on the box away from the rising water", in which case the two results combine as a conjunction true immediately after the event.

The scale allowed for each event type constrains the possible interpretations of the text. For example, in "I wiped the table clean", the table can change on the cleanliness scale and thus allows a RESULT construction here. In contrast, in "I wiped the table happy", the mood scale is not applicable to tables, ruling out the RESULT construction. Instead, one would prefer an alternative interpretation using a manner-like construction in which "happy" is used to qualify "I".

## 5   Integration with an Ontology

The second criterion we set for a theory of semantic roles is integration with an ontology. If semantic roles are to have an impact on deep language understanding and reasoning, they should be integrated with an ontology that supports that reasoning and stores commonsense knowledge. There is a large relevant literature concerning roles in semantic networks (e.g., Hayes, 1980; Thomason & Touretsky, 1991). Essentially such roles are functions from a type (for verbs an event type is one sense of the verb) to another object. If a word

| Type | New roles/ _Inherited roles_ | Verbs (often in sub-classes) |
|---|---|---|
| EVENT-OF-ACTION | AGENT | |
| EVENT-OF-AGENT-INTERACTION | AGENT1 _AGENT_ | Meet, collaborate, … |
| EVENT-OF-CREATION | AFFECTED _AGENT_ | Bake, establish, … |
| EVENT-OF-CAUSATION | AFFECTED, _AGENT_ | Push, control, … |
| MOTION | RESULT _AGENT_ _AFFECTED_ | Go, disperse, move, … |
| EVENT-OF-UNDERGOING-ACTION | AFFECTED | Die, inherit, … |
| EVENT-OF-STATE | NEUTRAL | |
| POSSESS | NEUTRAL1 _NEUTRAL_ | Own, possess, … |
| HAVE-PROPERTY | FORMAL _NEUTRAL_ | Be, seem, … |
| EVENT-OF-EXPERIENCE | EXPERIENCER _NEUTRAL_ | Appreciate, believe, … |

Table 3: Some Roles in the Event Ontology (Showing Role Inheritance)

sense $S$ has a semantic role $R$, then for all instances of $S$ there is an object that fills the $R$ role, i.e.,

$$\forall s.\ S(s) \Rightarrow \exists r.\ R(s,r)$$

This axiom captures what is sometimes called an *essential* role (Palmer, 2006), i.e., a semantic role that must exist even if not specified in the input sentence. Palmer notes that other roles are *obligatory*, that is, they are both essential and linguistically required, and still others are simply *optional* and may or may not be realized at either the linguistic or inferential levels.

A critical foundation of semantic networks is inheritance down type hierarchies, where a subtype "inherits" all properties of its parent types. When viewing this from the perspective of semantic roles, this means any role that is essential for a given type $T$ must also then be essential for all subtypes of $T$.

Surprisingly, the predominant models of semantic roles do not address such issues in any depth. VerbNet, for instance, creates its classes based on clustering by verb usage patterns, rather than semantic entailments, and has a very limited hierarchical structure defined in terms of extensions in allowed usages. Within these hierarchies, though, it does have inheritance of roles. Propbank, on the other hand, is word based and even semantically very similar verbs have different rolesets (e.g., compare the rolesets for the verbs *constrict, compress* and *squeeze*, which

one would expect would be clustered together in an ontology).

The TRIPS role set is fully integrated with a rich ontology. In this ontology, the concepts are organized both by entailment as well as the semantic roles the verb senses take. Table 3 shows a small part of the upper ontology for events and the roles that are defined for each type and inherited from ancestor types. Note TRIPS allows both essential and optional roles. Both are inherited down the hierarchy, and lower types can make an inherited optional role essential, but not vice versa. Whether a role is obligatory or not is not specified in the ontology, but rather in the lexicon where the words and the argument structures they allow are defined. Furthermore, it employs explanation closure techniques (Schubert, 1994) – if a role is *not* defined as possible in the ontology then the role is not possible.

The integration of roles with an ontology mutually constrains the assignment of roles to verb sense predicates *and* constrains the ontology itself. For instance, if we believe that the verb *disappear* takes the AFFECTED role and not an *AGENT* role, then its word sense cannot be under the EVENT-OF-ACTION or EVENT-OF-STATE hierarchies, but it looks like a good candidate for being under the EVENT-OF-UNDERGOING-ACTION category. Likewise, although you might think the verb *analyze* might take an EXPERIENCER role, if you believe that *analyze* falls under EVENT-OF-ACTION, then it should take the AGENT role instead.

One other aspect that relates to the ontology is selectional preferences. As in many semantic networks, one can also constrain the semantic type of the arguments that can fill a role. For instance, one might say that the event type EAT typically concerns an animate entity (as AGENT) and some comestible substance (as the AFFECTED role). Such knowledge is critical for driving semantic disambiguation during parsing. The TRIPS restrictions are soft constraints, i.e., the parser prefers interpretations that satisfy the constraints, but can construct interpretations that do not. As with roles, the selectional preferences are inherited down the hierarchy, with more specific event types accumulating all the constraints imposed on their ancestors in the hierarchy. We do not have the space to discuss this further here.

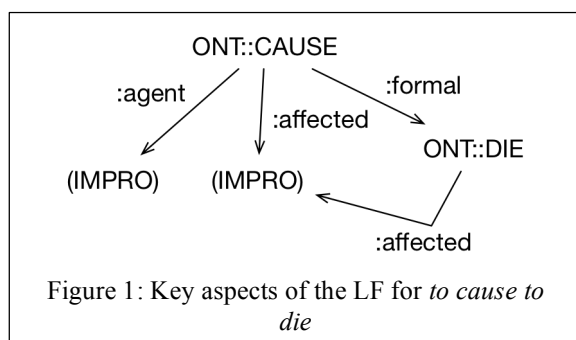| Role Chain | Role Identified | Example Word and Definition | Justification |
|---|---|---|---|
| V $\xrightarrow{R}$ Arg | R | *wolf*: ϕ *eat* ϕ | An unfilled direct role R of the head verb is R |
| V $\xrightarrow{RESULT}$ V1 $\xrightarrow{GROUND}$ Arg | NEUTRAL | *near*: ϕ *move towards* ϕ | The ground of the property is not affected by the verb |
| V$_{change}$ $\xrightarrow{FORMAL}$ V1 $\xrightarrow{AGENT}$ Arg | AFFECTED | *jump*: ϕ *cause [ ϕ to jump]* | The agent of the embedded event is changed by undergoing the cause event |

Table 4: A few sample rules for deriving roles from definitions.

## 6 Derivability in Definitions

The third criterion we explore is derivability. The motivation is as follows: If semantic roles have a semantics independent of their predicates/events, then the semantic properties of roles would not change between a predicate and its definition. For example, consider the word sense corresponding to the predicate for *kill*, defined as *cause to die*. According to our analysis above, *kill* would take two essential roles: AGENT and AFFECTED. The definition, on the other hand, involves a predicate *cause* that takes an AGENT, AFFECTED and a FORMAL role (to die). A highly abbreviated logical form for this definition is shown in Figure 1. While *cause* has three roles, only the FORMAL role is fixed by the definition (i.e., it is the *die* event). The unfilled essential roles in the definition are AGENT and AFFECTED, exactly the roles for *kill*.

Our hypothesis is that given a good definition of a word sense, the essential roles can be derived from that definition automatically. This, of course, has significant impact. If the definitions in WordNet (Fellbaum, 1998) are generally reasonable, we can bootstrap from the items predefined in the TRIPS ontology and lexicon, and derive the semantic roles for any verb in WordNet. Our preliminary evaluation described here indicates that this is a very feasible goal.

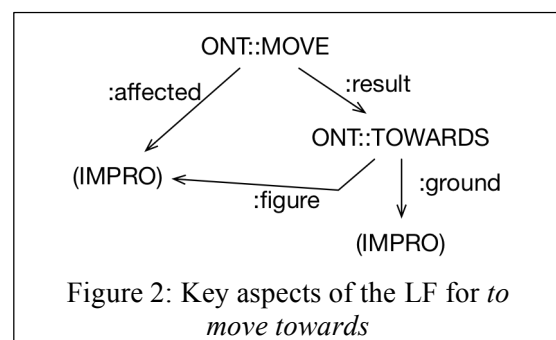To implement this we need a strategy for identifying unfilled semantic roles in a definition. The most common case we have encountered in WordNet definitions is that the unfilled roles are elided in the definition. The TRIPS parser instantiates such gaps in the logical form using its IMPRO construct. Other cases include indicating the unfilled roles by an indefinite pronoun, such as *someone* or *something*.

More complicated cases occur when the roles do not occur at the top level but in an embedded clause, as in one definition of *approach* (*move towards*), shown in Figure 2. The roles for the predicate *Towards* are FIGURE and GROUND, and an IMPRO fills the GROUND role of the predicate. Thus we have a role chain, from the head verb through RESULT and GROUND. We have analyzed such role patterns and created a mapping based on the semantics of the roles. In this case, it would indicate a NEUTRAL role, as the GROUND role is not changed by the event. Table 4 summarizes a few of the most common rules for identifying roles from role chains.

### 6.1 Experimental Evaluation of Derivability

To test this technique, we built a customized TRIPS system to parse definitions. The main customization was the addition of about a dozen top level syntactic rules that capture the common forms of definitions, which as we have seen contain much ellipsis. Otherwise, the grammar and lexicon are exactly the same as in all the other variants of the TRIPS parser. Our experiment also takes advantage of the fact that the TRIPS ontology has an extensive mapping to WordNet synsets, and uses the WordNet hypernym hierarchy to

Figure 1: Key aspects of the LF for *to cause to die*

Figure 2: Key aspects of the LF for *to move towards*

242

identify abstractions. Any WordNet synset can be mapped to its most specific type in the TRIPS ontology (Allen & Teng, 2017).

The experiment was set up as follows. We repeat the following to obtain 40 test cases:

1) Randomly choose a WordNet synset $S$ such that: (a) $S$ has a direct mapping to a TRIPS ontology type $T$, (b) the TRIPS lexicon has word $W$ with type $T$ that is also in $S$, and (c) the WordNet definition is not circular.

2) Remove the lexical entry for $W$ from the TRIPS lexicon after recording its essential roles as the gold standard answer.

3) Parse the definition for $S$ and extract the essential roles as described above.

4) Compare the roles from steps 2) and 3) to compute precision and recall.

For a baseline we assigned each verb the AGENT and AFFECTED roles. Table 5 shows the results of the experiment. We obtained 88% precision and 77% recall for our approach, versus 63% and 66% respectively using the baseline assignments. Because a large number of English verbs are simple transitive verbs describing change, the baseline did better than one might expect. Still our approach based on parsing definitions performed far better, lending strong support that our role set passes the derivability test.

Based on a manual analysis of the errors in this experiment, the errors arose from a combination of parse errors, ambiguities, and definitions that are terse and loose. For instance, one sense of *appear* is defined as *to come into sight*. This is ambiguous between entering into some state (the right interpretation here), and the acquisition reading where *sight* is acquired. The parser chose the latter, leading to an assignment of AGENT as the role rather than AFFECTED. We also missed a few roles because of parser errors. For instance, a sense of *pronounce* is defined as *to declare judgement on* (e.g., *they pronounced him unfit*). The parser failed to identify the second gap (i.e., the missing object of the preposition *on*). In addition, this is the one case we found where our strategy for identifying arguments was inadequate. There is an argument to this verb that is the judgement, but this is not signaled by an elided argument or an indefinite pronoun. Thus the gold answer is AGENT, NEUTRAL, FORMAL but our system was able to identify only the AGENT role.

There are several additional processing steps we could take to improve the performance. For instance, sometimes multiple definitions are presented in WordNet and we could process them all and try to combine them. Currently we only analyze the first one. Also, we could try to verify the roleset by attempting to parse examples that are given in the gloss. Often it seems that the lexicographer depends on the examples to supplement and disambiguate the definitions.

In our error analysis we did not find any example where if we had identified the correct parse of the definition, we would have identified an inappropriate role for the word being defined. This indicates that the semantic roles appear to be consistent across the lexicon and, furthermore, are identifiable by the semantic properties induced by the events in the definitions. In other words, they meet the criterion of derivability!

## 7 Discussion

We have presented three possible criteria for how one can produce and validate a semantics for a semantic role set. We have illustrated the techniques by looking at the roles in the TRIPS framework. Note we do not claim this indicates the TRIPS roleset is the only roleset that could be useful in linguistic theory and computational semantic models. However, we have shown that the TRIPS roleset is internally consistent and has a set of desirable semantic properties: (1) it is amenable to axiomatization in a temporal logic, (2) it is integrated into an ontology that supports inheritance, and (3) the roles are derivable in the sense that they can be derived for verb senses based on their definitions.

The TRIPS roleset was created by considering the properties of causality, temporality, existence, and sentience. Other researchers choose to create rolesets based on other criteria. We encourage those researchers to attempt to formalize their roles along the dimensions we have defined to create a firm theoretical foundation by which all theories can be compared.

### Acknowledgements

|  | Our Approach | Baseline |
|---|---|---|
| Precision | 88.4% | 63.4% |
| Recall | 77.2% | 65.8% |
| F1 score | 82.4% | 64.5% |

Table 5: Results in Deriving Rolesets

# 8 References

Allen, J. F. (1983). Maintaining Knowledge About Temporal intervals. *Communications of the ACM* 26(11):832-843.

Allen, J. F. (1984). Towards a general theory of action and time. *Artificial Intelligence*, 23(2):123-154.

Allen, J.F. (1995). Natural Language Understanding, Benjamin Cummings, Second Edition.

Allen, J. F. and G. Ferguson (1994). Actions and events in interval temporal logic. *Journal of Logic and Computation*, 4(5):531-579.

Allen, J. F. and C. M. Teng (2013). Becoming different: A language-driven formalism for commonsense knowledge. In the 11th International Symposium on Logical Formalizations of Commonsense Reasoning, Cyprus.

Allen, J., and C. M. Teng (2017). Broad coverage, domain-generic deep semantic parsing. *AAAI Workshop on Construction Grammar*, March, Stanford University.

Baker, Collin F., Charles J. Fillmore, and John B. Lowe. (1998). The Berkeley FrameNet Project. In *Proceedings of COLING-ACL'98*, pages 86–90, Montréal, Canada.

Bonial, C., et al. (2011). A Hierarchical Unification of LIRICS and VerbNet Semantic Roles. *5th IEEE Conf. on Semantic Computing (ICSC)*, Palo Alto, CA.

Bunt, H. C. and L. Romary (2002). Requirements on multimodal semantic representations. *Proceedings of ISO TC37/SC4 Preliminary Meeting*, pages 59–68.

Davidson, D. (1967). The Logical Form of Action Sentences. *The Logic of Decision and Action*. N. Rescher. Pittsburgh, University of Pittsburgh Press.

Dowty, D. (1991). Thematic Proto-Roles and Argument Selection, *Language*, Vol 67-3, pp547-619

Fellbaum, C. (1998, ed.) *WordNet: An Electronic Lexical Database*. Cambridge, MA: MIT Press.

Fillmore, Charles (1968) "The case for case", in E. Bach and R. Harms, (eds) *Universals in Linguistic Theory*, Holt, Rinehart, and Winston, New York, 1–90.

Hayes, Pat (1980). The Logic of Frames. In *Reading in Artificial Intelligence,* Margan Kaufman.

Kipper, K., A. Korhonen, N. Ryant, and M. Palmer (2008). A large-scale classification of English verbs. *Language Resources and Evaluation Journal*, vol. 42, pp. 21–40, 2008.

Palmer, M. (2006). *Semantic Processing for Finite Domains,* Cambridge University Press.

Palmer, M., Gildea, D. and Kingsbury, P. (2005). The Proposition Bank: A corpus annotated with semantic roles, *Computational Linguistics Journal*, 31:1.

Petukhova, V. and Bunt, H. (2008). LIRICS Semantic Role Annotation: Design and Evaluation of a Set of Data Categories. *LREC*.

Schiffrin A. and H. C. Bunt (2007). LIRICS Deliverable D4.3. Documented compilation of semantic data categories. http://lirics.loria.fr

Schubert, L.K. (1994). Explanation closure, action closure, and the Sandewall test suite for reasoning about change, *J. of Logic and Computation 4*(5), Special Issue on Actions and Processes, pp. 679-799

Thomason R. and D. Touretsky (1991). Inheritance Theory and Networks with Roles, in *Principles of Semantic Networks,* Morgan-Kaufmann.