

Double Trouble: The Problem of Construal in Semantic Annotation of Adpositions

Jena D. Hwang and Archna Bhatia
IHMC
{jhwang, abhatia}@ihmc.us

Na-Rae Han
University of Pittsburgh
naraehan@pitt.edu

Tim O’Gorman
University of Colorado Boulder
timothy.ogorman@colorado.edu

Vivek Srikumar
University of Utah
svivek@cs.utah.edu

Nathan Schneider
Georgetown University
nathan.schneider@georgetown.edu

Abstract

We consider the semantics of prepositions, revisiting a broad-coverage annotation scheme used for annotating all 4,250 preposition tokens in a 55,000 word corpus of English. Attempts to apply the scheme to adpositions and case markers in other languages, as well as some problematic cases in English, have led us to reconsider the assumption that an adposition’s lexical contribution is equivalent to the role/relation that it mediates. Our proposal is to embrace the potential for **construal** in adposition use, expressing such phenomena directly at the token level to manage complexity and avoid sense proliferation. We suggest a framework to represent both the scene role and the adposition’s lexical function so they can be annotated at scale—supporting automatic, statistical processing of domain-general language—and discuss how this representation would allow for a simpler inventory of labels.

1 Introduction

Prepositions and postpositions (collectively **adpositions**) are widespread in the world’s languages as grammatical markers expressing spatial, temporal, thematic,¹ and other kinds of semantic relations. Unfortunately for semantic processing, a handful of high-frequency types carry an immense payload by way of extreme polysemy. Thus, disambiguation of adpositional meaning is crucial to piecing together the interpretation of a sentence (§2).

A line of previous work (Srikumar and Roth, 2013a; Schneider et al., 2015, 2016, see §2) has developed a scheme for broad-coverage annotation

of adpositions with an eye toward building automatic disambiguation systems. Their most recent proposal consists of an inventory of 75 categorical labels known as **supersenses** that characterize the polysemy of English prepositions in a lexically-neutral and coarse-grained fashion. They envision disambiguation as assigning a single one of these supersenses to each preposition token.

While formalizing disambiguation via single-label classification works well for prototypical members of the categories, on closer examination, we argue that it is overly simplistic for many usages. This became particularly evident when we tried to adapt the English-centric supersense labels to other languages.

Here we advance a more nuanced view that an adposition can contribute a semantic perspective, or **construal**, over and above the scenario relation that its object participates in. We argue that it is essential to distinguish the contribution of the preposition itself, i.e., what the adposition **codes** for, from the semantic role or relation that the adposition mediates and that a predicate or scene **calls** for; and as a result, the label that would be most appropriate is underdetermined for many tokens (§3). In our view, the mismatch can be understood through the lens of **construal**, and this should be made explicit in corpora (§4).

To that end, we sketch an annotation approach that disentangles the two elements of the meaning while retaining the advantages of a broad-coverage (rather than lexicographic-sense-based) scheme. Preliminary analysis suggests that this scheme will work well not only for English, but also for the other languages examined. §5 surveys some of the phenomena that our new analysis addresses with examples from multiple languages; §6 suggests that this added flexibility at the token level removes the need for a great deal of complexity in the supersense inventory itself: i.e., we can get away with

¹In the sense of thematic roles (agent, patient, etc.).

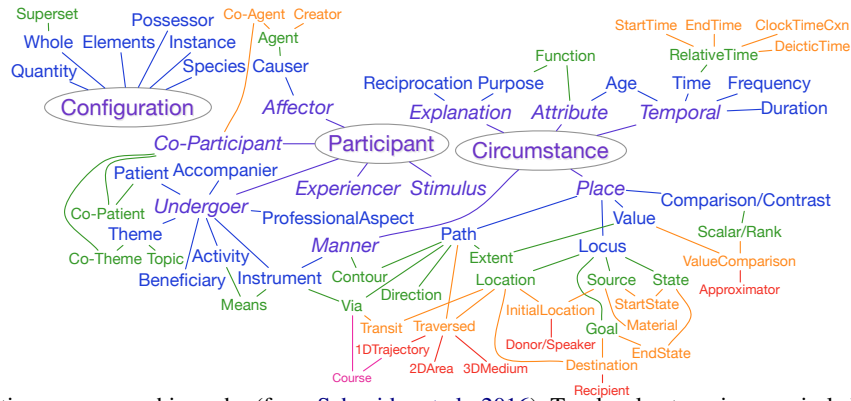


Figure 1: Preposition supersense hierarchy (from Schneider et al., 2016). Top-level categories are circled and subcategories radiate outward.

1/3 fewer semantic labels, reducing somewhat the practical concern of sparse data. Here we also discuss challenges and tradeoffs inherent in the proposed approach. We have begun testing the workability of this proposal empirically by annotating data in multiple languages, with disambiguation experiments to follow.

2 Approaches to Prepositional Polysemy

The most frequent English prepositions are extraordinarily polysemous. For example, the preposition **at** expresses different information in each of the following usages:

- (1) a. It is **at** 123 Main St. (**LOCATION**)
- b. We met him **at** 7pm. (**TIME**)
- c. Everyone pointed **at** him. (**GOAL**)
- d. She laughed **at** my acting. (**STIMULUS**)
- e. He held her **at** gunpoint. (**INSTRUMENT**)

When confronted with a new instance of **at**, NLU systems must determine whether it marks an entity or scene’s location, time, goal, or something else.

As lexical classes go, prepositions are something of a red-headed stepchild in the linguistics literature. Most of the semantics literature on prepositions has revolved around how they categorize space and time (e.g., Herskovits, 1986; Verkuyl and Zwarts, 1992; Bowerman and Choi, 2001). However, there have been a couple of lines of work addressing preposition semantics broadly. In cognitive linguistics, studies have examined abstract as well as concrete uses of English prepositions (e.g., Dirven, 1993; Lindstromberg, 2010). Notably, the polysemy of **over** and other prepositions has been explained in terms of sense networks encompassing core senses and motivated extensions (Brugman, 1981; Lakoff, 1987; Dewell, 1994; Tyler and Evans, 2001, 2003). The Preposition Project (TPP; Litkowski and Hargraves, 2005) broke ground in

stimulating computational work on fine-grained word sense disambiguation of English prepositions (Litkowski and Hargraves, 2005; Ye and Baldwin, 2007; Tratz and Hovy, 2009; Dahlmeier et al., 2009). Typologists, meanwhile, have developed *semantic maps* of functions, where the nearness of two functions reflects their tendency to fall under the same adposition or case marker in many languages (Haspelmath, 2003; Wälchli, 2010).

Preposition supersenses. Following Srikumar and Roth (2013b), Schneider et al. (2015) developed coarse-grained semantic categories of prepositions as a broader-coverage alternative to fine-grained senses, using categories similar to those appearing in semantic maps (**LOCATION**, **RECIPIENT**, etc.) rather than lexicalized senses. Schneider et al. (2015) refined their inventory of categories through extensive deliberation involving the use of dictionaries, corpora, and pilot annotation experiments. They call the categories **supersenses** to emphasize their similarity to coarse-grained classifications of nouns and verbs that go by that name (Ciaramita and Altun, 2006; Schneider et al., 2012).

The **at** examples in (1) are accompanied by the appropriate supersenses from the supersense scheme. Most supersenses resemble thematic roles (cf. Fillmore (1968)); a few others are needed to describe preposition-marked relations between entities. There are multiple English prepositions per supersense; e.g., “**in** the city” and “**on** the table” would join “**at** 123 Main St.” in being labeled as **LOCATIONS**. We understand the supersenses as prototype-based categories, and in some cases use heuristics like paraphrasability (“in order to” for **PURPOSE**) and WH-question words (“Why?” for **PURPOSE** and **EXPLANATION**) to help determine which tokens are instances of the category.

The 75 supersenses are organized in a taxonomy based on that of VerbNet (Bonial et al., 2011), with

PARTICIPANT, **CIRCUMSTANCE**, and **CONFIGURATION** at the top level.² The taxonomy uses multiple inheritance to account for subcategories which are considered to include properties of multiple supercategories. The full hierarchy is in figure 1.

The approach to preposition annotation is *comprehensive*, i.e., every token of every preposition type is given a supersense label. The supersenses were applied to annotate a 55,000 word corpus of online reviews in English, covering all 4,250 preposition tokens (Schneider et al., 2016). For each token, annotators chose a single label from the inventory. This is not an easy task, but with documentation of many examples in a lexical resource, **PrepWiki**,³ trained university students were able to achieve reasonable levels of inter-annotator agreement. Every token was initially labeled by at least two independent annotators, and differences were adjudicated by experts.

3 Problems with Preposition Supersenses

While the above approach works reasonably well for most English tokens, difficulties in directly applying the scheme to adpositions and case markers in other languages, as well as some of the persistent issues arising in English, have led us to conclude that perhaps the supersense hierarchy as it stands is too simplistic to provide a faithful account of the prepositions' semantic behavior. This has caused us to examine fundamental assumptions made by previous work and reevaluate what it means to semantically label an adposition.

3.1 Semantic Overlap in English

In the original English annotation (Schneider et al., 2016), a few phenomena caused much hand-wringing—not because there was no appropriate supersense, but because *multiple* supersenses seemed to fit. For example, it was observed that **TOPIC** and **STIMULUS** could compete for semantic territory. (2) evinces related usages of **about** with different governors:

²These loosely correspond to event arguments, adjuncts, and adnominal complements, respectively. However, supersense organization does not make any claims with regard to coreness or the argument/adjunct distinction, as there are many phenomena that do not conform to either of the prototypes for argument and adjunct (for a review of the literature on the argument/adjunct distinction, see Hwang, 2011). We are also not convinced that a firm distinction between lexical and non-lexical/functional adpositions (Rauh, 1993) can be established, though the relevance of this distinction in the context of the construal approach merits further investigation.

³<http://tiny.cc/prepwiki>

- (2) a. I read [a book **about** the strategy].
b. I read **about** the strategy.
c. I knew **about** the strategy.
d. I cared **about** the strategy.

Usages (2a–2c) could reasonably be labeled as **TOPIC**. This is because the **about**-PP indicates what is communicated (2a, 2b) and known (2c). The fourth example (2d), however, presents an overlap in its interpretation. On the one hand, traditional thematic role inventories include the category **STIMULUS** for something that prompts a perceptual or emotional experience, as in (3).

- (3) I was afraid **of** the strategy.

Surely, *cared* in (2d) describes an emotional state, so **about** marks the **STIMULUS**. However, much like examples (2a–2c), the semantics relating to **TOPIC** is still very much present in the use of **about**, drawing attention to the aspects of the caring process involving thought or judgement. This contrasts with the use of **for** in “I cared **for** my grandmother,” where the prepositional choice calls attention to the benefactive aspect of the caring act.

If we are constrained to one label per argument, where should the line be drawn between **STIMULUS** and **TOPIC** in cases of overlap? In other words, should the semantic representation emphasize the semantic commonality between all of the examples in (2), or between (2d) and (3)?

Observing that annotators were inconsistent on such tokens, Schneider et al. (2016) drew a boundary between **TOPIC** and **STIMULUS** in an attempt to force consistency, stating that “**TOPIC** should be used if the governor is a predicate of communication or of ‘higher-level’ cognition—i.e., primarily mental/intellectual rather than emotional/perceptual/bodily in nature”.⁴ This criterion seems artificial to us; at the very least, it splits hairs in a way that would be difficult to explain to annotators.

Below, we instead argue that the idea of construal/conceptualization offers a more principled answer; in our new analysis, the **TOPIC** suggested by **about** and the **STIMULUS** suggested by *cared* can coexist.

3.2 Applying the Supersenses to Other Languages

One of the premises of using unlexicalized supersenses was that the scheme would port well to other

⁴<http://tiny.cc/prepwiki/index.php/Category:SST-Topic>

languages (as the WordNet noun and verb supersenses have: Picca et al., 2008; Schneider et al., 2012, *inter alia*). To test this, we have begun applying the existing supersenses to three new languages, namely, Hebrew, Hindi, and Korean. Pilot annotation in these languages has echoed the fundamental problem discussed in the previous section.

Consider the Hindi examples below. In (4a), the experiencer of an emotion is marked with a postposition **kaa**, the genitive case marker in Hindi.

- (4) a. [Hindi]: **EXPERIENCER** vs. **POSSESSOR**
 bipaashaa **kaa** gussaa
 Bipasha **GEN** anger
 “Bipasha’s anger”
- b. [Hindi]: **EXPERIENCER**
 bipaashaa bahut gussaa hui
 Bipasha very angry became
 “Bipasha got very angry.”

The use of **kaa** strongly suggests possession (in (4), possession of an abstract quality). However, the semantics of the phrase also includes **EXPERIENCER**—thus, it seems inappropriate to force a choice between **EXPERIENCER** and **POSSESSOR** for this token. (The same problem is seen in a similar phrase “the anger **of** Bipasha” in English.) There are other ways to attribute anger to Bipasha—e.g., see (4b). Here Bipasha is not construed as a possessor when the postposition **kaa** is not used.

Our preliminary annotation of Hindi, Korean, and Hebrew has suggested that instances of overlap between multiple supersenses are fairly frequent.

4 The Construal Analysis

Why do “cared **about** the strategy” in (2d) and “anger **of** Bipasha” (cf. (4a)) not lend themselves to a single label? These seem to be symptoms of the fact that no English preposition prototypically marks **EXPERIENCER** or **STIMULUS** roles, though from the perspective of the predicates, such roles are thought to be important generalizations in characterizing events of perception and emotion. In essence, there is an apparent mismatch between the roles that the verb *care* or the noun *anger* calls for, and the functions that English prepositions prototypically code for. While **about** prototypically codes for **TOPIC** and **of** prototypically codes for **POSSESSOR**, there is no preposition that “naturally” codes for **EXPERIENCER** or **STIMULUS** in the same way. Thus, if a predicate marks an **EXPERIENCER** or **STIMULUS** with a preposition, the preposition

will contribute something new to the conceptualization of the scene being described. With “cared **about** the strategy,” it is **TOPIC**-ness that the preposition brings to the table; with “anger **of** Bipasha,” it is the conceptualization of anger as an attribute that somebody possesses.

Thus, we turn to theories in Cognitive Semantics to define the phenomenon of **construal** as a means of understanding the contributions that are emerging from the adpositions with respect to the expressed event or situation. Then, we propose a method to handle the problem posed by **construal** and to resolve the apparent semantic overlap which is pervasive across languages.

4.1 Defining Construal

The world is not neatly organized into bits of information that map directly to linguistic symbols. Rather, linguistic meaning reflects the priorities and categorizations of particular expressions in a language (Langacker, 1998; Jackendoff, 2002; Croft and Cruse, 2004, ch. 3). Much like pictures of a scene from different viewpoints will result in different renderings, a real-world situation being described will “look” different depending on the linguistic choices made by a speaker. This includes within-language choices: e.g., the choice of “John sold Mary a book” vs. “John sold a book to Mary” vs. “Mary bought a book from John.” In the process called **construal** (a.k.a. **conceptualization**), a speaker “packages” ideas for linguistic expression in a way that foregrounds certain elements of a situation while backgrounding others.

We propose to incorporate this notion of **construal** in adposition supersense annotation. We use the term **scene** to refer to events or situations in which an adpositional phrase plays a role. (We do not formalize the full scene, but assume its roles can be characterized with supersense labels from figure 1.) Contrast the use of the prepositions **by** and **of** in (5):

- (5) a. The festival features works **by** Puccini.
 b. I’m an expert on the works **of** Puccini.

While both prepositional phrases indicate works created by the operatic composer Puccini (i.e., **CREATOR**), the different choices of preposition reflect different **construals**: **by** highlights the agency of Puccini, whereas **of** construes Puccini as the source of his composition. Thus, “works **by** Puccini” and “works **of** Puccini” are paraphrases, but present subtly different portrayals of the relationship between

Puccini and his works. In other words, these phrases are not identical in meaning because the preposition carries with it different nuances of construal. In this paper, we focus on differences in construal manifested in different adposition choices, and the possibility that an adposition construal complements the construal of a scene and its roles (as evoked by the governing head or predicate).

For instances like “I read **about** the strategy” in (2b) that were generally unproblematic for annotation under the original preposition guidelines, the semantics of the adposition and the semantic role assigned by the predicate are congruent. However, for examples like “cared **about** the strategy” in (2d) and “anger **of** Bipasha” in (4a), we say that the adposition construes the role as something other than what the scene specifies. Competition between different adposition construals accounts for many of the alternations that are near-paraphrases, but potentially involve slightly different nuances of meaning (e.g., “talk **to** someone” vs. “talk **with** someone”; “angry **at** someone” vs. “angry **with** someone”).

Thus, the notion of construal challenges Schneider et al.’s (2015; 2016) original conception that each supersense reflects the semantic role assigned by its governing predicate (i.e. verbal or event nominal predicate), and that a single supersense label can be assigned to each adposition token. Rather than trying to ignore these construals to favor a single-label approach, or possibly create new labels to capture the meaning distinctions that construals impose on semantic roles, we adopt an approach that gives us the flexibility to deal with both the semantics coming from the scene and the construal evoked by the adpositional choice.

4.2 Formulating a Construal Analysis

We address the issues of construal by decoupling the semantics signaled by the adposition from the role expected by the scene. Essentially, we borrow from Construction Grammar (Fillmore et al., 1988; Kay and Fillmore, 1999; Goldberg, 2006) the notion that semantic contributions can be made at various levels of syntactic structure, beginning with the semantics contributed by the lexical items.

Under the original single-label analysis, the full weight of semantic assignment rested on the predicate’s semantic role, with the indirect assumption that the predicate selects for adpositions relevant to the assignment. Under the construal analysis, we assign semantics at both scene and adposition lev-

els of meaning: we capture what the scene *calls* for, henceforth **scene role** and what the adposition itself *codes* for, henceforth **function**. Both labels are drawn from the supersense hierarchy (figure 1). Allowing tokens to be annotated with both a role and a function accounts for the non-congruent adposition construals, as in (6).

- (6) a. The festival features works **by** Puccini.
scene role: CREATOR; function: AGENT
 b. I’m an expert on the works **of** Puccini.
scene role: CREATOR; function: SOURCE

We recognize that both of these sentences carry the meaning represented by the supersense **CREATOR** at the scene level, but also recognize the construal that arises from the chosen preposition: **by** is assigned the function of **AGENT** and **of** is assigned the function of **SOURCE**.⁵

5 Applying the Construal Analysis

In this section, we discuss some of the more productive examples of non-congruent construals in English as well as in Hindi, Korean, and Hebrew. Hereafter, we will use the notation **ROLE**→**FUNCTION** to indicate such construals. Adopting the “realization” metaphor of articulating an idea linguistically, this can be read as “**ROLE** is realized with an adposition that marks **FUNCTION**.”

5.1 Emotion and Perception Construals

Scenes of emotion and perception (Dirven, 1997; Osmond, 1997; Radden, 1998) provide a compelling case for the construal analysis. Consider the sentences involving emotion in example (7):

- (7) a. I was scared **by** the bear.
STIMULUS→**CAUSER**
 b. I was scared **about** getting my ears pierced.
STIMULUS→**TOPIC**

Comparing examples (7a) and (7b), we notice that there are two different types of stimuli represented in otherwise semantically parallel sentences.

⁵We also acknowledge that there is a level of construal contributed by the verb. For example, *Alex* in *Alex sent the package to Pam* can be **AGENT** or **SOURCE** depending whether the interpretation is focused on the agency of the argument or the spatial relation it has in reference to the action described by the verb. These verb-triggered construals have been previously explored, most notably by Jackendoff (1990). Perspective can also be evident in the choice of syntactic constructions, e.g., active vs. passive voice (*I made a mistake* versus *Mistakes were made*), which can be connected to sentiment (Greene and Resnik, 2009). We specifically focus on the construal that arises from the adposition in a given sentence.

The preposition **by** gives the impression that the stimulus is responsible for triggering an instinctive fear reflex (i.e., **CAUSER**), while **about** portrays the thing feared as the content or **TOPIC** of thought.⁶

In some languages, the experiencer can be conceptualized as a recipient of the emotion or feeling, thus licensing dative marking.⁷ In the Hebrew example (8a), the experiencer of bodily perception is marked with the dative preposition **l(e)-** (Berman, 1982). Similarly, in Hindi, the dative postposition **-ko** marks an experiencer in (8b).

- (8) a. [Heb.]: **EXPERIENCER**~>**RECIPIENT**
Koev l-i ha-rosh
Hurts **DAT**-me the-head
“My head hurts.”
- b. [Hindi]: **EXPERIENCER**~>**RECIPIENT**
mujh-ko garmii lag rahii hai
I-**DAT** heat feel PROG PRES
“I’m feeling hot.”

Contrast this with examples where scene role and adposition function are congruent:

- (9) a. I ate dinner **at** 7:00. **TIME**~>**TIME**
b. Let’s talk **about** our plan. **TOPIC**~>**TOPIC**

In (9a) and (9b), the preposition is prototypical for the given scene role and its function directly identifies the scene role. Because the semantics of the role and function are congruent, these cases do not exhibit the extra layer of construal seen in (7) and (8).⁸ In essence, our analysis helps capture the construals that characterize the less prototypical scene role and function pairings.

5.2 Professional Associate Construals

The online reviews corpus (Schneider et al., 2016) shows that, at least in English, professional relation-

⁶Interestingly, “scared **about**” seems to require an explicit or metonymic event/situation as the complement. Thus, “scared **about** the bear” would be felicitous to describe apprehension about some mischief that the bear might get up to. It would be less than felicitous to describe a hiker’s reaction upon being surprised by a bear.

⁷English displays this to a limited extent: “It feels/seems/looks perfect **to** me.”

⁸One might object that most or all adpositions impose a spatial construal—and thus, (9a) should be annotated as **TIME**~>**LOCATION**. We do not discount the possibility that such a metaphor can be cognitively active in speakers using temporal adpositions; in fact, there is considerable evidence that time-as-space metaphors are cross-linguistically pervasive and productive (Lakoff and Johnson, 1980; Núñez and Sweetser, 2006; Casasanto and Boroditsky, 2008). However, we do not see much practical benefit to annotating temporal **at** or topical **about** as spatial.

ships (especially employer–employee and business–client ones) are fertile ground for alternating preposition construals. The following were among the examples tagged as **PROFESSIONALASPECT**:

- (10) a. My dad worked **for** a record label.
PROFESSIONALASPECT~>**BENEFICIARY**
- b. Dr. S— **at** CVTS is not a good doctor.
PROFESSIONALASPECT~>**LOCATION**
- c. Nigel **from** Nidd Design has always been great!
PROFESSIONALASPECT~>**SOURCE**
- d. The owners and employees **of** this store ...
PROFESSIONALASPECT~>**POSSESSOR**

All of these construals are *motivated* in that they highlight an aspect of prototypical professional relationships: e.g., an employee’s work prototypically takes place at the business location (hence “work **at**”), though this is not a strict condition for using “work **at**”—the meaning of **at** has been extended from the prototype. Likewise, the pattern “*person* {**at**, **from**, **of**} *organization*” has been conventionalized to signify employment or similar institutional-belonging relationships.

The construal analysis equips us with the ability to use the existing labels like **BENEFICIARY** and **LOCATION** to deal with the overloading of the **PROFESSIONALASPECT** label, instead of forcing a difficult decision or creating several additional categories. This analysis also accounts for similar construals presented by adpositions in other languages. For example, the overlap of **PROFESSIONALASPECT** with **SOURCE**, as seen in English example (10c), occurs in Hindi and Korean as well.

5.3 Static vs. Dynamic Construals

Another source of difficulty in the original annotation came from caused-motion verbs like *put*, which takes a PP indicating part of a path. Sometimes the preposition lexically marks a source or goal, e.g., **into**, **onto**, or **out of** (11a). Often, however, the preposition is prototypically locative, e.g., **in** or **on** (11b), though the object of the preposition is interpreted as a destination, equivalent to the use of **into** or **onto**, respectively. This locative-as-destination construal is highly productive, so analyzing **on** as polysemous between **LOCATION** and **DESTINATION** does not capture the regularity. The PP is sometimes analyzed as a resultative phrase (Goldberg, 2006). In our terms, we simply say that the scene calls for a **DESTINATION**, but the preposition codes for a **LOCATION**:

- (11) a. Cynthia put her things **into** a box.
 DESTINATION~>DESTINATION
 b. Cynthia put her things **on** her bed.
 DESTINATION~>LOCATION

Thus, we avoid listing the preposition with multiple lexical functions for this regular phenomenon.

The opposite problem occurs with fictive motion (Talmy, 1996): a path PP, and sometimes a motion verb, construe a static scene as dynamic as seen in “A road runs **through** my property.” Rather than forcing annotators to side with the dynamic construal effected by the language, versus the static nature of the actual scene, we represent both: the scene role is **LOCATION** (static) and the preposition function is **PATH** (dynamic) (i.e., **LOCATION~>PATH**).

5.4 Metaphoric Scenes

Finally, our analysis gives us a way to handle metaphoric scenes (Lakoff and Johnson, 1980). In (12), the locative-as-destination construal (§5.3) is layered with the states-are-locations metaphor. We annotate the scene in terms of the governing predicate’s **target domain** (domain which we seek to describe), and the adposition function in terms of the **source domain** (domain from which we draw metaphorical expressions to conceptualize the target domain):

- (12) The election news put him **in** a very bad mood.
 ENDSTATE~>LOCATION

Our construal analysis can capture both source and target domains by assigning the source domain meaning to the function of the preposition and the target domain meaning to the scene role.

6 Toward a Revised Hierarchy

The annotation of both scene and function levels of semantics allows us to trade more complexity at the token level for less complexity in the label set. As discussed in §4, separating the scene role and function levels of annotation will more adequately capture construal phenomena without forcing an arbitrary choice between two labels or introducing further complexity into the hierarchy.

In fact, we intend to *simplify* the current super-sense hierarchy, by collapsing some of the finer-grained distinctions that can be accounted for with the construal analysis instead. Candidates for removal include the labels with multiple inheritance such as **CONTOUR** (inheriting from **PATH** and

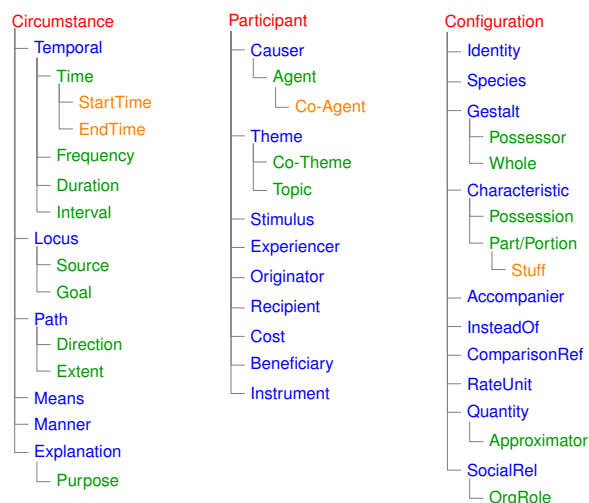


Figure 2: Preliminary revised hierarchy of 50 adposition super-senses.

MANNER; e.g., “The fly flew **in** zig-zags”) and **TRANSIT** (inheriting from **VIA** and **LOCATION**; e.g., “We traveled **by** bus”).

A preliminary proposal for a new hierarchy appears in figure 2.⁹ It weighs in at only 50 categories, a third fewer than the original 75. A significantly smaller inventory will both ease the cognitive burden on annotators and reduce the sparsity of labels in the data, which should facilitate better statistical generalizations with limited data.

The added representational complexity of construals seems justified to account for many of the phenomena discussed above, especially as the project grows to include more languages. But is the complexity worth it on balance? We consider some of the tradeoffs below.

6.1 Challenges in Function Assignment

We encountered several examples in which function labels are difficult to identify. Consider the following paraphrases:

- (13) a. [Korean]: **LOCATION~>LOCATION**
 Cheolsu-nun undongcang-**eyse** tallyessta.
 Cheolsu-NOM schoolyard-**at** ran.
 “Cheolsu ran in the schoolyard.”
 b. [Korean]: **LOCATION~>?**
 Cheolsu-nun undongcang-**ul** tallyessta.
 Cheolsu-NOM schoolyard-**ACC** ran.
 “Cheolsu ran in the schoolyard.”

⁹Apart from changes enabled by the construal analysis, a number of other simplifications and enhancements are incorporated into the proposal, which space does not allow us to enumerate here: for example, collapsing **LOCUS/LOCATION/STATE**, **SOURCE/INITIALLOCATION/INITIALSTATE**, **GOAL/DESTINATION/ENDSTATE**, and **TIME/RELATIVETIME/CLOCKTIME/CXN**; and replacing **PROFESSIONALASPECT** with **SOCIALREL** and **ORGROLE**.

In (13a), “schoolyard” is accompanied by a postposition **-eyse** (comparable to English **at**), which marks it as the location of running. This is the unmarked choice. On the other hand, in sentence (13b), the noun is paired with the accusative marker **-ul**, the marked choice. The use of **-ul** evokes a special construal: it indicates that the schoolyard is more than just a backdrop of the running act and that it is a location that Cheolsu mindfully chose as the place of action. Additionally, marking the location with the accusative marker, pragmatically, brings focus to the noun (i.e., he ran in a schoolyard as opposed to anywhere else). Such construals are not limited to locations, but may also include other scene roles such as **GOAL** and **ACCOMPANIER**, in alternation with postpositions that can express those functions. Since accusative case markers generally serve syntactic functions over semantic ones, it may be difficult to identify a semantic function the accusative marker carries.

A similar phenomenon can be found in Hindi:

- (14) a. [Hindi]: bare NP as **DESTINATION**
 maiN library jaa rahii thii
 I library go PROG PST
 “I was going to the library.”
 b. [Hindi]: **DESTINATION**~>?
 maiN library-ko jaa rahii thii
 I library-ACC go PROG PST
 “I was going to the LIBRARY.” [more emphasis on the library]

This suggests that, apart from spatiotemporal relations and semantic roles, adpositions can mark **information structural** properties for which we would need a separate inventory of labels.

In some idiomatic predicate–argument combinations, the semantic motivation for the preposition may not be clear (15).

- (15) a. Listen **to** the violin! **STIMULUS**~>?
 b. What’s he proudest **of**? **STIMULUS**~>?
 c. Unhappy **with** my meal! **STIMULUS**~>?
 d. I’m interested **in** politics. **TOPIC**~>?

While the scene role in (15a) and (15b) is clearly **STIMULUS**, the function is less clear. Is the object of attention construed (metaphorically) as a **GOAL** in (15a), and the cause for pride as a **SOURCE** in (15b)? Or are **to** and **of** semantically empty argument-markers for these predicates (cf. the “case prepositions” of Rauh, 1993)? We do not treat either combination as an unanalyzable multiword

expression because the ordinary meaning of the predicate is very much present. (15c) and (15d) are similarly fraught. But as we look at more data, we will entertain the possibility that the function can be null to indicate a marker which contributes no lexical semantics.

6.2 Challenges in Scene Role Assignment

There are complications which we are not yet prepared to fully address. First, if the PP is not governed by a predicate which provides the roles—such as a verb or eventive/relational noun—the preposition may need to evoke a meaning more specific than our labels. E.g., for “children **in** pajamas” and “woman **in** black,” **in** may be taken to evoke the semantics of wearing clothing.¹⁰ The label set we use for broad-coverage annotation is, of course, vaguer, and would simply specify **ATTRIBUTE** for the clothing sense of **in**. Copular constructions raise similar issues. Consider “It is **up to** you to decide,” meaning that deciding is the addressee’s responsibility: this idiomatic sense of **up to** is closer to a semantic predicate than to a semantic role or figure-ground relation.

6.3 Multi-Construal Analysis?

In rare instances, we are tempted to annotate a chain of extensions from a prototypical function of a preposition, which we term **multiple construal**. For instance:

- (16) a. Bob’s boss yelled **at** him for his mistake.
RECIPIENT~>**BENEFICIARY**~>**GOAL**
 b. Jane was angry **at** him for his mistake.
STIMULUS~>**BENEFICIARY**~>**GOAL**
 c. I was involved **in** the project.
THEME~>**SUPERSET**~>**LOCATION**

“Yelled **at**” in (16a) is a communicative action whose addressee (**RECIPIENT**) is also a target of the negative emotion (**BENEFICIARY**~>**GOAL**: compare the use of **at** in “shoot **at** the target”). (16b) is similar, except “angry” focuses on the emotion itself, which Bob is understood to have evoked in his boss.

With regard to (16c), the item “involved **in**” has become fossilized, with **in** marking an underspecified noncausal participant (hence, **THEME** as the scene role). At the same time, one can understand

¹⁰Indeed, this is the position adopted by version 1.7 of FrameNet, where **in** is listed as a lexical unit of the WEARING frame (<https://framenet2.icsi.berkeley.edu/fnReports/data/frame/Wearing.xml>).

the **in** here as motivated by the member-of-set sense (cf. “I am **in** the group”), which would be labeled **SUPERSET**→**LOCATION** because it conceptualizes membership in terms of containment. A similar logic would apply to “people **in** the company”: **PROFESSIONALASPECT**→**SUPERSET**→**LOCATION**. Effectively, the multiple construal analysis claims that multiple steps of extending a preposition’s prototypical meaning remain conceptually available when understanding an instance of its use. That said, we are not convinced that this logic could be applied reliably by annotators, and thus may simplify the usages in (16) to just the first and second or the first and third labels.

6.4 The Annotation Process

Annotators are generally capable of interpreting meaning in a given context. However, it might be difficult to train annotators to develop intuitions about adposition functions, which reflect prototypical meanings contributed by the lexical item that may not be literally applicable. These distinctions may be too subtle to annotate reliably. As we are approaching this project with the goal of producing annotated datasets for training and evaluating natural language understanding systems, it is an important concern.

We are currently planning pilot annotation studies to ascertain (i) the prevalence of the role vs. function mismatches, and (ii) annotator agreement on such instances. Enshrining role–function pairs in the lexicon may facilitate inter-annotator consistency: our experience thus far is that annotators benefit greatly from examples illustrating the possible supersenses that can be assigned to a preposition.

If initial pilots are successful, we would then need to decide whether to annotate the role and function together or in separate stages. Because the function reflects one of the adposition’s prototypical senses, it may often be deterministic given the adposition and scene role, in which case we could focus annotators’ efforts on the scene roles. Existing annotations for lexical resources such as PropBank (Palmer et al., 2005), VerbNet (Palmer et al., 2017; Kipper et al., 2008), and FrameNet (Fillmore and Baker, 2009) might go a long way toward disambiguating the scene role, limiting the effort required from annotators.

6.5 Linguistic Utility of Annotated Data

Assuming the above theoretical and practical concerns are surmountable, annotated corpora would facilitate empirical studies of the nature and limits of adposition/case construal within and across languages. For example: Is it the case that some of the supersense labels can only serve as scene roles, or only as functions? (A hypothesis is that **PARTICIPANT** subtypes tend to be limited to scene roles, but this needs to be examined empirically.) Which role–function pairs are attested in particular languages, and are any universal? Thus far we have seen that certain scene roles, such as **EXPERIENCER**, **STIMULUS**, and **PROFESSIONALASPECT**, invite many different adposition construals—is this universally true? As adpositions are notoriously difficult for second language learners, would it help to explain which construals do and do not transfer from the first language to the second language?

7 Conclusion

We have considered the semantics of adpositions and case markers in English and a few other languages with the goal of revising a broad-coverage annotation scheme used in previous work. We pointed out situations where a single supersense did not fully characterize the interaction between the adposition and the scene elaborated by the PP. In an attempt to tease apart the semantics contributed specifically by the adposition from the semantics coming from elsewhere, we proposed a construal analysis. Though many details remain to be worked out, we are optimistic that our analysis will ultimately improve broad-coverage annotations as well as constructional analyses of adposition behavior.

Acknowledgments

We thank the rest of our CARMLS team—Martha Palmer, Ken Litkowski, Omri Abend, Katie Conger, and Meredith Green—for participating in weekly discussions of adposition semantics; Michael Ellsworth for an insightful perspective on construal, Paul Portner for a helpful clarification regarding approaches to conceptualization in the literature, and anonymous reviewers for their thoughtful comments. We also thank the participants in the AAI Spring Symposium on Construction Grammar in NLU, held earlier this year at Stanford, where an early version of this work was presented.

References

- Ruth A. Berman. 1982. Dative marking of the affectee role: data from Modern Hebrew. *Hebrew Annual Review* 6:35–59.
- Claire Bonial, William Corvey, Martha Palmer, Volha V. Petukhova, and Harry Bunt. 2011. A hierarchical unification of LIRICS and VerbNet semantic roles. In *Fifth IEEE International Conference on Semantic Computing*. Palo Alto, CA, USA, pages 483–489.
- Melissa Bowerman and Soonja Choi. 2001. Shaping meanings for language: universal and language-specific in the acquisition of spatial semantic categories. In Melissa Bowerman and Stephen Levinson, editors, *Language Acquisition and Conceptual Development*, Cambridge University Press, Cambridge, UK, pages 475–511.
- Claudia Brugman. 1981. *The story of ‘over’: polysemy, semantics and the structure of the lexicon*. MA thesis, University of California, Berkeley, Berkeley, CA. Published New York: Garland, 1981.
- Daniel Casasanto and Lera Boroditsky. 2008. Time in the mind: using space to think about time. *Cognition* 106(2):579–593.
- Massimiliano Ciaramita and Yasemin Altun. 2006. Broad-coverage sense disambiguation and information extraction with a supersense sequence tagger. In *Proc. of EMNLP*. Sydney, Australia, pages 594–602.
- William Croft and D. Alan Cruse. 2004. *Cognitive linguistics*. Cambridge University Press, Cambridge, UK.
- Daniel Dahlmeier, Hwee Tou Ng, and Tanja Schultz. 2009. Joint learning of preposition senses and semantic roles of prepositional phrases. In *Proc. of EMNLP*. Suntec, Singapore, pages 450–458.
- Robert B. Dewell. 1994. ‘Over’ again: Image-schema transformations in semantic analysis. *Cognitive Linguistics* 5(4):351–380.
- René Dirven. 1993. Dividing up physical and mental space into conceptual categories by means of English prepositions. In Cornelia Zelinsky-Wibbelt, editor, *The semantics of prepositions: From mental processing to natural language processing*, Mouton de Gruyter, Berlin, pages 73–97.
- René Dirven. 1997. Emotions as cause and the cause of emotions. In Susanne Niemeier and René Dirven, editors, *The Language of Emotions: Conceptualization, expression, and theoretical foundation*, John Benjamins, Amsterdam, pages 55–86.
- Charles J. Fillmore. 1968. The case for case. In Emmon Bach and Robert Thomas Harms, editors, *Universals in Linguistic Theory*, Holt, Rinehart, and Winston, New York, pages 1–88.
- Charles J. Fillmore and Collin Baker. 2009. A frames approach to semantic analysis. In Bernd Heine and Heiko Narrog, editors, *The Oxford Handbook of Linguistic Analysis*, Oxford University Press, Oxford, UK, pages 791–816.
- Charles J. Fillmore, Paul Kay, and Mary Catherine O’Connor. 1988. Regularity and idiomatcity in grammatical constructions: the case of *let alone*. *Language* 64(3):501–538.
- Adele E. Goldberg. 2006. *Constructions at work: the nature of generalization in language*. Oxford University Press.
- Stephan Greene and Philip Resnik. 2009. More than words: syntactic packaging and implicit sentiment. In *Proc. of NAACL-HLT*. Boulder, Colorado, pages 503–511.
- Martin Haspelmath. 2003. The geometry of grammatical meaning: semantic maps and cross-linguistic comparison. In Michael Tomasello, editor, *The New Psychology of Language: Cognitive and Function Approaches to Language Structure*, Lawrence Erlbaum Associates, Mahwah, NJ, volume 2, pages 211–242.
- Annette Herskovits. 1986. *Language and spatial cognition: an interdisciplinary study of the prepositions in English*. Cambridge University Press, Cambridge, UK.
- Jena D. Hwang. 2011. Making verb argument adjunct distinctions in English. Synthesis paper, University of Colorado, Boulder, Colorado.
- Ray Jackendoff. 1990. *Semantic structures*. MIT press.
- Ray Jackendoff. 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford.
- Paul Kay and Charles J. Fillmore. 1999. Grammatical constructions and linguistic generalizations: the *What’s X doing Y?* construction. *Language* 75(1):1–33.
- Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. 2008. A large-scale classification of English verbs. *Language Resources and Evaluation* 42(1):21–40.
- George Lakoff. 1987. *Women, fire, and dangerous things: what categories reveal about the mind*. University of Chicago Press, Chicago.
- George Lakoff and Mark Johnson. 1980. *Metaphors We Live By*. University of Chicago Press, Chicago.
- Ronald W. Langacker. 1998. Conceptualization, symbolization, and grammar. In Michael Tomasello, editor, *The New Psychology of Language: Cognitive and Functional Approaches to Language Structure*, Lawrence Erlbaum Associates, Mahwah, NJ, pages 1–39.
- Seth Lindstromberg. 2010. *English Prepositions Explained*. John Benjamins, Amsterdam, revised edition.
- Ken Litkowski and Orin Hargraves. 2005. The Preposition Project. In *Proc. of the Second ACL-SIGSEM Workshop on the Linguistic Dimensions of Prepositions and their Use in Computational Linguistics Formalisms and Applications*. Colchester, Essex, UK, pages 171–179.
- Rafael E. Núñez and Eve Sweetser. 2006. With the future behind them: convergent evidence from Aymara language and gesture in the crosslinguistic comparison of spatial construals of time. *Cognitive Science* 30(3):401–450.
- Meredith Osmond. 1997. The prepositions we use in the construal of emotions: why do we say *fed up with* but *sick and tired of*? In Susanne Niemeier and René Dirven, editors, *The Language of Emotions: Conceptualization, expression, and theoretical foundation*, John Benjamins, Amsterdam, pages 111–133.

- Martha Palmer, Claire Bonial, and Jena D. Hwang. 2017. VerbNet: Capturing English verb behavior, meaning and usage. In Susan E. F. Chipman, editor, *The Oxford Handbook of Cognitive Science*, Oxford University Press, pages 315–336.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: an annotated corpus of semantic roles. *Computational Linguistics* 31(1):71–106.
- Davide Picca, Alfio Massimiliano Gliozzo, and Massimiliano Ciaramita. 2008. Supersense Tagger for Italian. In Nicoletta Calzolari, Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odjik, Stelios Piperidis, and Daniel Tapias, editors, *Proc. of LREC*. Marrakech, Morocco, pages 2386–2390.
- Günter Radden. 1998. The conceptualisation of emotional causality by means of prepositional phrases. In Angeliki Athanasiadou, Elżbieta Tabakowska, René Dirven, Ronald W. Langacker, and John R. Taylor, editors, *Speaking of emotions: conceptualisation and expression*, Mouton de Gruyter, pages 273–294.
- Gisa Rauh. 1993. On the grammar of lexical and non-lexical prepositions in English. In Cornelia Zelinsky-Wibbelt, editor, *The Semantics of Prepositions: From Mental Processing to Natural Language Processing*, Mouton de Gruyter, New York, pages 99–150.
- Nathan Schneider, Jena D. Hwang, Vivek Srikumar, Meredith Green, Abhijit Suresh, Kathryn Conger, Tim O’Gorman, and Martha Palmer. 2016. A corpus of preposition supersenses. In *Proc. of LAW X – the 10th Linguistic Annotation Workshop*. Berlin, Germany, pages 99–109.
- Nathan Schneider, Behrang Mohit, Kemal Oflazer, and Noah A. Smith. 2012. Coarse lexical semantic annotation with supersenses: an Arabic case study. In *Proc. of ACL*. Jeju Island, Korea, pages 253–258.
- Nathan Schneider, Vivek Srikumar, Jena D. Hwang, and Martha Palmer. 2015. A hierarchy with, of, and for preposition supersenses. In *Proc. of The 9th Linguistic Annotation Workshop*. Denver, Colorado, USA, pages 112–123.
- Vivek Srikumar and Dan Roth. 2013a. [An inventory of preposition relations](http://arxiv.org/abs/1305.5785). Technical Report arXiv:1305.5785. <http://arxiv.org/abs/1305.5785>.
- Vivek Srikumar and Dan Roth. 2013b. Modeling semantic relations expressed by prepositions. *Transactions of the Association for Computational Linguistics* 1:231–242.
- Leonard Talmy. 1996. Fictive motion in language and “ception”. In Paul Bloom, Mary A. Peterson, Nadel Lynn, and Merrill F. Garrett, editors, *Language and Space*, MIT Press, Cambridge, MA, pages 211–276.
- Stephen Tratz and Dirk Hovy. 2009. Disambiguation of preposition sense using linguistically motivated features. In *Proc. of NAACL-HLT Student Research Workshop and Doctoral Consortium*. Boulder, Colorado, pages 96–100.
- Andrea Tyler and Vyvyan Evans. 2001. Reconsidering prepositional polysemy networks: the case of ‘over’. *Language* 77(4):724–765.
- Andrea Tyler and Vyvyan Evans. 2003. *The Semantics of English Prepositions: Spatial Scenes, Embodied Meaning and Cognition*. Cambridge University Press, Cambridge, UK.
- Henk Verkuyl and Joost Zwarts. 1992. Time and space in conceptual and logical semantics: the notion of Path. *Linguistics* 30(3):483–512.
- Bernhard Wälchli. 2010. Similarity semantics and building probabilistic semantic maps from parallel texts. *Linguistic Discovery* 8(1):331–371.
- Patrick Ye and Timothy Baldwin. 2007. MELB-YB: Preposition sense disambiguation using rich semantic features. In *Proc. of SemEval*. Prague, Czech Republic, pages 241–244.