

Word Senses, Semantic Roles and Entailment

James F. Allen

University of Rochester and
Florida Institute of Human and Machine Cognition

james@cs.rochester.edu

Abstract

We describe an approach to classifying word senses and semantic roles that focuses on the entailments that can be drawn from sentences. A key observation is that entailments arise from sentences, not word senses, and by taking a compositional approach to deriving entailments, we can significantly simplify the number of senses that are required. In addition, we argue that most uses of semantic roles can be replaced by taking prepositional meaning seriously, resulting in a strongly compositional approach to computing sentence meaning.

1 Introduction

While there has been substantial effort in building lexicons and defining word senses, there is still no consensus on a criteria for distinguishing senses or semantic roles. As a result, existing lexicons, while containing significant amounts of extremely valuable information, fall short as we attempt to construct deeper reasoning systems that can compute entailments from text. One of the first issues is atomization of word senses. Significant effort has been made to explore the various contexts in which words are used and to identify the distinct cases by coining different word senses. As others have observed, however, word senses appear to be in the eye of the beholder, and classifications range from fine grained distinctions such as Wordnet (Fellbaum, 1998) to more coarser grained senses which are refined by their context and common-sense knowledge about the terms (e.g., Pustejovsky, 1995). Others have argued that viewing the problem as enumerating specific senses of words is fundamentally misguided (Hanks, 2000). In this paper we will explore these issues, arguing that many apparent word senses are simply encodings of sentence meanings that can be computed compositionally. We are exploring this in the context of reorganizing a substantial semantic lexicon developed for the TRIPS parser (Allen et al, 2008), with more than 15,000 words indexed into a 2000+ concept ontology, together with mappings from over 80% of

Wordnet into our ontology for computing meaning of words not explicitly in the lexicon.

It seems that an underlying assumption of most computational approaches is that the meaning of a sentence is primarily captured by a specific predicate (corresponding to one of the senses of the main verb). These representations often follow Davidson's (1967) proposal for reifying events, primarily for the convenience of allowing variable arity predicates. Semantic roles are viewed as a way to encode argument positions on the verb. But viewing semantic roles as merely indexes to predicate arguments is a major lost opportunity. If we allow PP and other modifiers to have significant meaning in their own right (as suggested by Dowty 1989, 1990 and Goldberg, 1995, 2006) we can reduce the need for different verb senses as well as concisely encoding information for entailment and common-sense inference.

There are many motivations and intuitions underlying the notion of word sense (cf. Palmer et al, 2007, Navagli, 2009). Given our interest in deep understanding, we will focus on one central criterion—different senses should correspond to different entailments. But it is important to remember that entailments are the result of sentence meanings, not word meanings. In determining necessary senses, we apply Occam's razor, and only introduce senses in cases where the entailments cannot be computed compositionally from the sentence structure.

2 Identifying Primary Word Senses

As an ontology for the predicates in this paper, we will try to use names of Frames in Framenet (Johnson & Fillmore, 2000) whenever possible. Consider the first two senses of the verb *move* in Wordnet, senses that are also distinguished in Framenet.

The plant moved.

Wordnet sense move.1

Framenet frame: MOTION

John moved the plant.

Wordnet sense move.2

Framenet frame: CAUSE-MOTION

In these two lexicons, the difference between these sentences is captured in the verb sense. It is possible, however, to capture the difference using a single sense with additional entailments added compositionally. To be concrete, we will express meanings in an explicit temporal logic (Allen, 1984, Allen & Ferguson, 1994) with liberal use of reified events (cf. Hobbs, 1985). The same points could be made with other formalisms. Given such a representation, the representation of *The plant moved* might be

$$\exists e, p, t. \text{Motion}(p, e) \ \& \ \text{Plant}(p) \ \& \ \text{Occurs}(e, t) \ \& \ \text{Past}(t)$$

This says there exists an event involving motion of an object p , which is a plant, and this event occurred at a time in the past. The predicate $\text{Motion}(p, e)$, named after the Framenet frame MOTION and corresponding to an intransitive sense of *move*, contains two arguments - the object undergoing motion (the plant) and the reified event of this object moving. Common sense knowledge associated with this predicate would provide the entailments that an object undergoing motion changes position.

If we follow Wordnet and Framenet, the meaning of the second sentence, *John moved the plant*, would correspond to the formula in our notation

$$\exists e, p, j, t. \text{Cause-Motion}(j, p, e) \ \& \ \text{Plant}(p) \ \& \ \text{John}(j) \ \& \ \text{Occurs}(e, t) \ \& \ \text{Past}(t)$$

Here CAUSE-MOTION is a ternary predicate encoding the meaning of a binary/transitive sense of the verb *move*. By having two senses of *move* in these examples, we have created a word sense disambiguation problem. In this case, it is not a great complication for the senses can be reliably disambiguated by the form of the sentence (one is intransitive and the other transitive). But if we are to support reasoning and computing entailments, we also need to define or learn a common sense axiom for both senses, or learn an axiom that relates these two predicates, such as

$$\forall e, p. \text{Cause-Motion}(p, e) \Rightarrow \exists e'. \text{Motion}(p, e')$$

Of course, if *move* just had two senses, this would not be much of a problem. But Wordnet has many different senses for the verb *move*, and the list is not complete. But many of these senses are related, which tends to complicate the disambiguation process (even for human annotators), and requires defining or learning a cross-product of common-sense axioms relating the senses.

The alternative I propose is to follow Dowty and others (e.g., Pustejovsky, 1995; Gawron,

2008) who suggest taking sentence structure seriously, and use the meaning of roles/arguments to reduce the number of verb senses required and the number of common sense entailment axioms one needs to encode or learn. To show this, let's reconsider the example above.

Assume that there is one sense of *move*, corresponding to the Framenet frame: MOTION. In other words, *move* means something undergoing motion. Rather than having a second sense with an additional argument (the "agent"), we define a general predicate, *A-cause*, that takes an agent and an event and creates a new event that consists of the agent causing an event (in an appropriate agent-like way). Under this analysis, the interpretation of the sentence *The plant moved* remains the same, but for *John moved the plant* we now have:

$$\exists e, e', p, t. \text{A-cause}(j, e, e') \ \& \ \text{John}(j) \ \& \ \text{Motion}(p, e) \ \& \ \text{Plant}(p) \ \& \ \text{Occurs}(e', t) \ \& \ \text{Past}(t)$$

Rather than needing specific axioms to relate *Cause-Motion* and *Motion* senses as we had above, we now need only general axiom based on the *A-cause* predicate that applies to all verbs that describe events that can be caused by agents:

$$\forall e, e', a, t. \text{A-cause}(a, e, e') \ \& \ \text{Occurs}(e', t) \Rightarrow \text{Occurs}(e, t)$$

With this alternative treatment, we have not only reduced the number of word senses throughout the lexicon, but also dramatically reduced the number of common-sense axioms we will need to encode or learn. This insight leads to the first rule of thumb in defining word senses.

Minimal Valence Sense Criteria

The senses of the verb should be based on uses involving the minimum number of arguments, unless senses of a larger valence cannot be constructed compositionally.

The idea is that the entailments of sentences involving higher valence verbs will generally be derived compositionally by applying the meanings of modifiers to the simple form. We only introduce other senses when this is not possible.

For convenience later on, we will introduce an entailed role predicate that allows us to identify the argument in such predicates, which, following Dowty (1991) we'll call the patient (aka proto-patient). In other words, we have an axiom like the following for all "unary" verb meanings, defined in terms of an abstract event UNDERGO:

$$\forall e, p. \text{Undergo}(p, e) \Rightarrow \text{Patient}(e, p)$$

All “unary” verbs would be classified under this class in the ontology, essentially encoding axioms such as

$$\forall e, p. \text{Motion}(p, e) \Rightarrow \text{Undergo}(p, e)$$

Of course, not all verbs support an intransitive form. Rather, many verbs have a core binary valence, as in *John hit the cat*, *He pushed the ball*, *I love you*, *The foundation supports the house*, *I believe him*, and so on.

These verbs describe a binary relation between two arguments, and with the reified event argument we have three place predicate. Thus, *He pushed the door* might have the meaning

$$\exists e, a, b, t. \text{Push}(a, b, e) \ \& \ \text{He}(a) \ \& \ \text{Ball}(b) \ \& \ \text{Occurs}(e, t) \ \& \ \text{Past}(t).$$

where PUSH would be a realization of the Framenet frame MANIPULATION.

For predicates involving one object acting on another in some way (like PUSH), we identify the arguments using Dowty’s proto-agent (which we’ll call Actor) and proto-patient:

$$\forall e, a, b. \text{Act-on}(a, b, e) \Rightarrow \text{Actor}(e, a) \ \& \ \text{Patient}(e, b)$$

One might consider it unnatural that two seemingly similar verbs, *move* and *push*, end up with different arity predicates for their core sense. However, there are significant semantic differences between these verbs. The eventuality of pushing inalienably requires two objects, one applying force to the other. One can see this in the difference between the entailment of the sentence *She moved*, which involves one object, and *She pushed*, which requires an implicit object to push against, i.e.,

$$\exists e, b, t. \text{Push}(S, b, e) \ \& \ \text{Occurs}(e, t)$$

where we use the constant S here for the denotation of *She* just to simplify the formulas. Even when not mentioned in the sentence, we know some other object exists.

Also note that just because someone pushes a door doesn’t entail that the door moves. The interpretation of pushing as a way of moving, as in *He pushed the ball to the corner*, will actually be constructed compositionally, combining the apply-force sense of *push* with the result/goal interpretation of the prepositional phrase, to be described in the next section.

3 Taking Prepositions Seriously

This section considers how we can construct sentence interpretations without having to proliferate the number of verb senses. To be concrete with simple examples, we will first look a range of modifiers that describe physical location, motion and containment (often collapsed into a few semantic roles with names such as “goal” and “direction”).

Consider the sentences in Table 1 using the verb *push*. Depending on the framework, we get different granularities of senses. From our viewpoint, these analyses are too detailed on one hand (as a way of classifying senses) while being too coarse-grained on the other (as a way of capturing entailments). In our approach, each of these sentences has distinct entailments, yet there is a single verb sense. It is the modifiers that provide the additional entailments.

In each case, the sense of *push*, $\text{PUSH}(a, b, e)$ entails that e is the event of a applying force to b . A new event arises when we apply the meaning of a modifier. Consider the first example. The meaning of *to the house* is a predicate $\text{To-Loc}(d, e, e')$, which defines the event e' that consists of event e resulting in the patient of e being at d . More specifically, this entailment is captured by an axiom involving *To-loc*:

$$\forall e, e', x, o, t. \text{To-Loc}(o, e, e') \ \& \ \text{Patient}(e, x) \ \& \ \text{Occurs}(e', t) \Rightarrow \sim \text{At-loc}(x, o, t) \ \& \ \exists t'. \text{At-loc}(x, o, t') \ \& \ \text{Meets}(t, t')$$

This one axiom essentially captures the meaning of *To-Loc*, namely that the event results in the Patient being at the location.

Thus, the direct meaning of the sentence *she pushed the box to the house* would be

$$\exists e, e'. \text{To-loc}(H, e, e') \ \& \ \text{Push}(S, B, e) \ \& \ \text{Occurs}(e', t)$$

[Where for rhetorical purposes we have simplified the analysis of the noun phrases, using H for the house, S for she, and B for the box] Note that the entailment that the box moved is not explicit in the logical form, but easily inferred. Since the box was not at the house during the event, but at the house at the end, it must have moved.

The second example, *She pushed against the wall*, does not entail any movement (although it

Example	Framenet	WordNet Sense	VerbNet
She pushed the box to the house	CAUSE-MOTION	1	Carry.11-4
She pushed against the wall	MANIPULATION	5	Push.21-1-1
She pushed the twigs apart	CAUSE-MOTION	1	Split.23-2
She pushed the ball towards the house	CAUSE-MOTION	8	Carry.11-4
She pushed me	MANIPULATION	5	Hold.15-1
She pushed the ball into the box	CAUSE-MOTION	1	Funnel.9-3

Table 1: Senses/Uses of Push

doesn't exclude it). The predicate *Against(o, e, e')* is true only if the event e' consists of some event involving a force (e) being applied to an object (o). The third example, *She pushed the twigs apart*, entails movement and the fact that some set of objects was scattered. This would be encoded in the meaning of the preposition *apart*. Specifically, *Apart(e, e')* is true only if e is an event that results in the Patient of e , which must be a set of objects, being separated. The meaning of *Towards(d, e, e')* would entail that patient of e moved in the direction of d , but not that it got there. The fifth example is simply the base entailments from PUSH, some force was applied to an object. Finally, the last example derives much of its meaning from *Into-Loc(c, e, e')*, which entails that the patient of e ends up inside a container c .

$$\forall e, e', l, t, x . \text{Into-Loc}(l, e, e') \ \& \ \text{Patient}(e, x) \ \& \ \text{Occurs}(e', t) \ \Leftrightarrow \ \exists t', t'' . \text{Ends}(t', t) \ \& \ \text{Meets}(t, t'') \ \& \ \sim \text{In}(x, l, t') \ \& \ \text{In}(x, l, t'')$$

There are other advantages of this approach, besides the fact that the number of senses for words has been reduced considerably. This approach also provides finer grained semantic entailments than usual in semantic role taxonomies. For instance, Framenet, like most other analysis, uses a semantic role *Goal* to capture the meaning of both expressions such as *to the house* and *into the house*. But one of these entails that the object that moved is now contained in the house while the other doesn't. To address this issue, some refer to developing an addition spatial model. My point here is that there is no need - the semantic role model and preposition sense model (in a spatial theory) are one and the same problem.

But perhaps the most significant advantage of this approach is the reduction in the number of common-sense axioms that need to be encoded or learned. Verbnet, for instance, is one the few resources that is attempting to capture entailments systematically for a large number of verbs. In Verbnet, a separate axiom needs to be specified for every verb pattern within every verb sense/cluster. For instance, probably half the verb classes in Verbnet explicitly encode that the agent role indicates a causality entailment between the agent and the event. We capture this with one general axiom. In Verbnet, entailments about motion and resulting location must be encoded for each verb cluster. In our approach, we need only encode such information for each sense of the prepositions. With general axioms relating the prepositional meanings to a theory of space, we are capturing the commonsense

entailments for a vast number of verbs, including all motion verbs.

Another advantage of the compositional approach is its generativity. Consider a sentence with a verb you might not expect to occur with a *to* modifier, such as *He extracted the pin to the corner*. This is an unusual sentence, but I think it is can be understood that the pin ended up in the corner. Our approach provides this interpretation without having to introduce a new (and dubious) sense of extract. All the original implications of the extracting event are still true, *and* the pin ended up in the corner.

Note also that by taking this approach, we have provided a good foundation for obtaining partial understanding of verbs for which we have the merest idea of their meaning. For instance, say we read the sentence *the ball zooped into the box* - while we may not know anything about this verb *zooped*, using the axioms about *into*, and we can posit that the ball is now in the box.

4 Capturing Entailments

Here we look at compositionally constructing entailments from the sentence structure, and show that it provides us considerable flexibility and allows us to avoid some dilemmas one often faces when trying to identify and classify word senses. Consider the following sentence, together with a few sentences that would be reasonable entailments:

He slid it from the corner into the box
 --> *He moved it into the box*
 --> *It entered the box*
 --> *He removed it from the corner*
 --> *It left the corner*

Framenet captures each of these verbs as different frames: CAUSE-MOTION (slide and move), ARRIVING, REMOVING, and DEPARTING, respectively. To capture these entailment relationships, we'd need to explicitly link these frames in another theory. In our approach, however, we can capture these conclusions with only a few common-sense entailments. To see this, consider first the meaning of *He moved it from the corner into the box*:

$$\exists e, e', e'', e''', t . \text{Into-loc}(B, e''', e''') \ \& \ \text{From-Loc}(C, e'', e''') \ \& \ \text{A-cause}(H, e', e'') \ \& \ \text{Slide}(I, e) \ \& \ \text{Occurs}(e, t) \ \& \ \text{Past}(t)$$

[where B is the box, C is the corner, I is the referent of it and H is the referent of He] To entail the first sentence, we only need the common-sense ontological knowledge that any sliding *sliding* event is a moving event

$$\forall e, o . \text{Slide}(o, e) \Rightarrow \text{MOTION}(o, e)$$

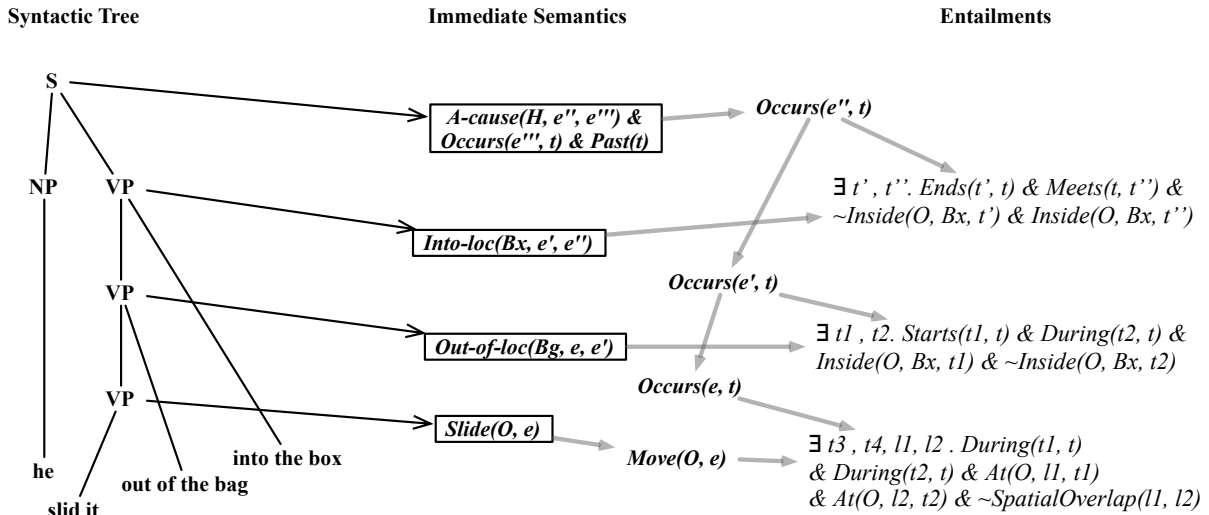


Figure 2: Compositionally Constructing Sentence Entailments

With this one axiom, we can entail all the variants of moving: *He moved it*, *he moved it from the corner*, *he moved it into the box*, and so on. For second example, we need to know the definition of the event *Enter*. This is one of those verbs where we can give a precise definition, following Allen (1983):

$$\forall e, o, loc, t . \text{Enter}(o, loc, e) \ \& \ \text{Occurs}(e, t) \\ \Leftrightarrow \exists t', t'' . \text{Ends}(t', t) \ \& \ \text{Meets}(t, t'') \ \& \\ \sim \text{In}(o, loc, t') \ \& \ \text{In}(o, loc, t'')$$

i.e., an entering occurs every time a containment relation is initiated. It is simple to see that any *Into-Loc* event entails an *ENTER* event.

Similarly, *REMOVING* events are defined in terms of an agent causing an object to not be in or at a location, which in turn would be entailed by the *A-cause* and *From-loc* parts of the sentence. And the fourth example, involving a *LEAVING* event, would be entailed using a definition of *LEAVING* that mirrors the *ENTER* definition.

Encoding all this knowledge is clearly a huge enterprise, but it is not nearly as bad as one might think looking at efforts like Verbnet. First, building meaning compositionally from the sentence structure allows us to define key entailments based on the meanings of the modifiers, rather than having to repeat them for each verb class. Note that the set of prepositions is a fixed class and manageable. The set of verbs is unlimited and continually growing. Second, we do not need to encode all knowledge before it becomes useful. In particular, encoding knowledge of core areas such as space and time, and abstractions of these, captures a critical body of knowledge about verbs. And, much of the key entailments can be obtained via inheritance. For instance, all the system needs to know about *SLIDE* to obtain the above entailments is that its

is a type of *MOTION* - information readily obtained in many existing ontologies and lexical resources. Of course, if the system knows more about sliding, then additional entailments can be drawn. And furthermore, once we have a basic understanding of sliding in terms of moving, we have a strong context which could help a system learn further properties of sliding from example sentences, or from observing instances of sliding in the actual world.

5 Building Meaning Compositionally

With this framework, we can develop a model of incremental deep parsing, in which each verb phrase corresponds to a new event, and forward chaining inference is performed as each event is constructed in order to compute possible entailments as parsing continues. Figure 2 shows a parse of the sentence *He slid it out of the bag into the box*, where as before, we omit the information about the referring expressions, using constants *O*, *Bg*, *Bx*, and *H*, for *it*, *the bag*, *the box* and *He* respectively.

As each verb phrase is constructed, a new event is created (generally all these events occur simultaneously). In general, only limited entailments can be drawn just from the VP interpretation until we can conclude that the event has occurred. This usually occurs when the top-level declarative S structure is produced. This triggers a chain of entailments derivable in a forward chaining fashion. The shaded arrows in Figure 2 connect the results of each derivation. When multiple arrows point to a conclusion, this means the conclusion follows from the conjunction of the antecedents. Thus, in this case, the fact that e''' occurs entails that e'' occurs (from an axioms involving *A-cause*), from which entails event e' occurs (from axioms involving

Into-loc), and so on. As each event is asserted to occur, we can conclude properties of the world describing the overall situation. Note in this example, we assumed the system knows nothing about sliding except that its a form of moving. Still we get a rich picture of the world described by this sentence.

6 Rethinking Semantic Roles

Semantic roles play a central role in many semantic theories, but generally there is no clear idea of the principles underlying the choices of roles. In many cases, the roles mainly seem to be mnemonic ways to indicate argument positions. Also, as described above, if we take prepositional meaning seriously, there is much less motivation for a separate theory of semantic roles independent of prepositional meaning. Here we explore this position.

As an example, consider the semantic role *Goal* (or *Destination* as it is called in VerbNet), which is commonly found in semantic role taxonomies and is generally taken to identify the end state (often a location) resulting from some verb of change. Two common realizations of the *Goal* role are PPs such as “to the house” and “into the box”. But we already have senses for these two prepositions so have no need for the semantic role. Furthermore, these two readings capture the different entailments that often seem to be ignored if we just label both of them as *goals*.

In a semantic role analysis, the fact that a “goal” type PP is allowed is specified as part of the verb frame. In this analysis, this would be captured by axioms involving the preposition, identifying the types of events that can be extended to have a resulting state. For example, say that a certain sense of *into* only applies to motion verbs. The we have an axiom

$$\forall o, e, e'. \text{Into-1}(l, e, e') \Rightarrow \exists o. \text{Motion}(o, e)$$

With a simple semantics, this prevents this sense of *into* from extending any non-motion verb, playing the role traditionally done by selectional restrictions on the verb.

But roles might be useful in verbs that encode the same entailments as these prepositions but are realized as a direct non-PP argument. For instance, consider the direct object of the verb *Enter*. It has entailments that seem equivalent to the predicate *Into*. Since this is easily captured in the definition of *Enter* presented earlier, as we need a definition for the verb in any case, adding a semantic role analysis here would not add any value.

Another place where the *Destination* role occurs in Verbnet are with verbs that involve

changing position or ownership and allow the indirect object form, as in *I sent her the book*, or *The train brought the city much needed supplies*. Some taxonomies would called this a *Recipient* role, but the name doesn’t influence our argument here at all. There are two approaches to this problem - the first would be to define a ternary predicate for these forms of verbs and directly encode the desired entailments. But this has the disadvantage of multiplying word senses and increasing the number of axioms needed to capture the phenomena. Much more attractive, given that the indirect object construction is highly generative, is to treat the indirect object like we are treating prepositional modifiers. Thus the interpretation of *I sent her the book* would have the form

$$\exists e, e'. \text{Involved}(H, e, e') \ \& \ \text{Sending}(I, B, e)$$

Where H is referent of her, *I for the speaker*, and *B for the book*, and the axiomatization of the predicate *Involved* will provide the appropriate entailments capturing our intuitions of the *Beneficiary*, *Destination* or *Recipient* depending on the event being modified.

For verbs like *send*, the interpretation is the same as the preposition *to*, so we might have an axiom that covers all the verbs involving transferring things (in Verbnet, the classes involving bring, carry, drive, send, and throw).

$$\forall o, x, y, e, e'. \text{Involved}(o, e, e') \ \& \ \text{Transfer}(x, y, e) \Rightarrow \text{To}(o, e, e')$$

These axioms would be expressed at the most abstract event type possible to avoid needing axioms for specific verb senses. We would have similar axioms to cover *Beneficiary* (one of the senses of the preposition *for*) and *Recipient* (the same or another sense of the preposition *to*).

The remaining cases of the *Destination* role in Verbnet occur in verbs such as *poke*. To explore these issues, we now focus on the semantic contribution of the directly subcategorized (non PP) arguments to a verb, namely the core arguments of subject and object. The semantic roles identify how these arguments relate semantically to the event described by the verb.

This verb *poke* allows a range of assignment of semantic roles to its core arguments. Consider the sentences below and what roles Verbnet would assign to the subject and direct objects (where A=agent, D=destination, I=instrument)

$$\begin{array}{ll} \text{He poked the ball.} & [A, D] \\ \text{The needle poked the ball.} & [I, D] \\ \text{He poked the needle into the ball.} & [A, I, D] \\ \text{He poked the ball with the needle} & [A, D, I] \end{array}$$

The challenge here is that we appear to have three binary senses of *poke* with different en-

tailments, and one of these senses, $[A, I]$, can only happen when we have a modifier that adds the destination. These are the kinds of examples that motivated semantic roles in the first place (Fillmore, 1966).

We would like to preserve a single sense of *poke* here. Considering the entailments, it seems that the core sense is that an object underwent some directed motion resulting in contact (with the destination object). So, we have this meaning for the second sentence:

$\exists e. \text{Poke}(N, B, e)$

The first example, involving the agent, would need to map to

$\exists e, e', x. A\text{-cause}(H, e, e') \ \& \ \text{Poke}(x, B, e)$

i.e., the person caused a poking event with the objects undergoing the motion (i.e., the instrument) unspecified, and may be the agent itself. And the third and fourth examples would specify all the players involved:

$\exists e, e'. A\text{-cause}(H, e, e') \ \& \ \text{Poke}(N, B, e)$

The challenge for our compositional framework is how the third example is constructed, as the VP *poked the needle* can only support the instrument interpretation in the context of the modifying PP *into the ball*. This remains an issue to be resolved and provides a strong argument for retaining a semantic role analysis for the core verbal arguments.

7 Using Semantic Role Abstractions

By using abstract roles and event types, we can concisely encode an even wider range of sentence uses with a single verb sense, albeit one more abstract. Say we redefine *Motion* to describe an event involving an abstracted sense of motion, which can be modified by abstracted senses of prepositions such as *To*, *From*, *Through*, and so on. In the abstract version, *From* relates an event with a state that was true at the start of the event, and likewise, *To* indicates a relation between its argument and the object involved in the event that is true at the end of the event. The more specific uses of these roles can be captured by entailments based on the prepositional relations. For instance, consider the meaning of *he moved from elation to depression*, which doesn't appear to be encoded in wordnet as a sense yet, involves a change in emotional state.

The predicate *To-loc* used above, is a specialization of *To* that involves physical locations. The entailments for *To-loc* remain just as before, and we have a rule that relates the two hierarchically:

$\forall x, e, e'. \text{To-loc}(x, e, e') \Rightarrow \text{To}(x, e, e')$

But more importantly, we have specialization rules that work in the opposite direction conditional on the type of object. For instance, if the object is a physical object (e.g., *to the house*) then there is a strong suggestion that *To* when applied to motion events has its *To-loc* sense, captured by the rule

$\forall x, y, e, e'. \text{To}(x, e, e') \wedge \text{MOTION}(y, e) \wedge \text{Phys-Obj}(x) \gg \text{To-loc}(x, e, e')$

Here, I introduced likely entailment operator (\gg) to indicate that this is likely but not logically necessary. For the purposes of this paper, we can assume an analysis such as the default conditional in Asher and Morreau (1991).

If the object of *To* is an emotion and the object undergoing the motion is cognitive entity, then a rule would specialize *To* to some relation indicating the cognitive entity being in the emotional state. Likewise, if the object of *To* is a legally-defined area (e.g., city, state), we'd have a rule suggesting that the relation of importance is the official residence (e.g., *He moved to Chicago*). While there might be many such rules for different cases of abstract motion, note that it is much easier to encode or learn such rules for a handful of prepositions rather than encode or learn an unlimited set of verb senses.

8 Discussion

This is a work in process. It has been motivated by the need to come up with a set of principles for determining an effective encoding of a broad-coverage lexicon to enable computing entailments from arbitrary texts and dialogue (i.e., deep understanding). Already we have found a few principles a great help in making sense of this task. Specifically, we are working to define a set of core guidelines, which include

- 1) Consider only senses of verbs in their base forms (i.e., intransitive forms, or transitive if not intransitive). Add additional senses only if no compositional approach can be found.
- 2) Replaced most semantic roles by predicates corresponding to the preposition senses, each of which is described with a set of axioms defining their entailments.
- 3) A more traditional semantic role analysis is only kept for interpretations of the core verbal arguments (i.e., subject and object).

The compositional approach appears to provide great benefit for reducing the number of senses and provides a useful criteria for any semantic role analysis.

Much of the entailments in the examples we have used can be found in Verbnet. For instance,

the class Push-12 has semantic annotations on each possible syntactic frame that captures the causation between the agent and the event, the contact of the agent and theme during the event, and the exertion of force during the event. As such, Verbnet provides a good resource for entailments. There are significant differences between the approaches, however. We have described a compositional approach that allows for a much more concise set of axioms with more uniform generality. In Verbnet, every syntactic frame of every verb class must have its own entailments specified. In addition, Verbnet classes, being based on Levin's (1993) analysis, do not always correspond to an intuitive classification of word senses. The verbs *enter* and *exit*, for instance, which appear as opposites along one dimension, occur in the same class for they support the same range of arguments.

There is also related work in trying to reduce the number of senses of verbs such as Ontonotes (Pradhan et al, 2007) and Propbank (Palmer et al, 2005). These efforts are highly compatible with our effort and will be a valuable resource in rethinking our verb senses. They do not, however, discuss issues of compositionality or entailment.

FrameNet has provided much inspiration for the identification of meaningful classes of verbs for our ontology. We, however, differ significantly in the importance we place on semantic roles. FrameNet uses a wide range of frame-specific roles making generalization across roles difficult. In our work, we have worked to reduce the use of roles just to the interpretation of direct arguments.

9 Acknowledgements

This work was supported in part by the National Science Foundation, grant #0748942, and the Office of Naval Research (N000140510314).

10 References

Allen, J. F. (1984). "Towards a General Theory of Action and Time." *Artificial Intelligence* 23.

Allen, J. F. and G. Ferguson (1994). "Actions and Events in Interval Temporal Logic." *Journal of Logic Computation* 4(5): 531-579.

Allen, J. F., Swift, M. and de Beaumont, W. (2008). *Deep Semantic Analysis for Text Processing. Symposium on Semantics in Systems for Text Processing (STEP 2008)*. Venice, September.

Asher, N. and M. Morreau. (1991) "Common Sense Entailment: A Modal Theory of Nonmonotonic Reasoning," in Proc. 12th Intl Joint Conference on Artificial Intelligence, Sydney, Australia.

Davidson, D. (1967). The logical form of action sentences. In Nicholas Rescher, (Ed.), *The Logic of Decision and Action*, U. of Pittsburgh Press.

Dowty, David (1989) "On the Semantic Content of the Notion of 'Thematic Role'." In Chierchia, Gennaro, Barbara H. Partee, and Raymond Turner. *Properties Types and Meaning Volume II: Semantic Issues*. Dordrecht: Kluwer Academic Publishers, pp.69-130.

Dowty, D. (1990) Thematic proto-roles and argument selection. *Language* 67.3: 547-619

Fellbaum, C (ed), (1998) *Wordnet, An Electronic Lexical Database*, MIT Press.

Fillmore, C. J. (1968). The Case for Case. *Universals in Linguistic Theory*. E. B. a. R. Harms. New York, Holt, Rinehart and Winston: 1-90.

Gawron, J. M. (2008) *Circumstances and Perspective: The Logic of Argument Structure*, San Diego Linguistics Papers, Issue 3,

Goldberg, Adele (1995) *Constructions. A Construction Grammar approach to argument structure*. Chicago: University of Chicago Press.

Goldberg, Adele (2006) *Constructions at Work: the nature of generalization in language*. Oxford University Press.

Johnson, C. and C. Fillmore (2000) The FrameNet tagset for frame-semantic and syntactic coding of predicate-argument structure. ANLP-NAACL, Seattle, WA.

Hanks, P. (2000). Do word meanings exist? *Computers and the Humanities*, 34(1-2):171-177.

Hobbs, J. (1985) *Ontological Promiscuity*, Proc. 23rd Mtg. of the ACL, Chicago, IL.

Kipper, K. Anna Korhonen, Neville Ryant, and Martha Palmer. (2006). *Extending VerbNet with Novel Verb Classes*. Fifth Intl Conf on Language Resources and Evaluation (LREC 2006). Genoa, Italy.

Levin, B. C. (1993). *English Verb Classes and Alternations: a Preliminary Investigation*, University of Chicago Press, Chicago, IL.

Navagli, Roberto. (2009). *Word Sense Disambiguation: A Survey*, ACM Computing Surveys, Vol. 41, No. 2.

Palmer M, Kingsbury P, Gildea D (2005). "The Proposition Bank: An Annotated Corpus of Semantic Roles". *Computational Linguistics* 31 (1)

Palmer, M., H. Dang and C. Fellbaum, (2007). *Making fine-grained and coarse-grained sense distinctions, both manually and automatically*, *Natural Language Engineering*, Volume 13, Issue 02.

Pradhan, S., E.H. Hovy, M. Marcus, M. Palmer, L. Ramshaw, and R. Weischedel (2007). *OntoNotes: A Unified Relational Semantic Representation*. in Proc. of the First IEEE International Conference on Semantic Computing (ICSC-07). Irvine, CA.

Pustejovsky, J. (1995) *The Generative Lexicon*, MIT Press.