Mechanisms of Sense Extension in Verbs

James Pustejovsky
Dept. of Computer Science
Brandeis University
Waltham, MA USA
jamesp@cs.brandeis.edu

Anna Rumshisky
Dept. of Computer Science
Brandeis University
Waltham, MA USA
arum@cs.brandeis.edu

Abstract

In this paper, we examine some of the mechanisms at work which relate distinct senses of a predicate. Independent of whether one adopts a lumping or splitting approach to sense differentiation and granularity in word meaning, the issue of how senses relate to one another is relevant for both approaches. While we differentiate between the selection and the coercion of an argument by a predicate, type coercion is not the mechanism responsible for most of the sense extensions discussed here. We claim that the different degree of meaning extension (degree of metaphoricity), results from a number of formal processes operating on the predicate. These include: generalizing the type of the argument; changing the argument structure and relative prominence of arguments; and finally, abstracting the core meaning of the verb itself. We demonstrate such sense extensions with several diverse verb classes, including motion predicates and locative relations, We examine the hypothesis that metaphorical interpretations are structured and scalar in nature. This would have the consequence that abstract metaphoric readings for concrete predicates should be accompanied by intermediate senses, exhibiting less metaphoricity.

1 Introduction

The creative manipulation of lexical content seems a hallmark of language use. The meanings of words in context can be modulated in many different ways, ranging from minor accommodations of the selectional preferences expected by a predicate, to major abstractions associated with the core semantic value of the predicate. In this paper, we look at the manner in which senses of a predicate appear to modulate in the language. We explore several independent generative mechanisms that we believe are responsible

for the creation of extended senses for predicates. These devices are distinct from the *type-preserving* coercion mechanisms introduced in Generative Lexicon Theory to account for type mismatches that frequently occur in argument selection (Pustejovsky, 1995). We claim that the different degrees of meaning extension result from a number of formal processes operating on the predicate. Those we have identified include the following:

(1) a. generalizing the type of the argument of the predicate;
b. changing the argument structure and relative prominence of arguments;
c. and finally, abstracting the core meaning of the verb itself.

We demonstrate such scalar-like sense extensions with several diverse verb classes, including motion predicates, locative and spatial relations, and change-of-state predicates. We examine the hypothesis that metaphorical interpretations are structured and scalar in nature. This would have the consequence that abstract metaphoric readings for concrete predicates should be accompanied by intermediate senses, exhibiting less metaphoricity.

2 Related Work

2.1 Relations between existing senses

Defining a set of senses available to a lexical item is one of the tasks in lexical semantics which has eluded practitioners in all the impacted fields from theoretical linguistics to lexicography to computational work on word sense induction and disambiguation. For polysemous verbs especially, constellations of related meanings make this task even more difficult.

In lexicography, "lumping and splitting" senses during dictionary construction – i.e. deciding when to describe a set of usages as a separate sense – is a well-known problem (Hanks and Pustejovsky, 2005; Kilgarriff, 1997; Apresjan, 1973). It is often resolved on an ad-hoc basis, resulting in numerous cases of "overlapping senses", i.e. instances when the same occurrence may fall under more than one sense category simultaneously. This problem has also been the subject of extensive study in lexical semantics, addressing questions such as when the context selects a distinct sense and when it merely modulates the meaning, what is the regular relationship between related senses, and what compositional processes are involved in sense selection (Pustejovsky, 1995; Cruse, 1995; Apresjan, 1973). A number of syntactic and semantic tests are traditionally applied for sense identification,

such as examining synonym series, compatible syntactic environments, coordination tests such as *cross-understanding* or *zeugma* test (Cruse, 2000); but none of these tests tend to be conclusive or universally applicable.

For verbs, one of the main confounding factors is co-composition, since semantics of the arguments in many cases contributes significantly to determining the overall meaning of the verb. Several current annotation efforts attempt to enumerate the context elements that accompany different senses of polysemous predicates. In particular, FrameNet and PropBank are two lexical resources that specify semantic interpretation ascribed to the verb's arguments in terms of semantic role labels. However, neither records the information regarding the relations between different senses, either in terms of semantic interpretation ascribed to the arguments, or in terms of the actual sense of the verb proper.

PropBank (Palmer et al., 2005) defines verb senses in terms of framesets, where each frameset consists of a set of semantic roles that are for the arguments of the verb. Semantic roles are defined on a verb-by-verb basis, with a separate set of roles specified for each sense. With the exception of the standard Agent and Patient/Theme assigned to Arg0 and Arg1, respectively, no relation between the roles for different senses is assumed.

FrameNet (Ruppenhofer et al., 2006) organizes lexical information in terms of script-like semantic frames, with semantic and syntactic combinatorial possibilities specified for each sense (the frame-evoking *lexical unit*). FrameNet uses Fillmore's case roles to represent semantics of the arguments. Context specification for each lexical unit contains such case roles (e.g. Avenger, Punishment, Offender, Injury, etc. for the *Revenge* frame) and their syntactic realizations, including grammatical function (Object, Dependent, External Argument (= Subject)), etc.), and phrase type (e.g. NP, PP, PPto, VPfin, VPing, VPto, etc.). Core frame elements represent semantic requirements of the target lexical unit, some of which may not be actually expressed in the sentence.

Case roles (*frame elements*) in FrameNet are again derived on ad-hoc basis for each frame, and since the analysis proceeds frame-by-frame, often only one or two main senses of frequent polysemous verb have any frame interpretation associated with them. One FrameNet mechanism that is relevant to the analysis of relations between different senses is the mechanism of between-frames relations. For example, FrameNet contains only three senses of the verb *drive*. The first sense, "convey in a car", is associated with the *Bringing* frame which **uses** the *Motion* frame. The second sense, "move under its own power or directed by a driver (of a vehicle)", is associated with the *Self-motion* frame which **inherits** from the *Motion* frame. However,

this type of specification does not provide insight into relations between the elements filling argument positions for different senses beyond the simple mapping.

In this paper, we explore more formally two of the factors that contribute to sense extension for verbs (1) changes in the semantic type of the arguments (2) the changes in argument structure. Mechanisms which relate predicate senses can be viewed as strategies employed in language for extending the meaning of a predicate, or as static formal mappings from one sense to another. The analysis we propose here applies equally to both diachronic and synchronic sense modification.

2.2 Treatments of metaphor

Metaphor and metonymy as mechanisms for sense extension in language has been explored extensively in the context of diachronic semantic change, with a number of linguists over the past century proposing typologies of semantic shift, including Bloomfield (Bloomfield, 1933), Stern (Stern, 1931), Ullmann (Ullmann, 1962), and earlier works by Jespersen (Jespersen, 1917), Bréal (Bréal, 1899), and others. Bloomfield (Bloomfield, 1933), for example, supplies the following examples for these mechanisms: Germanic 'biting' $>$ 'bitter' (metaphor, based on similarity); Old French 'joue' $>$ jaw (metonymy, based on spatial or temporal contiguity).

In more recent work, Zharikov and Gentner (Zharikov and Gentner, 2002), for example, look at the diachronic development of figurative meanings in nouns, using the data from Oxford English Dictionary to show that in these cases initial literal meaning is often followed over time by figurative uses marked by similes and other overt comparisons, which are in time replaced by strictly metaphoric uses which eventually turn into secondary meanings.

Another recent study relevant to the line of analysis we will pursue here was done by Hanks (Hanks, 2006). Looking mostly at the nominal metaphor, he proposes that there are different degrees of metaphoricity, depending on the 'frame' into which the metaphoric expression is embedded. The notion of the relative strength of the 'resonance' between the source and the target domain, is what he refers to as gradable metaphoricity. The notion that there are different degrees of metaphoricity, or, as he puts it, "some metaphors are more metaphorical than others" is explored by Hanks in a context of noun-based metaphor. The main claim he advances is that the more properties are shared between the target and the source somain, the *less* metaphorical is the metaphor. Here is an illustration of different de-

grees of metaphoricity in the metaphorical use of the noun *sea*, going from less metaphorical to more metaphorical:

(2) a. sea of [[Substance]]: mud, acid, snow, fire, foliage
b. sea of [N-Plural[Physical Object]]: faces, hats, umbrellas, hands
c. sea of [[Color — Light]]: light, darkness, blue, green, red and white
d. sea of [[Abstract]]: sensation, data, debt, misery, uncertainty, heartbreak, poverty, pleasure
e. sea of [N-Plural[Event]]: market transactions, activity, lies, troubles

Metaphors can be extended, so "the resonance between the target and the source is amplified and extended when other, related terms and concepts (significant collocates) of the 'secondary subject' are brought into play." Hanks proposes that less metaphorical cases are *less likely* to be extended further in the text.

Most of the work on metaphor in language has not dealt specifically with metaphoric transfer as the generative mechanism for creating verbal sense extension and has not addressed the relevant phenomena in sufficient detail. Metaphor is typically analyzed in terms of a mapping between the 'target' and 'source' conceptual domain (alternatively referred to as 'topic'/'tenor' and 'vehicle' of the metaphor, or 'primary' and 'secondary' subject of the metaphor), where the elements of the target domain are re-analyzed in terms of the elements of the source domain.

Some of the classic work on metaphoric transfer as it relates to the predicate sense extension seems to apply this general mapping device too broadly. For example, consider Lakoff's analysis (Lakoff, 1993) of some of the relevant data in (3). Two metaphoric mappings are postulated and several examples are given to illustrate each:

(3) a. Starting an Action is Starting out on a Path
We are just starting out.
We have taken the first step.

b. Success Is Reaching The End of the Path
We've reached the end.
We are seeing the light at the end of the tunnel.
We only have a short way to go.
The end is in sight.
The end is a long way off.

In these cases, however, strong metaphorical interpretation is confused with what appear to be compositional aspectual interpretations. There are two things to notice about these examples: first, they involve starting and ending activities. Second, the nouns are general functional expressions, that are arguably not even path denoting expressions. Finally, unless we have a case of metaphorical blending in the "The end is in sight", this is actually not a path metaphor at all, but rather a perceptual metaphorical extension.

This idea of metaphoric transfer between the target and the source domain is sometimes supplemented by the analysis of context, viewed as the 'frame' (Black, 1962) into which the word serving as the focus of the metaphor is embedded. In the examples below, for instance, the verb *plow* is seen as the *focus* of the metaphor, as it is the only word that is used metaphorically, while the other words in the sentence which form the *frame* of the metaphor are used in their literal senses:

(4) a. The chairman plowed through the discussion.
    b. I like to plow my memories regularly.

Black points out that the metaphoric use of *plow* in the two sentences above may or may not be considered instances of the same metaphor, but provides no formal mechanism for resolving this question.

Some discussion of the mechanisms involving argument type mismatches is provided in the work of Fass (Fass, 1988) which proposes the mismatch of the verb's preference for a particular argument type as the source of tension that has to be resolved through property matching between vehicle and topic of the metaphor. For example, in (5), *drink* has the Agent preference of for *animal*, which creates a tension due a type mismatch (a *car* is not an *animal*):

(5) The car drank gasoline.

This tension is resolved when an analogical match discovered between the salient property of an element from the source domain and a property of an element from the target domain: animals consume drinkable liquids as cars consume gasoline. This salient property serves as the *ground* for the metaphor. The mapping is then created between the two domains using other matched features (both animals and cars are bounded three-dimensional solids, etc.) The higher the proportion of differences to similarities between the topic and the vehicle, the "better" the metaphor.

Krishnakumaran and Zhu (Krishnakumaran and Zhu, 2007) use similar considerations to identify possible metaphoric uses in text using distributional information. They propose to identify metaphoric usages such as the one given in (6) by looking at the conditional probability of WordNet hyponym of a noun, given a particular verb. This, of course, would identify infrequent uses (including formerly literal ones that fell out of use) as possibly metaphoric.[1]

(6) He planted good ideas in their minds.

They consider such examples to be metaphoric usages of nouns, however, depite the fact that in this case, the predicate provides the features for the salient mapping between the target and source domain.

3 Sense Extension Mechanisms

While it is impossible to say how many meanings we create for a particular word in normal language use, we can reasonably ask how many meanings we have stored for that word in our mental lexicon. This is where linguists differ broadly in assigning responsibility for whether meaning shifts occur at all and, if so, how they occur. As a result of this divide, the role that compositionality and mechanisms of sense extension play in structuring not only the grammar but also the lexicon is significant.

For example, in conventional models of language meaning, an ambiguous verb is thought to have several distinct word senses. For each sense, the verb acts on its parameters (its arguments in syntax) in a compositional manner. This means that the semantics of the result of application of the verbal function to its argument is determined by the semantics of the function itself. Consider, for example, the way in which the verbs *throw* and *kill* each have several distinct senses.

(7) a. Mary threw the ball to John.
PROPEL(HUMAN, THING)
b. They threw a party for Bill.
ORGANIZE(HUMAN, EVENT)
c. Mary threw breakfast together quickly.
CREATE(HUMAN, THING)

[1]See Rumshisky, 2008 for further discussion of caveats associated with using distributional information for sense detection.

The use of *throw* in each sentence above illustrates a true verbal ambiguity, one that requires separate senses, each with specific subcategorization and semantic selection as illustrated. Likewise, the verb *kill* as used in (8) below, demonstrates a systematic sense distinction as well.

(8) a. John killed the plant.
CAUSE-TO-DIE(THING, ANIMATE)
b. Mary killed the conversation.
TERMINATE(HUMAN, EVENT)
c. John killed the evening watching TV.
SPEND(HUMAN, TIME)

As with the verb *throw*, each of these senses has a regular and productive distribution in the language, exemplified below.

(9) a. Mary killed the fish.
b. The President killed any attempt at dialogue with Cuba.
c. John killed the time reading.

Verb senses like these are distinct, semantic units, perhaps related to each other, but stored separately in the lexicon. Because they have distinct subcategorization and type selection frames, the semantic computation involving these senses in the syntax can be performed compositionally. These examples with the verbs *throw* and *kill* illustrate that lexical forms may be truly ambiguous, and as such, can be modeled adequately by a sense enumerative lexical (SEL) model (Pustejovsky, 1995).

Such is not the case, however, with more closely related (and overlapping) senses that are arguably computational extensions of a single sense. This appears to be the case with the different uses for the verb *arrive*, as illustrated below.

(10) a. The car arrived at the store.
b. The plane arrived at a cruising altitude of 30,000 feet.
c. Mary has not arrived at her goal weight of 125 lbs. yet.
d. The scientist finally arrived at a solution to the problem.

These senses exhibit a *sense extension cline*, one that takes us from the literal meaning for *arrive*, to different degrees of metaphorically extended senses for that verb.

In order to describe the sense extension mechanisms, we must assume, in a set of senses for a predicate, that one sense is an anchor, and the others

are related through some transformation(s) or not. True ambiguity exists, and not all senses for a predicate will be related, and unrelated senses may be treated as homonyms. Nevertheless, we are interested in identifying the formal processes through which anchor sense(s) are extended to produce related senses.

Using data from corpora, we illustrate the relation between the senses of several diverse predicates. We start with the mechanisms of pure selection from the anchor sense, and then radially branch out by applying semantic transformations to the predicate. This includes: subtyping; several types and levels of coercion; non-metaphoric sense extensions; "weak metaphors"; and finally "strong metaphor" readings.

By identifying specific sense extension mechanisms for verb meanings, we hope to be able to arrive both at a qualitative understanding of sense modulation, as well as a quantitative assessment of the notion of "degree of metaphoricity" of a predicate.

Given the notion of a predicate sense as a semantic feature matrix, we can imagine sense extensions as mutations over the anchor sense. Minor modifications would correlate with a small number of permutations, perhaps over peripheral features, while a major shift in sense would correspond with major feature modifications. The purpose of our investigation is to first see what features permit mutations, and secondly, to see if there are dependencies between these changes.

From our initial study, it appears that some "extended senses" are in fact associated with minor modifications in the predicate body. In particular, we identify an operation called *a sense-preserving type weakening*, where the modified typing on the arguments still satisfies the constraints imposed by the predicate in the qualia structure. A sense resulting from this operation we call a "non-metaphoric sense extension". This can be contrasted with more substantial type modifications in the arguments, giving rise to greatly generalized interpretations of the relation between them. For this reason, we call such operations *sense-extending type weakening*. This class will be further differentiated into "weak metaphors" and "strong metaphor", depending on whether the core meaning of the predicate undergoes generalization as a result of abstracting on the argument typing.

From our preliminary data analysis, it appears that sense extensions can be ordered according to these kinds of operations, giving rise to a basic scalar notion of sense extensibility. This would substantiate the view that metaphoric interpretation is best analyzed in terms of degrees of metaphoricity (Hanks, 2006).

In this section, we will illustrate this process of sense extension within

several verb classes: motion predicates, locative and spatial relations, and change of state predicates. These examples will support our view that sense extensions relate to the anchor sense through identifiable scalar modulations of the information structure associated with the predicate. First, however, we summarize theoretical assumptions made throughout the rest of the paper.

3.1 GL Assumptions

For our discussion, we will assume the model of Generative Lexicon (GL), as presented recently in Pustejovsky, 2006. Research in GL has traditionally focused on overcoming the static nature of lexical knowledge as deployed by linguistic theories, which traditionally have failed to account for creative word use in novel contexts. Rather than taking a snapshot of language at any moment of time and freezing it into lists of word sense specications, operations in GL do not preclude sense extensibility: the compositional processes are open-ended in nature and account for the novel, creative, uses of words in a variety of contexts by positing procedures for generating semantic expressions for words on the basis of particular contexts. Both lexical items and phrases encode the following four types of information structures: a. Lexical Type Structure: b. Argument Structure: c. Event Structure: and d. Qualia Structure.

The qualia values embody the modes through which we categorize our concepts. When certain features (qualia) are present or absent, we can abstract away from the representation, and generalize lexemes as belonging to one of three conceptual categories Pustejovsky (2001, 2006): Naturals types: Natural kind concepts consisting of reference only to Formal and Constitutive qualia roles; e.g., tiger, river, rock; b. Artifactual types: Concepts making reference to Telic (purpose or function), or Agentive (origin); e.g., knife, policeman, wine; and c. Complex types: Concepts integrating reference to the relation between types from the other levels; e.g., book, lunch, exam.

This enriched inventory of types for the language is motivated by the need for semantic expressiveness in lexical description. We also need, however, richer interpretive operations to take advantage of these new structures. Following [Pustejovsky, 2006], we argue that there are four ways a predicate can combine with its argument:

(11) a. *Pure selection* (type matching), where the type a function requires is directly satisfied by the argument;

b. *Accommodation*, where the type a function requires is inherited by the argument;

c. *Type coercion by exploitation*, where part of the argument's type is used to satisfy the type required by the function;

d. *Type coercion by introduction*, where the argument is wrapped (embedded) with the type required by the function.

The *argument structure* for a predicate is a specification of the number and type of the parameters associated with the predicate. For example, the verb *die* can be represented as a predicate taking one argument, *kill* as taking two arguments, while the verb *give* takes three arguments.

(12) a. **die**(x)
b. **kill**(x,y)
c. **give**(x,y,z)

In GL, arguments are differentiated into different types (cf. Pustejovsky, 1995, 2001 for discussion).

(13) a. *True Argument*: syntactically realized argument of the lexical item;

b. *Default Argument*: argument which participates in the logical expressions in the qualia, but which is not necessarily expressed syntactically;

c. *Shadow Arguments*: argument which is semantically incorporated into the lexical item. They can be expressed only by operations of subtyping or discourse specification; e.g.,

d. *True Adjunct*: argument which modifies the logical expression as part of the situational interpretation, but is not an argument to the relation proper.

As we will discover in the next section, the above argument types are exploited and modulated in different ways by distinct sense extension mechanisms. We begin with the case of metaphorical extensions involving motion predicates.

3.2 Case Study I: Motion Predicates

In this section, we will examine the sense extensions possible for motion predicates in English. Our thesis is that these senses can be ordered along a scale, one which measures the number of modifications to argument type restrictions on the predicate.

Motion predicates do not constitute one undifferentiated semantic class. Talmy, 1985 is perhaps the first to systematize the observation that languages have distinct strategies for expressing concepts of motion. He noticed that there are two basic constructions associated with the expression of motion: *verb-framed* and *satellite-framed* patterns (subsequent work on this includes Jackendoff, 1983; Talmy, 2000; Choi and Bowerman, 1991). This is also referred to as the *path* verb vs. *manner-of-motion* verb distinction. The latter strategy (satellite-framing) can be seen in sentences such as:

(14) a. *John hopped*$_{manner}$ *out of the room*$_{path}$.
b. *Mary crawled*$_{manner}$ *to the window*$_{path}$.

The path (verb-framed) construction is illustrated with the following examples:

(15) a. *John arrived*$_{path}$ *by foot*$_{manner}$.
b. *John descended*$_{path}$ *the stairs running*$_{manner}$.

We can split languages broadly into the two classes. MANNER CONSTRUCTION LANGUAGES encode *path* information in directional PPs, particles, and other adjuncts, while the main verb encodes the *manner-of-motion*; examples include English, German, Russian, Swedish, and Chinese. PATH CONSTRUCTION LANGUAGES encode *path* information in the matrix verb, while adjuncts optionally specify the *manner-of-motion*; examples include Modern Greek, Spanish, Japanese, Turkish, and Hindi.

As observed in (14) and (15) above, English allows both constructions, and these are common in everyday language. For example, *bike* is a manner verb used in a path PP-construction to indicate direction and path information. The verbs *arrive* and *leave* are both inherently path verbs and give no information regarding the manner-of-motion without further context.

We now return to our examples of sense extensions for the verb *arrive*, as well as other path predicates such as *reach*.

(16) a. The plane arrived in New York on time.
b. The plane arrived at a cruising altitude of 30,000 feet.

c. Mary has not arrived at her goal weight of 125 lbs. yet.
d. The scientist finally arrived at a solution to the problem.

(17) a. The plane reached the runway.
b. The hikers reached the top of the hill.
c. The two sides reached an agreement.

Consider the uses of *arrive* above. These senses exhibit a *sense extension cline*, one that takes us from the literal meaning for the verb in (16a), to increasingly extended senses. The lexical representation for the path predicate *arrive* makes reference to three arguments, two of them realized, and one defaulted:

(18) $\lambda y \lambda x \exists z \left[\begin{array}{l} \textbf{arrive} \\ \text{ARGSTR} = \left[\begin{array}{l} \text{ARG1} = x : phys \\ \text{ARG2} = y : location \\ \text{D-ARG1} = z : path \end{array} \right] \text{QUALIA} = \left[\begin{array}{l} \text{F} = at(x, y) \\ \text{A} = move(x, z) \end{array} \right] \end{array} \right]$

For the present discussion, we will abstract away some of the lexical structure of We can view this predicate as a *type template*, as illustrated in (19).

(19)

x:Phys	path_motion(x,y)	y:Location

The literal sense in (16a) is characterized by pure type selection on all argument positions to the verb: **the plane**: $phys$, **New York**: $location$. The denotation of a complex entity such as *New York* will, of course, involve more than typing as a location or geo-location, but it is defined at least by this type. [2]

Notice, however, that the location argument position in motion predicates is frequently filled by an NP denoting a physical object that is not a location, as in (22).

(20) a. The plane arrived at the gate on time.
b. Mary arrived at the store.

For both of these examples, the two entities can be thought of as denoting physical objects *where* something prototypical happens. For example, a gate can be construed as "a location for things dealing with planes", e.g.,

[2] in Pustejovsky, 2006, such entities are defined as Complex Types (dot objects), which have a more complex denotative status than normal entities. We will not concern ourselves with the issue of dot object selection and coercion in this discussion.

boarding, disembarking, and so forth. We will call this a *attribute functional coercion*, where an attribute of the argument satisfies the type of the selecting predicate. That is, for a sentence like (22a), the predicate's type is selected by virtue of a functional coercion over the NP in object position, viz., *the gate*; hence, the coercion returns the location of the gate, satisfying the type selection from *arrive*.

(21)

x:Phys	path_motion(x,y)	y:Location(w)

Observe that the literal sense is also seen when the location position is not inherently an entity of type *location*, but can be construed as "having a location."

(22) a. The musician arrived at the piano, ready to play.
b. We arrived at the door late.

In both these cases, the physical object in the locative-PP is type coerced to a location interpretation (REF TO GL PAPERS), without, however, changing or embellishing the core sense of the predicate *arrive*. Hence, in (22a), for example, the motion of the musician ends at a location, here construed as the location of the piano.[3]

Similarly, the sentences in (23) exhibit coercion in the locative-PP position, but no additional sense extension.

(23) a. The guests arrived at the concert.
b. Mary arrived at the talk.

In both these cases, the internal argument is an event nominal (*concert* and *talk*), being construed as the location for that event. This is also a fairly conventional type coercion (REF), and interestingly, does not bring about any sense extensions (as defined here) to the predicate itself. Hence, the type template above is being applied, within which the coercion is applied:

(24) a. The guests (PHYS) arrived at the concert (EVENT).

b.

x:Phys	path_motion(x,y)	y:Location(w:Event)

[3]For composition with the preposition *at*, things are actually more complicated, since it exploits any inherent functionality of the object being construed as having a location. In other words, there is really only one configuration relative to the piano, if one is going to play it conventionally. Similar remarks hold for other cases of functional *at*.

Now consider examples of locations that refer to scalar values of the location attribute itself. This includes nouns such as *location*, *height*, *altitude*, and so forth.

(25) a. The plane arrived at a cruising altitude of 30,000 feet.
b. The climbers arrived at a height of 3,000 meters.

Following from this behavior, we see that other scalar attributes can be modulated to fit into this selectional pattern. That is, nouns such as *weight* and *temperature* , which are scalar stative attributes, seem to fit easily into an extended sense of the meaning of the verb.

(26) a. The room finally arrived at a comfortable temperature.
b. Mary has not arrived at her goal weight of 125 lbs. yet.

We will refer to these examples as "weak metaphoric" extensions, licensed by a sense-extending type weakening of the LOCATION argument, to its immediate supertype, i.e., SCALAR_ATTRIBUTE. This gives rise to the following type template for the weak metaphoric reading above in (26).

(27)

x:Phys	path_change(x,y)	y:Scalar _Attribute(x)

The final example we look at is a further modulation of the sense, where other stative predicates are modulated to fit into a scalar-like interpretation. These are illustrated below.

(28) a. The women arrived at an agreement after much discussion.
b. The scientist finally arrived at a solution to the problem.

In these examples, which we will refer to as "strong metaphoric" extensions resulting from further type abstraction, the state is actually a result of some process or activity, and this is the activity which is construed as the path in the verb argument structure.

(29)

x:Phys	path_change(x,y)	y:State(x)

While the specific realizations of these sense extension processes vary, a similar sense mutation pattern is easily seen in sense inventories of other motion predicates. Consider the manner of motion verb *crawl*, with one of its anchor senses exemplified in (30).

(30) The baby crawled to its mother.
"To move slowly in a prone position, by dragging the body along close to the ground."

The weak metaphoric extension of this sense retains the slow motion component, while relaxing the type requirements on the external argument to $phys$.

(31) The cars were crawling along the road.
"To walk, go, or move along with a slow and dragging motion."

Now consider the strong metaphoric sense extension shown in (32) below, which results from the type weakening of both internal and external arguments.

(32) The book crawled to a conclusion.

Here we have the manner expressed without the core component of the motion from the senses given in (30) and (31). The movement is abstracted while retaining the manner of a slow rate of progression.

A specific syntactic pattern shown in (33) is associaed with another metaphoric extension of the same anchor sense.

(33) The field is crawling with policemen.

Here, the constraint is that the extended sense requires the spatial location of the activity as the subject, and the acting moving objects must be in the PP object position.

3.3 Case Study II: Locative Relation Predicates

In this section, we examine sense extension patterns observed for location relation predicates, such as *glue, bind, attach, fasten*, and *anchor*. For the present discussion, we will use the verb *anchor* for illustration.

In the CPA verb lexicon (CPA, 2009), the verb *anchor* is associated with four main patterns. The first two patterns correspond to the literal, primary sense, as in (34a) and (34b), with the causative and inchoative alternations accounting respectively for 21% and 15% of all instances. Examples (34c) and (34d) illustrate the other two patterns, which account for 20% and 40%, respectively.

(34) a. The boat was anchored several hundred yards offshore.
b. After several attempts at finding a landing spot, the vessel gave up and anchored in midstream.[4]
c. The lid was anchored to the sides by screws.
d. A written constitution must be anchored in the idea of universal citizenship.

Consider the argument structure specification for the literal sense, as given in (39). $boat\_location \sqsubseteq location$

(35) $\lambda z \lambda y \lambda x \exists w \left[\begin{array}{l} \textbf{anchor} \\ \text{AS} = \left[\begin{array}{l} \text{ARG1} = x : human \\ \text{ARG2} = y : boat \\ \text{D-ARG1} = z : location \\ \text{S-ARG1} = w : anchor \end{array} \right] \text{QUALIA} = \left[\begin{array}{l} \text{F} = connected(y, z) \\ \text{A} = attach(x, w, z) \end{array} \right] \end{array} \right]$

where w is understood as a *partitive connector* between y and z.[5] As with other alternating verbs, the above semantic representation corresponds to both the causative and the inchoative realizations of the verb *anchor*, as seen in (34a) and (34b), respectively.

The argument structure for this sense specifies four arguments:

(37) x: crew, y: boat, z: location, w: anchor

Notice that only the true arguments, x and y, allow coercion-derived sense extensions, as illustrated below.

(38) a. coerce(x: *country* $\rightarrow$ *human*) :
The US(x) anchored its Pacific Fleet in the Gulf.

b. coerce(y: *organization* $\rightarrow$ *phys*) :
The US anchored its Pacific Fleet(y) in the Gulf.

So far we have seen no modification in the argument structure, and the sense of the predicate has remained the same. Let us now examine the

[4]Note that in this case, the unaccusative carries a causative interpretation.

[5]Note that if the default argument $z : location$ is used, it has to more informative than its default type, by virtue of either subtyping on the default type $boat\_location$, or by virtue of supertyping, e.g.

(36) a. The boat was anchored in the *first* mooring.
b. The ship was anchored in Lyon.

extended senses corresponding to CPA patterns three and four (cf. (34c) and (34d)), which exhibit the following variation.

Notice how the verb sense of *anchor* is being slightly extended in sentence (34c). Here, the literal interpretation of the predicate is generalized slightly by abstracting on the shadow argument and weakening the typing requirements on the remaining arguments witin the same domain of interpretation. At the same time the default argument $z : location$ is promoted to a true argument. This results in a *sense-preserving type weakening*, whose lexical representation is illustrated below:

(39) $\lambda w \lambda z \lambda y \lambda x \left[\begin{array}{l} \textbf{anchor} \\ \text{AS} = \left[\begin{array}{l} \text{ARG1} = x : \top \\ \text{ARG2} = y : phys \\ \text{D-ARG1} = z : phys \\ \text{S-ARG1} = w : connector(y) \end{array} \right] \text{QS} = \left[\begin{array}{l} \text{F} = connected(y, z) \\ \text{A} = attach(x, w, z) \end{array} \right] \end{array} \right]$

The modified argument structure for this sense can be summarized as follows:

(40) x: TopType, y: phys, z: phys, w: connector

where z is interpreted as a region, i.e. a $phys$ interpreted as a $location$, and w is realized as glue, screws, sealant, and so forth, as illustrated below:

(41) The base boards were screwed into position through the side, in order to anchor the sides to the base.

Notice that while the argument structure changed and the typing requirements have generalized, the qualia structure of the predicate hasn't changed and and specifies the change in connectedness (cf. (39)).

We now turn our attention to the metaphoric extensions that modify the core meaning of the predicate. We identify two types of metaphoric extension that correspond to the fourth CPA pattern (cf. (34d)), weak and strong metaphoric extension, as illustrated in (42) and (43).

(42) Weak metaphoric extension:
a. Germany is now firmly anchored within the European Union.
b. A community is a social whole which is usually anchored in space.

(43) Strong metaphoric extension:
a. His behavior is no longer anchored in reality.
b. Constitution must be anchored in the idea of universal citizenship.

These metaphoric extensions are produced from the literal anchor sense by further weakening y and z, and again promoting $z : location$ to a true argument. In the weak metaphoric extension, both y and z are of the same type, and are not $abstract$, while in the strong metahporic extension case, both arguments are promoted to $abstract$.

The distinction made between the weak and the strong metaphoric extension can be illustrated through paraphrasing, e.g.

(44) a. Germany is now firmly *situated* within the European Union.
b. Constitution must be *grounded* in the idea of universal citizenship.
c. *Germany is now firmly *grounded* within the European Union.
d. *Constitution must be *situated* in the idea of universal citizenship.

3.4 Extensions Over Other Verb Classes

The sense-extension clines similar to the ones we have discussed above for several verb classes are easily seen in all areas of the lexicon. For example, for the verb *cultivate* below, (46b) shows non-metaphoric sense extension in due to the changes in relative prominence of the arguments, while a strong metaphoric extension shown in (46c) is produced by a *sense-extending type weakening* on the arguments.

(45)
$$\begin{bmatrix} \textbf{cultivate} \\ \text{ARGSTR} = \begin{bmatrix} \text{ARG1} = x : human \\ \text{ARG2} = y : land \end{bmatrix} \end{bmatrix}$$

(46) *cultivate.v*
object
a. LAND: garden, land, soil
b. PLANT (non-metaphoric sense extension): plants, flowers, shrubs, food, cabbage, coconut
c. ATTRIBUTE/ABSTRACT (metaphoric extension): talent, aura, reputation, image, impression; relationship, link

Some further generalizations can be made with respect to *sense-extending type weakening* mechanisms. For example, all the opposition predicates, as well as almost any change-of-state predicates, when metaphorized, acquire an aspectual reading. Consider the opposition and change-of-state predicates such as *open, close, fill, empty, kill, die, grow*. When the typing restriction on the direct object is relaxed, any event-denoting noun in the object position invokes the aspectual reading for the predicate, producing a strong metaphoric sense extension:

(47) a. We have to open/close these negotiations.
b. We were obliged to kill the project.
c. There is no way to kill this odor.
d. The conversation died.
e. That remark killed the conversation.

Physical objects and their gating predicates are typed according to a rich variety of semantic classes. Change of state verbs act in many different ways on the different types of objects. When applied to event nominals, they tend to act on the aspectual properties of the event as an interval, producing metaphoric readings that are aspectual in nature.

(48) a. *awaken*: The letter awakened strong feelings for John.

4 Concluding Remarks

In this brief paper, we have touched on the problem of how to analyze distinct but related senses of a predicate through a number of lexical semantic modulations. Specifically, we examined: generalizing the type of the argument; changing the argument structure and relative prominence of arguments; and finally, abstracting the core meaning of the verb itself. We briefly examined sense extensions within two verb classes, motion predicates and locative relations. The results of these preliminary thoughts suggest that degrees of metaphoricity may be modeled, in part, by focusing on specific semantic modulations within the predicate. Further work needs to to be done, including extensions to a broader range of verb classes, as well as other sense modulation phenomena.

References

Apresjan, J. (1973). Regular polysemy. *Linguistics*, 142(5):5–32.

Black, M. (1962). *Models and metaphors: Studies in language and philosophy*. Cornell University Press.

Bloomfield, L. (1933). *Language*. New York: Allen & Unwin.

Bréal, M. (1899). *Essai de smantique*. Paris: Hachette.

Choi, S. and Bowerman, M. (1991). Learning to express motion events in english and korean: The influence of language-specific lexicalization patterns. *Cognition*, 41(1-3):83 – 121.

CPA (2009). Corpus Pattern Analysis.

Cruse, D. A. (1995). Polysemy and related phenomena from a cognitive linguistic viewpoint. In Dizier, P. S. and Viegas, E., editors, *Computational Lexical Semantics*, pages 33–49. Cambridge University Press, Cambridge, England.

Cruse, D. A. (2000). *Meaning in Language, an Introduction to Semantics and Pragmatics*. Oxford University Press, Oxford, United Kingdom.

Fass, D. (1988). *Collative semantics: A semantics for natural language*. PhD thesis, Doctoral dissertation, University of Essex.

Hanks, P. (2006). Metaphoricity is gradable. *Corpus-Based Approaches to Metaphor and Metonymy*, pages 17–35.

Hanks, P. and Pustejovsky, J. (2005). A pattern dictionary for natural language processing. *Revue Française de Linguistique Appliquée*.

Jackendoff, R. (1983). *Semantics and Cognition*. MIT Press.

Jespersen, . (1917). *Negation in English and other language*. Copenhagen: A. F. Hst.

Kilgarriff, A. (1997). I don't believe in word senses. *Computers and the Humanities*, 31:91–113.

Krishnakumaran, S. and Zhu, X. (2007). Hunting elusive metaphors using lexical resources. In *Proceedings of the Workshop on Computational Approaches to Figurative Language*, pages 13–20. Association for Computational Linguistics.

Lakoff, G. (1993). The contemporary theory of metaphor. In Ortony, A., editor, *Metaphor and thought*, pages 202–251. Cambridge University Press.

Palmer, M., Gildea, D., and Kingsbury, P. (2005). The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.

Pustejovsky, J. (1995). *Generative Lexicon*. Cambridge (Mass.): MIT Press.

Pustejovsky, J. (2006). Type theory and lexical decomposition. *Journal of Cognitive Science*, 6:39–76.

Rumshisky, A. (2008). Resolving polysemy in verbs: Contextualized distributional approach to argument semantics. *Distributional Models of the Lexicon in Linguistics and Cognitive Science, special issue of Italian Journal of Linguistics / Rivista di Linguistica*.

Ruppenhofer, J., Ellsworth, M., Petruck, M., Johnson, C., and Scheffczyk, J. (2006). *FrameNet II: Extended Theory and Practice*.

Stern, G. (1931). *Meaning and change of meaning with special reference to the English language*. Göteborg: Elander.

Talmy, L. (1985). Lexicalization patterns: Semantic structure in lexical forms. In Shopen, T., editor, *Language typology and semantic description Volume 3: Grammatical categories and the lexicon*, pages 36–149. Cambridge University Press.

Talmy, L. (2000). *Towards a cognitive semantics*. Massachusetts Institute of Technology.

Ullmann, S. (1962). *Semantics: An introduction to the science of meaning*. Oxford: Blackwell.

Zharikov, S. and Gentner, D. (2002). Why do metaphors seem deeper than similes. In *Proceedings of the twenty-fourth annual conference of the cognitive science society*, pages 976–981.