Identification and Representation of Caused Motion Constructions

by

Jena D. Hwang

B.A., Calvin College, 2002

B.S., Calvin College, 2002

M.A., University of Colorado at Boulder, 2008

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Linguistics and Institute of Cognitive Science

2014

This thesis entitled:
Identification and Representation of Caused Motion Constructions
written by Jena D. Hwang
has been approved for the Department of Linguistics and Institute of Cognitive Science

Dr. Martha Palmer

Dr. Annie Zaenen

Date

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Hwang, Jena D. (Ph.D., Linguistics and Cognitive Science)

Identification and Representation of Caused Motion Constructions

Thesis directed by Dr. Martha Palmer

Advances in syntactic parsing and semantic role labeling have been a boon to Natural Language Processing. However, they perform poorly with sentences that do not conform to expected syntax-semantic patterning behavior. For example, in the sentence "The crowd laughed the clown off the stage", a verb of non-verbal communication laugh is coerced into the semantics of a caused motion construction (CMC) and gains a motion entailment that is atypical given its inherent lexical semantics. Accurate semantic role labeling for such sentences requires that NLP classifiers accurately identify these coerced usages in data. Given accurate semantic role labels, the sentence would also require a semantic interpretation with appropriate representations that include the semantics of the CMCs.

This thesis focuses on the definition, identification, and representation of the CMCs. We expand on the work from Construction Grammar to develop the semantic types and varieties of CMCs for corpus annotation. Utilizing the annotation as the training and test data, we train automatic CMC classifiers and demonstrate that CMCs can be reliably identified in the corpus data. Furthermore, we develop a new set of semantic predicates in VerbNet for the semantic representation of CMCs. These predicates will provide for the representation of CMC sentences, but also give VerbNet a more consistent explicit representation for paths of motion. Finally, we demonstrate that CMC representation can help give the proper semantic representation to sentences even when the verb in the sentence does not include the semantics of CMC.

The overall contribution of this work is the establishment of the processes involved in identifying and representing constructions in an empirical setting. This work assesses the necessary steps to define and annotate constructions in a corpus setting, train classifiers for constructions, and represent the semantics of constructions through VerbNet predicates. While we have focused on the

identification and representation of caused motion constructions, a similar corpus-driven study can be conducted for other constructions whose sentence representations would not be possible with the semantics of the verb alone.

Acknowledgements

First and foremost, I would like to thank my advisor, Martha Palmer, for her knowledge, insight, friendship and support without which this dissertation would never have been a reality. I am especially grateful to my dissertation committee, Annie Zaenen, Laura Michaelis, Jim Martin, and Bhuvana Narhashimhan, for their advice and guidance throughout this process. I would also like to thank Tim O'Gorman, Shumin Wu, Claire Bonial, Susan Brown, James Gung, Jared Schwartz, Jenette Preciado, and the many past and present folks at CompSem lab who have contributed to this work in various ways. To my fellow master's cohort and my doctoral colleagues and friends: I thank you for the friendships that made my years at CU more complete and enjoyable.

I give a thousand thanks to Aous Mansouri, Ashwini Vaidya, Jinho Choi, Susanne Stadlbauer, Nina Jagtiani, Alison Hilger and Jill Duffield, each of whom I have been able to count on for feedback, unconditional support and friendship (with the accompaniment of great food and much laughter, of course!). I would like to thank my family: my parents, my sisters and Abi, whose steady love and care has made this work possible. Finally, I give my most deepest gratitude to my husband Ryan, whose hours of detailed copy editing and undivided attention to my ramblings and rantings are only outdone by his patience, love and encouragement, and his unwavering belief in my ability to complete this work.

Contents

Tables

Table

Figures

Figure

Chapter 1

Introduction

1.1 Overview

When it comes to representation and identification of event expression, natural language processing (NLP) has traditionally focused on the semantics provided by the syntactic relationship between the verb and its arguments. Semantic role labeling, for example, centers around the recognition that there is a relationship between the semantics of the verb and the syntactic context in which the verb is realized; a recognition that has improved performance in NLP's semantic understanding systems (Guildea and Palmer, 2002). However, event representations that center solely around the semantics of the verb might not provide a complete picture. Sentences where the verb does not conform to the expected syntax-semantic patterning behavior remain problematic. Consider the following sentences:

(1) a. The boy threw the ball into the field.

b. The crowd hissed the clown off the stage.

c. The market tilted the economy into recession.

These sentences are semantically related – an entity causes a second entity to go along the path described by the prepositional phrase: in 1a, the boy causes the ball to go into the field, in 1b, the crowd causes the clown to go off the stage, and in 1c, the market causes the economy to go into recession.

While only the verb in the first sentence is generally identified as a verb of motion that can appear in a caused motion context, all three are examples of caused motion constructions (CMCs). The verb *hiss* in sentence 1b is normally considered an intransitive sound emission verb (e.g. *The snake hissed at the clown*), but in this sentence, the verb is coerced into the caused motion interpretation and the semantics of the verb gives the manner in which the movement happened (e.g. *the crowd caused the clown to move off the stage by means of hissing*). The verb *tilt* is a verb of spatial configuration normally taking, as its object, the inclined item (e.g. *He tilted the bottle*). In 1c, the verb is not only coerced into the caused motion reading, the coerced meaning is also abstract rather than physical as in *He tilted the liquid into his mouth and swallowed*. Whether the motion is physical or abstract, the semantics of the sentences is parallel: all three sentences have a causal argument responsible for the event, an argument in motion, and a path that specifies the initial, middle, or final location, state or condition of the argument in motion.

Despite the semantic similarities expressed in these sentences, a sentence is most likely identified as a a case of caused motion if the verb includes the semantics of motion in its prototypical sense. Consider the following two sentences, which are first translated into Spanish using Google Translate, then hand translated back into English[1] :

(2) a. They threw him out of the university.
Ellos le echaron fuera de la universidad.
Tr. They threw him out of the university.

b. They hissed him out of the university.
Le silbó fuera de la universidad.
Tr. They whistled to him outside the university.

The verb *throw* in sentence 2a is an example of a verb that most commonly appears in a caused motion usage, while the verb *hiss* in sentence 2b would be considered atypical in this particular semantic context. Presented with these sentences, a human reader would have no difficulty

[1] Results from Google translate (http://translate.google.com/) run in August 2014.

in judging the overlap in meaning. When the verb *throw* is used, the results are correct. However, the translation for 2b involving the verb *hiss*, which does not include semantics of motion in its prototypical sense, is incorrect and has absolutely no implication that a group of people did something to cause another person to leave the university. Moreover, the resulting Spanish does not at all show the semantic similarities seen in the English sentences.

The reason these coerced usages are challenging for natural language systems is because the current methods for natural language processing rely heavily on the lexical aspects of the language. New and unusual usages of lexical items are problematic, because they are often infrequent in the data and the frequency of their usage is far outnumbered by that of more conventional verbs in the same semantic context. This is especially significant in discriminative learning methods employed by semantic classification tasks, because their predictive capability is dependent on the examples they have seen or experienced during the training stage. If examples for usages like 1b are sparse (i.e., dwarfed in frequency) or simply do not appear in the training data, the classifier cannot gather the necessary statistics about the verb *hiss* used in a caused motion context to be able to correctly predict its meaning when later encountering the usage. It is a well-known problem: various studies have shown that verbs of ambiguous meaning or verbs that do not have training examples are identified as causes in classifier performance drop (Giuglea and Moschitti, 2006; Carreras and Marquez, 2005).

Studies have noted that the semantic problem posed by lexical sparsity can be remedied by recognizing the relationship between semantic and syntactic behavior in verbs. That is, if we can characterize verbs according to not only their meaning but also their syntactic realizations, we can group verbs that are both syntactically and semantically coherent (Jackendoff, 1990; Levin, 1993), thus allowing for generalizations beyond the individual lexeme. For example, if we can associate *hiss* and *throw* based on their shared syntactic and semantic behavior, such knowledge would help NLP systems with the issue of unseen predicates. Conversely if we know that the verbs *hiss* and *throw* are in the same class, one verb could be used to predict the syntactic and semantic behavior of the other verb. Indeed, verb classifications based on both meaning and syntactic realization have been shown to help in various tasks such as machine translation, word sense disambiguation, semantic role labeling, acquisition of subcategorization frames, and discourse parsing (Dorr, 1997; Korhonen, 2002; Subba and Eugenio, 2009; Merlo and Stevenson, 2002).

The recognition of the importance of classifying verbs according to their syntactic form and their semantic value has led to the creation of the resource VerbNet (Kipper et al., 2006, 2008). Following the work by Levin 1993, the verb classes in VerbNet are grouped by their semantic and syntactic similarities, and thereby VerbNet allows the user to abstract away from individual verb types to capture more general features based on the classes' syntactic behavior or semantic types. VerbNet has been used in various tasks in NLP including semantic role labeling (Swier and Stevenson, 2004; Giuglea and Moschitti, 2006; Merlo and van der Plas, 2009; Das and Smith, 2009; Bauer and Rambow, 2011), semantic parsing (Shi and Mihalcea, 2005), creation of conceptual graphs (Hensman and Dunnion, 2004), and natural language generation (Habash et al., 2003). Furthermore, the detailed semantic predicate information that VerbNet provides in its verb classes has been used to derive appropriate inferences from natural texts (Zaenen et al., 2008).

Even so, the challenge that VerbNet and all other hand crafted lexically based resources face is that they can never be complete enough to handle all possible usages and combinations that human language can throw at them. For example, VerbNet is capable of providing NLP systems with a way to generalize the semantic and syntactic behavior of a verb in an atypical usage as seen in 1b, if *hiss* is included in the same VerbNet class as the verb *throw*. Thus, the problem is circular: if we want to use VerbNet to correctly interpret the unfamiliar use of the verb *hiss* as used in the sentence, VerbNet needs to have already associated it with the behavior of verbs like *throw*. However, in order for VerbNet to make such an association, the resource creators would already have to be familiar with the use so that *hiss* could be included in the same class as *throw*. Thus, we are left with the question: how do we interpret unseen words and novel lexical usages in the context of supervised learning?

The premise that drives this dissertation, therefore, is that the challenge of generalizability could possibly be addressed by linking semantics more directly with syntax, as theorized by Construction Grammar (Fillmore, 1988; Goldberg, 1995; Kay and Fillmore, 1999; Michaelis, 2004; Goldberg, 2005). Construction Grammar is in agreement with VerbNet's assumption of a relationship between the meaning of word and its syntactic realization, but it goes a step further in proposing that particular syntactic structures can themselves carry meaning, in the same way a lexical item is thought to bear meaning. This theory suggests that the meaning of a sentence arises not only from the lexical items but also from the morpho-syntactic structures or *constructions* the lexical items sit in. In other words, meaning must be interpreted at both the lexical and constructional levels of sentences. This is especially important for sentences in which the lexical semantics of the verb is at odds with the semantics of the sentence.

Thus, if the semantic interpretation is strictly based on the expected semantics of the verb and its arguments, in some cases it will fail to to include the relevant information from the CMC. An accurate semantic role labelling for such sentences requires that NLP classifiers accurately identify these coerced usages in data. Furthermore, once the CMCs are identified and the semantic roles are properly assigned, the sentence would also require an accurate semantic interpretation with appropriate representations that include the semantics of the CMCs.

1.2 Improving VerbNet

This thesis is part of a greater effort at equipping VerbNet with systematic ways for dealing with coercive usages of verbs. Through coercion, a verb is allowed to violate its typical selectional restriction and to be used in atypical syntactic contexts (Pustejovsky and Jezek, 2008; Michaelis, 2005; Goldberg and Jackendoff, 2004; Pustejovsky, 1993)[2] . For example, in sentence 1b, a typically intransitive verb *hiss* (e.g. *The snake hissed at the boy*) is coerced into caused motion context. In the same way, an intransitive action verb like *blink* (e.g. *John blinked at me*) gains a caused motion reading when used with an object of blinking and a path prepositional phrase in *She blinked the snow off her eyelashes*.

[2] This argument structure level of coercion is also discussed in literature in terms of valence extension or valence augmentation.

VerbNet is currently useful at providing an analysis of the meaning of a sentence that is predictable given the semantics of the verb. However, when a verb is used in a syntactic context that is atypical of the verb, VerbNet does not have a good analysis for the 'extra' meaning the verb gains through coercion. Consequently, VerbNet's current treatment of the coerced instances is not consistent. VerbNet handles these coercive usages either by inserting the verbs into syntactically relevant classes at the expense of semantic unity or by including the CMC as a frame for the verb's class even if the CMC is not compatible with the inherent meanings of the member verbs. Consider the following three sentences:

(3) a. John slouched himself into the chair.

b. The crowd laughed the clown off the stage.

c. Cynthia blinked the snow off her eyelashes.

The verb *slouch* in example 3a is a posture verb appearing in the ASSUMING_POSITION-50 class along with other verbs like *hunch, lean* and *slump*. VerbNet currently interprets the CMC usage of this verb by including it in the RUN-51.3.2 class along with verbs of directed motion such as *canter*, *run*, *trot* and *walk*, although the semantics of the verb does not indicate directed motion in its typical usage. The verb *laugh* is a member of the NONVERBAL_EXPRESSION-40.2 class appearing with other verbs such as *moan*, *smirk*, and *weep*. In this case, rather than including *laugh* as a member of the RUN-51.3.2 class as it was done for *slouch*, VerbNet includes the syntactic frame directly into the NONVERBAL_EXPRESSION-40.2 class definitions[3] . Finally the verb *blink* is in the HICCUP-40.1.1 class, which includes other verbs of bodily movements and processes such as *burp*, *snore*, and *burp*. Unlike the other two instances, this particular usage of a CMC is not at this time addressed by VerbNet.

[3] The inclusion of the caused motion frame in the NONVERBAL_EXPRESSION-51.3.2 class comes directly from Levin (1993). This class has since been reanalyzed and the CMC frame was removed. While this change serves to strengthen the semantics unity of the member verbs, much like the *blink* example in sentence 3c, the sentence representation for 3b becomes unattainable.

A better way, we believe, is introducing constructional definitions – VerbNet constructions, if you will – that will interact with the current VerbNet classes to project the constructional meaning on to the sentence where the inherent semantics of the verb does not include it. With these constructional definitions, VerbNet should give a uniform semantic treatment for all sentences of caused motion regardless of the specific lexical meaning of the verb. Figure 1.1 is a visualization of a CMC instantiated by a verb typical of motion.

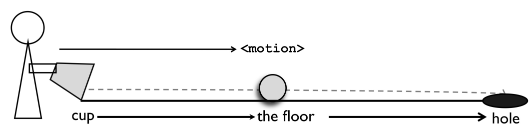

Figure 1.1: *"Cynthia rolled the ball out of the cup, across the floor, and into a hole."*

VerbNet would ideally represent the semantics expressed in the above sentence through the following semantic predicates:

```
cause(Cynthia, E)
motion(E, the ball)
rel_to_path(start(E), the ball, the cup)
rel_to_path(during(E), the ball, the floor)
rel_to_path(end(E), the ball, a hole)
```

Cynthia causes an event **E** in which *the ball* is put into motion. At the beginning of **E**, *the ball* is in the cup. During **E**, *the ball* moves across the floor to eventually to be located in *a hole* at the end of the event.

Because the verb *roll* and the members of the ROLL-51.3.1 class are verbs of motion, VerbNet would include this syntactic frame within the class. Verbs such as *blink* belonging to HICCUP-40.1.1 class, would not take caused motion syntax as one of their typical frames as the members of the class are neither verbs of motion or verbs of locational change. Nonetheless, VerbNet should give these verbs a similar semantic representation when they appear in a caused motion context. Figure 1.2 shows the visualization of the caused motion event in example 3c and its desirable sentence representation.

Since the class to which the verb *blink* belongs does not include caused motion frame, the

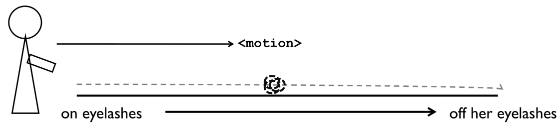

cause(*Cynthia*, **E**)
motion(**E**, *the snow*)
body_process(**E**, *Cynthia*)
rel_to_path(start(**E**), *the snow*, *her eyelashes*)
rel_to_path(end(**E**), *the snow*, **not**(*her eyelashes*))

Figure 1.2: *"Cynthia blinked the snow off her eyelashes."*

semantics of motion and path information is projected by the semantics of the CMC. More specifically, the `cause()` predicate, the `motion` predicate and the `rel_to_path` predicates is supplied by the semantics of the caused motion construction. One notable difference from a standard VerbNet sentence representation is that the *blink* sentence representation that includes the CMC carries an additional semantic component, namely, the act of blinking itself. Thus, what the VerbNet HICCUP-40.1.1 class would then provide is the specialized semantics that are particular to the lexical meaning that instantiates the CMC. In the above example, this lexical information is captured with the `body_process()` predicate which comes from the class HICCUP-40.1.1 to which the verb *blink* belongs[4] .

The semantic interaction between caused motion constructions and the ROLL-51.3.1 class or other classes of motion such as THROW-17.1 with members like *kick, lauch*, and *toss* will produce the same results that they would have produced by themselves, since the CMC frame is inherent in the member verbs (e.g. *She carelessly tossed her book on the table*). However, when a caused motion construction is paired with the HICCUP-40.1.1 class, which includes verbs such as *blink*, and *sneeze*, it produces the correct analysis for sentences like *She blinked the snow off her eyelashes* or *Pat sneezed the foam off the cappuccino*.

[4] In this thesis, we rely on the semantic predicates provided by the individual VerbNet classes to supply the necessary semantic representation that describe the semantics projected by the verb. This means that just like the example 1.2 receives lexically specific meaning through its verb's association with a particular class, we can distinguish the differences in the semantics expressed in *He shoved the ball into the hole* from *He rolled the ball into the hole* by the fact that *shove* belongs to the PUSH-12 class that specifies the exerted force involved in the event. This also means that the differences in the lexically specific semantics between verbs in the same class (e.g. *He rolled the ball into the room* vs. *He bounced the ball into the room*) is not addressed in this thesis.

1.3 Objectives

The objective of this thesis is to facilitate the semantic representation of CMCs for the conventional usages as well as the coerced ones. This objective breaks down into four general stages.

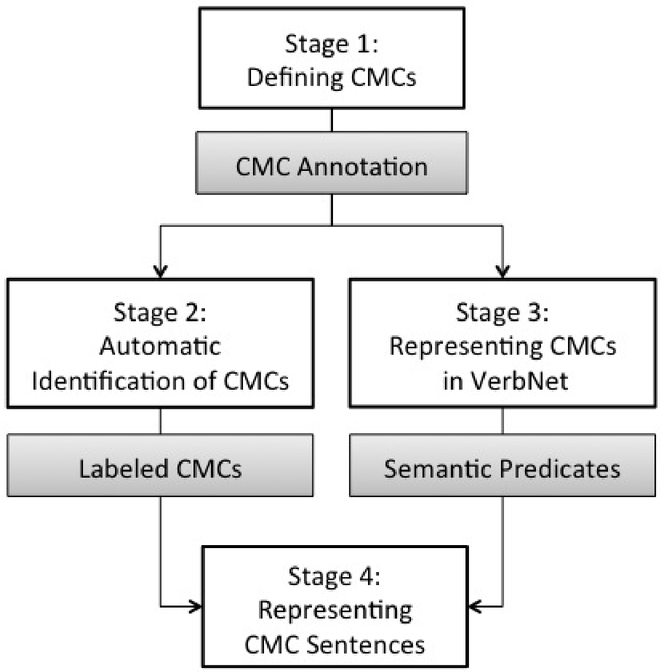

Figure 1.3: Stages.

The **first stage** focuses on our efforts at defining the semantic types and varieties of CMCs through an iterative annotation process and establishing annotation guidelines based on these criteria to aid in the production of a consistent and reliable annotation. During this stage, semantic types of CMCs are identified and detailed criteria for identifying CMCs are developed. The semantic types and the criteria aid in the production of consistent annotation with high inter-annotator agreement. Through the iterative annotation and revision process, CMC semantic type classification is established. The CMC annotation produced at this stage feeds into the next two stages of the study.

The **second stage** focuses on the training of a CMC classifier for automatic identification. In a pilot study, we determined that CMCs can be automatically identified with high accuracy

(Hwang et al., 2010). The pilot study was conducted in a highly controlled environment over a small portion of Wall Street Journal data. This dissertation further expands on the pilot study with a larger set of high-quality annotated data for the further training and testing of CMC classifiers.

The **third stage** focuses on establishing semantic predicate representations to represent CMCs. For semantic representation, we turn to the lexical resource VerbNet and the semantic predicates it provides for sentence representation. As described, VerbNet groups verbs according to their typical semantic and syntactic behaviors and is built to best handle instances where the verb is used in its typical syntactic context like the one seen in example 1a. VerbNet does not currently have effective ways of handling instances in which the verb is used in an atypical syntactic context, through which the verb gains new meaning. That is, because VerbNet classes are designed to focus on the protypical or conventional behavior of verbs, semantic coercions that often cross-cut through the established class boundaries are problematic. In order to improve our ability to use VerbNet, we seek to augment VerbNet with the information necessary to provide a unified treatment and consistent representation of both the coercive and the conventional usages of verbs.

The **fourth stage** focuses on the semantic representation of sentences, with a special focus on the instances belonging to VerbNet classes that are not normally associated with caused motion usages. This stage pulls together the results in the automatic classifier from stage two and the semantic predicates developed in stage three to give representation to caused motion instances.

1.4 Organization

This thesis is organized as follows. Chapter 2 provides background on Construction Grammar, VerbNet, semantics in NLP, and the resources used throughout this dissertation. Chapter 3 presents stage one of our objectives. It discusses the defining characteristics of caused motion constructions and presents our classification of caused motion constructions based on the semantic types found in the corpus data. It also presents our efforts on the corpus annotation of caused motion constructions. Chapter 4 presents stage two of our objectives. It describes our experiments on the automatic classification of caused motion constructions, using the annotated corpora developed in Chapter 3. Chapter 5 presents stage three of our objectives. It introduces a new semantic representation for caused motion constructions that will provide a more consistent representation for verbs that denote the semantics of path of motion. It also discusses the two stages of implementation process to update VerbNet to the new representation. Using the results from the classification efforts presented in Chapter 4 and the semantic predicates developed in Chapter 5, Chapter 6 presents stage four of our objectives by presenting our investigation on the sentence representation of coercive instances of caused motion constructions. Chapter 7 summarizes this thesis, presents its conclusions and contributions, and describes future research.

Chapter 2

Background

2.1 Construction Grammar

In the traditional view of semantics, meaning is considered to be the responsibility of individual lexical items and the morphemes that make up a word. Expanding on this notion, Construction Grammar recognizes a more inclusive category called *constructions* as carriers of meaning. Constructions are defined as linguistic forms that are assigned meaning in the same way that words are – via convention rather than composition (Fillmore, 1988; Goldberg, 1995; Kay and Fillmore, 1999; Michaelis, 2004; Goldberg, 2005; Michaelis, 2012). As expected, constructions include lexical items and morphemes, but they also include fixed, productive and semi-productive syntactic structures. These structures include regular syntactic devices (e.g. wh-questions, relative clause constructions, passive constructions) and more idiosyncratic yet productive expressions (e.g. *What's X doing Y?* (Kay and Fillmore, 1999), ditransitive (Goldberg, 1995), or *let alone construction* (Fillmore, 1988)). The meaning of a given phrase, a sentence, or an utterance, thus, is a combination of the meaning from the lexical items and from the morpho-syntactic structure in which they are found. Construction Grammar suggests that semantics can be interpreted at the level of phrase and sentence level structures such as caused motion constructions, resultative constructions, or ditransitive constructions. These constructions are associated with a distinct meaning independently from the semantics of the individual lexical items.

There are several subgroups of theories within the Construction Grammar framework. Perhaps the three best known groups are Sign-Based Construction Grammar (Sag et al., 2012) – a unification-based formalization of syntax, ‘Goldbergian’ Construction Grammar (Goldberg, 1995, 2005) known for its description of argument-level constructions and studies in psycholinguistics and cognitive linguistics, and Cognitive Grammar (Langacker, 1987) known for its conceptualization of semantics using image schemas. The subgroups in Construction Grammar may differ with respect to the details of their conceptualization and formalization of syntax and semantics, however, they all share the same underlying understanding of constructions as discussed above. A grammar for any Construction Grammarian is an “intricate network of overlapping and complementary patterns ‘blueprints’ for encoding and decoding linguistic expressions of all types” (Fried, http://www.constructiongrammar.org/).

2.2 Caused Motion Constructions

Semantic descriptions of caused motion constructions have a long history in linguistics, most often receiving attention in literature dealing with verbs of motion (Levin and Rappaport-Hovav, 1988, 1995; Talmy, 1991; Gruber, 1962), paths of motion (Slobin, 2005; Jackendoff, 1983), event causation (Levin and Rappaport-Hovav, 2005; Croft, 1991; Langacker, 1987; Talmy, 1976), resultative constructions (Boas, 2003; Goldberg and Jackendoff, 2004). In Goldberg (1995), CMC receives a focused attention, independent of other related semantic or syntactic phenomena. Goldberg proposes the existence of a single caused motion construction (often labeled as CAUSE-MOVE construction). The existence of the construction is motivated by the observation that in many of the attested usages of CMCs, the sentences retain meaning that cannot be attributed to the semantics of the main verb alone:

(4) a. They laughed the poor guy out of the room.

b. Frank sneezed the tissue off the table.

c. Mary urged Bill into the house.

d. Sally threw a ball into the net.

In examples 4a-4c, the semantics of the main verbs *laugh*, *sneeze*, and *urge* do not directly predict the caused motion meaning. For example, the verb *laugh* is a sound emission verb often licensing an agent, who does the laughing, and a theme, who is laughed at in intransitive contexts such as *They laughed at the poor guy*. Only in example 4d is the meaning of the verb *throw* consistent to the caused motion meaning of the sentence. For sentences 4a-4c, it is the merging of the semantics of the verb with the construction that gives the verb its additional arguments and coerces it into the meaning not projected by its lexical meaning. In the case of 4d, the constructional meaning completely overlaps with the sense of the verb. The merge does not cause the verb to receive additional semantics.

The lesson we would like to draw from this theoretical framework is the concept that the interpretation of meaning in a sentence does not have to be strictly restricted to the meaning drawn from the verb's lexical meaning. In fact, if we base our interpretation of a sentence strictly on the typical semantic and syntactic behavior of the verb (e.g. *laugh* as a sound emission verb or *sneeze* as a bodily process verb), we are likely to miss out on the proper interpretation for sentences that do not fully conform to their verbs' meaning. Thus, this thesis follows the general premise of Construction Grammar that, in addition to assigning labels at the level of lexical items and predicate arguments as a way of piecing together the meaning of a sentence, we recognize the non-compositional meaning provided by constructions.

By abstracting away from the individual verb types we can get at some of the semantic generalizations that would not be possible if we only focused on the individual words. Consider the following three sentences, having the same syntactic structure, each taken from different genres of writing available on the web.

(5) a. Blogger arrested - blog him out of jail! [Genre: Blog]

b. Someone mind controlled me off the cliff. [Genre: Role Playing Game]

c. He clocked the first pitch into center field. [Genre: Baseball]

Each of these sentences makes use of words, especially the verb, in ways particular to the genre. In example 5a, readers are asked to help the release of their fellow arrested blogger from prison by posting protests in their individual blogs. In example 5b, a gamer relates how he or she fell as a victim to the whims of another gamer with the in-game capacity to control the movement of other gamers. In example 5c, a player hits the ball into center field. Even if we are unfamiliar with the specific jargon used, as a human we can infer the general meaning intended by each of the three sentences: person X causes entity Y to move in the path specified by the prepositional phrase. In a similar way, if we can assign a meaning of caused motion at the sentence level and an automatic learner can be trained to accurately identify the construction, then even when presented with an unseen word, a useful semantic analysis is still possible.

2.3 VerbNet

VerbNet (Kipper et al., 2008) is a lexical resource that expands on Levin's (1993) verb classification. In accordance with Levin's work, VerbNet's classification of verbs is based on the hypothesis that verbs realized in similar syntactic environments will share in their semantics. That is, VerbNet class membership is determined by shared meaning and shared syntactic alternation. Thus, a VerbNet class is characterized by a set of semantic roles shared by all members in the class, syntactic frames in which the verbs occur, and the semantic representation of the event (designated by *E*). For example, the verbs *loan* and *rent* are grouped together in class 13.1 with roughly a "give" meaning, and the verbs *deposit* and *situate* are grouped into 9.1 with roughly a "put" meaning.

Additionally, VerbNet's class definition is hierarchical in nature; meaning that a member can either be associated at the most general level of class description or at one of the more semantically specific subclass levels. Take for example the *pocket-9.10* class shown in example 6. The *pocket-9.10* class includes verbs of putting, such as *bag* and *pocket* and is associated with the thematic roles: AGENT, THEME, DESTINATION and the frame NP V NP PP.DESTINATION, which is defined by both syntactic/word-order representation and semantic predicate representations.

(6)

class: pocket-9.10	
Roles	AGENT, THEME, DESTINATION
Members	bag, imprison, pocket etc.
Frame	NP V NP PP.DESTINATION
Ex(Example)	*I pocketed the coins in the side pocket.*
Syn(Syntax)	AGENT V THEME DESTINATION
Sem(Semantics)	cause(Agent,E) motion(during(E),Theme) not(Prep(start(E),Theme,Destination)) Prep(end(E),Theme,Destination)

The subclass *pocket-9.10-1* seen in 7 and its member verbs (i.e., *dock*, *land*) inherits all of the roles and frames associated with its parent class *pocket-9.10*. Additionally, the subclass defines a more specific frame (i.e., NP V PP.DESTINATION) that does not apply to the parent class.

(7)

subclass: pocket-9.10-1	
Members	dock, land, lodge, etc.
Frame	NP V PP.DESTINATION
Ex	*I landed in Russia.*
Syn	THEME V DESTINATION
Sem	motion(during(E),Theme) not(Prep(start(E),Theme,?Destination)) Prep(end(E),Theme,?Destination)

2.3.1 Why VerbNet?

We consider VerbNet as the most appropriate resource for expressing constructions, because unlike FrameNet and WordNet that deal primarily with semantics, VerbNet cares about both the syntactic and semantic level of description. The theoretical basis behind VerbNet is the recognition that there is a systematic relationship between semantics and syntax: a verb's syntactic behavior is indicative of the verb's meaning since "particular syntactic properties are associated with verbs of a certain semantic type" (Levin, 1993, p.5; c.f., Pinker, 1989; Jackendoff, 1990; Gleitman, 1990; Levin and Rappaport-Hovav, 2005). Take an example from Levin (1993, p.5):

(8) The sailors gallied the whales.

Even without knowing the meaning of the verb *gally*, an archaic whaling term, a speaker of English may make an assumption about its meaning: the verb could possibly mean something synonymous with verb *see* (e.g., *The sailors saw the whales*) or verb *frighten* (e.g., *The sailors frightened the whales*). If the speaker knew the verb could be used in the middle construction (e.g. *Whales gally easily*), they would likely associate *gally* more strongly with the meaning of *frighten* which allows for the middle construction (e.g. *Whales frighten easily*), than the verb *see* which does not (e.g. **Whales see easily*).

Thus, the essential difference between what we see in VerbNet (and Levin's verb classification) and Construction Grammar's constructions is the role of the verb types. A VerbNet class is a grouping of verb types that retain similar syntactic behavior and the semantic coherence they have is something of a byproduct of grouping verbs displaying similar syntactic alternations. A construction, however, is defined by a single syntactic form which is associated with a meaning, regardless of the verb types that appear in the construction.

However, VerbNet is conducive to Construction Grammar because in both syntax and semantics play a role. Constructions are the pairing of form and meaning. VerbNet pairs a group of syntactic realizations to a meaning by way of grouping verb types. Thus, VerbNet is a natural fit for an additional layer of constructional definitions in which each class is associated with one or more constructions.

2.4 Supervised Learning

The objective of any statistical classification system based on machine learning in natural language processing is to predict the correct and relevant label for a piece of text it is shown. For example, in the case of a word sense disambiguation (WSD) task, we want the system to be able to label a verb in a sentence with the proper semantic sense of the verb. Amongst a number of favored methods of learning used by natural language processing, supervised machine learning is currently the most dominant approach. Supervised learning relies heavily on manually labeled data for training and testing purposes.

In supervised learning, a classifier is given training data that is already tagged with the proper labels. This training data comes in the form of features that characterize the raw text. Features can be encoded based on the words themselves, the characteristic and linguistic aspects of the text, relevant information that can be associated with the words in the text (e.g., gathered from resources like VerbNet) or any other information the researcher sees as being representative of the annotated text. Features are extracted from the raw text and handed to the classifier for training. The classifier then gathers statistics about the features over the whole of the training data and builds a model of what it sees. The system uses this model to predict the labels for a new and unlabeled sentence. The accuracy of the system is then evaluated by comparing the classifier's predicted label against the manually annotated "gold" standard.

Because of its high reliance on manual annotation, the performance of a supervised system crucially depends on both the amount and the quality of the annotation. The supervised classifier does best if there are plenty of exemplars of the phenomenon it is asked to predict for, and it does poorly when the exemplars are few or entirely missing from the training data. For example, a WSD classifier is likely to incorrectly label *hiss* in *They hissed him out of the university* if it has only seen examples for the predominant sense of the word as exemplified in *The snake hissed at the man*. This assumes that there was a correct label (i.e., "expel with contempt") available for the classifier to choose. If the data is entirely missing the correct label, then the classifier has zero chance of getting the sense right – it will default to the label with the best chance of being 'right' given the statistics it gathered. As shown in the CoNLL-2005 Semantic Role Labeling (SRL) shared task (Carreras and Marquez, 2005), system performance numbers drop significantly when a classifier, trained on the Wall Street Journal corpus, is tested on the Brown corpus. This is largely due to "highly ambiguous and unseen predicates (i.e. predicates that do not have training examples)" (Giuglea and Moschitti, 2006).

2.5 Semantics with Natural Language Processing

In natural language processing, the issue of semantic analysis in the presence of lexical and syntactic variability is often perceived as the purview of either word sense disambiguation (WSD) or semantic role labeling (SRL) or both. In the case of WSD, the issue is often tackled through the use of large corpora tagged with sense information for trainining a classifier to recognize the different shades of meaning of a semantically ambiguous word (Ng and Lee, 1996; Agirre and Edmonds, 2006). In the case of SRL, the goal is to identify each of the arguments of the predicate and label them according to their semantic relationship to the predicate (Gildea and Jurafsky, 2002).

There are several corpora available for training WSD classifiers such as WordNets SemCor (Miller, 1995; Fellbaum et al., 1998) and the GALE OntoNotes data (Hovy et al., 2006). However, most, if not all, of these corpora include only a small fraction of all English predicates. Since WSD systems train separate classifiers for each verb, if a particular verb does not occur in the sparse training data, a system cannot create an accurate semantic interpretation. Even if the predicate is present, the appropriate sense might not be. In such a case, the WSD will again be unable to contribute to a correct overall semantic interpretation. This is the case in example (1), where even the extremely fine-grained sense distinctions provided by WordNet do not include a sense of hiss that is consistent with the caused motion interpretation rendered in the example.

Available for SRL tasks are efforts such as PropBank (Palmer et al., 2005; Kingsbury and Palmer, 2003) and FrameNet (Fillmore and Petruck, 2003) that have developed semantic role labels (based on differing approaches) and have labeled large corpora for training and testing of SRL systems. PropBank identifies and labels the semantic arguments of the verb on a verb-by-verb basis, creating a separate frameset that includes verb specific semantic roles to account for each subcategorization frame of the verb. Much like PropBank, FrameNet identifies and labels semantic roles, known as Frame Elements, around a relational target, usually a verb. However, unlike PropBank, Frame Elements are less verb specific, but rather are defined in terms of semantic structures called frames evoked by the verb. That is, one or more verbs can be associated with a single semantic frame. Currently FrameNet has over 2000 distinct Frame Elements.

Although differing in the nature of their tasks, WSD and SRL systems both treat lexical items as the source of meaning in a clause. In WSD, for every sense we need a new entry in our dictionary to be able to interpret the sentence. With SRL, we need the semantic role labels that describe the predicate argument relationships in order to extract the meaning.

2.6 Resources Used

2.6.0.1 English PropBank

The English PropBank, is a semantic resource that was developed for the task of semantic role labeling (SRL), which seeks to automatically identify and extract the different semantic relationships between words in a given text (Palmer et al., 2005; Kingsbury and Palmer, 2003). Resources that PropBank provides include two things: annotated data and a repository of lexical information. PropBank's annotated data includes various types of text including newspaper, magazine, and weblog texts. The data is annotated with the lexical repository that defines verbs' argument structures.

It is in the definitions of argument structures that the argument-adjunct treatment is found. The definitions of verbs are contained in what is called a Frame File that can include one or more Frameset for a given verb, and each Frameset corresponds to a different subcategorization frame of the verb. Here is an example:

```
Frameset id: give.01, transfer
Arg0: giver
Arg1: thing given
Arg2: entity given to
```

A verb entry includes a unique Frameset id, a descriptive text of the word, and the verb's argument structure. For the verb *give*, the argument structure includes the *numbered arguments* (i.e. ARG0, ARG1 and ARG2) as described in the Frameset. The Frameset id makes a distinction between this use of the verb as in "John *gave* cookies to Mary" and the use in "I am *given* to wonder what is so great about it!". In addition to the numbered arguments, PropBank defines *modifier labels* starting with ARGM. Accompanying the label ARGM will be one of the many function (semantic) labels that describes the modifier (e.g. TMP for temporal, LOC for locative). For example, in a sentence such as "John left cookies for Mary on Thursday", *on Thursday* would be labeled as ARGM-TMP, with the temporal function tag, TMP, which defines the semantics of the modifier.

2.6.0.2 OntoNotes Verb Sense Groups

OntoNotes Verb Sense Groups (Duffield et al., 2007) provide an inventory of English verb semantic senses. The senses in the OntoNotes groupings are finer-grained than PropBank's Framesets, but coarser-grained than WordNet verb senses (Miller, 1995). In fact, the OntoNotes sense grouping project started out as the grouping of WordNet's finer grained senses. This effort led to improved inter-annotator agreement (IAA) and system performance.

As an example of verb sense granularity, take the transfer of possession meaning of the verb *give*. For PropBank's single *give.01*, OntoNotes currently provides 3 senses:

```
Sense 1: transfer of possession
         (e.g. He gave her a bouquet of roses)
Sense 2: relinquishing of personal resource, attribute or knowledge
         (e.g. He gave his life for his children)
Sense 3: intentional grant or decree
         (e.g. Will she give him a divorce?)
```

These three senses are far more coarse-grained than the WordNet senses. They represent 27 of the WordNet senses.

2.6.0.3 Syntactic Parses

In this thesis, we use two styles of syntactic parses: phrase structure trees and dependency trees. The phrase structure trees we use come from the English Penn Treebank (Taylor, 1996; Taylor et al., 2003). It is a syntactic resource that provides manually parsed corpora with phrase structure parses. Here is an example of a Penn Treebank-style phrase structure tree.

```
(TOP (S (NP-SBJ (NP (NNP John)))
        (VP (VBD gave)
            (NP (NNP Mary))
            (NP (NP (DT a)
                (NN cookie))))
        (. .)))
```

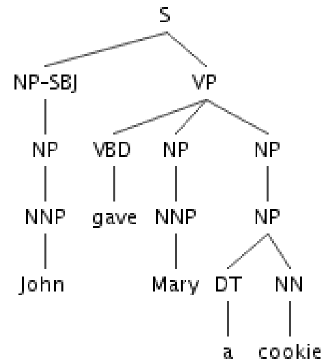

Like the syntactic trees from transformational grammar, Penn Treebank trees use empty categeories, traces, and indexed movements of the constituents. It provides its own set of part of speech tags and phrase structure labels, and includes a small set of semantic function tags when appropriate.

We also use CLEAR dependency parses (Choi, 2012). Unlike the phrase structure trees where constituents are recursively grouped under phrase structure labels, dependency parses consists of trees which represent all of the dependency relations to the main verb of the sentence (whether the dependency is direct or indirect). The following is an example of a dependency parse for the sentence “John gave Mary a cookie”.

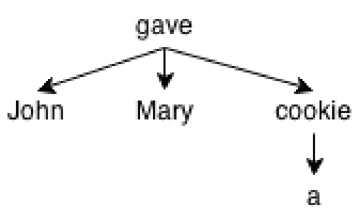

The verb is the top most head node, and *John*, *Mary* and *cookie* are dependents of the verb. The determiner *a* is a dependent of *cookie*.

Chapter 3

Defining and Annotating Caused Motion Constructions

3.1 Introduction

This chapter focuses on our efforts to define the semantic types and varieties of CMCs through an iterative annotation process and to establish annotation guidelines based on these criteria to aid in the production of consistent and reliable annotation. The annotated data will serve two major purposes. First, it will serve as the training and testing data for the automatic classifier experiments presented in Chapter 4. Second, the typological classification of CMCs will serve as a guideline for establishing consistent semantic predicates for CMCs in VerbNet as presented in Chapter 5.

We take an iterative approach to the development of semantic categories and guidelines for CMCs. Rather than making use of a single set of pre-determined annotation guidelines for all our annotation efforts, we allow the guidelines to be updated and revised based on the feedback we received from the annotation process. This process of annotation, evaluation, and guideline update is iterated over a number of cycles until we arrive at a set of criteria for identifying CMCs that yield a reasonably high inter-annotation agreement rate.

In Section 3.2, we present Goldberg's (1995) definition on CMCs. In Section 3.3, apply the linguistic theories to the creation of annotation guidelines. We start with the definitions of a CMC as proposed by Goldberg (1995). Through the course of this section we will describe the motivations for the constraints and guidelines we set into place for the identification of CMCs exploring how these decisions compare and contrast to the the constraints as presented by Goldberg

(1995). As we will see, the final set of guidelines both expand and restrict on Goldberg's analysis, specifying the necessary conditions for CMC identification that are best suited to streamlining the annotation process. Finally, we present our present our annotation efforts in CMC's identification in Sections 3.4, and a pilot study, which involves the assignment of a concreteness rating to the CMC instances, is presented in Section 3.5.

3.2 Goldberg (1995) Definition

Caused motion constructions can be defined as having the following coarse-grained syntactic structure: (NP (V NP PP)). The construction is characterized by three argument roles: a causal argument, an undergoer or theme argument, and one or more path arguments. The construction roughly means "the causal argument (NP) directly causes the undergoer (NP) to move along a path specified by the prepositional phrase". Figure 3.1 represents Goldberg's CMC.

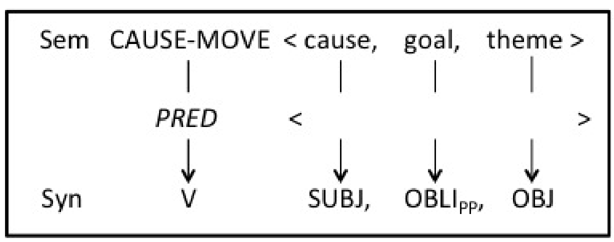

Figure 3.1: Diagram for uninstantiated caused motion constructions (Goldberg, 1995, 160).

The diagram includes three levels of description. The top level, labeled as *Sem*, specifies the semantics and argument structure specific to the CAUSE-MOVE construction. The middle level specifies the information coming from the predicate, *PRED*, which currently is unspecified. When the construction is instantiated with a verb, the verb's argument structure is merged with that of the construction. The lowest level, labeled as *Syn*, describes the syntactic realization of the sentence. The arrows designate the merging of the construction's argument roles with the verb's argument roles. The following two examples have been diagrammed in Figures 3.2 and 3.3, respectively.

(9) a. Sally **threw** a ball into the net.

b. They **laughed** the poor guy out of the room.

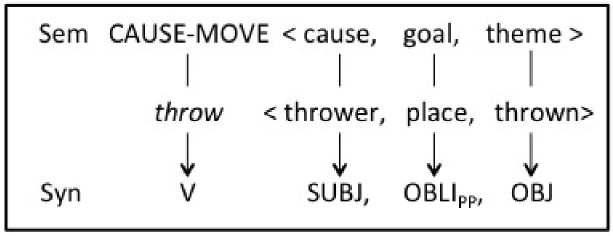

Figure 3.2: Diagram for caused motion construction instantiated with the verb *throw*.

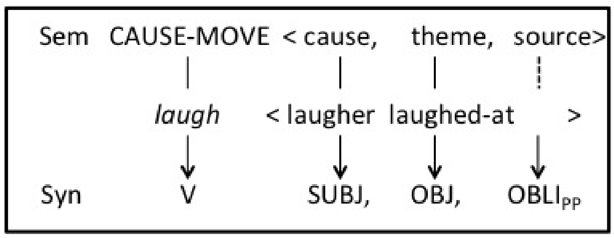

Figure 3.3: Diagram for caused motion construction instantiated with the verb *laugh*. The verb is coerced into a caused motion reading

For the construction instantiated by the verb *throw* as diagrammed in Figure 3.2, solid lines link the arguments from the construction to the arguments of the verb: there is a one to one mapping between the *cause* and the *thrower* argument, between the *path* argument and the *place* to which the item was thrown and the *theme* and the *thrown* arguments. For a construction that is merging with a verb whose argument structure does not have a one to one correspondence with the arguments in the construction, the verb receives the missing arguments from the construction. For example, in9b, diagrammed in Figure 3.3, the *path* argument does not have a corresponding argument in the semantic argument structure of the verb. Thus, the sentence receives the path argument from the construction, which is represented by a dotted line going from the construction's *path* argument to the empty slot in the verb's argument structure.

Nevertheless, caused motion is limited as to the verbs or lexical items with which it can merge. Thus, constructional approaches make use of semantic constraints to explain why certain verbs cannot merge with the construction. Constraints on the construction determine if the construction and the lexical items are semantically compatible, thus providing an explanation as to why certain sentences are not observed in the English language, or why certain verbs sound odd in the context of the construction.

(10) a. *Steve opened Melissa into the room. (Boas, 2003, pg. 99)

For example, the verb *open* in the caused motion context renders the sentence nonsensical. While we readily accept verbs like *laugh* to be coerced in the same context, the construction does not equally accept all verbs. Goldberg proposes a series of constraints on caused motion constructions. We examine these constraints to identify the ones with the most descriptive power with the intent of adapting them for the annotation process. On the whole we find that Goldberg's analyses are helpful in characterizing CMCs. We also find the analysis is limited for one basic reason: the constraints are more helpful in making sense of those instances that have already been judged as CMC or non-CMC rather than in making the discriminative decision in the first place.

While it is not the intent of this dissertation to provide a more comprehensive set of constraints for caused motion constructions, we do examine the existing constraints from the viewpoint of developing annotation guidelines that will provide guidance for annotators when distinguishing CMCs from non-CMCs. One theme that surfaces throughout our study is that the basic needs of the annotation process of the corpus data inevitably causes our analysis of CMCs to diverge from that of Goldberg's. Our decisions on what semantic aspects we use to constrain CMCs are based on the need to make categorical decisions at a sentence level. While discerning the semantics that the construction provides to the sentence versus the semantics that come from the lexical items is not the primary focus of this thesis, our work on characterizing CMCs and providing constraints for their identification in corpus data is a contribution to understanding the function of CMCs in the English language.

3.3 From Theory to Practice

One of the challenges of adapting theoretical work to a practical annotation setting lies in identifying the existing discrepancies in the approaches and goals, and appropriately reconciling them with the practical needs of the annotation process. Goldberg's (1995) analysis provides a specific account of caused motion constructions, primarily of the concrete and physical type, for the purposes of defining their semantic constraints and providing motivation for existence. The analysis is conducted over prototypical examples that clearly represent the constraints at play. Unfortunately for machine learners, exemplifying a linguistic phenomenon through the sole use of prototypical examples is not enough. A learning system needs to know how to handle all caused motion constructions, including the abstract/metaphorical usages and marginal cases found in the actual data, so that it can draw the boundaries between the caused motion usages and non-caused motion usages. In other words, to make automatic identification possible, we need to establish distinctive characteristics for identifying CMCs so that these distinctive features are reflected in the annotation used for training automatic systems. As we move from theory to practice we seek to strike a balance between being faithful to linguistic theories and being cognizant of the practical needs of their application. Therefore, we refine, expand, and detail certain aspects of the linguistic analysis of CMCs according to the guiding principles that motivate this study:

- **Meaningful Labels:** Because the annotation process is to be used for the automatic identification of CMCs, annotated labels should describe meaningful and coherent semantic phenomena. That is, if the descriptive features of a construction are too broad, then automatic identification is less likely. If the label is semantically too general, no meaningful inferences can be made from its identification.

- **Expedited Annotation:** Annotation guidelines should expedite a quick decision process and minimize ambiguity and long deliberations. Annotation should not require a high degree of linguistic knowledge.

- **Explicit Guidelines:** In order to expedite annotation, definitions and guidelines for CMCs must be explicit. This means, the guidelines should provide labels that are specific enough for quick decision-making and semantic criteria that are general enough to allow annotators to make decisions over all corpus instances presented to them.

We examine the characteristics of CMC at two layers of description. At the first layer of description, we analyze the theories concerning semantic constraints over CMCs. Through the course of this discussion, we develop three semantic criteria that generalize over all instances of CMCs, including both the physical and abstract types. At the second layer of description, we examine the metaphorical extensions of CMCs to recognize salient categories that naturally surface in the data. We separate these categories out and assign them distinct labels for the ease of the identification and annotation process.

3.3.1 Constraint-Based Distinctions

3.3.1.1 Direct Causation

A central constraint that drives most of Goldberg's more fine-grained semantic constraints is the constraint of direct causation. Direct causation can be generally defined as the pairing of a causal act and the caused effect without the interference of an external third cause (e.g. in "She killed my cactus", the death is directly caused by the specified subject). This is distinguished from the more indirect or periphrastic causation where the effect is not directly paired with any particular cause (e.g. in "She caused my cactus to die" , no particular act is directly responsible for the death of said cactus). The constraint of direct causation for CMC, requires that the cause in the event must be directly responsible for the effect, namely, the undergoer's motion.

A generalization that is evident in Goldberg's analysis is that a CMC's compatibility with the constraint of direct causation is strongly based on (1) whether or not the syntactic subject[1] is a causal entity, (2) whether or not the syntactic object[2] can be subjected to a causal force, and (3)

[1] Syntactic subject in the canonical CMC syntactic form: (NP-SUBJ (V NP PP))
[2] Syntactic object in the canonical CMC syntactic form: (NP (V NP-OBJ PP))

what the semantics of the verb specifies about the subject and the object.

Goldberg's cause argument constraint specifically seeks to allow only the subjects wielding causal force as the cause arguments of a CMC. The constraint states that only an *agent* or a *natural force* can be the causal argument. This effectively rules out instruments as potential cause arguments. In the following sentence, the cane is the instrument of aid for the man as he is moving himself into the car. A cane cannot wield the force necessary to be the cause argument that will cause the motion of the man.

(11) ?His cane helped him into the car. (1995, pg. 165)

Most of the rest of Goldberg's constraints on direct causation are directed towards establishing whether or not the syntactic object can be subjected to a causal force. The undergoer would have no decision on the matter of where it goes. Only the cause argument can decide the path on which the undergoer will travel. This, Goldberg notes, is the reason that the specification of path of motion is odd in the following sentence:

(12) *Pat shot Sam across the room.

This sentence is unacceptable on the interpretation that Pat shot Sam with a gun and the bullet forced him across the room (1995, pg. 170)[3] . The shooting event has an effect that does not coincide with the effect described by the rest of the sentence. Specifically, the effect caused by Pat in the shooting event is that Sam is hit and penetrated by a moving bullet, and not the effect that describes Sam's trajectory across the room. Example 12 is unacceptable because for CMCs, the constraint of direct causation must hold, the cause argument must set the undergoer moving over a path determined by the cause argument. If there can be a conventional sequence of events that can connect the cause event to the effect of motion, then, Goldberg notes that the scenario must be "cognitively packaged as a single event".

[3] It was brought to my attention that *Pat shot the can across the room* is acceptable. If the *can* in the sentence is the projectile, then this sentence does not violate Goldberg's constraint. If this sentence could be interpreted such that Pat shot the can with a gun and the bullet from the gun shooting event, in turn, caused the can to go across the room, then we have a counter example to Goldberg's constraint. Further investigation would be needed to see if both readings are available for native speakers.

(13) They company flew her to Chicago for an interview.

While the company isn't the one who is doing the actual flying, this sentence is deemed acceptable because, in a conventionalized situation such as this, the flying event is inclusive over all of the processes that occur when arranging for someone to be flown somewhere (e.g. purchasing tickets).

The proposed constraints are useful in describing generalities of CMCs. Unfortunately, there are naturally occurring examples of sentences that should be rendered unacceptable, consequently, should not occur in data given these constraints. For example, as Boas (2003) points out, an instrument can be construed as the cause argument if the pragmatic knowledge allows it. In each of the examples below, Boas notes that given the proper contextual background information, "the subject's force-emitting properties [...] make it possible to construe it as a causer" (2003, pg. 103).

(14) a. The GPS system guided Christian through the city.

b. The espresso maker blew steam into the milk.

A GPS system's ability to provide navigational guidance causes Christian to move through the city. In a similar way, the force exerted by the espresso maker causes steam to move into the milk. These examples stand in contrast to Goldberg's sentence in example 11, which is unacceptable because there is no perceivable context that can attribute a 'force-emitting property' or proto-agent-like properties (Dowty, 1989) to a walking stick.

The problem is not specifically in the general constraint of direct causation. The examples in 14 still suggest that the subjects of the two sentences are forceful entities responsible for the undergoer arguments' motion. We are attributing agent like abilities to otherwise non-agentive entities. Rather, the problem seems to be in Goldberg's articulation of the details of the constraint. Let us take another example. This one concerns the nature of the undergoer argument.

Goldberg (1995) observes that verbs like *encourage*, *persuade* and *convince* do not appear in the construction due to the direct causation constraint the construction imposes on the sentence.

(15) a. Sam **coaxed** Bob into the room.

b. *Sam **encouraged** Bob into the room.

Goldberg points out that *coax* and other verbs of psychological states such as *frighten* and *lure* can appear in the construction because they do not “entail the existence of a cognitive decision” (1995, pg. 166). This is contrasted with verbs like *persuade* that denote a cognitive decision on the part of the undergoer argument. That is, if the sentence 15b were allowable, then Bob would cause himself to go into the room by Sam’s encouragement. Consequently, if Bob is the causal entity of Bob’s own movement, then the ”mediating cognitive decision” violates the semantic constraint on direct causation, thereby rendering the sentence non-felicitous for CMC.

Despite this, there are attested cases, though infrequent, where verbs like *persuade* or *encourage* appear in a caused motion-like context, as exemplified in the following sentences. This requires the annotators to decide whether or not such instances are cases of caused motion.

(16) a. Wade’s sheer enthusiasm [...] **persuaded** me into a renewed fascination. (BNC)[4]

b. For the second day in a row, I have **convinced** myself out of biking to work (enTenTen)[5]

The constraint that rules out verbs that denote the existence of a cognitive decision from being instantiated in a caused motion construction serves to explain why the construction favors certain verbs (e.g. 15a) over others (e.g. 15b). In fact, it gives us insight into why cognitive decision verbs, even when they occur in data, do so rarely as CMCs.[6] Nevertheless, the existence of such sentences in natural data suggests that Goldberg’s cause argument constraint limited as to which sentences it allows as CMC.

[4] British National Corpus

[5] English TenTen web corpus

[6] For the English TenTen corpus, there were approximately 60 instances of caused motion like the example 16b out of 464K+ instances of the verb *convince*.

Here is another example that involves the cause argument. Sentences that involve a body part as the undergoer argument of the sentence can also lead to false positives. Compare the cause of the following two sentences:

(17) a. In frustration, Mary hit her head against the wall.

b. When the train came to an abrupt stop, Mary hit her head against the wall.

In both sentences, the undergoer of the event is Mary's own body part. In the first sentence, Mary intentionally wills her head to move along a path directed against the wall. In the second sentence, however, the cause that drives Mary's head to hit the wall is a force beyond Mary's control. In other words, Mary is the direct cause of the motion of her head in sentence 17a, making the sentence an instance of CMC. The second sentence, however, cannot be a CMC as the causation is indirect and accidental.

Defining a series of detailed constraints that spell out what types of instances are barred by direct causation may be the problem. That is, saying that certain sentences are not supposed to occur is not very helpful when the annotators come across these 'unacceptable' sentences, requiring us to be more general with our description for the purposes of annotation.

We opt for a simple approach to our formulation of direct causation for our guidelines by expressing the constraint of direct causation as a semantic criterion that will generalize over various potential sentence types:

> **Criterion 1**: The cause argument must exert a force over the undergoer, resulting in the undergoer's motion. The causal entity must be either agentive or be capable of wielding the force necessary for a direct effect on the undergoer.

This criterion correctly informs the annotators to select for instances in which the causal entity has a direct effect on the undergoer entity, regardless of the identity of the verb (e.g. 16), and to discount instances that lack direct causation (e.g. 11 and 17b). It also correctly rules out instances in which the undergoer moves as a result of an accident – a force external to the subject of the sentence.

3.3.1.2 Entailment of Motion

If the previous constraints on direct causation were too narrow, the one dealing with entailment of motion is too general for our purposes. Goldberg recognizes that depending on the verb that instantiates the CMC, the construction can suggest a different set of inferences when it comes to motion. Consider the following examples:

(18) a. Harry locked Joe into the room.

b. Harry allowed Joe into the room.

c. Harry invited Joe into the room.

d. Harry helped Joe into the room.

In sentence 18a, Harry prevents Joe's motion by introducing a barrier in Joe's potential path of motion, "causing the patient to stay in a location despite its inherent tendency to move" (1995, pg. 162). Goldberg argues that the path argument in the sentence codes for a 'complement' of the potential motion as Harry's potential path would have been to move *out of* the room. The sentence 18b has approximately the opposite effect: Harry makes Joe's motion possible by removing any prohibitions that may have kept Joe out of the room. Sentence 18c's semantics are similar to that of 18b, but entailment of motion is dependent on the condition of Joe accepting Harry's invitation. The motion is implied, but not entailed. In both 18b and 18c, the motion can be negated as shown in 19a and 19b, respectively.

(19) a. Harry allowed Joe into the room, but Joe never went into it.

b. Harry invited Joe into the room, but Joe never went into it.

c. *Harry helped Joe into the room but Joe never went into it.

Finally, sentence 18d is likely the only example here where Joe's motion is unconditionally entailed: it is not possible to negate Joe's movement into the room (e.g. 19c). Goldberg identifies these distinctions as four senses of CMC. We will refer to them as prevented motion, enabled motion, implied motion and accompanied motion, respectively.

The instances of prevented motion by definition do not entail motion. In fact, the opposite is true: the entailment states that Joe's motion was restricted. Enabled and implied motion are trickier because of their potential for motion. For either of the sentences, if the motion does take place, the motion of the undergoer can be directly attributed to the causal force of the subject – the constraints for direct causation are satisfied. In other words, provided that Joe's acceptance of Harry's permission or invitation is presupposed, we can say that Harry's motion is implied.

Goldberg's constructional definition of CMCs does not have to require a strict entailment and, in her framework, the implication of motion along a path is sufficient to identify the sentences as CMC. This analysis essentially gives the construction the responsibility of specifying the preconditions necessary for the undergoer's motion and the burden of whether or not the motion is truly entailed is left to the individual verbs with which the constructions merge. If the sentence does not violate the constraints of direct causation in the event that the preconditions are met, it is considered acceptable as a case of CMC.

We take a narrower view of what defines CMCs. Specifically, we focus on the subset of Goldberg's definitions for which the motion of the undergoer is strictly entailed. In the annotation process, we look at individual sentences to decide whether or not the instance is a case of caused motion. If we classify CMC according to Goldberg's general rule of accepting any instance that satisfies the necessary preconditions for both restriction and potential motion and leave much of the entailment to the lexical items themselves, we allow too much semantic variety under the caused motion label. By restricting CMCs to instances in which motion is strictly entailed, we seek to enable the inference that the object is in motion (physical or abstract) during the course of the event with a corresponding change of location or state.

What we seek are general rules by which we can make specific decisions over corpus sentences. Furthermore, such rules should identify a meaningful and coherent semantic phenomena. Thus, the key specification in the issue of entailment of motion is that we not only need the sentences to imply the undergoer's motion, but they also strictly entail it for the event.

> **Criterion 2**: The undergoer must be in motion during the event. Only when the motion of the undergoer is strictly entailed for the event is a sentence judged as a caused motion construction.[7]

Effectively, this criterion excludes 18a, 18b, and 18c from CMCs.

3.3.1.3 Existence of Path of Motion

Caused motion constructions are not uniquely definable by their syntactic form. Not all syntactic structures of the form (NP-SBJ (V NP PP)) belong to CMCs. Consider the following sentences.

(20) a. We saw the bird in the shopping mall.

b. The boy bounced the ball in front of the house.

c. I considered Ben as one of my brothers.

d. Mary finally kicked the ball, to my relief.

The prepositional phrases in 20 do not specify paths of motion. In sentences 20a and 20b they do denote locative information, but, more specifically, they are static locations where the events occur. In 20a, the prepositional phrase indicates the location where the seeing event happened, not a path along which the subject caused the bird to move. In sentence 20b, the boy is causing the ball to bounce, however, the prepositional phrase only indicates where the bouncing event occurred, not where the ball was caused to move.

---

[7] We note here that the entailment of motion does not test for whether or not the motion proceeds uninterrupted. As Rappaport-Hovav and Levin (2008) note, certain verbs like *throw* do not entail the attainment of a goal, while for others verbs like *drag* the goal cannot be denied:

- I threw the ball to Julian, but it fell short of him.
- *I dragged the box to the door, but stopped before I got there.

Our criterion is specifically for the entailment of motion, rather than the attainment of goal.

While no human annotator would have trouble recognizing that these prepositional phrases do not carry path information, there are still cases where it is necessary to specify to the annotators that the cause motion constructions must have a path. We specify the existence of a path as a requirement for identifying CMCs. For specification of the criterion, we adopt from Goldberg's observation that the path of motion must fully be determined by the cause argument.

> **Criterion 3**: The prepositional phrase must denote the undergoer's path of motion. The path of motion must exist and be fully determined by the cause argument.

We present three areas where this criterion has a direct relevance.

Semantically symmetrical arguments: Consider the pairs of examples in 21. For each example, the first sentence is a case of CMC, while the second one is not. We observed that these types of sentences are particularly tricky for annotators. The reasons are not clear, but it is likely because these instances normally take verbs (e.g. *tape*, *tie*, *paste*, or *split*) that can appear in CMC usage, being used as non-CMCs in a CMC syntactic form.

(21) a. I added an egg to the batter.
I added 5 to 6.

b. The company introduced a superconcentrated Lemon Cheer in Japan.
I introduced my fiance to my mother.

c. He later separated a single cell from a 16-celled embryo.
How do you separate oil from water?

What we have in the above examples are verbs being used in caused motion and non caused motion contexts. The prepositional phrases in the first sentence in each pair can be construed as a path: the egg put into the batter in 21a, the detergent is brought into Japan as a product in 21b, and a single cell is removed from an embroyo in 21c. However, the semantics of path does not exist in the second sentence of each pair: the direct object of the verb and the object of the preposition are separate entities that are being either joined or separated. For example, in the first sentence 21a, the batter is the whole to which the egg is being incorporated. In the second sentence, however, 6 is not the whole to which 5 is being added; rather 6 and 5 are added together to make up a larger whole, 11. In the case of 21b, the second sentence indicates that the finance and the mother are being introduced to each other.

The non-CMC instances are able to undergo Levin's reciprocal alternation (Levin 1993; also see Dougherty (1974) for *each other* test) while the CMC instances cannot.

(22) a. ?I {added/mixed} an egg and the batter (together).[8]
I added 5 and 6 (together).

b. *The company introduced a superconcentrated Lemon Cheer and Japan (to each other).
I introduced my fiancee and my mother (to each other).

c. ?He separated a single cell and a 16-celled embryo (from each other).
We separated oil and water (from each other).

Applying Levin's reciprocal alternation as in example 22 suggests that there is a semantic symmetry (Fillmore, 1966; Dowty, 1991) between the direct object of the verb and the complement of the preposition in the non-caused motion sentences. In other words, these sentences cannot be CMCs as the semantic value of the two arguments in reference to the verb are equivalent and, therefore, the argument expressed in the prepositional phrase cannot be a path.

Thus, the annotators need to be mindful not to make the CMC selection based on the verb alone. In addition to Levin's reciprocal alternation, we provide a second test to the annotators to evaluate if they are dealing with symmetrical arguments. In this test we have them evaluate if it is semantically allowable to swap the lexical contents of the direct object and the object of the prepositions:

[8] An interesting side point: on a judgment test, 2 out of 5 speakers judged *"I mixed an egg and the batter together"* unacceptable, but reordering the mention of the two ingredients, *"I mixed the batter and an egg together"* made it acceptable.

(23) a. ?I added the batter to an egg.

I added 6 to 5.

b. *The company introduced Japan in a super concentrated Lemon Cheer. (WSJ)

I introduced my mother to my fiance.

c. *He later separated 16-celled embryo from a single cell.

How do you separate water from oil?

The idea is that if the arguments are symmetrically used then they should be able to exchange places in the syntactic template without greatly altering the semantics of the sentence. In examples above, the swapping of the arguments of CMC renders the sentences unacceptable. While the first sentence in 23a is understandable, it would be odd to order the arguments in this way. In the case of example 23b, the first sentence of the pair could be rephrased to read *The company introduced Japan to a super concentrated Lemon Cheer*. However, if we do this, then the semantics that this sentence invokes is the social encounter meaning much like the second sentence of the pair.

Creation event: The following examples are instances of the creation event. Each of them has the appropriate syntactic structure and a locative prepositional phrase to pass for a caused motion sentence. However, earlier in the annotation studies, a decision was made to exclude cases of creation events as instances of CMCs. This decision is based on the observation that there is no perceptible path indicated in a creation event.

(24) a. Yet, Sierra's hard-line stance has created something of a rift in the organization.

b. He composed a new series of events into his ever-growing novel.

c. The IRS builds another factor into its secret computer formula for selecting returns for audit.

d. Create an empty folder on the hard drive which you can point iTunes at.

e. FirstBank opened a new branch in Denver.

In the above sentences, the entity described by the direct object is coming into being as a result event designated by the verb. In sentence 24a, a rift comes into being as a result of creation. Similarly, in sentence 24b, a new series is being created through composition. The verb *open* in example 24e is not normally associated with verbs of creation unlike verbs such as *create*, *build*, or *compose*. However, the semantics of the sentence shows parallels with the rest of the examples in that a new branch of FirstBank is caused to come into being in Denver. Compare these sentences to the following:

(25) I installed a new SSD into my Lenovo G580.

The semantics of sentence 25 resemble those in 24 in that we can say that by the end of the event the direct object now exists in the location described by the prepositional phrase. Compare example 25 to 24e:

(26) a. In 25: *I installed a new SSD into my Lenovo.*
→ At end of event: *A new SSD is now in my Lenovo*

b. In 24e: *FirstBank opened a new branch in Denver.*
→ At end of event: *A new branch of FirstBank is now in Denver*

However, that's the extent of their parallels. Because of the creation event, the created object, in this case, the direct object, cannot exist before the beginning of the event, making further inferences based on the two sentences different for the beginning of the event:

(27) a. In 25: *I installed a new SSD into my Lenovo.*
→ At start of event: *I have a new SSD, but it is not (yet) in my Lenovo.*
→ At end of event: *A new SSD is now in my Lenovo*

b. In 24e: *FirstBank opened a new branch in Denver.*
→ At start of event: **FirstBank has a new branch, but it is not (yet) in Denver.*
→ At end of event: *A new branch of FirstBank is now in Denver*

Creation events cannot have a starting location or state as the created item does not exist before the moment of creation. In other words, creation events do not have an identifiable path. The only place these instances share semantics with path of motion instances are in the end-of-event inference that the item in question now exists in the location designated by the prepositional phrase.

Value assignment: As a result of using data from the Wall Street Journal corpus for the creation of guidelines and annotation, there were financial sentences using verbs like *value* and *price* to indicate a value assignment event. Here are a couple examples:

(28) a. They valued its new proposal at $2.29 billion. (WSJ)

b. The notebook, the TI Model 12, will be priced at $4,199. (WSJ)

The reason these sentences came to our attention is because verbs like *put*, which prototypically participate in CMCs, were appearing in these sentence types (e.g. "The state assessor had put the value at $5.7 billion"). The question that arose was: can we liken these sentences to a putting event where the undergoer of the movement is the value and the *at* phrase specifies a location in a scalar path? It turns out that with these instances we can apply the same analysis used on creation events:

(29) In 24e: *They valued its new proposal at $2.29 billion.*

→ At start of event: **They had a value for the new proposal, but it is not (yet) placed at the value of $2.29 billion.*

→ At end of event: *They now place the proposal at the value of $2.29 billion.*

The only way we could interpret these sentences as cases of movement along a scalar path is if we can say that the value moved from one point to another. However, because the direct object does not gain a value until the moment of the assignment, the valued object had no price at the start of the event. Therefore, we cannot say that these instances express path information.

3.3.1.4 Summary

The following are the three criteria we have discussed in this section.

- The cause argument must exert a force over the undergoer, resulting in the undergoer's motion. The causal entity must be either agentive or be capable of wielding the force necessary for a direct effect on the undergoer.

- The undergoer must be in motion during the event. Only when the motion of the undergoer is strictly entailed for the event is a sentence judged as a caused motion construction.

- The prepositional phrase must denote the undergoer's path of motion. The path of motion must exist and be fully determined by the cause argument.

Thus far, we have implicitly touched upon the issues concerning abstract and metaphorical realizations of CMCs without an explicit reference to them. For the rest of this section, we look at the metaphorical distinctions we make in our definitions of CMC.

3.3.2 Metaphorical Distinctions

It should be of no great surprise that sentences expressing abstract entities in their arguments or instances that deal with metaphorical extensions of CMCs tend to be more difficult for the annotators. We found that clear-cut motion (e.g. *He kicked the ball into the bin*) or even coerced instances involving concrete motion (e.g. *The crowd laughed him off the stage*) were easier for human annotators to categorize as CMCs. However, the more abstract the meaning, the more difficult was the evaluation. Consider the following examples of CMCs:

(30) a. They plan to recycle them into fresh sealing clay. (WEB)

b. Ron Brierley raised its stake in the company Friday to 15.02% from about 14.6% Thursday. (WSJ)

c. Exxon [...] sold its 53-floor Rockefeller Center skyscraper to a Japanese company. (WSJ)

The difficulty of the sentences in example 30 hang on the interpretation of the nature of the path invoked in these sentences. Where does the stake in the company go in sentence 30b? Is there any motion implied in 30a and 30c? This is especially difficult when the criteria as presented earlier in Section 3.3 are based on physical motion rather than the abstract.

We currently recognize three major metaphorical extensions to the physical instantiations of CMCs. These categories are (1) *cause-transform-cx*, which expresses changes of state as a metaphorical extension to changes of location; (2) *cause-change-scale-cx*, where the path of motion is conceived of as a linear scale; and (3) *cause-change-possession-cx*, which includes CMCs denoting the transfer of possession or information. We also recognize a fourth category, *cause-displacement-cx*, which includes the physical CMCs. However, unlike the first three categories, which are motivated by the metaphorical usages of CMCs, this fourth category is semantically broader and includes instances that were identified as CMCs by the criteria specified in Section 3.3.1.4, and yet do not fit into any of the metaphorical categories. This last category, as we will see, could merit a closer attention as it requires further semantic disambiguation.

3.3.2.1 Changes of State

Various studies have recognized the metaphorical link between changes of state and changes of location (Gruber, 1962; Jackendoff, 1983; Lakoff, 1993; Goldberg, 1995; Goldberg and Jackendoff, 2004). According to the metaphor, states map to physical locations and change of state maps to motion. Thus, the transition from an initial state to a resultant state is conceived of as a path of motion.

That said, our category of *cause-transform-cx* is slightly narrower in its semantics. It is exemplified in 30a and in the following examples:

(31) a. Joe broke the bottle into pieces. F

b. Rival gangs have turned cities into combat zones. (WSJ)

This category describes CMCs whose undergoer argument goes through a transformational change in the event described by the verb. More specifically, this category indicates a change in the property or attribute of the undergoer argument (c.f. *identificational transition* in Gruber 1962). For example in 31a, the bottle undergoes a transformational change and the path prepositional phrase specifies the resultant state of the bottle. In the same way, the cities are changed to become combat zones in 31b.

Goldberg's diagram for the metaphorical extension involving changes of state is shown in Figure 3.4. The diagram for *cause-transform-cx* is instantiated with the sentence in 31a. We have already seen the diagram for the general caused motion sense in Section 3.2. The diagram below it shows that the goal role is specified as the result, and that the two constructions are related by the metaphorical extension link: I_M.

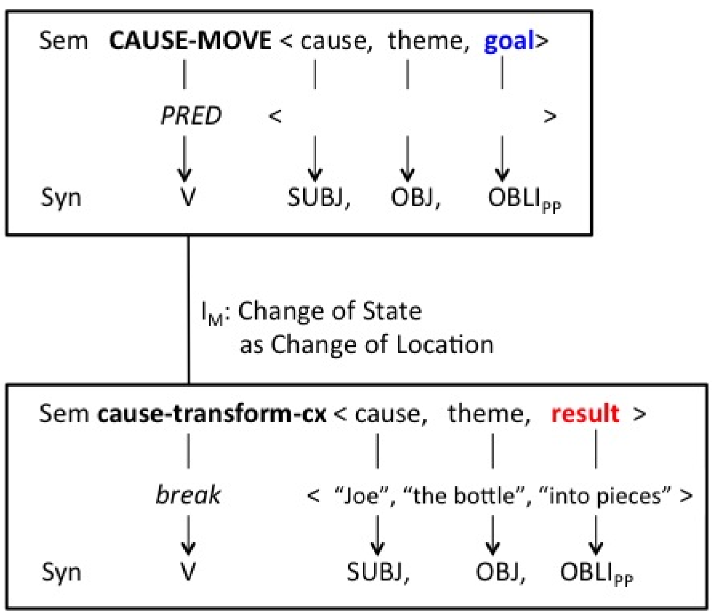

Figure 3.4: Diagram for cause-transform-cx.

3.3.2.2 Changes in Scale

Among the well-recognized metaphors is Lakeoff's 'LINEAR SCALES are PATHS' (Lakoff and Johnson, 1980; Jackendoff, 1983; Lakoff, 1993). Under this metaphor, values are mapped to locations, and changes in value are mapped to motion. In sentence 30b, once the path of motion is conceived of as a linear scale, the motion can be described as the movement from one point to another point in that scale. The following sentence shows an instance of a CMC that expresses the complete path (i.e., source, goal, and traversed path locations):

(32) The cash injection boosted Zeta's capital to 8.47 billion pesetas from 1.82 billion pesetas. (WSJ)

The above sentence is diagrammed in Figure 3.5 using the Goldberg-style representation. The *cause-transform-cx* definition has been instantiated with sentence 32. As you can see, both the goal and the source arguments are realized in this sub-construction, and the I_M link specifies the mapping of the change in scale to the change of location.

Since the verb *boost* normally specifies for one participant role (e.g. the boosted entity as in *'Borat' boosted tourism to Kazakhstan*), the sentence receives the semantics of the source/initial and goal/final values directly from the construction, which is represented through the dotted lines connecting the top semantics level to the middle verb's argument structure level.

3.3.2.3 Transfer of Ownership and Transfer of Information

We also recognize the *cause-change-possession-cx* category. The metaphor at work in this category maps possessions to the entity in motion and the act of being possessed to the spatial location (Gruber, 1962; Lakoff and Johnson, 1980; Jackendoff, 1983; Lakoff, 1993; Goldberg, 1995; Rappaport-Hovav and Levin, 2008). Consequently, under this metaphor, the transfer is mapped to motion and the path of transfer conceptualized as the path of motion. Additionally, by way of the conduit metaphor (Reddy, 1979; Lakoff and Johnson, 1980), we include transfer of information in this category. Thus, information or ideas are objects that are transferred from the

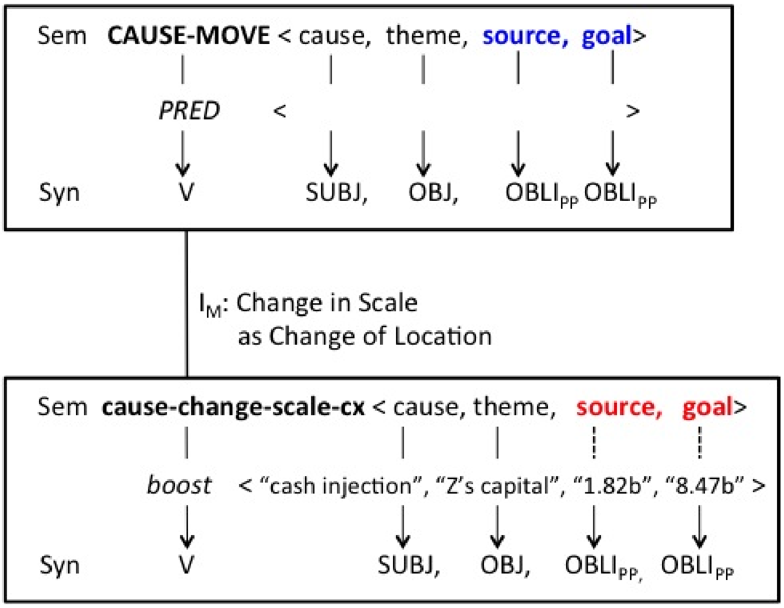

Figure 3.5: Diagram for cause-change-scale-cx.

speaker/donor to the recipient. Like changes of possession, we conceptualize the transfer of ideas from person A to person B as movement of an object from location X to location Y.

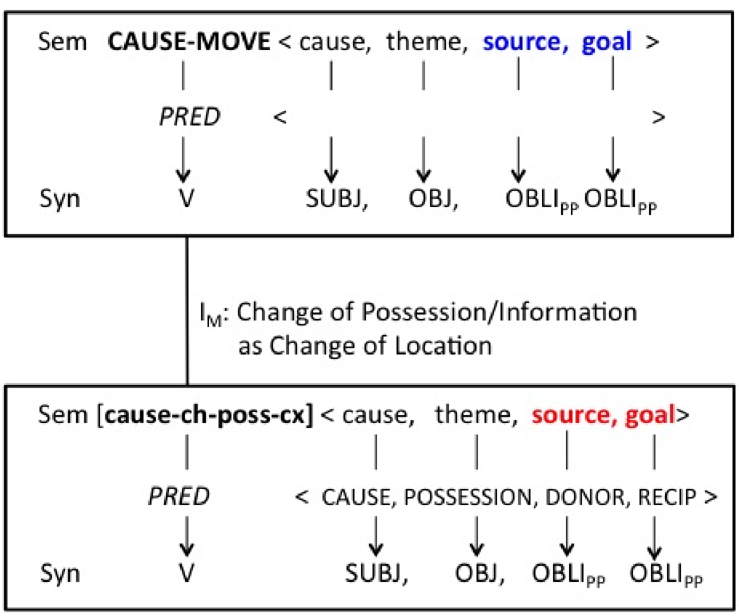

Figure 3.6: Diagram for cause-change-possession-cx.

Figure 3.6 diagrams the change of possession category. The possession is linked to the theme in motion, and the source and goal arguments map to the predicate's donor and recipient arguments. Unlike the previous diagrams, the particular specifications of this construction calls for both the donor and recipient, which cannot exist in English. Consider the following pairs of sentences.

(33) a. Mary bought the coat from the Salvation Army.

Mary bought the coat from the Salvation Army *to her sister*.

b. Mary give her coat to the Salvation Army.

Mary give her coat *from her sister* to the Salvation Army.

The only way that the second sentence in 33a can be made acceptable is to rephrase the to-prepositional phrase as a for-prepositional phrase (i.e. *Mary bought her coat from the Salvation Army for her sister*). However, this would indicate a benefactive role rather than a recipient roal for the buying event; whether or not her sister would receive the item would be underspecified. The second sentence in 33b is unacceptable in the reading where "from her sister" is an argument of the verb. In other words, the buying event cannot instantiate the recipient role as an oblique argument, and the giving event cannot instantiate the donor role as an oblique argument. For this reason, we introduce a *cause-have-cx* category for inclusion of instances like 33a along with the *cause-receive-cx* for the inclusion of instances like 33b.

Goldberg's (1995) analysis provides the metaphorical mapping seen in Figure 3.6, the *cause-receive-cx* construction, which only includes a recipient argument. In light of the thematic constraints on oblique arguments, we reformulate Goldberg's definitions slightly. We illustrate both constructions with Goldberg-style representations in Figure 3.7.

Please note the coindexation marked in red ink between the cause and donor arguments in *cause-receive-cx* and between the cause and recipient arguments in *cause-have-cx*. The lexical layers of the diagrams receive all four arguments from the constructional layers, but the syntactic layers only realize one of the two arguments. Effectively, the cause argument takes on a second role for both constructions: the cause arguments take on the donor role in *cause-receive-cx* and the cause argument takes on the recipient role in *cause-have-cx*.

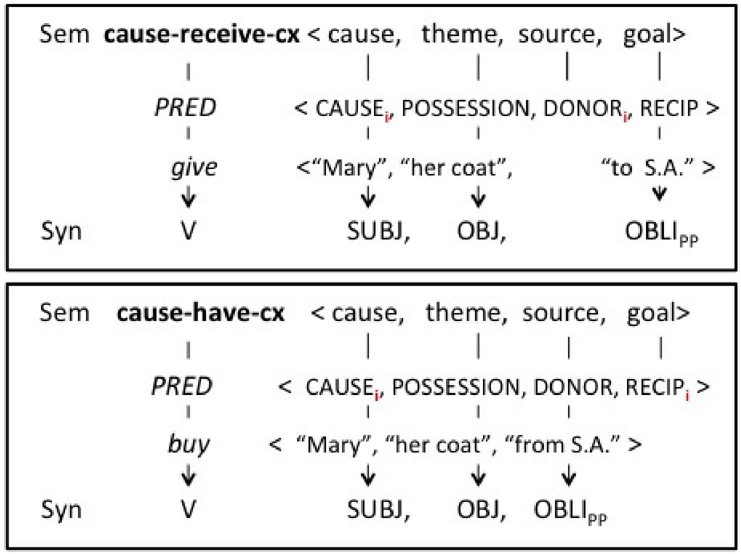

Figure 3.7: Diagram for cause-receive-cx and cause-have-cx.

3.3.3 The *cause-displacement-cx* Category

Now we turn to our current representation's most general category: *cause-displacement-cx*. IIt is the most prototypical of the CMC categories as it includes the concrete motion on which all the other metaphorical extensions are based. In this category we also make a distinction between the *cause-remove-cx* and *cause-put-cx* types. The *cause-remove-cx* category includes instances that can only take the goal location, and the *cause-put-cx* category includes instances that only takes the source locations. Sentences that can instantiate both goal and source locations are considered instances of the cause-displacement-cx category. Following are a few examples:

(34) a. He kicked the ball into the bin. [*cause-displacement-cx*]

b. She put the book on the table. [*cause-put-cx*]

c. She wiped the dirt from her hands. [*cause-remove-cx*]

However, it is also the most varied of the categories as it includes abstract instances, judged by the criteria in Section 3.3.1.4, as CMCs, that do not neatly fit in to the metaphorically motivated categories we have seen so far. The following are some sentences that show the degree of variety under this label. The verb that heads the construction in each sentence is bold-faced for ease of reading.

(35) a. I [..] will stitch together three sides and **add** a matching zipper to the fourth side.

b. He **bullies** Kate into a dance that consists of drooling on her while trying to break her ribs.

c. Occam's Razor should **incline** us towards the more mundane explanations, which almost invariably prove to be correct.

d. Since then he has expanded his fleet and can now **bring** his furs to the front door of retailers as far away as the Midwest.

e. President Bush **called** his attention to the matter during the Italian leader's visit.

f. Can anyone **shed** some light on this please?

The first example, sentence 35a, is highly concrete, as it deals with attaching a zipper to a fabric of some kind. Sentences 35b and 35c have abstract paths, the first of which feels more concrete as the dance the sentence is referring to describes a setting that can be described through physical qualities (e.g. dance floor). Example 35d speaks not of the actual bringing of furs to a physical front door but of an importing or a business event that is being described using physical verbiage. And finally the last two examples could be considered idiomatic expressions.

The problem with lumping these semantically varied sentences into the same category is in the drawing of inferences. For example, given sentences 35a and 35f we can say that by the end of the event the zipper and the light move to the specified goal location, but the nature of the movement obviously is very different – the first expresses a physical movement and the latter expresses a metaphorical movement. For the example in 35a, we can say that the zipper is now physically at the location of the fabric. The same inference does not hold for 35f. Thus, for tasks where the distinction is important we require further disambiguation. During the early stages of CMC type analysis, we have discussed a possible inclusion of finer grained subcategorization of CMC types, which did not make into the final classification. These categories are listed in Appendix A.

As a first step to distinguishing other metaphorical extensions represented in the *cause-displacement-cx* category, we chose to annotate the data that was labeled as CMC with a concreteness rating. Since the metaphorical extensions of CMCs are, in fact, the abstract usages of CMCs, a concreteness evaluation that would serve to show the different levels of abstractness represented in the data might be a helpful first initial effort. The annotation process, guidelines and results are presented in Section 3.5 The study presented is in its early stages of investigation.

3.4 Annotation of the Caused Motion Constructions

3.4.1 Annotation Labels

Our final semantic type classification is shown in Figure 3.8. It includes 4 main categories of CMCs, one of which does not have any instances of its own (i.e. *cause-change-possession-cx*), and 4 semantically finer grained sub-categories. Each of these categories detail the particular characteristics and semantic types that they lend to consistent semantic inferences. Bracketing around the category label specifies that the category is uninstantiable.

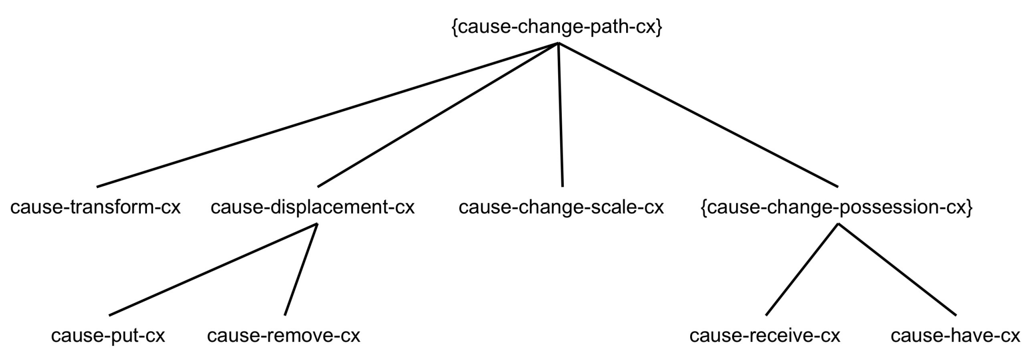

Figure 3.8: CMC Classification

3.4.1.1 cause-displacement-cx

The category *cause-displacement-cx* (DISPLACE) describes the most prototypical of the CMC types. The labels in this category are given to sentences in which the theme undergoes a change of location (physical or abstract) without an attributive or transformational change. DISPLACE has two finer grained subcategories: *cause-put-cx* (PUT), which includes instances that syntactically realize only the GOAL arguments and *cause-remove-cx* (REMOVE), which includes instances that syntactically realize only the SOURCE arguments. Unlike the two sub-categories, DISPLACE can occur with either of the path arguments. The following are examples in these sub-categories:

(36) a. (PUT) She put the book *[from the shelf]-SOURCE [on the table]-GOAL.

b. (REMOVE) She removed the book [from the shelf]-SOURCE *[onto the table]-GOAL.

3.4.1.2 cause-change-scale-cx

The category *cause-change-scale-cx* (SCALE) identifies CMCs in which the path of motion is mapped to a linear scale. It is specifically reserved for instances where the patient argument moves along a path that is expressed as a linear scale, and includes sentences like the one seen in example 30b. This movement can be in any direction. Following are further examples of this category:

(37) a. I marked the price down to 2 dollars.

b. The heavy downpour raised the level of the lake to 1000 ft.

3.4.1.3 cause-change-possession-cx

The category *cause-change-possession-cx* (TRANSFER) includes caused motion instances that denote transfer of possession. In this category the path of motion (physical or abstract) is defined as the path on which the patient argument moves from the SOURCE entity's possession to GOAL entity's possession.

There are two sub-categories under TRANSFER: *cause-receive-cx* (GIVE), which includes instances that syntactically realize only the GOAL (or recipient) argument and *cause-have-cx* (RECEIVE), which includes instances that syntactically realize only the SOURCE (or giver). Instead of the SOURCE argument for GIVE and GOAL argument for RECEIVE, the cause argument (in this case an AGENT) takes on the role of the GOAL and SOURCE arguments, respectively. Following are some examples:

(38) a. (GIVE) Mary gave my coat [to the salvation army]-GOAL.

b. (RECEIVE) Mary bought my coat [from the salvation army]-SOURCE.

3.4.1.4 cause-transform-cx

The category of *cause-transform-cx* (TRANSFORM) identifies CMCs in which the object of the verb undergoes a transformational change in the event described by the verb. As discussed earlier, this category includes CMCs for which the path is conceived of as a change of state, and the motion is described as the movement from an initial state/attribute to a final state/attribute as exemplified in sentence 30a. The following are further examples of this category.

(39) a. I broke the vase into little pieces.

b. This company renders crayon scribblings of toddlers into some kind of abstract art.

In example 39a, the PATIENT argument *vase* undergoes a breaking event which results in the attribute described in the prepositional phrase *into little pieces*. In the same way, in example 39b, the company transforms – presumably by re-interpretation – the attribute or the condition of the crayon scribblings of toddlers into abstract art. As seen in the examples, this label can be given to both physical (39a) and abstract (39b) sentences.

3.4.1.5 Summary

The classification, thus, introduces a total of 7 CM labels for annotation according to the classification seen in 3.4.1.

- **DISPLACE** for cause-displacement-cx
- **PUT** for cause-put-cx
- **REMOVE** for cause-remove-cx
- **SCALE** for cause-change-scale-cx
- **TRANSFORM** for cause-transform-cx
- **GIVE** for cause-have-cx
- **RECEIVE** for cause-receive-cx

In addition to these 7 labels, annotators were provided with a miscellaneous OTHER category under which they could classify sentences they judged as CMC but did not feel fit readily fit into the 7 classified categories. Furthermore, the classification specifies standards for disambiguating ambiguous instances and default labels in cases where the decision is made impossible (e.g. lack of context, specialized jargon that is difficult to understand).

3.4.2 Annotation Process

The data for this study was pulled from the Wall Street Journal (WSJ) [9] , Web Text (WEB), and Broadcasting News (BN) corpora of the OntoNotes project (Weischedel et al., 2012). Verbal phrases matching the syntactic form “Subject Verb Object PrepPhrase” (i.e. (NP-SBJ (V NP PP))) were selected from the above corpora using the Penn Treebank annotation. Note that because Penn Treebank includes traces, the selection includes instances that are marked as underlyingly (NP-SBJ (V NP PP)). In other words, the selection also includes passive sentences (e.g. *Coffee was shipped*

[9] This corpus contains over 846K words selected from the non “strictly financial (e.g., daily market reports) portion of the Wall Street Journal included in the Penn Treebank II Marcus et al. (1994)).

from Colombia by Gracie.) and instances with traces in the object NP or PP including questions and relative clauses (e.g. *What did Gracie ship from Colombia?*).

Corpus	Instances Annotated	Double Annotated	Annotation Labels	Annotator's Previous CMC Knowledge
WSJ	15,209	10%	1 CMC, 1 NON-CMC	Strong
WEB	1,824	100%	7 CMC, 1 NON-CMC	Strong
BN	3,776	100%	7 CMC, 1 NON-CMC	Very little

Table 3.1: Corpus of annotation, annotated instances, and annotator knowledge.

For the first two passes of annotation, the sentences from the WSJ (totalling 15209 instances) were manually judged as either a CMC or a NON-CMC with 10% of the data double annotated by two annotators, both of whom were deeply familiar with the linguistic literature on CMCs. In the first of the two passes, the annotators were asked to mark categories that they felt were distinct from one another in addition to specifying the binary CMC or NON-CMC labels. In this process, they were given general directions to look out for instances of changes of possession and transformation usages with the ability to specify different types of instances as they saw fit. The double annotated data was then adjudicated and the disagreements were analyzed.

In an effort to systematize the annotation and raise the inter-annotator agreement, we took a closer look at the annotated data, with special focus to the categories the annotators came up as they were doing the annotation. Based on these annotations and a series of discussion, we carried out a detailed analysis of the semantic types and entailments of CMCs. In turn, the analyses lead to the the establishment of the CMC classification and the annotation guidelines. In the second part of the two passes devoted to WSJ, the data annotated as CMC were re-annotated under the new classification.

The third and fourth passes of the annotation were conducted under the newly established guidelines. For the third pass of the annotation, instances matching the CMC syntactic forms were selected from the WEB corpus (totalling 1824 instances). These instances were double annotated by the same two expert annotators. For the fourth pass of the annotation, instances matching the

CMC syntactic forms were selected from the BN corpus (totalling 3776 instances). For this final pass of the annotation, two new annotators unfamiliar with CMC were given training in the guidelines before the annotation of BN. The annotation of the third and fourth passes was adjudicated for disagreements.

3.4.3 Results

3.4.3.1 Annotation Label Distribution

Table 3.2 shows a breakdown of the labels according to the annotated data. The columns labeled "Annotation Labels" display the counts for each label with the relative frequency (i.e. count/total annotated instances) in parenthesis. The labels DISPLACE, PUT and REMOVE, and the labels RECEIVE (i.e., cause to have) and GIVE (i.e., cause to receive) are placed together in their respective semantic groups. The count and relative frequencies of these semantic groups are also provided. The sum of each of the "Annotation Labels" rows equals the total number of CMCs found in the data. The total CMC and NON-CMC counts are shown in the far right columns.

	Annotation Labels							**Totals**	
	DISP.	PUT	REM.	REC.	GIVE	SCALE	TRANSF.	CMC	NON-CMC
WSJ	643 (0.042)	511 (0.034)	107 (0.007)	257 (0.017)	446 (0.029)	202 (0.013)	84 (0.006)	2250 (0.148)	12959 (0.852)
	1261 (0.083)			703 (0.046)				15209	
WEB	199 (0.109)	202 (0.111)	11 (0.006)	25 (0.014)	62 (0.034)	12 (0.007)	22 (0.012)	533 (0.292)	1291 (0.708)
	412 (.226)			87 (0.054)				1824	
BN	249 (0.066)	199 (0.053)	63 (0.017)	30 (0.008)	121 (0.032)	16 (0.004)	25 (0.007)	703 (0.186)	3073 (0.814)
	511 (.135)			151 (0.040)				3776	
CMC Type	cause-displace-cx			cause-ch-possession-cx		cause-ch-scale-cx	cause-transf.-cx		

Table 3.2: Label distribution given the total number of annotated instances in the corpus.

The number of CMCs found in the corpora seems to vary from corpus to corpus in relative proportions. In the case of the WEB corpus, CMCs account for almost 30% of the syntactically similar instances, while in the WSJ corpus, CMCs account for less than 15% of the annotated data. Table 3.3 shows the distribution of the individual CMC labels within the data identified as CMCs. Again, the counts are accompanied by their respective relative frequencies given the total number of CMC instances in the corpus.

<table>
<tr><th>Corpus</th><th>DISP.</th><th>PUT</th><th>REM.</th><th>REC.</th><th>GIVE</th><th>SCALE</th><th>TRANSF.</th></tr>
<tr><td rowspan="2">WSJ</td><td>643
(0.286)</td><td>511
(0.227)</td><td>107
(0.048)</td><td>257
(0.114)</td><td>446
(0.198)</td><td rowspan="2">202
(0.090)</td><td rowspan="2">84
(0.037)</td></tr>
<tr><td colspan="3">1261
(0.560)</td><td colspan="2">703
(0.312)</td></tr>
<tr><td rowspan="2">WEB</td><td>199
(0.373)</td><td>202
(0.379)</td><td>11
(0.021)</td><td>25
(0.047)</td><td>62
(0.116)</td><td rowspan="2">12
(0.023)</td><td rowspan="2">22
(0.041)</td></tr>
<tr><td colspan="3">412
(0.773)</td><td colspan="2">87
(0.163)</td></tr>
<tr><td rowspan="2">BN</td><td>249
(0.354)</td><td>199
(0.283)</td><td>63
(0.090)</td><td>30
(0.043)</td><td>121
(0.172)</td><td rowspan="2">16
(0.023)</td><td rowspan="2">25
(0.036)</td></tr>
<tr><td colspan="3">511
(0.727)</td><td colspan="2">151
(0.215)</td></tr>
</table>

Table 3.3: Label distribution given the total number of instances annotated as CMC.

The *cause-displace-cx* labels make up the majority of all CMC labels. The WSJ corpus shows the most change of possession type instances, which is likely attributable to the semantics of purchase and transaction that is common in the financial newspaper genre. The fact that the frequency of the SCALE label is almost 4 times higher in the WSJ corpus could also be attributed to the the financial usages found in the corpus.

3.4.3.2 Verb Distribution

Table 3.4 shows the 20 most frequent CMC verbs in each corpus. 6 verbs that appear in all three corpora have been italicized.

The verbs from the WSJ and BN corpora consist of verbs we often expect to find in caused motion usages. The verb list for the WEB corpus is a little more interesting: we find coercive usages of CMC for verbs like *drum*, *ram*, *incline*, or *shine*. Here are a couple example sentences:

WSJ		WEB		BN	
put	provide	shed	post	*bring*	sell
sell	buy	*put*	spend	*put*	move
give	pay	*send*	add	*send*	fire
bring	*get*	plug	incline	*take*	carry
send	place	ram	*take*	*give*	deliver
take	turn	*bring*	drum	*get*	submit
receive	transfer	*give*	shine	place	pass
increase	add	*get*	apply	turn	drive
raise	remove	rip	buy	throw	receive
move	invest	pump	sweep	lead	push

Table 3.4: Verb Distribution in CMC Data.

(40) a. We'll **drum** these criminals out of office in the shame they deserve, even if it takes until January 2008 .

b. If your ideas are any good you'll have to **ram** them down people's throats.

The verbs in Table 3.4 are members of 43 VerbNet classes[10] . The five most frequent class series represented include **Verbs of Change of Possession** (series 13 classes, including *Give-13.1*, *Get-13.5.1*, and *Contribute-13.2*), **Verbs of Sending and Carrying** (series 11 classes, including *Send-11.1*, *Bring-11.3*, and *Carry-11.4*), **Verbs of Putting** (series 9 classes, including *Put-9.1* and *Funnel-9.3*), **Verbs of Removing** (series 10 classes, including *Remove-10.1*, *Wipe_Manner-10.4.1* and *Steal-10.5*), and **Verbs of Motion** (series 51 classes, including *Roll-51.3.1*, *Run-51.3.2* and *Accompany-51.7*).

accuse	create	have	open	reach
base	do	hold	play	rise
be	fall	keep	prevent	see
build	file	leave	produce	use
charge	find	offer	quote	view

Table 3.5: Top 25 NON-CMC verbs

[10] The count is made by strict membership count. That is, each VerbNet class in which these verbs were members were counted.

Table 3.5 lists the top 25 highly frequent verbs that do not have CMC usages in the corpora. This list includes verbs like *be*, *do*, or *have* that cannot appear in caused motion usages, and verbs like *keep*, *leave*, or *prevent* that are contrary to the semantics of CMCs as discussed in this chapter. This list also includes verbs that could potentially be found in CMC usages such as *build* (e.g. *Was it Dr Madnar or Doktor that built him into a cyborg?*) and *file* (e.g. *Kay filed the paper into a crisp new manilla folder*).

3.4.3.3 Inter-Annotator Agreements

Table 3.6 reports on the inter-annotator agreement scores. One notable aspect of the passes conducted after the establishment of the guidelines is that the annotators were asked to label instances with one of the 9 potential labels (i.e. 8 CMC labels and 1 NON-CMC label; see Section 3.4.1.5). The "Overall IAA" reports the annotator agreement over the overall caused motion label (CMC vs. NON-CMC)[11] .

	CM Expertise	Corpus	Overall IAA	F-Score	
Pre-Guidelines	expert	WSJ	.883	0.667	
Post-Guidelines	expert	WEB	.881	0.764	*
	beginner	BN	.839	0.606	

Table 3.6: Inter-annotator agreement rates for WSJ, WEB, and BN

The agreement numbers show 88.3% and 88.1% agreement between the expert annotators for the pre-guideline and post-guideline annotations, respectively. The annotators newly trained on CMC guidelines also showed a high agreement rate of 83.9%.

Because the negative label NON-CMC makes up the majority class (85.2% in WSJ, 70.1% in WEB, 81.4% in BN), the true negative labels outnumber the CMC instances. This means that while the values represented by the overall IAA are indicative of the general annotator agreement over both CMC and NON-CMC labels, they do not indicate how the annotators performed on the

[11] The 8 labels of CMC counts were collapsed under a single CMC label.

CMC labels in particular. We find that the F-score, which only takes into account the positive instances in its calculation, better represents the inter-annotator agreement for the caused motion label (Hripcsak and Rothschild, 2005). The F-score differences were also evaluated via a chi-squared test at a significance level of $p<0.05$ (Yeh, 2000). They are reported in Table 3.6.

For the expert annotators, we found that the F-score agreement after the introduction of the CMC guidelines was significantly higher ($p = 0.018$; $chi^2 = 5.59$, DF = 1) than the score obtained from the annotation done before the guidelines were introduced. Additionally the annotators qualitatively reported that the guidelines helped with problematic or ambiguous cases, resulting in an improved annotation experience. The newly trained annotators' agreement, however, was lower than that of the expert agreement (F-score of 0.606). However, this difference was not statistically significant ($p = 0.192$; $chi^2 = 1.70$, DF = 1) from the expert annotation of the WSJ data.

The inter-annotator rates in the different passes of annotation show a comparable agreement of over 83%. However, the F-scores suggest that there was a significant difference in the annotation agreement when the agreement is calculated based on the positive labels only. The results in Table 3.6 suggest that the guidelines were useful for the expert annotators. Despite the fact that more labels for annotation normally means a higher cognitive load, it is encouraging to find that the usefulness of the guidelines in annotation outweighed the cognitive load to produce higher performance.[12] For the newly trained annotators, who were previously unfamiliar with caused motion literature, we were able to achieve comparable annotation agreement rates as that of the pre-guideline efforts by expert annotators.

3.5 Concreteness Rating: A Pilot Study

Currently, the *cause-displacement-cx* category is a fairly coarse-grained group. It includes physical and abstract instances that may lead to differing inferences. Here are a couple examples of the we discussed earlier:

[12] The IAA rates of the individual labels from the post-guideline passes show 79.5% and 79.2% agreement for WEB and BN corpora annotations, respectively. However, as it was for the overall IAA rates, we expect that these numbers will have to be studied a little more closely to understand what the agreements truly signify.

(41) a. I [..] will stitch together three sides and **add** a matching zipper to the fourth side.

b. President Bush **called** his attention to the matter during the Italian leader's visit.

The first sentence is an instance of a truly physical CMC, while in the second instance the *attention* is conceptualized as the undergoing entity in motion. Unlike the first sentence that requires no additional interpretation to say that the zipper is caused to be attached at a certain location, in the second sentence we only get the caused motion reading via metaphorical extension (e.g. Attention is a metaphorical object: e.g. "You lost my attention!"; Topics or issues are metaphorical locations: e.g. "Let's move on to the next matter at hand").

As a first step in distinguishing what other metaphorical extensions are represented in the *cause-displacement-cx* category, we chose to annotate the data that was labeled as CMC with a concreteness rating. Specifically, the 2250 instances from the WSJ corpus that were identified as CMCs in the previous section were annotated with one of the following four labels:

- **Fully Concrete**: The undergoer is a physical entity, and the path is a physical place in this world. The sentence must have an interpretation where the event describes a completely physical event. The following are examples from the corpora:

  (42) a. I *scooted* them into the dog run and stood in the doorway myself.

  b. Thieves *stole* a 12th century fresco from an abandoned church in Camerino, Italy.

- **Partially Concrete, Partially Abstract**: At least one argument (undergoer or path) is a concrete or tangible entity. If the undergoer argument is a physical entity, then the path argument cannot be a physical place. If the path argument is concrete, then the undergoer argument cannot be a physical entity. The abstract arguments are bolded in the following examples:

  (43) a. The company *plugged* itself right into **Carter campaign rhetoric**.

  b. Cathay also is *moving* **some of its [...] operations** outside Hong Kong .

- **Fully Abstract**: Both undergoer and path arguments are abstract, and the sentence cannot have a physical interpretation.

(44) a. That could *shove* a weak economy into recession

b. The Associated Press 's earthquake coverage *drew* attention to a phenomenon that [...]

c. This expected blow has *cast* a pall over the economy's prospects.

- **Idiomatic**: These are the instances that an annotators may not consider to be CMCs because they feel highly idiomatic. The sentence retains a meaning beyond that denoted by the caused motion interpretation.

(45) a. Arkansas Democrat David Pryor *spilled* his guts on the Senate floor the other day.

b. We could *sweep* it under the rug and hide it.

The 2250 instances of CMCs from the WSJ corpus were annotated by one annotator, and 783 instances (approximately a third of the data) were double annotated by a second annotator. The double annotated portion of the data shows an agreement rate of 79%. What we seek here is to capture the concreteness of the event by determining whether or not the undergoer or the path refer to physical entities/places. By design we keep the "Fully Concrete" label fairly fine grained in terms of its concreteness. However, the distinctions between the "Fully Abstract" and "Partially Abstract" instances seems to be the largest source of the disagreements. Here are three examples of the disagreements:

(46) a. The brokers' hope has been that they could soon coax investors into *shifting* some of their hoard into the stock market.

b. The 1989 act simply *extends* these guarantees to the private sector.

c. We scuttled along for a few feet before he *plunged* us into the drink again.

Most of the disagreements hinge on the interpretation of money or organizations. For example, in sentence 46a, it is likely that the dispute is in the concreteness of *the hoard*. If this sentence had a path that was clearly physical (e.g. *The thieves shifted their hoard into the house*), it would force a concrete reading out of the direct object. As the path is not a physical one, the concreteness of the object is not as certain. If the object is abstract, then this would be an instance of “Fully Abstract”. Otherwise, we would have a “Partially Abstract” instance. We have a similar issue with sentence 46b. The annotation depends on the annotator’s interpretation of the path prepositional phrase.

The last example, 46c, is one that confuses the “Fully Concrete” vs “Fully Abstract” labels. Here we have a physical *us* and a *drink* that can have a very real and concrete form. The event also indicates a highly tangible setting, calling up images of people scuttling along a street, very likely drunk. However, this sentence can have an abstract reading where the *us* refers to the self that is capable of inebriation and *drink* refers to the state of drunkeness. In other words, under this interpretation, a physical reading of the sentence would have the cause argument literally shove a group of people into a large container full of drink. Given this, it is likely that the annotator who chose to label this as a fully concrete instance was relying on the setting of the event, while the second annotator was cuing in on the abstract usage that stands in contrast to the alternate physical usage.

One result of this effort is the understanding that these types of distinctions are possible with relatively high level of agreement. However, the approach used in this study is still at an early stage, and there is much potential for further improvement. The agreement rate could certainly use improvement. One way to do so would be to fill in an obvious gap: this study’s current lack of a theoretical backing. The study was borne out of the general need for a way to distinguish more abstract usages from physical ones that retain very different semantic inferences and semantic consequences. However, with the feasibility of such annotation established, the next step would be to look at the literature from linguistics, psycholinguistics or other lexical analysis studies to better understand what has been said about the nature of concreteness and abstractness in order to address some of the issues we are seeing in the annotators' disagreements between the fully and partially abstract labels. There are some examples that may warrant multiple labels, given their ambiguity.

3.6 Final Considerations and Future Work

In this chapter we presented our work on characterizing and defining caused motion constructions by expanding and refining the definitions proposed by Goldberg (1995). We carried out a detailed analysis and annotation of CMCs, systematizing their defining characteristics into a typological classification and creating annotation guidelines that will improve corpus annotation of CMCs. The outcome of this effort is the semantic typing of CMCs and the development of the categorical types of CMCs as they occur in corpus data. Our classification of CMCs gives insight into the types and varieties of semantic inferences entailed by CMCs in corpus data. We used categorical labels of CMCs for the annotation of corpus data. The annotation, thus, provides data for the training and testing of automatic CMC classifiers. In the next chapter, we present our classification efforts.

There are several directions we could take with the future work. First, we only recognize a few of the metaphorical distinctions in our current analysis. Additional distinctions would call for a further analysis and refinement of the current groupings based on corpus analysis, particularly the *cause-displace-cx* category. We take our first stab at making this distinction through the concreteness ratings, but, as noted earlier, this approach requires further investigation. Secondly, this work does not address the status of verb particle constructions where the particle expresses directional information in the event described by the verb. For example, sentences like *I threw the pill up about an hour after taking it*, where the particle *up* can be analyzed as encoding the path of the pill's movement, have been excluded from this study. In fact, future studies will give specific attention to these verb particle constructions that encode directional expressions to determine their relevance and relatedness to caused motion constructions.

Finally, a similar corpus driven classification could be conducted for other constructions. Future work might include extending the constructional analysis to additional construction types including, but not limited to, adjectival resultatives (e.g. *Mary hammered the metal flat*), conative constructions (e.g. *Brian wiped at the counter with a damp rag*), and ditransitive or “cause to give” constructions (e.g. *John lent me a bicycle*). Definitions of these constructions will aid in producing sentence representations that would not be possible with the semantics of the verb alone.

Chapter 4

Automatic Identification of Constructions

4.1 Introduction

In this chapter, we present our studies on the automatic classification of caused motion constructions (CMCs). Our study makes use of a supervised learning approach, where an automatic classifier is trained to identify a phenomenon based on exemplars in the training data. At the decoding stage, the classifier is asked to predict the label for the phenomenon in a new set of data. The results are then evaluated against the gold standard annotation to gauge the classifier's performance. We utilize the CMC annotation we produced in Chapter 3 as training and testing data to explore the semantic and syntactic features that would best help characterize caused motion constructions in an automatic learning setting.

First we discuss a small-scale pilot study we conducted to determine the feasibility of detecting CMCs in corpus data. The rest of the chapter is devoted to further experiments expanding beyond this exploratory study. Thus, Section 4.2 will present the pilot study conducted over a small subset of the WSJ corpus. Section 4.3 describes our experimental setup. The rest of the chapter discusses the experiments we conducted to establish the most predictive and generalizable features for classifying caused motion constructions.

4.2 Pilot Study

As an initial step in determining the usefulness of construction grammar for interpreting semantics in computational linguistics, we conducted a pilot study aimed at ascertaining if a classifier can be taught to identify CMCs (Hwang et al., 2010). This section reports on on our investigations into which features were the most useful in the classification of CMCs.

The results presented show that a classifier can be trained to automatically identify the semantics of constructions; at least for CMCs, and can do this with high accuracy. Furthermore, we have determined that for the pilot the preposition feature is the most useful feature for identifying CMCs. Additionally we found that using constructions to inform semantics resulted in equally high performance on the out-of-vocabulary predicates. This serves as evidence that semantic analysis of novel lexical combinations and unseen verbs can be improved by enriching semantics with a construction-level analysis.

It would be important to mention that the pilot study was conducted before the establishment of CMCs' semantic types and categories. Therefore, the decisions on what should be considered caused motion, at this stage, was based on Goldberg's (1995) definitions. CMC definitions included both literal usages (e.g. *Well-wishers stuck little ANC flags in their hair*) and non-literal usages (e.g. *Producers shepherded 'Flashdance' through several scripts*) of caused motion.

4.2.1 Data Selection

The data for this study is pulled from the Wall Street Journal part of Penn Treebank II (Marcus et al., 1994). From this corpus, all sentences with the syntactic form (NP-SBJ (V NP PP)) were selected. Treebank syntactically treats caused motion sentences in the same manner as the sentences in which the PP is an adjunct phrase to the verb. The following two sentences would both be parsed in Treebank as (NP-SBJ (V NP PP)).

(47) a. Bill rolled the ball down the hill. [Caused Motion]

b. Bill ate an apple for lunch. [Non-Caused Motion]

The selection allowed for intervening adverbial phrases (e.g. *Sally threw a ball* ***quickly*** *to him*) and additional prepositional phrases (e.g. *Sally threw a ball to him* ***on Tuesday*** or *Sally threw a ball* ***in anger*** *into the scorer's table*). A total of 14.7k instances were identified in this manner.

To reduce the size of the corpus to be labeled to a target of 1800 instances, we first removed instances containing traces as parsed by the Penn Treebank. These included passive usages (e.g. *Coffee was shipped from Colombia by Gracie*) and instances with traces in the object NP or PP including questions and relative clauses (e.g. *What did Gracie ship from Colombia?*). Secondly, we removed the instances of sentences that can be deterministically categorized as NON-CMCs: instances containing ADV, EXT, PRD, VOC, or TMP type object NPs (e.g. *Cindy drove* ***five hours*** *from Dallas* or *You listen,* **boy**, *to what I say!*). Because we can automatically identify this category, keeping these examples in our data would have resulted in even higher performance. We also considered reducing the size by removing certain classes of verbs such as verbs of communication (e.g. *reply*, *bark*), psychological state (e.g. *amuse*, *admire*), or existence (e.g. *be*, *exist*). While it is reasonable to say that these verb types are highly unlikely to appear in a CMC, if we were to remove sets of verbs based on their likely behavior, we would also be excluding interesting usages such as *The stand-up comedian joked me into a state of uncontrollable giggles* or *The leader barked a command into a radio*.

After filtering these sentences, 8700 remained. From the remaining instances, we randomly selected 1800 instances for the experiments presented.

4.2.2 Labels and Classifier

The 1800 instances were hand-labeled with one of the following two labels: Caused Motion (CMC) and Non Caused Motion (NON-CMC). After the annotation, the corpus was randomly divided into two sets: 75% for training data and 25% for testing data. The distribution of the labels in the test data is 33.3% CMC and 66.7% NON-CMC. The distribution in the training set is 31.8% CMC and 68.2% NON-CMC. For our experiments, we used a Support Vector Machine (SVM) classifier with a linear kernel. In particular we made use of LIBSVM (Chang and Lin, 2001) as training and testing software.

4.2.3 Features

The baseline consisted of a single conceptual feature - the lemmatized, case-normalized verb. We chose the verb as a baseline feature because it is generally accepted to be the core lexical item in a sentence, which governs the syntactic structure and semantic constituents around it. This verb feature was encoded as 478 binary features (one for each unique verb in the dataset), where the feature value corresponding to the instance's verb was 1 and all others were 0. In the experiments, we utilized gold-standard values for two of the PP features for a proof of feasibility. In addition to the baseline verb feature (feature 1), our full feature set consisted of 8 additional types for a total of 334 features. The full set of features is listed in Table 4.1.

4.2.4 Results

For the baseline system, the model was built from the training data using a linear kernel and a cost parameter of C=1 (LIBSVM default value). When using the full feature set, the model was also built from the training data using a linear kernel, but the cost parameter was C=0.5, the best value from a 10-fold cross validation on the training data. We report our results in Table 4.2.

The results show that the addition of features 2-9 shown in Table 4.1 resulted in a significant increase in both precision and recall, which boosted the f-score from 0.624 to 0.857, an increase of 0.233.

4.2.4.1 Feature Performance

In order to determine the usefulness of the individual features in the classification of caused-motion, we evaluated the features in two ways. In the first study, we compared the performance of each of the features to a majority class baseline (i.e. 66.7% accuracy). A useful feature was expected to show a statistically significant increase over this baseline. The significance of each feature's performance was evaluated via a chi-squared test ($p < 0.05$). Our results show that features 1, 2, 3 and 5 performed significantly better than the majority class baseline. The features

#	Name	Count	Description
Baseline Features			
1	Lemma	478	The lemmatized, case-normalized verb encoded as binary features
PP Features			
2	Preposition	76	The preposition heading the prepositional phrase was encoded as 76 binary features, one per preposition type in the training data. For instances with multiple PPs, preposition features were extracted from each of the PPs.
3	Function Tag on PP	11	Penn Treebank encodes grammatical, adverbial, and other related information on the PPs POS tag (e.g. PP-LOC). The function tag on the prepositional phrase was encoded as 10 binary features plus an extra feature for PPs without function tags. .
4	Complement Category to P	19	Normally a PP node consists of a P and a NP. However, there are some cases where the complement of the P can be of a different syntactic category (e.g. *So, view permanent insurance [[for]*P ***[what it is]*****SBAR***]*PP.). Thus, the phrasal category tags (e.g. NP, SBAR) of the prepositions sister nodes were encoded as 19 binary features. For instances with multiple PPs, all sister nodes of the prepositions were collected.
VerbNet Features: The following features were automatically extracted from VerbNet classes with frames matching the target syntactic structure, namely "NP V NP PP".			
5	VerbNet Classes	123	The verbs in the data were associated with one or more of the above VerbNet classes according to their membership. If a verb belongs to multiple matching classes, each corresponding feature was set.
6	VerbNet PP Type	27	VerbNet frames associate the PP with a description (e.g. NP V NP PP**.location**). The features represented the union of all PP types (i.e. if a VerbNet class included multiple PPs, each of the corresponding features was assigned a value of 1). If a verb was associated with multiple VerbNet classes, the features were set according to the union over both the corresponding classes and their set of PP types.
Noun Entity (NE) Features: These features were automatically annotated using BBNs IdentiFinder (Bikel et al., 1999).			
7	Subj NP	23	The union of all NE under the NP-SBJ node was encoded as 23 binary features.
8	Obj NP	27	The union of all NE under the object NP node was encoded as 27 binary features.
9	PPs Obj NP	28	The union of all NE under the NP under the PP node was encoded as 28 binary features.

Table 4.1: Full description and examples of features used (Part 2).

4, 7 and 8 were unable to distinguish between the caused-motion constructions and the non caused-motion usages. Their precision values could not be calculated because these features resulted in zero positive (CMC) classifications.

Features	P%	R%	F	A%
Baseline Set	78.0	52.0	0.624	79.1
Full Set	87.2	86.0	0.857	91.1

Table 4.2: System performance: precision (P), recall (R), F1 score, and accuracy (A)

In a second study, we evaluated the performance of the system when each feature was removed individually from the full set of features. The removal of a useful feature was expected to show a statistically significant drop in performance compared to that of the full feature set. The significance in this performance degradation when compared against the full set of features was evaluated via a chi-squared test (McNemar, $p < 0.05$). When features 1, 3 and 8, were removed, we observed a statistically significant performance drop. The rest of the features were not shown to have a statistically significant effect on the performance. Our results show that the preposition feature (feature 2) is the single most predictive feature and the feature with the most significant effect in the full feature set. These results are encouraging: unlike the purely lexical features (i.e. the named entity features 6, 7, and 8) that are dependent on the particular expression used in the sentence, prepositions are function words. Like syntactic elements, these function words also contribute to the patterned structures of a construction as discussed in Chapter 3. Furthermore, unlike the semantics of features that are dependent on content words that are subject to lexical variability, prepositions are limited in their lexical variability, which make them good general features that scale well across different semantic domains.

In addition to the preposition feature, the verb feature (feature 1) was found to affect performance at a statistically significant level in both cases. Based on numerous studies in the past that have shown the usefulness of the verb as a feature, this is not an unexpected result. Interestingly, our results seem to indicate interactions between features. This can be seen in two different instances. First, while feature 8 (NEs for Object NP) on its own was not found to be a predictive feature, when removed, it resulted in a statistically significant drop in performance compared to that of the full feature set. The opposite effect can be seen with the VerbNet class feature. While it showed a statistically significant boost in performance when introduced into the system by itself, when dropped from the full feature set, the drop in the system performance was not found to be significant. This seems to indicate that the NEs for Object NP and the VerbNet Classes features have strong interactions with one or more of the other features.

4.2.4.2 Out-of-Vocabulary Verbs

Additionally, we separately examined the performance on the test set verbs that were not seen in the training data (i.e., out-of-vocabulary/OOV items). Just over a fifth of the instances (92 out of 450 constructions) in the test data had unseen verbs, with a total of 83 unique verb types. The results show that there was no decrease in the accuracy or F-score. In fact, there was a chance increase, not statistically significant, in a two-sample t-test ($t = 1.13$; $p > 0.2$).

We carried out the same feature studies detailed in the previous section for the OOV verbs. The performance of both studies reflected the results Section 4.2.4.1, with one expected exception. The verb feature was found to be of no value to the predictor. What is interesting here is that the verb feature did perform at a significant level for the full test data. From this observation, it would be expected that the overall performance on the OOV verbs would be negatively affected since there is no available verb information. However, this was not the case.

4.3 Experimental Setup

For the rest of the chapter we present the experiments that expand on the pilot study. The following experimental setup applies to all of the experiments presented in the rest of this chapter.

4.3.1 Corpora

We have briefly mentioned the relevant corpora in Section 3.4.2 of Chapter 3. Here are further details on each corpus. The data comes from the latest version of OntoNotes, version 5.0, (Weischedel et al., 2012). Gold annotations for Penn Treebank, PropBank, and Verb Sense

Annotation are available for all of OntoNotes corpora. As we did for the pilot study, we use the Wall Street Journal (WSJ) corpus. This corpus contains over 846K words selected from the non "strictly" financial (e.g., daily market reports) portion of the Wall Street Journal included in the Penn Treebank II (Marcus et al., 1994). We also pull from the smaller of the two WebText (WEB) data sets published in OntoNotes. This corpus contains 85K words selected from English weblogs. This portion of the data is not to be confused with the the larger 200K word web data, which is a separate corpora in OntoNotes. The third corpora used in our experiments is the 200K word Broadcast News (BN) data. OntoNotes' BN data contains news texts from broadcasting sources such as CNN, ABC, and PRI (Public Radio International).

4.3.2 Data Selection

In order to narrow the data down to a more manageable size for annotation, we exclude instances that can be deterministically categorized as NON-CMCs using the gold Penn Treebank annotation of the corpora. To do this we first select all sentences with the base syntactic form (NP-SBJ (V NP PP)) based on the Penn Treebank gold annotation. Note that the Penn Treebank includes traces. In the pilot study, we had excluded instances with traces, opting to match sentences by their surface structure. In other words, the data selection for the pilot excluded other syntactic selections that match caused motion constructions such as passive sentences (e.g. *Coffee was shipped from Colombia by Gracie.*) and instances with traces in the object NP or PP including questions and relative clauses (e.g. *What did Gracie ship from Colombia?*). For the current study, our data selection includes instances that retain an underlying syntactic form (NP-SBJ (V NP PP)). In effect, we extend the syntactic variability in the data.

Additionally, we use a set of heuristics (a smaller set than the pilot) to further select instances of potential CMCs. Instances which satisfy the following three conditions are extracted for annotation:

(1) An NP exists in the verb phrase.

(2) At least one PP exists in the verb phrase.

(3) The NP precedes the PP in the verb phrase.

For the remaining data, already annotated instances from the pilot study are separated out for double-checking. We also set aside instances that can be deterministically categorized as non-caused motion: instances with the function tags ADV, EXT, PRD, VOC, or TMP. These sentences are kept for a quick verification at the annotation stage that they indeed are cases of non-caused motion and labeled as such.

4.3.3 Labels and Classfiers

All of our experiments are geared towards the identification of the CMC label and the DISPLACE label. We build two binary classifiers: one for each of the two labels. The CMC label includes all of the instances labeled with one of the 7 caused motion types described in Chapter 3. The DISPLACE label includes the instances labeled with the *cause-displacement-cx* type, which also includes the PUT (i.e. *cause-put-cx*) and REMOVE (i.e. *cause-remove-cx*) instances. Table 4.3 shows the CMC type and the annotation labels included in the two classification labels. Table 4.4 shows the classification label distribution across the three corpora.

Classification Label	**CMC Type**	**Annotation Labels**
CMC	all types	DISPLACE, PUT, REMOVE TRANSFORM, SCALE, GIVE, RECEIVE
DISPLACE	cause-displacement-cx	DISPLACE, PUT, REMOVE

Table 4.3: Classification Labels

For all our experiments, 80% of annotated data is randomly selected as the training/development data and the remaining 20% is set aside as the test/evaluation set. For our experiments, we use a Support Vector Machine (SVM) classifier with a linear kernel. In particular, we use LIBSVM (Chang and Lin, 2001) as our training and testing software. We use a 5-fold cross-validation process for the development stage.

	WSJ		WEB		BN	
CMC	2250	14.8%	533	29.2%	703	18.6%
NONCMC	12959	85.2%	1291	70.8%	3073	81.4%
DISPLACE	1261	8.3%	412	22.6%	511	13.5%
NONDISPLACE	13948	91.7%	1412	77.4%	3265	86.5%

Table 4.4: CMC and DISPLACE label distribution in training and test data

4.3.4 Features

The features encode syntactic and semantic information that targets four elements in the sentence: (1) the verb, which expresses the event or the situation of the sentence, (2) the preposition, which instantiates the path information in a caused motion sentence, (3) the complement of the preposition, which covers the rest of the prepositional phrase, (4) the cause argument, which is recovered from the subject of the sentence or the prepositional by-phrase in a passive sentence, and (5) the undergoer argument, which is recovered from the direct object position of the sentence or from the subject position in a passive sentence. We will discuss the cause and undergoer argument recovery in further detail later.

Furthermore, features that encode **SEMANTIC** information include:

- **Nominal Entity** features which are automatically generated using BBNs IdentiFinder (Bikel et al., 1999). The IdentiFinder annotates relevant noun phrases with labels such as “Persons”, “Time”, “Location”, or “Organization”.

- **PropBank Frameset** features specify the verb’s sense based on its subcategorization frame. This is extracted from the gold annotation provided by Ontonotes.

- **Ontonotes Verb Sense** features which specify the verb’s sense. The semantics of these features are generally finer grained than what the PropBank framesets encode. These features are also provided as gold annotation in OntoNotes.

- **VerbNet Class** features that encode each of the VerbNet classes in which the verb is a member. A verb can be a member of one or more VerbNet classes.

- **Preposition Type** features obtained from the automatic preposition labeller developed in a recent study by Srikumar (2013). The labeller introduces a set of 32 roles that semantically disambiguate a preposition to distinguish between the different semantic sense of the preposition as expressed in a sentence (e.g. the preposition *from* in *Poor care lead to her death* ***from*** *pneumonia.* (Cause) vs. *She copied the lines* ***from*** *the film.*(Source))

Features encoding **SYNTACTIC** information include the:

- **Part of Speech Tag** of the lexical item in the syntactic parse.

- **Dependency Relation Tag** of the lexical item in a dependency parse.

The following is a complete listing of all the feature types used in the experiments. The verb lemma feature is the baseline feature for all our experiments. Anytime we use the terms "Full Set" or full feature set, we referring to a set of features that includes all of the 16 feature types below.

Verb Feature Sets:

(1) **verb lemma [Baseline]:** lemmatized and case-normalized verb

(2) **verb pos:** part of speech of the verb

(3) **verb deprel:** dependency relation of the verb to its head

(4) **VerbNet class:** verb's associated VerbNet classes. If a verb belongs to multiple matching classes, each corresponding feature is set.

(5) **PB frameset:** verb's associated PropBank frames item **verb sense:** verb's associated Ontonotes verb sense

Preposition Feature Set:

(6) **preposition:** the preposition feature is set for each preposition found in the sentence

(7) **preposition type:** the preposition type feature is set for each preposition found in the sentence

Complement of Preposition (pcomp) Feature Sets:

Note: The features are set for each prepositional phrase in the sentence.

(8) **pcomp pos:** part of speech of the complement of the preposition

(9) **pcomp deprel:** complement's dependency relation to the preposition

(10) **pcomp ne:** the pcomp ne feature is set for each nominal entity found as a complement of the preposition.

Cause Argument Feature Sets:

(11) **cause pos:** part of speech of the head of the cause argument

(12) **cause deprel:** head of the cause argument's dependency relation to the verb

(13) **cause ne:** the cause ne feature is set for each nominal entity found in the cause argument.

Undergoer Argument Feature Sets:

(14) **undergoer pos:** part of speech of the head of the undergoer argument

(15) **undergoer deprel:** head of the undergoer argument's dependency relation to the verb

(16) **undergoer ne:** the undergoer ne feature is set for each nominal entity found in the undergoer argument.

4.3.4.1 On Syntactic Features

One major difference from the pilot study is that all syntactic information gathered after the data selection process comes from dependency parses, rather than the phrase structure parses. The decision to encode syntactic features from the dependency parses rather than for phrasal parses was based on the flexibility and the amount of additional information we gain through the dependency parse type. After a series of experimental runs with features from both parse types, it was determined that further syntactic features based on the phrase trees produced relatively similar performance to that of its counterpart labels on the dependency trees. However, the dependency labels are functionally finer grained than phrase structure labels for those syntactic elements that are most relevant to the CMCs. For example, dependency parsing subdivides the prepositional usages by their function. Here are three examples:

(48) a. They don't break teeth [as-MARK they thrash around up there].

b. It must suck to be hated [by-AGENT everyone].

c. [It] will punch straight [through-PREP all platters] and cause enough shock [...]

The phrasal parses would give prepositional phrase (PP) treatment to both of these phrases. Because dependency parse is interested in the relationship between two words, thus, the prepositions are labeled with a dependency label that is relevant to their relationship to the verb (e.g. MARK for word introducing a finite clause subordinate to another clause and AGENT complement of a passive verb which is introduced by the preposition). This gives us greater specificity in the selected features and flexibility for their use.

In our experiments, we employ the CLEAR dependency parses (Choi, 2012). In particular, these parses have been automatically converted from the Penn Treebank phrasal trees. In the pilot study, the two features encoded from the phrasal tree were the syntactic category of the complement of the prepositional phrase and the semantic function tag (part of the gold annotation) of the prepositional phrase (e.g. in "PP-LOC" the function tag "LOC" describes the semantic value of the phrase given the verb). The information encoded by the complement category is generally captured from the preposition by the dependency relation feature. We no longer include the semantic function tag in our experiments.

Beyond what phrasal trees can provide for us, dependency parses are able to provide the dependency relations such as *nsubjpass* and *agent*, which are useful in capturing the proper causal argument and the undergoer argument in passive sentences. The following examples show simplified phrasal and dependency labels for an active and a passive sentence:

(49) a. [Phrasal: Active] (The dog)-NP-SUBJ chased (the cat)-NP.

[Phrasal: Passive] ((The cat)-*i*)-NP-SUBJ was chased (*-*i*)-NP (by the dog)-PP

b. [Dependency: Active] (The dog)-NSUBJ chased (the cat)-DOBJ.

[Dependency: Passive] (The cat)-NSUBJPASS was chased (by the dog)-AGENT.

The phrasal parse handles a passive sentence by introducing a trace after the verb, where the direct object (moved to the subject position) would be found in the underlying representation. This trace is coindexed with the subject of the passive sentence to indicate the movement from the object position to the subject position. The oblique that represents the prepositional phrase is marked as a prepositional phrase, which can be indistinguishable from other *by*-phrases (e.g. *The cat had been chased* ***by the dog*** *for several minutes* ***by the time I arrived***.). The dependency tree lacks the trace, but the dependency relations are formulated to reflect the relationship of the arguments to the verb. The *nsubjpass* label specifies that the constituent is the nominal subject of a passive sentence and the *agent* label indicates that the constituent is the agent of the passive sentence. So in dependency trees, what information we lose by not having the trace information for the passive sentences, we gain back in a more obvious label that distinguishes the *by*-phrase from other prepositional phrases.

4.3.4.2 Cause and Undergoer Argument Recovery

Another difference reflected in these features is the treatment of the cause and undergoer arguments. For this study, we make a pre-processing pass of the data to recover these arguments when possible. The recovered arguments are as following:

- **Passive Sentences:** For passive sentences, the complement of the *by*-prepositional phrase is recovered as the cause argument and the subject is recovered as the undergoer argument.
- **Conjunctions:** Given two verbal conjuncts sharing the subject, as in *"He cut the peppers and diced the tomatoes"*, the phrasal tree places the two verb phrases and the conjunction

as sister constituents, while the CLEAR dependency parse places the conjunction and the second conjunct as dependents of the first verb. The following shows a simplified phrasal tree on the left and a dependency tree on the right.

```
(S                                 cut
 +- (NP-SUBJ He)                     +- He
 +- (VP cut the peppers)             +- the peppers
 +- (CONJ and)                  vs.  +- and
 +- (VP diced the tomatoes)          +- diced
)                                         +- the tomatoes
```

This means, that while in phrasal trees we could simply access the verb's NP-SUBJ for the cause argument, in dependency trees the two conjuncts' access to the cause argument is not symmetrical. The argument *He* is accessible to the verb *diced* via the verb *cut*, as the argument is a direct dependent of the verb *cut* and not the verb *diced*. To recover the arguments of the second verb conjunct we reach for the dependent on the first conjunct as necessary.

- **Subordinate clauses:** For verbs that are found in subordinate clauses whose head node is a verb (also called matrix verb) such as an infinitival clause (e.g. *He [plans]*-HEAD *to* ***cut*** *the peppers into pieces*), or a relative clause (e.g. *Joe [cut]*-HEAD *the tomatoes Mary* ***washed***.), we reach for the head node's arguments to fill in the missing cause and theme arguments. If there is an intervening relative pronoun (e.g. *Joe cut the tomatoes* ***that*** *Mary washed*), the relative pronoun is retrieved as the argument (either as cause or theme depending whether or not the subordinate clause is a passive), instead.

4.3.4.3 Part of Speech Tags & Dependency Relation Tags

After a series of experimental runs, it was determined that the part of speech and the dependency relation features might be too fine grained to provide useful information to the classifier. More coarse groupings did improve the f-score by a few hundredths of points. Therefore, all of the features expressed by the part of speech and the dependency relation are featurized in the following manner.

Part of Speech Tags[1]

- Cardinal numbers (CD), pronouns (PRP), and gerundial (VBG) and participial (VBN) forms of verbs are featurized as found (one feature per tag).
- Rest of the verb forms are mapped to the base tag VB.
- Plural nouns are mapped to their singular counterparts.
- Adjectives and adverbs are all mapped to the base tag JJ and RB, respectively.
- Rest are given the tag: OTHER

Dependency Relation Labels

- Relations specifying subjects, direct object, and agent (oblique of a passive sentence), and relations specifying the object of the preposition, complement clauses, and relative clauses are featurized as found (one feature per tag).
- Complement clauses (e.g. *pcomp*, *acomp*) are grouped under a single *comp* label.
- Modifiers (e.g. *partmod*, *advmod*) are grouped under a single *mod* label.
- Rest are given the tag: Other

4.4 Classifier Experiments

4.4.1 Syntactic vs. Semantic Features

Tables 4.5 and 4.6 show the precision and recall percentages and the f-score values for our experiments. Here we show results for three feature combinations: the Baseline set encoded from the verb's lemma, the Baseline plus the preposition feature set (Baseline+P), and the Full Set that includes all of the features listed in Section 4.3.4. The best performance values are bold-faced. The significance of a feature set's performance was evaluated via a chi-squared test (McNemar, $p < 0.05$). Statistically significant change from the Baseline feature set is marked with a †. Additionally, for the CMC classification we show the inter-annotator agreement (Gold) f-score from

[1] A complete listing of the part of speech tags with their description is found in Appendix B.

Chapter 3. Our best performances in CMC classification as measured by the f-score are comparable or higher than the inter annotator agreement f-score.

	WSJ			WEB			BN		
	P	R	F	P	R	F	P	R	F
Baseline	61.23	37.56	0.4656	75.6	55.7	0.641	71.4	53.6	0.612
Baseline+P	**75.00**	74.67	0.7483†	78.0	**80.2**	**0.791**†	**84.8**	75.7	0.800†
Full Set	74.00	**77.78**	**0.7584**†	**79.0**	78.3	0.787†	84.1	**82.9**	**0.835**†
Gold			0.667			0.764			0.606

Table 4.5: System performance on CMC label classification.

Statistically significant change from the Baseline feature set is marked with a †.

	WSJ			WEB			BN		
	P	R	F	P	R	F	P	R	F
Baseline	66.80	63.89	0.6531	72.7	58.5	0.649	71.3	55.9	0.626
Baseline+P	**76.33**	74.21	**0.7525**†	73.4	70.7	0.720	80.0	70.6	0.750†
Full Set	72.52	**75.40**	0.7393†	**76.5**	**79.3**	**0.778**†	**80.6**	**77.5**	**0.790**†

Table 4.6: System performance on DISPLACE label classification.

Statistically significant change from the Baseline feature set is marked with a †.

With the exception of the DISPLACE classifier on the WEB corpus, both the Baseline+P and the full set of features perform significantly better than the Baseline in both sets of experiments. What is interesting here is that the Baseline+P set performs just as well and sometimes better than the full set of features fairly consistently across the corpora, though the differences in the values are not statistically significant.

In order to gain a better understanding of the performance on the full set of features, a series of feature experiments were carried out. For these experiments, the full feature set was divided into syntactic features and semantic features as described in Section 4.3.4. As a means of control, both the syntactic and semantic feature sets also include the features for the verb lemma and the preposition. Out of the different feature combinations examined, the distinction between semantic and syntactic features is the most salient. Table 4.7 shows the system performance values for the syntactic and semantic features. We also show the performance of the Baseline+P plus VerbNet class (Baseline+P+VNCLS) feature set, as it gives better insight into the semantic feature performance.

CMC Classification:									
	WSJ			WEB			BN		
	P	R	F	P	R	F	P	R	F
Syntactic	63.79	41.11	0.5000	76.6	55.7	0.645	72.4	54.3	0.620
Semantic	71.02	72.44	0.7173	77.3	64.2	0.701	80.5	76.4	0.784
Baseline+P+VNCLS	71.78	76.89	0.7425	78.8	77.4	0.781	85.9	82.9	0.844
DISPLACE Classification:									
	WSJ			WEB			BN		
	P	R	F	P	R	F	P	R	F
Syntactic	66.80	63.89	0.6531	73.8	58.5	0.653	72.3	58.8	0.649
Semantic	72.94	73.81	0.7337	76.3	70.7	0.734	74.3	79.4	0.768
Baseline+P+VNCLS	74.81	76.59	0.7569	78.7	72.0	0.752	82.8	75.5	0.790

Table 4.7: System performance on semantic and syntactic features.

The numbers indicate that the semantic features have a consistently higher performance than the syntactic features. The syntactic feature sets, perform significantly lower than the full feature sets and they barely pass the Baseline features in performance. In fact, the syntactic features are significantly lower than the Baseline+P features, despite the fact that, just like the semantic features, they include the verb lemma feature and the preposition feature. This suggests, that the syntactic features even in the presence of the lexical features are not strongly predictive of caused motion constructions. Moreover, these numbers seem to indicate that the performance on the full set of features likely comes from the semantic feature performance.

Amongst the semantic features, the Baseline feature, the Baseline+P feature, and the feature for VerbNet class membership of the verb (i.e. Baseline+P+VNCLS) give the highest results. With the exception of the CMC classifier on the BN corpus, the numbers for the Baseline+P+VNCLS set are not significantly different from either the semantic feature or the full feature set performance. Other semantic combinations were also tested, but they did not result in any particular change from the semantic feature set and the full feature set.

In conclusion, the semantic features come out of these experiments as the most predictive features. This finding makes intuitive sense. Recall that during the data selection stage, we selected for instances that show syntactic compatibility with CMCs. Although syntactic variability still exists in the selected data (e.g. relative clauses and passive sentences), because of the data selection stage based on syntax, the task of identification comes primarily down to the semantic distinction between existing sentences. Additionally, some of the existing syntactic differences are neutralized by the cause and undergoer argument pre-processing stage described in Section 4.3.4.2. Thus, it stands to reason that most of the useful contributions come from the lexical items themselves and the semantics of the verb and its arguments.

Finally, the baseline system of the DISPLACE classification shows either a similar or improved performance over the CMC classifier. The overall performances across the different feature sets show similar values. Given that DISPLACE label makes up a smaller percentage of the total data as shown in Section 4.3.3 (e.g. DISPLACE label for WSJ accounts for just under 9% of the total test and training data), the comparable performance is likely indicative of the fact that the DISPLACE label represents a more semantically coherent phenomenon than the CMC label.

4.4.2 Cross-Genre Experiments

To investigate how the features perform for out-of-genre texts, the classifiers trained on WSJ data were evaluated on WEB and BN data. Table 4.8 shows the f-score results of these experiments. For ease of reading, we include in the tables the difference between the out-of-genre and the in-genre values, from previous tables (e.g. the Baseline feature for CMC classification when the WSJ model is tested on WEB data shows a decrease of 0.148 from the WEB classifier model when tested on WEB data). The † symbol marks a statistically significant change from the in-genre system performance (McNemar, $p < 0.05$). The best performances for the tested corpora are in bold-face.

In the cross-genre classification task, the Baseline+P+VNCLS set surfaces as the best performer, closely followed by the full feature set. As expected the classifiers show a general drop in

WSJ Model	Test on WEB				Test on BN			
	CMC		DISPLACE		CMC		DISPLACE	
Baseline	0.494	(-0.148)†	0.443	(-0.206)†	0.535	(-0.077)	0.553	(-0.073)
Baseline+P	0.655	(-0.136)	0.632	(-0.089)	0.787	(-0.013)	0.667	(-0.083)
Baseline+P+VNCLS	**0.663**	(-0.118)	**0.663**	(-0.109)	**0.817**	(-0.027)	0.716	(-0.074)
Semantic	0.578	(-0.123)†	0.487	(-0.247)†	0.757	(-0.027)	0.620	(-0.148)
Syntactic	0.481	(-0.163)†	0.443	(-0.210)†	0.552	(-0.068)	0.556	(-0.093)
Full Set	0.645	(-0.142)†	0.622	(-0.156)	0.811	(-0.024)	**0.737**	(-0.053)

Table 4.8: F-scores on CMC and DISPLACE labels across genre.

Statistically significant change from the in-genre system performance is marked with a †

performance across genre. The significance testing shows that none of the drops in performance seen for test on BN data are significant. This suggests that the statistics the classifiers were able to gather from WSJ, given the features, were more like the BN data than the WEB data. This certainly makes logical sense: Broadcast News data comes from a different genre/type of news than WSJ, but it is a new text, nevertheless.

What is also interesting here is that although the semantic feature set fares just as well as the full feature set on in-genre performance, the significant drop in performance suggests that semantic features alone may not be capable of scaling to other genres. Given the better performance displayed by the full feature set versus the semantic set, it might indicate that even if syntactic features are by themselves not as strong predictors for caused motion construction, they may contribute more when dealing with new genres. Different genres of text (e.g. newspaper versus web text) will be variable in literary and stylistic usages, which may contribute to differing preferences in syntactic choice. The improved performance suggests that there is a syntactic consistency across the corpora, allowing the syntactic features to have a positive impact on the classifier.

4.4.3 Removing Frequent NON-CMC Verbs

The following is a list of the top 25 highly frequent verbs that do not appear in a CMC usage. In fact, the semantics of most of the verbs are not compatible with the established definitions of CMCs. For example, verbs like *be*, *do*, or *have* cannot have caused motion usages, and verbs like *keep*, *leave*, or *prevent* are contrary to the semantics of CMCs as discussed in Chapter 3.

accuse	create	have	open	reach
base	do	hold	play	rise
be	fall	keep	prevent	see
build	file	leave	produce	use
charge	find	offer	quote	view

Table 4.9: Top 25 NON-CMC verbs

In this experiment, we remove the verbs in the above list from both the training and testing data. Because these verbs have no instances of CMCs or DISPLACEs, only the negative label was reduced in size. Effectively, the removal of the verbs increases the proportion of the positive labels in the corpora. The numbers are shown in Table 4.10.

	NON-CMC Instances			Positive Instances			
	Removed	Remaining		CMC		DISPLACE	
Corpus	Count	Count	% of Total	Before	After	Before	After
WSJ	2888	10071	77.7	14.8%	18.3%	8.86%	10.2%
WEB	216	1075	83.3	29.2%	33.1%	24.2%	25.6%
BN	521	2552	83.0	18.6%	21.6%	14.3%	15.7%

Table 4.10: Removed lemma count and its effect on NON-CMC label

Tables 4.11 and 4.12 show the precision and recall percentages and the f-score values when the instances of the most frequent NON-CMC verbs are removed from the training and testing data.

There is a general improvement in performance after the removal of the verbs from the data. The most marked improvement is in the WEB models (both CMC and DISPLACE) and the BN model's DISPLACE label classification. In particular the recall value shows improvement in these classifier models. As we have seen before, the Baseline+P+VNCLS set and the full feature set show the best predictions. There is no noticeable improvement in the WSJ classifiers except for a slight (statistically insignificant) increase in the baseline values.

	WSJ			WEB			BN		
	P	R	F	P	R	F	P	R	F
Baseline	63.32	40.67	0.4953	69.0	54.7	0.611	75.7	60.0	0.669
Baseline+P	71.71	71.56	0.7164	80.7	86.8	0.836	79.2	81.4	0.803
Baseline+P+VNCLS	70.97	73.33	0.7213	81.6	87.7	0.845	79.6	83.6	0.815
Semantic	69.37	68.44	0.6890	74.6	80.2	0.773	77.1	84.3	0.805
Full Set	73.88	76.67	0.7525	76.2	87.7	0.816	79.5	82.9	0.811

Table 4.11: System performance on CMC label classification with frequent NON-CMC verbs removed.

	WSJ			WEB			BN		
	P	R	F	P	R	F	P	R	F
Baseline	63.25	58.73	0.6091	70.3	63.4	0.667	71.1	57.8	0.638
Baseline+P	72.77	67.86	0.7023	74.1	76.8	0.754	79.4	75.5	0.774
Baseline+P+VNCLS	74.89	69.84	0.7228	76.1	81.7	0.788	79.8	81.4	0.806
Semantic	71.81	64.68	0.6806	73.8	75.6	0.747	74.5	77.5	0.760
Full Set	73.60	73.02	0.7331	76.7	84.1	0.802	81.4	81.4	0.814

Table 4.12: System performance on DISPLACE label classification with frequent NON-CMC verbs removed.

4.4.4 Random Downsampling of Negative Labels

As we have seen in Section 4.3.3, the CMC and the DISPLACE instances in WSJ are outnumbered by the negative, NON-CMC labels. The previous experiment on removing NON-CMC verbs effectively brought up the percentage of positive labels for the CMC and DISPLACE labels to 20% and 11%, respectively. However, label proportions of 20-80 or, worse, 11-89 are still highly unbalanced. Several studies have shown that in cases of training size imbalance, downsampling data can help with the performance of supervised classifiers (Weiss and Provost, 2001; Kubat and Matwin, 1997). Thus, for this experiment, we randomly downsample the negative labels in the WSJ training data to increase the percentage of positive labels For the sake of simplicity, we base the downsampling proportions on the CMC label: we cut the negative label so that the CMC label makes up 25% and 30% of the total data. The proportions of the DISPLACE labels are, therefore, 14.0% and 16.8%, respectively. Table 4.13 shows the percentages of the CMC labels and the DISPLACE labels after each cut. The column labeled as "Previous" shows the percentages at the beginning of this experiment.

	Previous		Downsample 1		Downsample 2	
CMC	1800	18.3%	1800	25.0%	1800	30.0%
NONCMC	8057	81.7%	5394	75.8%	4195	70.0%
DISPLACE	1009	10.2%	1009	14.0%	1009	16.8%
NONDISPLACE	8848	89.8%	6185	86.0%	4986	83.2%

Table 4.13: CMC and DISPLACE label distribution in the training set

Table 4.14 shows the performance of the WSJ models on the downsampled training set. The results indicate that the downsampling of the negative labels in the training data leads to increased performance. We have also tested the semantic feature set and the Baseline+P feature set as well. Their performances are approximately equal with no significant difference from the Baseline+P+VNCLS, so we do not include those numbers.

CMC Classification:						
	Downsample 1			Downsample 2		
	P	R	F	P	R	F
Baseline	63.20	55.33	0.5900	57.09	68.00	0.6207
Baseline+P+VNCLS	72.47	86.00	0.7866	70.72	89.11	0.7886
Full Set	75.76	88.89	0.8180	73.92	91.33	0.8171

DISPLACE Classification:						
	Downsample 1			Downsample 2		
	P	R	F	P	R	F
Baseline	65.17	69.05	0.6705	61.89	75.40	0.6798
Baseline+P+VNCLS	71.43	85.32	0.7776	69.59	88.10	0.7776
Full Set	76.29	88.10	0.8177	72.56	91.27	0.8084

Test data: WSJ test with frequently NON-CMC verbs removed (from Section 4.4.3)

Table 4.14: Classification performance with downsampled training data

Again, the downsampling was only applied to the training set, altering the distribution of labels only for the training data. The test set remains identical from its previous distribution in Section 4.4.3.

Interestingly, with the random downsampling of the training data, we see a boost in the full feature set's performance far more than the Baseline+P+VNCLS set's performance. In fact, in all cases we observed that the full features now show a significantly higher performance than the other features (McNemar, $p < 0.05$). The observed results for the two downsampled classifiers are not statistically distinct from one another.

To ascertain if this improved performance in downsampling carries over to the classification of out-of-genre texts, we test the WSJ model on the WEB data. As you will recall from the previous out-of-genre experiment, while the classification on the BN corpus showed no statistically significant degradation when using the WSJ model, tests on the WEB corpus showed mixed results. In this experiment we want to see if this significant degradation can be remedied by making the WSJ training data more balanced. Table 4.15 shows the f-score results of two downsampled classifiers' performance on the WEB test data and compares them to the results from Section 4.4.2.

	CMC		DISPLACE	
Full Set system from 4.4.2	0.645	(-0.142)†	0.622	(-0.156)
Downsample 1	0.711	(-0.076)	0.653	(-0.125)
Downsample 2	0.732	(-0.055)	0.688	(-0.091)

Test data: Original WEB test data with
original label distributions (from Section 4.4.2)

Table 4.15: Full feature set performance of downsampled WSJ model tested on WEB

The classifiers trained on downsampled data not only shows improvement in the f-score value, but now also show that the WSJ model does not cause a statistically significant drop in performance when compared to the in-genre results (i.e. WEB trained, WEB tested). As shown in the table, the Downsample 1 set shows slightly better performance than the Downsample 2, but the difference between the two values is statistically insignificant. In conclusion, the results show that the downsampling of the training data not only improves the in-genre performance of the classifiers, but it also improves on the prediction of the WEB data, which was previously not as robust.

Table 4.16 shows a breakdown of the feature performance. In the classification of CMC and DISPLACE, the best performance is observed for the full set of features. The Baseline+P+VNCLS set does perform just as well, but with a caveat.

	CMC		DISPLACE	
	Down 1	Down 2	Down 1	Down 2
Baseline+P+VNCLS	0.6562	0.7219	0.7006	n/a
Semantic	0.6298	0.7156	0.5915	0.6329
Full Set	0.711	0.7317	0.653	0.6875

Table 4.16: Feature performance of downsampled WSJ model tested on BN

For the Downsample 2 set for the DISPLACE classification, we found that the Baseline+P+VNCLS system gives a fully zero positive (DISPLACE) classification. In other words, the system predicts a negative label for every instance in the test data. This is a perplexing finding given that this feature set has been exhibiting good results across the previous experiments. Further investigation is required to ascertain the cause behind the Baseline+P+VNCLS set's apparent failure to classify; but the finding suggests that there simply wasn't enough information in the downsampled training data to be able to make a proper decision at the testing stage.

In the end, the full feature set gives the highest performance when trained on the Downsample 2 data. The performance of semantic features are good, but the full feature set, which includes the syntactic features as well as the semantic and the lexical ones, scales better than the other features across genre.

4.5 Final Considerations and Future Work

In this chapter, we have presented our work on the automatic classification of CMCs in corpus data using the annotated data produced in Chapter 3. Our studies have shown that we can achieve the identification of caused motion instances at a higher rate than the inter-annotator agreement scores, the best performance that can be realistically expected. We have also shown that semantic information is highly indicative of the caused motion phenomenon, confirming our general intuition that the caused motion construction is a semantic phenomenon. However, we also find that syntax provides scalable features that generalize well across different types of text, producing better results in cross-genre experiments. We have also shown that the downsampling of the negative label has a positive impact on the classification of the labels.

There are several directions we can go from here. First, this work has made use of various gold annotations for the purposes of feature extraction. The most obvious next step in this investigation will involve experimentation with automatically obtained features. Secondly, in this work we have focused on the classification of CMCs as a single phenomenon and the classification of a single subtype, the *cause-displacement-cx* CMC type. Future work will look at expanding the classification to the other subtypes of the CMCs, as described in Chapter 3.

Additionally, we hope to examine the impact of further features. As the experiments have shown, the lexical and semantic features (lemma, preposition, VerbNet classes) surface as strong predictors of CMCs. It follows from this, that we should expand the feature search to other semantic information. One particular set of features that might be interesting, would be based on FrameNet frames. Since FrameNet's frames represent different conceptual semantic domains, features from FrameNet may be instrumental at capturing and highlighting the semantics of CMCs that are spread across VerbNet classes of differing semantic types. Moreover, it would also be interesting to expand on the lexical features: lexical features can be extended to not just the verb of the sentence but also to the noun phrases. Further investigation into using resources like WordNet (Miller, 1995; Fellbaum et al., 1998) might be needed to remedy sparse data issues that lexical features based on words from the noun phrases might create.

Chapter 5

Representing Caused Motion Constructions in VerbNet

5.1 Introduction

VerbNet currently supplies a large array of semantic predicates for the representation of meaning, making it a natural choice for the representation of CMCs. What's more, verbs of directed motion (e.g. *Put-9.1, Run-51.3.2*), verbs of transfer of possession or information (e.g. *Give-13.1*, *Transfer_mesg-37.1.1*) and verbs of changes of transformation (e.g. *Turn-26.6.1*), whose semantics directly lend to the description of caused motion events, are currently represented in VerbNet. As we seek to leverage VerbNet's semantic predicate representations for CMC representation, we are faced with the issues of generalizability: VerbNet currently has no single set of predicates that can consistently represent the different meaning components of CMCs.

Because VerbNet's semantics were conceived on a class by class basis, the predicates are inevitably honed to the particular semantics of the class. So, while useful in describing the class where they are used, the predicates do not generalize very well across verb classes of similar meaning, let alone classes related in meaning through metaphorical extensions. There are some consistent generalizations that can be made from the semantic predicates: for example, each of the classes that fall under the series 13, the *Verbs of Change of Possession*, are described by a similar set of semantic roles and predicates. However, these generalizations are limited to the classes Levin (1993) identified as being related.

A similar issue is noted in a study by Zaenen et al. (2008), where the authors seek to extract such inferences relating to the pre- and post-states of change of location from VerbNet's predicate information. Their general conclusion is that even though many of the VerbNet classes contain information leading to a path interpretation, VerbNet is limited in that not all of the VerbNet classes of change of location code for path information, and the path information is inconsistent in its representation across classes. That is, when comparing classes that do contain the desired path information for inference making, the information is not necessarily consistent across all of VerbNet, making the inferencing difficult. What is clear is that whether it is for text-specific semantic representations or for inferencing, it is important that VerbNet's representation be explicit, unambiguous and consistent. A proper representation of a CMC will require a semantically informed representation of the path of motion. For the automatic identification and classification of constructions, we need a unified representation that can be associated with the caused motion usages.

Consider the following examples:

(50) a. The horse jumped *in the paddock*.

b. The horse jumped *into the neighbor's garden*.

In both sentences, the italicized prepositional phrases refer to a location of a sort. The unmistakable difference here is that the phrase in the first sentence refers to a location that remains static over the course of the event, while the phrase in the second sentence refers to the location where the horse has arrived by the end of the event. Additionally, the implication present in the second sentence, though lacking in the first, is of the change of location: at the beginning of the event the horse is located in an unknown location and by the end of the event it has moved to the neighbor's garden. VerbNet's representations should not confuse the meaning sentences like these have if the sentences are to be represented in a way that expedites the proper inferences.

The purpose of this chapter is to present our efforts at establishing and implementing systematic predicates for the representation of paths of motion. Since the focus of this thesis is the caused motion construction, CMCs will receive close attention. However, the predicates will have to be formulated so that in the future they can be extended for other constructions. If we focus only on the CMC usages, we may produce a useful representation for CMCs that is too narrow to be extended to other usages of paths of motion. Thus, for the purposes of this chapter, we entertain slightly broader notions of semantics of paths of motion and changes of location (physical or abstract) that should be applicable related usages.

In Section 5.2, we evaluate the current status of VerbNet's representation of the different semantic components of CMCs to identify the existing inconsistencies and limitations. In Section 5.3 we present a new and semantically explicit representation that can better generalize across multiple classes in VerbNet. In Section 5.4, we discuss methods for tackling the implementation of the new predicates.

5.1.1 Approach

For the representation of caused motion constructions, it is necessary that the semantic predicates consistently and explicitly capture the semantics of causation, the motion event, and the path of motion. Causation is fairly straightforward – we need to be able to specify the argument that is responsible for the motion in the sentence. For the motion event, we need to be able to capture the semantics of the undergoer's movement across the path, whether that motion is physical or metaphorical (e.g. transformation or transfer). For the path of motion, we want to be able to account for all the elements expressed in a path of motion: *source* – where the moving entity is at the beginning of the event, *path* – the location over which the moving entity travels, and *goal* – where the moving entity is at the end of the event.

Furthermore, we aim for a representation that can be extended to other path designations outside of CMCs. Consider the following two sentences:

(51) a. The dog jumped *the fence*.
The trainer jumped his dog *over the fence*.

b. Why would you want to fry my hair some more after you already burnt it *crisp*?
Jennifer baked the potatoes *to a crisp*.

The phrases in italics, despite their differing syntactic functions, semantically play the same role in each of the sentence pairs. In example 51a, both phrases refer to the path over which *the dog* travels. In example 51b, we have an adjective expressing the resultant states of the event. A successful semantic representation of the path of motion should treat the two phrases in a similar manner, such that the same path inference available to one is available to the other.

Finally, what we seek is an explicit representation that supports proper inferencing and reasoning based on what we know about motion and change of location (Zaenen et al., 2008, 2010). For example, given sentence 50b we want to be able to conclude that the horse is outside *the neighbor's garden* at the beginning of the event and is located inside *the neighbor's garden* at the end of the event. Such a representation should distinguish between a static location and a path of motion. That is, the representations must distinguish arguments denoting the path of motion (see ex. 50b) from constituents that encode locative information about where an event occurs (see ex. 50a).

We adopt the following guiding principles throughout this study. First, we stay faithful to the semantics of CMCs as we have defined them in Chapter 3. We want a representation that will unify the semantics of CMCs across their semantic types, and still allow each type its distinguishing feature. Secondly, we do our best to alter VerbNet as little as possible. When possible, we pull what we can from existing predicates, and suggest new ones only in cases where the current predicates cannot handle the necessary semantics. Finally, we stay mindful of other constructions that use path of motion or signal changes of location or state that do not fall under the CMC categories. Examples of these are:

- Directed motion (e.g. *Eric went from Paris to Lyon*, *The dog jumped over the fence.*)
- Fictive motion (e.g. *The meeting went from 3pm to 5pm*, *The road meanders through the forest.*)
- Ditransitives (e.g. *Jennifer gave Mary a book.*)
- Adjectival resultatives (e.g. *The water froze solid*, *A harsh winter froze the lake solid.*)

- Path verbs (e.g. *The plane arrived*, *We landed in Russia.*)

5.1.2 Terminology and Conventions

To avoid any potential confusion in terminology, here are the definitions of terms as they will be used in this chapter.

- Unless otherwise specified, we will generally refer to the event or situation expressed by the change of location/state, transformation, or transfer as the **motion event**.

- We will generally use the term **undergoer** to refer to the object or entity that undergoes the motion event. UNDERGOER, THEME and PATIENT are the usual thematic roles in VerbNet that describe this entity in relation to the motion event.

- The term **source** refers to the location (physical or otherwise) that describes the region where the undergoer is located at the start of the event. The thematic roles SOURCE and INITIAL_LOCATION are often used to describe the source location in VerbNet.

- The term **goal** refers to the location (physical or otherwise) that describes the region where the undergoer is located at the end of the event. Thematic roles GOAL, DESTINATION, RESULT are often used to describe the goal location in VerbNet.

- **Path** refers to the 'region' the motion traverses. It is the connecting region between the source and goal locations. We will sometimes use the term **trajectory** to refer to the path when the motion event expresses a change of location or state. In VerbNet, thematic roles PATH and TRAJECTORY are often used to describe the path.

- **Path of Motion** refers to the complete region the undergoer traverses through the motion event. Source, path, and goal are points along the path of motion.

Additionally, here is a list of conventions used in this chapter for VerbNet components. This is how they will be presented in the paragraph text.

- VerbNet class will be *italicized*: *Calibratable_cos-45.6*.
- VerbNet roles will be in SMALL CAPS: TRAJECTORY, UNDERGOER
- VerbNet predicates will be `typewriter` font: `cause()`, `Prep()`.
- VerbNet variable for event will always be presented as *E*.

5.2 Current VerbNet Semantic Representations

In earlier work (Hwang et al., 2013), we conducted a study to ascertain the current status of the VerbNet representation of path of motion for physical changes of location. This section extends on this study by including an exploration of VerbNet's current representation of the semantics of changes of location, changes of state, transfer of possession, and transfer of information.

5.2.1 Causation

In VerbNet, the semantics of causation is represented using a `cause()` predicate as seen in the example below. The predicate specifies two arguments: the causal argument, AGENT, and the argument *E*, the event variable for which the predicate holds true. The use of the predicate `cause()` for the semantics of causation is, for the most part, highly consistent across VerbNet.[1]

(52)

Roll-51.3.1	
Roles	AGENT, THEME, LOCATION, RESULT
Frame	NP V NP PP
Ex	*Bill rolled the ball down the hill.*
Syn	AGENT V THEME LOCATION
Sem	cause(Agent,E) motion(during(E), Theme) Prep(E, Theme, Location)

[1] There are still a small number of cases where the `cause()` predicate is not used although the semantics suggest it. For example, this can be seen in the *Hit-18.1* class where the sentence *The stick hit the door open* does not have causation represented (i.e., The stick caused the door to open). There are two possible reasons the causative statement is missing. The first could be simply the result of a human error rather than a general tendency or inconsistency in representation. The second reason may not be unlike the example for *Roll-51.3.1* class, here we have an INSTRUMENT as the causer. It could be that VerbNet needs a better analysis for cases in which the causer is not the AGENT

5.2.2 Changes of Location

5.2.2.1 Motion

The semantics of motion is represented, in general, with the `motion()` predicate, which takes two arguments: the event *E*, specifically a subpart of the event `during(`*E*`)` and the THEME of the event. Much like the `cause()` predicate, the `motion()` predicate is used fairly consistently across VerbNet to represent motion. While the predicate represents motion, it does not necessarily indicate change of location. This is exemplified for the two frames in the *Run-51.3.2* class below, where the former frame simply denotes the existence of motion, while the second is a case of caused motion with an associated path of motion.

(53)

Run-51.3.2		
Roles		AGENT, THEME, LOCATION, RESULT
Frame		NP V
	Ex	*The horse jumped.*
	Syn	THEME V
	Sem	motion(during(E), Theme)
Frame		NP V PP.LOCATION
	Ex	*The horse jumped over the fence.*
	Syn	THEME V LOCATION
	Sem	motion(during(E), Theme) Prep(E, Theme, Location)

While the use of the `motion()` predicate is mostly consistent, there are still cases where the motion event is missing the predicate and the representation is left up to the location predicate.

(54)

Banish-10.2		
Frame		NP V NP PP.DESTINATION
	Ex	*The king deported the general to the isle.*
	Syn	AGENT V THEME DESTINATION
	Sem	cause(Agent, E) location(start(E), Theme, ?Source) location(end(E), Theme, Destination)

The causation element of the semantics is represented. The `location()` predicates indicate that the THEME entity was at the source location (question marked because it is presupposed but uninstantiated) at the start of the event and at the DESTINATION at the end of the event. In such cases the motion event must be introduced to the frame.

5.2.2.2 Path of Motion

When it comes to the representation of path associated with changes of location, there are no set standards for representation in VerbNet. However, we observe that there are a number of general tendencies in how VerbNet currently handles the path of motion. We will focus on specific predicates (i.e., `location()`, `Prep()`, `via()` and `direction()`) and what we perceive as inconsistencies in their current usage.

Predicate `location()`: In general, for classes in which the realized piece of the path refers to either endpoint of the change of location, VerbNet makes use of the `location()` predicate in conjunction with the `motion()` predicate to represent path information. Consider the *Slide-11.2* class:

(55)

Slide-11.2		
Roles		AGENT, THEME, SOURCE, DESTINATION
Frame		NP V PP.SOURCE
	Ex	*The book slid from the table onto the floor.*
	Syn	THEME V SOURCE DESTINATION
	Sem	motion(during(E), Theme) location(start(E), Theme, Source) location(end(E), Theme, Destination)

The specification of path in terms of `location()` allows for VerbNet to specify that the THEME is at the specified location at the start of the event (i.e. *the table*), is set in motion during (i.e. *during(E)*) the event, and is located at the DESTINATION (i.e. *the floor*) by the end of the event. However, it is not always the case that only change of location type events will have `location()`

predicates marked as the start or end of an event. In the example 56, the predicate `location()` is used as a means of indicating that the decoration (i.e. THEME; *the name*) is made on the ring (i.e. DESTINATION) and does not indicate that there was a change of location that occurred over the course of the event.

(56)

Illustrate-25.3	
Ex	*The jeweler decorated the ring with the name.*
Syn	AGENT V DESTINATION THEME
Sem	cause(E, Agent) created_image(result(E), Theme) location(end(E), Theme, Destination))

Even if we were to say that change of location happens only for classes that have the `motion()` predicate along with a `location()` predicate, it is not a reliable measure as there are classes such as *Banish-10.2* (e.g. *The king deported the general to the isle.*) that fail to specify the predicate `motion()` in their frames.

Predicate `Prep()`: If the sentence expresses the trajectory of the motion, rather than the starting or ending location, VerbNet currently tends to favor the use of the `Prep()` predicate accompanied by the `motion()` predicate as seen in example 53. The most obvious problem is that unlike the `location()` predicate that represents the location of an event, the `Prep()` predicate, on its own, does not represent a specific meaning. That is, given the `Prep()` predicate alone, it is difficult to determine what relationship THEME and LOCATION have with each other and how they relate to event *E*. Even if we set this problem aside, this representation of path is somewhat problematic since the same predicate `Prep()` is used to represent a static location of an object such as *in the corner* in the sentence *"The statue stood in the corner"* (*Spatial_configuration-47.6*), which is unlike the trajectory `Prep()` predicate that was used in example 53.

The way to distinguish such instances is seemingly to look for the `motion()` predicate along with the `Prep()`. However, such a heuristic is not fail-safe either. Consider example 57 in a comparison with example 53, in which the representation (e.g. the use of the `Prep()` in the presence of `motion()`) seems to look much like the representation we see in 53. However, unlike 53, the `Prep()` in the *Swarm-47.5.1* class refers to a static location where the motion occurs.

(57)

Swarm-47.5.1	
Ex	*Bees are swarming in the garden.*
Syn	THEME V LOCATION
Sem	exist(during(E), Theme) motion(during(E), Theme) Prep(during(E), Theme, Location)

Finally, unlike the `location()` predicate, which we have seen in the previous section and the `via()` and `direction()` predicates, which we will see in the following section, `Prep()` is a variable predicate that gets instantiated by a specific preposition. Our intent is to replace this with predicates such as `via()` or `direction()` that have argument slots for prepositions, thus avoiding second order logic.

Predicates `via()` and `direction()`:

Less frequently used to represent the path of motion are the `direction()` and the `via()` predicates. Like `location()` and `Prep()`, they are used in conjunction with the `motion()` predicate. The following is an example of the use of the `via()` predicate:

(58)

Run-51.3.2	
Roles	AGENT, THEME, LOCATION
Frame	NP V NP.LOCATION
Ex	*The horse jumped the stream.*
Syn	THEME V LOCATION
Sem	motion(during(E), Theme) via(during(E), Theme, Location)

Although the LOCATION role examples 53 and 58 specify the trajectory along which the THEME moves, the creators of this class chose to treat 58 as a case of `via()` instead of `Prep()` as seen in 53. The motivation behind this treatment was likely a syntactic one – the LOCATION in

53 is found in a prepositional phrase, while the same role appears as the noun phrase in 58. The predicate `via()` appears in two other classes, *Vehicle-51.4.1* (e.g., *She skated the canals.*) and *Nonvehicle-51.4.2* (e.g., *She rowed the canals.*); both of which are also classes where the syntactic frame in which the path of motion appears is taken into account in the creation of the semantic frames.

The `direction()` predicate is generally found to represent the trajectory in a given path of motion as seen in example 59. This representation is comparable to the one for 53.

(59)

Escape-51.1	
Roles	THEME, LOCATION
Frame	NP V PP.OBLIQUE
Ex	*The prisoners advanced across the field.*
Syn	THEME V OBLIQUE
Sem	motion(during(E),Theme) direction(during(E),Prep_Dir[2] ,Theme,Oblique)

One advantage this representation has over the use of `Prep()` is that while with `Prep()` the actual preposition is instantiated from the context, the `direction()` predicate specifies that the THEME is directed towards the LOCATION during the event. Outside of this factor, it is not clear the VerbNet needs to distinguishes the path information in 59 from that in 53.

5.2.3 Changes of State

Now we turn to the representation of the changes of state. The `motion()` predicate represents the motion event for changes of location and they apply fairly consistently across all of the relevant VerbNet classes. Change of state, however, does not have a semantic representation that is equivalent to the motion event for caused motion constructions that can be uniformly applied to all instances of change of state. This is as expected: in stative changes the transition of the undergoer from one state to another is conceptualized as the motion. Thus, the semantics of change of state are represented in VerbNet through class-specific predicates.

Consider the two classes shown as examples 60 and 61. In each, the specific semantic predicates that represent the change of state are in boldface. For the verb *baked*, the semantic representation for the change of state is represented through `cooked()`, and for the verb *break*, the predicates are specific about the degradation the glass goes through and the change to its physical form.

(60)

Cooking-45.3	
Members (e.g.)	bake, boil, fry, roast
Frame	NP V NP PP
Ex	*Jennifer baked the potatoes to a crisp.*
Syn	AGENT V PATIENT RESULT
Sem	cause(Agent, E) **apply_heat(during(E), ?Instrument, Patient)** cooked(result(E), Form, Patient) Pred(result(E), Patient)

(61)

Break-45.1	
Members (e.g.)	break, crack, rip
Frame	NP V NP PP
Ex	*Tony broke the glass to pieces.*
Syn	AGENT V PATIENT RESULT
Sem	cause(Agent, E) contact(during(E), ?Instrument, Patient) **degradation_material_integrity(result(E), Patient)** **physical_form(result(E), Form, Patient)** Pred(result(E), Patient)

The predicate `Pred()` is consistently used in VerbNet to represent the end state as seen in both examples above. The `Pred()` predicate is found in 20 VerbNet classes and is almost[3] exclusively used for resultative phrases.

There is, however, one glaring weakness to this representation. Note that the first argument to `Pred()` uses `result()` rather than `end()` to predicate the event *E*. The specification of the

[3] The *Stimulus_subject-30.4* class is the only exception to this. The sentence *That pea soup tasted delicious to me* is described by the semantic predicate *perceive(during(E), Experiencer, Stimulus)* ***Pred(during(E), Stimulus)*** *in_reaction_to(E, Stimulus).*

`result()` predicate for *E* suggests that the `Pred()` predicate is relevant to the resulting state of the event, but, unfortunately, all this achieves is to specify that there is some resulting state involved in the event without specifying what that resulting state is. In other words, the predicate only includes as variables the Event *E* and the `Patient` role but leaves out the RESULT role that would instantiate the exact nature of the resultant state.

Aside from the `Pred()` plus `result()` predicate combination, there is one other predicate that seems to be responsible for the semantics of changes of state. Consider the *Turn-26.6.1* class:

(62)

Turn-26.6.1	
Members (e.g.)	convert, turn, transform
Frame	NP V PP
Ex	*He turned into a frog.*
Syn	PATIENT V RESULT
Sem	not(state(start(E), Result, Patient))
	state(result(E), Result, Patient)

The *Turn-26.6.1* class uses the *state()* predicate to specify the state at the `start()` of the event and the state at the `result()` (resulting state) of the event. Note that just like the example from *Banish-10.2* in 54 the semantics of transformation are not at all represented. It is left up to the `state()` and the negation of the earlier state to imply that there was a state change. Classes like *Turn-26.6.1* will need an explicit predicate to handle the semantics of transformation.

5.2.4 Transfer of Possession

The current representation for changes of possession has probably the most consistent and explicit predicates for representation out of those we have looked at so far. Changes of possession are represented through the `transfer()` predicate in conjunction with the `has_possession()` predicate. The following representation is highly consistent across all changes of possession classes (i.e. the 13.1-13.7 classes) and the removal of possession class *steal-10.5*.

(63)

Give-13.1	
Members (e.g.)	lend, give, pass, sell
Roles	AGENT, THEME, RECIPIENT
Frame	NP V NP PP.recipient
Ex	*They lent a bicycle to me.*
Syn	AGENT V THEME RECIPIENT
Sem	cause(Agent,E) transfer(during(E), Theme) has_possession(start(E), Agent, Theme) has_possession(end(E), Recipient, Theme)

The semantic predicates show that the AGENT causes the event *E* during which there is a transfer of the THEME. At the start of the event the AGENT has possession of the THEME, and at the end of the event, the THEME is in the possession of the RECIPIENT.

The transfer of possession classes whose verbs take the source rather than the goal information in the oblique position, specify the SOURCE and other predicates that parallel the *give-13.1* class shown above. In *Obtain-13.5.2*, the `has_possession()` predicates specify that the possession of the THEME is at the SOURCE location at the start of the event and at the AGENT location at the end of the event.

(64)

Obtain-13.5.2	
Members (e.g.)	accept, obtain, receive
Roles	AGENT, THEME, SOURCE
Frame	NP V NP PP.goal
Ex	*Carmen obtained a spare part from Diana.*
Syn	AGENT V THEME SOURCE
Sem	cause(Agent,E) transfer(during(E), Theme) has_possession(start(E), Source, Theme) has_possession(end(E), Agent, Theme)

The representation for changes of possession generally reflects the way the change of location verbs are represented: the `transfer()` predicate is, here, counterpart to the `motion()` predicate. It specifies that the possession is being transferred during the course of the event, in the same way that the `motion()` predicate specifies the entity in motion during the event. The `has_possession()` predicate captures the path semantics of possession transfer. This is an improvement over the motion verbs where the path representation can be represented through any of the three potential predicates (i.e. `location()`, `via()`, and `direction()`).

5.2.5 Transfer of Information

The predicates for the transfer of information, use one predicate, `transfer_info()`, to represent both the transfer and the path information expressed in the sentence. The following is an example from the *Transfer_mesg-37.1.1* class.

(65)

Transfer_mesg-37.1.1	
Members (e.g.)	explain, narrate, recite
Roles	AGENT, TOPIC, RECIPIENT
Frame	NP V NP PP.recipient
Ex	*I explained the matter to them.*
Syn	AGENT V TOPIC RECIPIENT
Sem	cause(Agent,E) transfer_info(during(E), Agent, Recipient, Topic)

The semantic predicates show that the AGENT causes the event *E* during which there is a transfer of information. The `transfer_info()` predicate indicates that the TOPIC is transferred from AGENT to RECIPIENT. However, unlike the semantics of the change of possession, the AGENT does not relinquish possession of the information that is transferred. By the end of the event, the transferred information is in the possession of both the AGENT and the RECIPIENT.

5.2.6 Changes in Scale

The closest we come to a change in scale type representation is found in one class: *calibratable_cos-45.6*. Here is an example.

(66)

calibratable_cos-45.6	
Members (e.g.)	balloon, decrease, increase, jump, lower
Roles	Patient, Attribute, Extent
Frame	NP.attribute V PP.recipient
Ex	*The price of oil increased by ten percent.*
Syn	Attribute (of) Patient V Extent
Sem	change_value(during(E), Direction, Attribute, Patient) amount_changed(during(E), Attribute, Patient, Extent)

The calibratable change of state class was originally designed by Levin (1993) to include verbs that exhibit the possessor-attribute alternation (e.g. *The price of oil increased.* vs *The oil increased in price*). Consequently, the verbs in the class may be relevant, but the class lists only those frames that exhibit one of the two possessor-attribute alternations, without reference to causative usages.

In accord with Levin's goals for this class, the `change_value()` predicate specifies that the Patient's Attribute changes in value. And the value moves in the direction specified by the role Direction, which presumably would be made available through an external method based on the semantics of the verb. Interestingly, the creators of VerbNet included the Extent role along with the frame that requires it in example 66. This role shows up in the `amount_changed()` predicate that specifies the amount by which the Patient's Attribute was changed during the course of the event. In effect, `change_value()` represents the Theme's move from one value to another, and `expresses_changed()` the path of change.

There exists one other class where verbs describing changes along a linear scale could potentially fit: the *other_cos-45.7* class. This class was created as a part of the VerbNet extension project (Kipper et al., 2006). Unfortunately, this class is currently a miscellaneous category that includes approximately 340 verbs denoting a telic event (e.g. *americanize*, *gladden*, *metabolize* and *yellow*). While it includes many verbs of scalar change (e.g. *accelerate*, *increase*,or *slow*), the class will, in the future, require further subdividing to make it more semantically coherent.

5.2.7 Summary

The table below displays a summary of the current predicates and roles that are relevant in representing the entire range of CMCs in VerbNet. The motion of the event (physical or metaphorical) and the path of motion describing a CMC are captured through the use of predicates. The entity in motion and the locations (physical or metaphorical) along the path are captured through the use of the semantic roles.

	Event	Entity in Motion	Path of Motion	Locations Along Path
Change of Location	`motion()`	THEME	`location()` `via()` `Prep()` `direction()`	INIT_LOCATION DESTINATION TRAJECTORY
Change in Scale	`change_value()`	PATIENT	`amount_changed()`	EXTENT
Change of State	specific to class	PATIENT	`Pred()`	RESULT
Transfer of Poss.	`transfer()`	THEME	`has_possession()`	SOURCE RECIPIENT
Transfer of Info.	`transfer_info()`	TOPIC	`transfer_info()`	RECIPIENT

Table 5.1: Current predicates and roles relevant for CMCs in VerbNet classes

For the most part, the semantic roles as they are currently used in VerbNet fit for the role they play in describing caused motion semantics. All that the semantic roles would require is a way to capture the relationship between those that appear as entities of motion versus those that represent physical or abstract locations along the path. This is, to an extent, already available through the semantic role hierarchy (Bonial et al., 2011a) developed in recent years. We now leverage the hierarchy for the representation of CMCs. For the semantic predicate representation, we revamp some of the current predicates and introduce new ones to systematize how path is being represented. In the next section, we discuss changes to the current predicates and roles, and present new ones that will improve the consistency and informativeness of the CMC representations.

5.3 New VerbNet Semantic Representation

Table 5.2 shows a summary of the changes to the predicates and roles. The changes are marked in blue. The bold-faced predicates and roles indicate the new additions to VerbNet. The

	Event	**Entity in Motion**	**Path of Motion**	**Locations Along Path**
Change of Location	`motion()`	THEME	**`path_rel()`**	INIT_LOCATION DESTINATION TRAJECTORY
Change in Scale	`change_value()`	PATIENT		EXTENT
Change of State	specific to class **`transform()`**	PATIENT		RESULT
Transfer of Poss.	`transfer()`	THEME		SOURCE RECIPIENT
Transfer of Info.		TOPIC		RECIPIENT

Table 5.2: New predicates and roles relevant for CMCs in VerbNet classes

most significant change to the current predicate semantics is the introduction of the `path_rel()` predicate to handle the representation for path of motion of both the physical and the metaphorical usages. We also introduce the `transform()` predicate for those classes that express changes of state, but entirely lack a motion event representation as seen in Section 5.2.3. Additionally, `transfer_info()` is replaced by the `transfer()` predicate for the representation of transfer of information. Finally, we make a revision to the `change_value()` predicate to parallel the `motion()` and `transfer()` predicates. Changes to the existing semantic roles are not as extensive. We introduce one new role, TRAJECTORY, highlighted in the table above, and consider the addition of a second role, as will be discussed in the following subsections.

5.3.1 Semantic Predicates

Table 5.3 is another perspective on the information in Table 5.2, displaying the information relevant to semantic predicates, in particular. This table includes a new column: the second column displays the CMC types associated with the categories specified in the leftmost column.

	CMC Type	Event	Path of Motion
Change of Location	cause-displace-cx	`motion()`	**`path_rel()`**
Change in Scale	cause-chng-scale-cx	`change_value()`	
Change of State	cause-transform-cx	specific to class **`transform()`**	
Transfer of Poss.	cause-chng-poss-cx	`transfer()`	
Transfer of Info.			

Table 5.3: New predicates and roles relevant for CMCs in VerbNet classes

For a semantically informed VerbNet representation of path of motion, we introduce the new `path_rel()` predicate accompanied by the consistent inclusion of a motion event predicate (column 3) that describes the physical or abstract motion of the undergoer. The name of the `path_rel()` predicate reflects the fact that the information contained in the predicate is expressed in **rel**ation to the **path**. The predicate is defined in Table 5.4:

`path_rel`(<Event>,<Undergoer>, <Location>, <LocType>, <Prep>)	
where:	
<Event>:	Event *E* with the time function predicates (i.e. `start()` `during()` and `end()`) specifying at which point in time in the event the path of motion is relevant
<Undergoer>:	The semantic role of the entity in motion.
<Location>:	The semantic role for the physical or metaphorical/abstract location where the Undergoer is at the time specified by the Event
<LocType>:	A constant that specifies the path's semantic type
<Prep>:	The preposition that instantiates the path predicate, if available

Table 5.4: Definition of the `path_rel()` predicate

By allowing the Event argument to specify that the predicate is applicable for the different points in time along the progression of the event, we can use the `path_rel()` predicate to represent the source, path, and goal elements of the path. For example, the `path_rel()` predicate can be demonstrated in the following. The first set in the example represents the uninstantiated predicates. The second set shows the predicates instantiated with the words in the sentence.

(67)

Cynthia rolled the ball out of the cup, across the floor, and into a hole.
cause(Agent,E) motion(during(E), Theme) path_rel(start(E), Theme, Initial_Location, *ch_of_loc*, prep) path_rel(during(E), Theme, Trajectory, *ch_of_loc*, prep) path_rel(end(E), Theme, Destination, *ch_of_loc*, prep)
cause("Cynthia",E) motion(during(E), "the ball") path_rel(start(E), "the ball", "the cup", *ch_of_loc*, "out of") path_rel(during(E), "the ball", "floor", *ch_of_loc*, "across") path_rel(end(E), "the ball", "a hole", *ch_of_loc*, "into")

The `path_rel()` predicate, thus, includes the event *E*, the same *E* found in the `motion()` predicate, which specifies the THEME's motion. Additionally, the first `path_rel()` predicate indicates the position, INITIAL_LOCATION, of the THEME at the start of the event, the second `path_rel()` predicate indicates location over which the THEME travelled during the event, and finally the third `path_rel()` predicate indicates the location where the THEME is by the end of the event. The `path_rel()` predicate, thus, allows for the consistent representation of all possible combinations of the expression of path, which so far has not been possible in VerbNet.

The `path_rel()` predicate also uses the choice of the LocType constant to distinguish between the different semantic types of a CMC. The following is our current list of constants.

- **ch_of_loc**: specifies a path that participates in a change in location
- **ch_of_state**: specifies the path of motion from one state to another
- **ch_of_scale**: specifies a change in value on a linear scale
- **ch_of_poss**: specifies a transfer in possession
- **ch_of_info**: specifies a transfer in information

Compare the above example of physical change of location, a case of *cause-displace-cx*, with the following example, which is an instance of *cause-change-scale-cx*, taking the constant *ch_of_scale*. Note that the `change_value()` predicate takes the place of the `motion()` predicate. We will examine this predicate in further detail in Section 5.3.1.2.

(68)

The heavy downpour raised the level of the lake to 1000ft.
cause(Cause,E) change_value(during(E), Patient) path_rel(end(E), Patient, Destination, *ch_of_scale*, prep)
cause("The heavy downpour",E) change_value(during(E), "the level of the lake") path_rel(end(E), "the level of the lake", "1000ft", *ch_of_scale*, "to")

Note in this example that not all roles in the `path_rel()` predicate have to be instantiated in the sentence. Uninstantiated instances are allowed to be unspecified in VerbNet's semantic representation. These instantiated roles or definite null complements are syntactically optional but, nonetheless, semantically essential (Fillmore, 1986; Palmer, 1985). For example, even though an initial value for the level of the lake is not expressed, the meaning of these roles is still implicit in the sentence. The lake must have been at a value that is not 1000ft and, therefore, had a different initial value at the beginning of the event. The advantage of the representation of path of motion through the predicate `path_rel` is that we get the desired uninstantiated inferences relating to path for free. In other words, even though VerbNet does not make the implicit arguments explicit in the frames as in the example above, the inferences can be made explicit for the purposes of sentence representation. Recognizing the existence of uninstantiated roles allows us to make semantically implicit information explicit in the sentence representation (Gerber and Chai, 2013, 2010; Palmer et al., 1986).

On a related note, with the exception of the transfer of information CMC type, there is an implicit understanding in these CMC events that the THEME can only exist in a single place for any given point in time of the event. In other words, in example 67, the inference we want to retrieve is that the ball is in the cup before the event begins and is no longer there at the end of the event. In the same way, the ball is not in the hole before the event, but is there at the end of the event. The advantage of a consistent and explicit representation of path of motion is that we get these inferences for free.

For the most part VerbNet handles this inference, albeit inconsistently, by negating the locative predicates. Take for instance the *banish-10.2* class:[4]

(69) a. *"The king banished the general from the army."*

cause(Agent, E)

location(start(E), Theme, Source)

not(location(end(E), Theme, Source))

b. *"The king deported the general to the isle."*

cause(Agent, E)

location(start(E), Theme, ?Source)

location(end(E), Theme, Destination)

In 69a, only the SOURCE location is marked as being true at the beginning of the event but not at the end of the event. In 69b, both SOURCE location and DESTINATION are represented, but neither of them are marked as not being true at the other points in time. For a complete representation of all possible inferences we would have had to specify:

(70) location(start(E), Theme, Source)

location(end(E), Theme, Destination)

not(location(start(E), Theme, Destination))

not(location(end(E), Theme, Source))

In the new representation, it is assumed that for any given path of a CMC, the semantics of all points in the path are available even if they are not syntactically realized. Thus, we can make the generalization that for every instance of a `path_rel()` predicate the entity in motion will be (1) located at the source at the beginning (and *only* at the beginning) of the event, (2) located over the specified trajectory during (and *only* during) the event, and (3) located at the destination at

[4] It is not by omission that the `motion()` predicate is not present in this example. This class does not currently include it, although it should as a class that expresses change of location.

the end (and *only* at the end) of the event. The only exception to this generalization is the *cause-possession-cx* type of CMC for the transfer of information. For information transfer, only the (2) and (3) of the generalization will hold.

Finally, we noted in the introduction that the path of motion predicate should be available not only for the CMC instances but also for instances of other syntactic forms that express directed motion. Here is the revised view of the *Run-51.3.2* class. Compare this to example 53 in Section 5.2.2.

(71)

Run-51.3.2 *(revised)*
Roles: AGENT THEME SOURCE TRAJECTORY DESTINATION
Frame: NP V PP.TRAJECTORY
The horse jumped over the river. motion(during(E),Theme) path_rel(during(E), Theme, Trajectory, *ch_of_loc*, prep)
Frame: NP V NP.TRAJECTORY
The horse jumped the river. motion(during(E),Theme) path_rel(during(E), Theme, Trajectory, *ch_of_loc*, prep)
Frame: NP V PP.SOURCE PP.DESTINATION
The horse jumped from the rocks onto the shore. motion(during(E),Theme) path_rel(start(E), Theme, Initial_Location, *ch_of_loc*, prep) path_rel(end(E), Theme, Destination, *ch_of_loc*, prep)

The `path()` predicate, thus, would replace both the `Prep()` predicate (see ex. 53) and the `via()` predicate (see ex. 58). With the replacement of `Prep()`, we are achieving a more semantically explicit representation of the path. The replacement of `via()` allows us to represent the semantics of the trajectory independently of the constituent's syntactic expression.

5.3.1.1 On *cause-transform-cx*

For the *cause-transform-cx* type of CMC, we effectively replace the `Pred()` predicate with `path_rel()`. Example 60, taken from *Cooking-45.3* has been revised in example 72 to use the new predicate. The change has been bold-faced.

(72)

Cooking-45.3 *(revised)*		
Members (e.g.)		bake, boil, fry, roast
Frame		NP V NP PP
	Ex	*Jennifer baked the potatoes to a crisp.*
	Syn	AGENT V PATIENT RESULT
	Sem	cause(Agent, E)
		apply_heat(during(E), ?Instrument, Patient)
		cooked(result(E), Form, Patient)
		path_rel(end(E), Patient, Result, *ch_of_state*, prep)

Additionally, we introduce the new `transform()` predicate. It is applied only to the current classes that do not have a predicate to describe the motion event and are purely relying on the path information to imply that there was a change of state. We see this behavior for the classes that include verbs of transformation. The following is the revised semantic representation for *Turn-26.6.1*:

(73)

Turn-26.6.1 *(revised)*		
Members (e.g.)		convert, turn, transform
Frame		NP V PP
	Ex	*He turned into a frog.*
	Syn	PATIENT V RESULT
	Sem	**transform(during(E),Patient)**
		path_rel(end(E), Patient, Result, *ch_of_state*, prep)

Much like the `motion()` predicate we discussed earlier and the `transfer()` predicate we will see in later sections, the `transform()` predicate takes two arguments: the event *E* and the PATIENT role undergoing the change. The `transform()` predicate specifies that the transformation of the PATIENT holds true during the event *E*. The `path_rel()` predicate specifies that by the end of the event the PATIENT has reached the resulting state specified by the RESULT role.

5.3.1.2 On *cause-change-scale-cx*

Several changes are in effect for the representation of the *cause-transform-cx* type of CMC, which at the present time is restricted to the *Calibratable_cos-45.6* class. Example 74 shows the old and new semantic representation of a sentence with semantics of changes along a linear scale.

(74)

Calibratable_cos-45.6 *revised*	
Members (e.g.)	balloon, decrease, increase, jump, lower
Roles	PATIENT, ATTRIBUTE, EXTENT
Frame	NP.attribute V PP.recipient
Ex	*The price of oil increased by ten percent.*
Syn	ATTRIBUTE (OF) PATIENT V EXTENT
Old Sem	change_value(during(E), Direction, Attribute, Patient) amount_changed(during(E), Attribute, Patient, Extent)
New Sem	change_value(during(E), Patient, Direction) property(Patient, Attribute) path_rel(during(E), Patient, Extent, *ch_of_scale*, prep)

First, the `path_rel()` predicate replaces the `amount_changed()` predicate, allowing the complete path expression beyond that of the EXTENT role alone. Secondly, the `change_value()` predicate is reanalyzed as two separate predicates. The change_value() predicate is redesigned to be specifically used to represent the motion the PATIENT undergoes and it mirrors the current design of the `motion()` and `transfer()` predicates. We keep the DIRECTION role in this predicate as it is designed to express in which direction the motion tends. Additionally, the ATTRIBUTE role is removed from `change_value()` and reassigned to the `property()` predicate, which is already in existence in VerbNet (found in, for example, *function-105.2*, *coil-9.6*). The `property()` predicate is currently used to designate a role describing the quality or property of another role. The first argument specifies the role to which the property is attributed, and the second argument specifies the attribute of that role.

What is still wanting is the representation of causation. Many of the verbs in this class allow for a cause argument (e.g. *The recent developments increased the price by ten percent*). In the recent weeks this class has undergone a division that separates those verbs that subcategorize a causal entity as their argument from those that do not. For the representation of the verbs that allow for a cause argument, the CAUSE role is now being considered.

5.3.1.3 On *cause-change-possession-cx*

Although the CMC semantic type classification unites the transfer of information with the transfer of possession, the VerbNet classes do not. VerbNet distinguishes verbs that deal with information (class 37 and 40 series) from the verbs of possession (class 13 series). Here we evaluate them separately. We will begin with an analysis of the representation of change of possession:

(75)

Give-13.1 *(revised)*		
Members (e.g.)		lend, give, pass, sell
Roles		AGENT, THEME, RECIPIENT
Frame		NP V NP PP.recipient
	Ex	*They lent a bicycle to me.*
	Syn	AGENT V THEME RECIPIENT SOURCE
	Old Sem	cause(Agent,E) transfer(during(E), Theme) has_possession(start(E), Agent, Theme) has_possession(end(E), Recipient, Theme)
	New Sem	cause(Agent,E) transfer(during(E), Theme) path_rel(start(E), Theme, ?Source, *ch_of_poss*, prep) path_rel(start(E), Theme, Recipient, *ch_of_poss*, prep) equals(Agent,?Source)

With the verbs of possession, the `has_possession()` predicate is replaced with the more general purpose `path_rel()` predicate, with the LocType constant specifying it as an instance of a change of possession event. One distinct feature of the representation of the *cause-change-possession-cx* type of CMCs is that while syntactically a sentence cannot realize both the SOURCE role and the RECIPIENT role as an oblique argument (the two are mutually exclusive), both roles are specified in the class. The reason being that, as you will recall from an earlier discussion, the AGENT itself takes on the role of the RECIPIENT in the *cause-receive-cx* type of CMCs and SOURCE in the *cause-have-cx* type of CMCs. We represent the relationship between the AGENT and the SOURCE or RECIPIENT by uniting the two arguments under the `equals()` predicate. This predicate is already in use in VerbNet to equate two events or roles in an event representation.

For transfer of information, the `transfer_info()` predicate is replaced by the `transfer()` and `path_rel()` predicates. The LocType constant in the `path_rel()` predicate specifies that this is an instance of transfer of information.

(76)

transfer_mesg-37.1.1 *(revised)*	
Members (e.g.)	explain, narrate, recite
Roles	AGENT, TOPIC, RECIPIENT
Frame	NP V NP PP.recipient
Ex	*I explained the matter to them.*
Syn	AGENT V TOPIC RECIPIENT
Sem	cause(Agent,E) transfer_info(during(E), Agent, Recipient, Topic)
New Sem	cause(Agent,E) transfer(during(E), Topic) path_rel(start(E), Topic, ?Source, *ch_of_info*, prep) path_rel(start(E), Topic, Recipient, *ch_of_info*, prep) equals(Agent,?Source)

As a counterpart to the THEME role in transfer of possession, the TOPIC role in transfer of information specifies the entity undergoing the transfer. We recognize that the inferences of this movement differ from the transfer of possession in that the TOPIC does not completely depart from the AGENT who initiates the transfer. However, the distinction can be easily made on the basis of the LocType constant that specifies that this is a transfer of information rather than a transfer of possession. As we saw earlier, the `equals()` predicate again captures that the AGENT is the SOURCE of the information.

While we are specifically interested in the representation of CMCs, the representation is also made available to semantically related but syntactically distinct frames like the following instances from *Tell-37.2*. All of these sentences receive the same semantic analysis as specified in example 76.

(77) a. Ellen told a story to Helen (CMC)

b. Ellen told Helen a story (Ditransitive)

c. Ellen told Helen that the party would be tonight. (Sentential Complement)

d. Ellen told Helen, ”Leave the room.” (Direct Speech)

5.3.2 Semantic Roles

So far in this chapter we have discussed the arguments relevant to describing CMCs, without addressing the issue of the differences in how they are realized as thematic roles across different VerbNet classes and CMC semantic types. Individual classes realize roles that are semantically specific for the given needs of that class. What we need to capture are the semantic equivalences across classes expressed by the semantic roles. Whether the undergoer argument is realized as either a PATIENT or a THEME, we want to generalize across them to find a common representation. VerbNet's semantic role hierarchy (Bonial et al., 2011a) captures the necessary semantic generalization of the semantically specific roles by assigning them to more coarse-grained supertypes (see Appendix C for the complete hierarchy). At the finer-grained level, the semantic role hierarchy meets the necessary semantic specificity of the individual classes; at the coarser-grained level it provides roles to which the finer-grained roles can be generalized.

The original hierarchy includes the SOURCE and GOAL roles, but the path of motion (i.e. the location traversed through motion) is missing from the hierarchy. We introduce the PATH role to fill this gap, and the TRAJECTORY role as its subcategory for use in the change of location classes. Additionally, we establish a connection between the the existing EXTENT role and the PATH role. We also consider a potential new role VIA. Table 5.5 shows the breakdown of the current roles by their semantic types, with the supertype of each column displayed at the bottom of the table. Figure 5.1 also shows the new roles and their place in the semantic hierarchy. The new roles are marked in blue in both Table 5.5 and Figure 5.1.

Using this hierarchy, the classes have the flexibility to pull back to the coarser-grained super roles if the finer-grained roles are not appropriate. For example, a class expressing semantics of change of location might want to use the TRAJECTORY role. If the semantics of the class requires the role to be semantically underspecified as to the type of motion event it describes (i.e. allow for abstract usages of path of motion), the PATH role may be more acceptable for it.

Additionally, allowing for coarse-grained roles as supertypes to the finer grained roles helps

CMC Type	Cause Entity	Entity in Motion	Locations in the path of motion		
			start	during	end
cause-displace-cx	AGENT	THEME	INIT_LOCATION	TRAJECTORY	DESTINATION
cause-chng-scale-cx	CAUSE	PATIENT		EXTENT	
cause-transform-cx	AGENT	PATIENT	SOURCE	n/a	RESULT
cause-chng-poss-cx	AGENT	THEME / TOPIC	SOURCE	**?VIA**	RECIPIENT
Super Roles:	**ACTOR**	**UNDERGOER**	**SOURCE**	**PATH**	**GOAL**

Table 5.5: New predicates and roles relevant for CMCs in VerbNet classes

us do away with the creation of multiple semantically specific roles for the needs of the individual classes. During the initial stages of this study, we had considered the creation of subtype roles for the specific semantics they represent (e.g. INITIAL_STATE, INITIAL_VALUE, FINAL_VALUE). However, keeping in with our efforts to reuse as many of the resources that VerbNet already provides, we have chosen to not include these roles at this time. Creating semantically specific roles might be useful in meeting the specific semantics of a class, but it also creates a degree of unnecessary sparsity (e.g. INITIAL_VALUE role would only be used once in the whole of VerbNet). Finer-grained roles can in future be added as they become necessary to distinguish a large subset of the semantics represented by the roles.

Finally, in a recent series of discussions on semantic roles, it has been suggested that the VIA role could be added as a subtype of the PATH role to describe the portion of the path that connects the SOURCE and the RECIPIENT. For transfer of possession, the VIA role would describe the method of transaction (e.g. *We purchased it **on/via e-bay***). For transfer of information, the VIA role would describe the method of communication (e.g. *They acknowledged **by/via email***). The unique property that sets this role apart from the other PATH roles is that its semantics includes a manner of transfer or it is sometimes read as an instrument of transaction or communication, while none of the current semantic roles in VerbNet represent this type of information (c.f. Schneider (2014)). This role is currently also under discussion for inclusion in the semantic hierarchy.

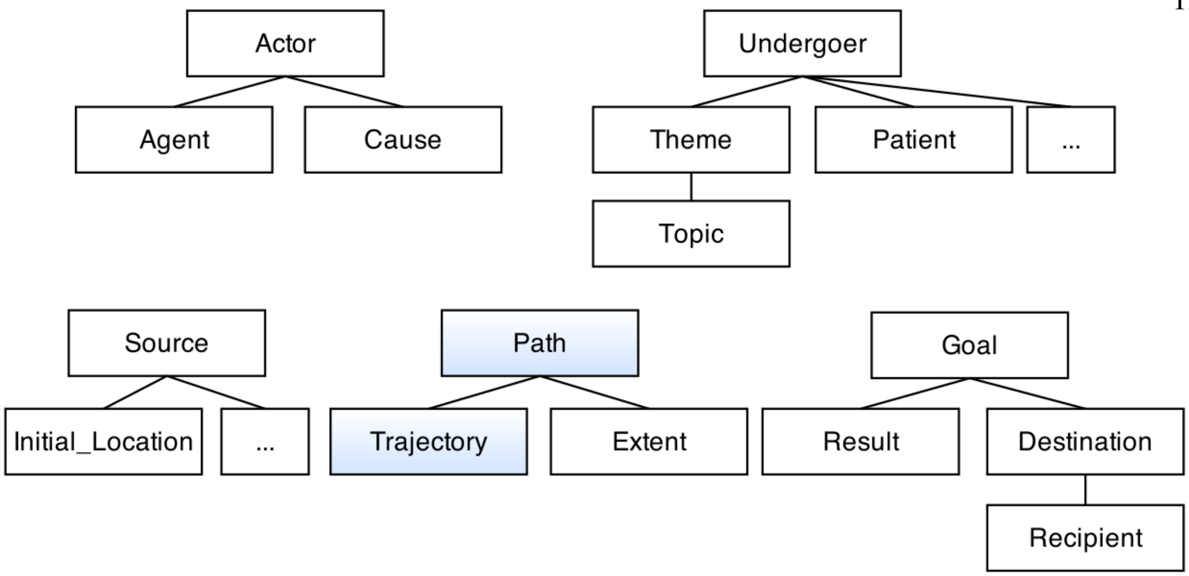

Figure 5.1: Semantic roles relevant in the representation of CMCs

5.4 Implementing Changes to VerbNet

In Section 5.2, we looked at how the current VerbNet semantic representations handle the semantic components of CMCs, pointing out inconsistencies and problems. In Section 5.3, we have discussed in detail the new predicates that allow for a VerbNet wide semantic representation, in particular, the semantics of path of motion, addressing the issues specified in Section 5.2. We now turn to the implementation of the changes detailed in Section 5.3. Because VerbNet's semantics were analyzed on a class by class basis with some consistency across classes of the same series (e.g. verbs of putting, for example, takes up the class 9 series), an approach to the implementation of the VerbNet wide changes requires elaboration. Although there are some clear one-to-one correspondences between the current and the new representations, the current representations are generally too heavily conditioned on individual class semantics to accommodate VerbNet-wide generalizations on how exactly the current predicates must be updated. We divide up the task of implementation to suit the semantic needs of the individual VerbNet classes.

Currently, the changes to predicate representation are scheduled for late in the Summer of 2014, and we expect they will continue into the Fall of 2014. While there is no simple way of anticipating every problem we might face as the implementation makes its headway, we suggest the groups of classes according to their needs and their expected challenges.

We envision the process of implementation to take place in two general stages. In the first stage, we tackle the VerbNet classes that already specify one or more frames that include the motion event and the path of motion. At this stage, we will use generalizations given what we know of the semantics of these VerbNet classes to carry out a sweeping implementation pass over a large number of classes. New predicates will replace the existing predicates and fill in any semantic gaps we see in the representation, and new roles will be introduced as needed to complete the semantic descriptions of the classes.

During the second stage of implementation, trickier classes that could not be handled during the first stage will get a closer look. Additionally, the second stage will introduce semantic representation of the path of motion to any verb classes that should include it but do not.

5.4.1 Stage 1: Large Scale Updates for Current Classes

The purpose of this stage is to fix the problematic representations. The goals of this stage include separating the static locations from the paths of motion, adding motion event representations to classes that should include it but do not, and addressing the issue of problematic predicates, with a specific focus on the `Prep()` predicate.[5] Each of the VerbNet class representations including any of the predicates discussed in Section 5.2 were evaluated to see if there are any emergent semantic groupings that would help us categorize and prioritize the implementation process.

Out of the 288 classes that are currently in VerbNet, 123 classes include one or more of the predicates discussed in Section 5.2. The breakdown of class counts by predicate type is shown in Table 5.6.[6]

[5] As you will recall from Section 5.2, the `Prep()` predicate does not have a semantic value. It simply specifies that whatever information is contained within the predicate, was expressed in the prepositional phrase.

[6] Please note that the numbers of the counts do not add up to 123 because a class can include more than one of these predicates. Also, the count for `motion()` in this table includes one instance of `rotational_motion()` which appears in *Coil-9.6*. This class is listed as one of Levin's *Verbs of Putting*. Given the semantics of this class, it seems that this particular predicate might have been intended to be used as a subtype of `motion()`.

Predicate	Count	Predicate	Count
`amount_changed()`	1	`Pred()`	21
`change_value()`	1	`Prep()`	38
`direction()`	5	`transfer()`	10
`has_possession()`	14	`transfer_info()`	20
`location()`	45	`state()`	8
`motion()`	32	`via()`	4

Table 5.6: Count of VerbNet classes using our predicates of interest.

We segment the 123 classes into three semantic groups: (1) classes that express a motion event with a path of motion, (2) classes that express no motion event and no path of motion, and (3) classes that express no motion event, however, there is a path of 'motion' encoded in the semantics. Each of these categories requires distinct ways of handling the predicates. We expect the first stage to cover the updating of 113 out of the 123 classes, leaving 10 classes from this particular list for stage 2 of the implementation.

5.4.1.1 +Motion Event; +Path of Motion

This is the most prototypical semantic grouping that includes all of the CMC categories, and it also includes the verb classes that express a motion event (whether that is change of location/state/scale or transfer of possession or information). One point that should be made early on is that these groups are not exclusive to CMC usages and also include directed motion and path verbs. Consider the following three sentences. None of these can be analyzed as caused motion instances as they do not abide by CMC's syntactic structure.

(78) a. The ship appeared on the horizon. (*Appear-81.1.1*)

b. John left. (*Leave-51.2*)

c. The team reached the hill. (*Reach-51.8*)

A complete listing of the verb classes that fall under this category are found in Tables D.1 through D.4 in Appendix D. The changes and updates, as expressed in Section 5.3, apply fairly neatly to the classes in this group, and there are some sweeping generalizations we can make about classes found in this group:

- **`direction()` and `via()` → `path_rel()`**

  All instances of the predicates `direction()` and `via()` map directly to the `path_rel()` predicate. In each of the classes, where they appear, these predicates are used to describe a location along the path.

- **Path `location()` and `Prep()` → `path_rel()`**

  All instances of `location()` or `Prep()` happening at `start(`*E*`)`, `end(`*E*`)`, and `result(`*E*`)` will map directly to the `path_rel()` predicate. However, a `location()` or `Prep()` predicate happening `during(`*E*`)` will have to be examined to make sure that it is a path of motion (e.g. *The horse jumped* ***over the fence***) rather than a static location (e.g. *The horse jumped* ***in the paddock***).

- **Static `Prep()` → `location()`**

  All instances of `Prep()` that express a static location map to `location()`.

- **`Pred()` → `path_rel()`**

  All instances of `Pred()` used for resultatives map to `path_rel()` taking place at the `end()` of the event.

Additionally, each of these classes must specify a motion predicate relevant for the class. Consider the following sentences:

(79) a. Mary put the book on the table. (*Put-9.1*)

b. Doug removed the plates from the table. (*Remove-10.1*)

In sentence 79a, there is a motion associated with the event: Mary causes the book to be moved onto the table. In the same way, in 79b, Doug causes the plates to move away from the table. The event indicates the motion of the plates, just like the motion of the book, in the previous sentence. Interestingly, only the first example receives the `motion()` predicate, while the second does not. We observe that classes identified as Levin's *Verbs of Putting* (class 9 series), *Verbs of Send and Carry* (class 11 series), and *Verb's of Motion* (class 51 series), consistently include the `motion()` predicate to describe the motion event in the sentence. However, classes like *Verbs of Removing* and various verbs for which one would expect the semantics of motion, do not consistently include the `motion()` predicate.

The class *Appear-48.1.1* is another example. This class includes verbs like *appear*, *come*, and *flow* that describe an event in which an entity becomes visible in the scene. In other words, the verbs encode a path of motion that indicates that the entity, which was at the beginning of the event somewhere off the scene, appears in the scene by the end of the event. The class *Appear-48.1.1* tries to code this path of motion in example 78a by assigning two `location()` predicates at the start and end of the event, both indicating the horizon as its destination, and negating the `location()` predicate relevant for the start of the event with the `not()` predicate (see example 54 in Section 5.2.2). By including the `motion()` predicate in the semantic predicates, along with the `location()` predicates that are replaced by the more explicit path predicate, `path_rel()`, we can represent the motion event that the verb is encoding for. The new predicates for example 78a are the following:[7]

(80) A ship appeared on the horizon.

`motion`(*during(E)*, THEME)

`path_rel`(end(E), THEME, GOAL, *ch_of_loc*, *prep*)

In the same way, we introduce the `transfer()` predicate to the verbs of transfer of possession/information and `transform()` to VerbNet classes missing motion events in their representation:

[7] Currently the *Appear-48.1.1* class specifies a LOCATION role for the final location in which the THEME argument is found at the end of the event. With the addition of the `motion()` predicate, the role will have to be updated accordingly.

(81) a. The thief stole the painting from the museum. (*Steal-10.5*)

`cause`(AGENT, *E*)

`manner(during`(E), *illegal*, AGENT)

`has_possession(start`(E), SOURCE, THEME)

`has_possession(end`(E*),* AGENT, THEME)

`not(has_possession(end`(*E*), SOURCE, THEME))

b. The witch turned the prince into a frog. (*Turn-26.6.1*)

`cause`(AGENT, *E*)

`not(state(start`(*E*), RESULT, PATIENT))

`state(result`(*E*), RESULT, PATIENT)

In both of these classes, note that the path over which the possession moves and states involved in the transformation event are certainly represented, but much like *Appear-48.1.1*, they are missing the predicate that describes the event. The final predicates appear as follows:

(82) a. The thief stole the painting from the museum. (*Steal-10.5*)

`cause`(AGENT, *E*)

`transfer`(*during(E)*, THEME)

`manner(during`(E), *illegal*, AGENT)

`path_rel(start`(E), THEME, SOURCE, *ch_of_loc*, *prep*)

b. The witch turned the prince into a frog. (*Turn-26.6.1*)

`cause`(AGENT, *E*)

`transform`(*during(E)*, THEME)

`path_rel(end`(E), PATIENT, RESULT, *ch_of_state*, *prep*)

The list also includes other classes that could receive a motion event predicate given the current semantic representation, but could also require further discussions. Here are a couple examples:

(83) a. The dogs ate their food off the floor. (*Eat-39.1*)

b. Mary hired two secretaries as helpers. (*Hire-13.5.3*)

The reason that the *Eat-39.1* class is in this list is for the examples like 83a that take on a caused motion-like reading: the sentence approximately describes an event where the food is caused to be removed from the floor by the means of eating. It is uncertain why this particular frame was included in this class, given that this is a consumption class. But if we are to include instances like this in the class, we also need to introduce the `motion()` predicate to express the movement that the food undergoes in the event. A better way to deal with this situation is to simply remove the problematic frame from the class. In other words, let the *Eat-39.1* class give the verb the meaning of eating or ingesting. Let the sentence receive the relevant semantic predicates of caused motion from the CMC itself.

The *Hire-13.5.3* class includes verbs that express employment like *hire*, *recruit*, *enlist*. This class, is a later addition to VerbNet that likely was analyzed as belonging to the class 13 series, *Verbs of Change of Possession*, as a parallel to the *Get-13.5.1* class (e.g. *Mary got two secretaries as helpers*; the verb *hire* is a member in Levin's original *Get-13.5.1*). Again, if we are to analyize this sentence as a case of change of possession, we would have to give the sentence the motion event predicate to indicate that there is a movement or obtaining of the two secretaries during the event. It seems, though, that the semantics indicated by *Hire-13.5.3* are different from those of *Get-13.5.1*; in the first sentence we have a clear indication of a beginning of an employment event by an organizational entity, for which the *Get-13.5.1* class is underspecified. One better solution would be to give it the semantics of hiring and employment, much like FrameNet's *Hiring* frame, and move this class out of the class 13 series.[8]

[8] As you will see, the *Hire-13.5.3* class is included in Table D.1 along with `motion()` taking classes, rather than in the change of possession list, where it belongs since it is a part of class 13 series. The reason it was included in this class is to draw a parallel with the *Fire-10.10* and *Resign-10.11* classes that are semantically related but sit in the *Verbs of Removing*, class 10 series.

5.4.1.2 -Motion Event; -Path of Motion

The next group to be tackled are the events in which there is no motion event specified for an undergoer argument, and what locative information that is specified is static. A complete listing of the verb classes that fall under this category are found in Table D.5 in Appendix D. The following are the generalizations for this group of classes:

- **Reconsider the `motion()` predicate**

  3 of the 16 classes in this group currently have the `motion()` predicate specified. However, unlike the `motion()` predicate as we have defined it, the motion in these classes denotes existence of movement rather than a movement along a path. The two types of motion should be disambiguated. We will look at this in a little more detail shortly.

- **Revisit `location()` and `has_possession()` predicates**

  Since these two predicates represent static places in which an object or a possession is located, the temporal function denoting the predicate's applicability with respect to the event (e.g. `start(E)`, `end(E)`, `during(E)`) should be unspecified like *own* in *Own-100.1* (e.g. *He owns two books.* → `has_possession`(***E***, *He*, *two books*)).

- **`Prep()` → `location()`**

  All instances of `Prep()` that express a static location, map to `location()`. Currently there aren't any `Prep()` predicates that map to `has_possession()` predicates.

The three classes whose use of the `motion()` predicate does not line up with the semantics of a motion event are exemplified in the following sentences:

(84) a. [A flag]-THEME fluttered over the fort. (*Modes_of_being_with_motion-47.3*)
`motion(during(`*E*`)`, THEME)

b. The garden is swarming [with bees]-THEME. (*Swarm-47.5.1*)
`motion(during(`*E*`)`, THEME)

c. [The dog]-$_{\text{AGENT}}$ flopped in the corner. (*Assuming_position-50*)

`motion(during(`*E*`),` AGENT)

The first two examples are from the class 47 series, *Verbs of Existence*. The `motion()` predicates, here, represent motion that "characterizes the existence of these entities" (Levin, 1993). In the example for *Assuming_position-50*, the predicate describes the manner in which the entity assumes a spatial configuration. And this manner is encoded in the verb. Here is a more detailed view of example 84a.

(85)

Modes_of_being_with_motion-47.3		
Frame		NP V PP.location
	Ex	*The flag fluttered over the fort.*
	Syn	THEME V
	Sem	motion(during(E), Theme) exist(during(E), Theme) Prep(E, Theme, Location)

The above sentence syntactically resembles the intransitive example of jump seen in example 53 in Section 5.2.2.1 (e.g. *The horse jumped over the fence*). Unlike the example 53, however, the `motion()` predicate is used here to denote the existence of movement without denoting motion along a path. The predicate here is representing a *manner* of motion: the motion of the fluttering flag that is taking place at a static location. Besides the `motion()` predicate, we have the `exist()` predicate that encodes the existence of the flag during the event. Its location over the fort is represented with the `Prep()` predicate.

It is obvious that the predicate for the path of motion needs to be changed to distinguish the directed motion instance from the static location seen in 85. That is, once we replace the `Prep()` predicate with the `location()` predicate, we can indicate that the LOCATION is a static location where the fluttering flag exists. Even so, we are still left with a confusing predicate if the verbs in the *Modes_of_being_ with_motion-47.3* are used in a directed motion reading via coercion. Consider the following sentence in comparison to the intransitive example in 53. What would `motion()` really represent in such a case?

(86) (I pulled open the locker and) a note fluttered into my hands.

When the verb *jump* is used in a directed motion sense, the jumping event is interpreted as the means by which the horse moves (e.g. *The horse went over the fence by jumping*, Levin and Rappaport-Hovav (1988); Jackendoff (1990)). The jumping event may as well be described as a manner of motion much like the fluttering, but when used in the change of location sense, there is no ambiguity since jumping will be interpreted as the means by which the horse went over the fence. On the other hand, when *flutter* is used in a directed motion context, the fluttering is interpreted as a manner (e.g. means:**The note went into my hands by fluttering* vs. manner:*The note went into my hands (while) fluttering*). Since we still use `motion()` for both instances, the predicate effectively codes for two different types of movements. Thus, there might be a need for predicate disambiguation.

A possible way to remedy this ambiguity is by in making use of the `manner()` predicate we find in the class 18 series, *Verbs of Contact by Impact*, and class 21 series, *Verbs of Cutting*. The `manner()` predicate in these classes is used to represent the motion relating to the manner in which the event was conducted. Here is an example from *Cut-21.1*:

(87) Carol cut the bread

`cause`(AGENT, *E*)

`manner`(`during`(*E*), *Motion*, AGENT)

`contact`(`during`(*E*), ?INSTRUMENT, PATIENT)

`degradation_material_integrity`(`result`(*E*), PATIENT)

In this example, rather than using the `motion()` predicate, the `manner()` predicate is used to indicate that during the course of the event, a manner of motion, specified by the constant *Motion*, is associated with the AGENT. Thus, if we make a similar interpretation of the three classes, we can indicate that the flag exists in the location described by the `location()` predicate and that its existence can be unambiguously characterized by the fluttering motion:

(88) [A flag]-$_{\text{THEME}}$ fluttered over the fort. (*Modes_of_being_with_motion-47.3*)

`manner(during(`*E*`),` *Motion*, THEME)

Another way to handle the manner of motion is to introduce a new motion predicate such as `manner_of_motion()`. Currently the `manner()` predicate is also used to represent other manners of events such as designating *illegal* events in *Steal-10.5* and *forceful* events in *Bump-10.4*. Thus, rather than overloading the semantics on the `manner()` predicate, a separate predicate can be dedicated to the representation of the manner of motion.

5.4.1.3 -Motion Event; +Path of Motion

This group includes three general categories of verb meaning: the aspectual verbs, the fictive motion verbs (as defined by Talmy (2000); c.f. pseudo-motional locatives in Dowty (1979)), and the creation verbs (see Table D.6 in Appendix D). These verbs allow for the expression of a path of motion without an expression of a motion event. Consider the following examples:

(89) a. The party continued until 8PM. (*Continue-55.3*)

b. A river runs through the valley. (*Meander-47.7*)

c. His property begins at the fence. (*Terminus-47.9*)

d. Saul jotted down readings on a notepad. (*Scribble-25.2*)

Admittedly, calling this group a case of "-motion event" is something of a misnomer. That is, labelling this group as lacking a motion event does not account for those instances where we can establish a metaphorical extension based on motion and changes of location. For example, Jackendoff (1983) defines a conceptual field, called a *circumstantial* field, that maps to the spatial semantics . In it, the "things" involved in the event are the entities in motion, and events or states are conceptualized as locations in a path. Thus, in example 89a, the party *circumstantially goes* to 8PM. Similarly, by Lakoff's conceptual metaphor of "Creating Is Moving To A Location" (Lakoff et al., 1991), the created object is conceptualized as the entity in motion, and by this account we could possibly say, in example 89d, that Sam causes the readings to be moved to the location of the notepad. This group of verbs is considered "- motion event" in so far as that they do not fit in our current definitions of motion events.

Nevertheless, what is important here is that these classes do express a path. Much like 89a and 89d, the river extends over a path that runs through the valley in example 89b, and, in 89c, the property has a boundary established at the point of the fence and extends from this point. The existence of a path, suggests that the `path_rel()` predicate may be an appropriate representation for these verb classes.

There is one generalization that can be made for updates to existing predicates. For all classes in this group, `Prep()` and `location()` should be reanalyzed as `path_rel()`. Additionally, in order to complete the `path_rel()` predicate, we would need a constant to capture the type of path the semantics describes, which we will call *extend*, for the purposes of this example. The following is an example of what the final representation of 89a would look like:

(90) [The party]-$_{\text{THEME}}$ continued until [8PM]-$_{\text{TIME}}$.

Old: `continue`(*E*, THEME) `Prep`(*E*, THEME, TIME)

New: `continue`(*E*, THEME) `path_rel`(end(*E*), THEME, TIME, *extend*, *prep*)

We expect that we will have to give a particular attention to the constant we labeled *extent* in the above example. Note that the semantics expressed in the path phrases differ in examples 89a through 89d. For example, the extent of time in 89a is contrasted from the location over which a river stretches in 89b. In fact, compare them to the following sentence:

(91) The temperature went from 60F to 15F (in a matter of 3 hours!) (*Meander-47.7*)

This example is an instance of fictive motion as the example 89b. However, unlike 89b where the path does not indicate change of location (i.e. the river is not changing its location from one point of a valley to another – it simply extends over the valley), the path in example 91 is an effective change in temperature, which might be more aptly captured through the use of the *ch_of_state* constant (see 5.3.1). Thus, the assignment of an appropriate constant will require further attention.

5.4.1.4 Summary

In this section we discussed the generalizations we can make about VerbNet classes and the changes that will affect classes that already represent changes of motion, motion events, and/or path of motion in some capacity. This process will replace usages of `direction()`, `via()`, `Prep()`, `state()`, and `transfer_info()`, leaving the predicates `location()` and `has_possession()` predicates specifically for static locations and possessions, respectively.

Throughout this discussion, we make sweeping generalizations about the different groups we see as semantically coherent for the purposes of predicate implementation. However, we expect that there will be cases that require closer attention. For example, Table D.4 only lists 3 classes out of the 9 classes that are a part of the class 26 series. We expect that we will need to analyze the other classes in the series in order to see if the same motion event predicate, `transform()`, will apply to them as well. In a similar way, the *Continue-55.3* class is one of the 6 classes that describe aspectual verbs. The rest of the class 55 series might deserve an examination to see if the same analysis we apply to *Continue-55.3* will also apply to them.

Moreover, as we have specified, we expect Stage 1 to handle 113 classes out of the 123 classes that make use of a predicate relevant to verbs of motion and paths of motion. The remaining 10 classes and any other classes that are identified as needing extra attention will roll over to Stage 2 of the implementation. The second stage will give closer attention to the classes that need extra analysis the most, and also explore the needs of the classes that should describe paths of motion, but currently do not (i.e. false negatives for paths of motion classes).

5.4.2 Stage 2: Close-up Evaluation of Remaining Classes

In a previous study (Bonial et al., 2011b), we examined each of the VerbNet classes[9] to determine their compatibility with or coerciblity into caused motion constructions. Using this work as a starting point, we reevaluated the classes under the new CMC guidelines to produce a list of all classes that are compatible with caused motion usage or coercible into CMC. Through this process, we were able to identify 85 compatible or coercible classes, not overlapping with the previous 123 classes. Out of the 85 classes, we have identified 16 classes where CMCs should be considered for inclusion with their conventional frames.

The second stage of the implementation process , thus, will focus on those 16 classes that are suitable for the inclusion of CMCs and the remaining 10 classes from the previous stage. A complete listing of the verb classes is found in Table D.7 in Appendix D. The classes in this group, especially the 16 newly-identified classes, have not been specifically analyzed for the appropriateness of their representation of the motion event and the path of motion. They will have to be put through the same analysis as the classes in Stage 1. The following are a couple of initial observations about these classes, and the types of challenges they might present.

5.4.2.1 Missing Frames or Incomplete Semantics

The newly identified classes are either missing a frame for caused motion instances or they already include appropriate frames but do not have the necessary semantic representation. For example the *Ferret-35.6* class specifies the semantic roles AGENT, THEME and SOURCE and a frame that exemplifies CMC usage, "I ferreted the secret out of him." What is lacking is not the semantic roles or the appropriate semantic frame, but it is the semantic representations for the motion event and the path of motion that are entirely missing:

(92) I ferreted the secret out of him.(*Ferret-35.6*)

`search(during(`*E*`)`, AGENT, SOURCE, THEME)

[9] At the time there were a total of 270 classes in existence in VerbNet. 18 new classes have been added since.

The semantic predicate captures the aspect of the event in which the AGENT is searching for the THEME, but the sentence still needs the right semantics to express that the AGENT was successful in attaining the secret, and that the secret was pulled out of the SOURCE.

5.4.2.2 Semantic Role Mismatch

Ideally all of the classes that represent path of motion and caused motion constructions will have thematic roles that reflect the prototypical thematic roles, namely, Actor, Theme, Source, Goal, and Path. The classes represented in Stage 1 either use semantic roles that match roles in this list or use roles that belong as a sub role to one of these roles (e.g. Agent, Destination, Trajectory) in the thematic hierarchy. However, in Stage 2 we will come across classes that do not have this clean semantic role lineup. Consider the following example:

(93) [Linda]-AGENT taped [the picture]-PATIENT to [the wall]-CO-PATIENT. (*Tape-22.4*)

The *Tape-22.4* class includes verbs of attaching such as *anchor*, *glue*, and *screw*, and the semantic roles we see in the example are very specific to the features of the class. The above example is an instance of CMC in that the picture is caused to be on the wall by the means of taping. That is, the taped item is the undergoer and the wall to which it is taped is a destination. However, this class additionally allows for the direct object and the object of the prepositional phrase to be semantically symmetrical arguments (c.f. discussion in 3.3.1.3) exemplified by sentence like "Linda taped the pen to the pencil". This sentence, unlike example 93, can undergo the *together* reciprocal alternation: "Linda taped the pen and the pencil together (Levin, 1993, pg. 163). Because of this class feature, the thematic roles for the direct object and the prepositional phrase are PATIENT and CO-PATIENT, respectively.

Because the thematic roles represent the specific semantic characteristics of the verb class we cannot simply replace the CO-PATIENT with GOAL, without violating the semantic coherence of the class. One way of tackling classes like this would be segregating the classes into those that allow symmetrical arguments from those that do not. Barring that, we would have to introduce rules that specify an alternate mapping for the semantic roles in specific frames.

5.4.2.3 Non Motion Classes

Some of the classes listed in this group are not actually verbs of motion nor verbs that express paths of motion. There are three classes of this type: *Stimulus_subject-30.4*, *Initiate_communication-37.4.2*, and *Become-109.1*. The reason we cannot simply remove these classes from the analysis is that they make use of semantic predicates such as `state()` and `transfer_info()` that Stage 1 removes from usage. These classes will have to be examined to determine the predicates that should replace `state()` and `transfer_info()`.

5.5 Final Considerations and Future Work

In this chapter, we have presented a new set of semantic representations that will handle the path of motion in a more consistent and explicit manner. This remedies the issues seen in the current VerbNet semantic representation, which has no consistent way of handling verbs expressing the semantics of path of motion. We have focused on the development of semantic predicates to represent the cause motion constructions, but we gave particular attention to making sure these new semantic predicates are applicable to other usages of path of motion (e.g. motion events without cause, directed motion, and fictive motion). Moreover, in this chapter, we have discussed in some detail the process of implementing the new predicates as part of VerbNet, pointing out a number of challenges and issues that we might face during the implementation process. We have broken down the process into two stages: the first stage expedites the update of a large number of classes based on a number of semantic generalizations, and the second stage gives closer attention to classes that require further considerations. In the following chapter, we use the new predicates discussed in this chapter to give a semantic representation to sentences that have been automatically identified as CMCs in Chapter 4.

In terms of future work, we would like to point out two specific issues that could be addressed. As it was discussed a number of times in this thesis, the `cause-displacement -cx` CMC type is too semantically varied. Consequently, the range of semantics that the `motion()` predicate covers is rather wide. We have dealt with this problem by assigning appropriate values to the LocType constant in the `path_rel()` predicate that the `motion()` predicate accompanies, to restrict the semantic nature of the path to narrower meanings like change of location (i.e. *ch_of_loc*), change of state (i.e. *ch_of_state*), or change of possession (i.e. *ch_of_poss*). Even so, especially during the implementation stage, there may be a need to pay attention to the motion type a class specifies to see if there are disparities that warrant a new motion event predicate type.

The second issue that future work can address lies in bringing a measure of structure to the current semantic predicates and their use in VerbNet. To our knowledge, this is the first effort that has looked at updating VerbNet to bring consistency in representation to a particular semantic phenomenon that affects a large number of VerbNet classes, and it will not likely be the last. If we are to expedite future endeavors to bring more semantic consistency across VerbNet representations, a helpful first step may be to carry out an analysis of the current semantic predicates in VerbNet in order to identify semantic or functional similarities and differences. For example, this chapter presented predicates like `motion()`, `transfer()` and `transform()`, which are related to each other in that they represent the motion event. The `motion()` predicate, moreover, is related to other predicates that represent a motion type like `body_motion()` (from e.g. *Wink-40.3.1*) or `manner_of_motion()` (see Section 5.4.1.2). In turn, the `body_motion()` predicate is related to `body_process()` in that they describe something about the human body. Assessing the similarities and differences in the semantic predicates can help identify semantic redundancies and gaps. By doing this, we could guide the semantic predicates in the direction that the semantic roles have gone: recognizing relationships or possible hierarchies to bring more structure and usability to the semantic predicates. This will serve to aid any future efforts of carrying out a VerbNet wide semantic update, because if we know the relationships between the predicates, it would be easier to discern what related changes VerbNet might require if a single predicate is updated.

Chapter 6

CMC Sentence Representations

6.1 Introduction

As discussed in Chapter 1, this thesis is part of a greater effort to equip VerbNet with systematic ways for dealing with coercive usages of verbs. VerbNet currently provides meaning representation for the predictable semantics of the verb, but it is not as well-equipped to give proper representation to verbs when they are used in syntactic and semantic contexts that are atypical for the verb.

Equipping VerbNet with the capacity to provide a semantic representation for coercive usages of CMCs requires that the resource should be consistent and explicit in its representation of CMC and its semantic components. Thus, in Chapter 3, we have defined the semantic types of CMC, providing a categorical distinction that aids in the identification of CMCs in naturally occurring data. In Chapter 5, we have revised the semantic predicate representation in VerbNet to provide a consistent representation for all verbs expressing paths of motion. We have also outlined a two-stage implementation process that will update VerbNet with the newly designed semantic predicates. Once VerbNet is updated with the new predicates and there is consistency in the representation of the motion event and the path of motion, we can then utilize the prior semantics of VerbNet together with the semantics of caused motion to give coercive instances the representation they need.

The purpose of this chapter is to investigate how the coercive representations will work, and what issues and challenges we may expect in the process of generating representations. This chapter pulls from the results of the previous two chapters to give sentence representations to caused motion instances in the data. The classifier results in Chapter 4 give us automatically labeled CMC instances. The semantic predicates established in Chapter 5 and the VerbNet analysis carried out in the chapter, provide us with the necessary semantic representations for the caused motion sentences. Moreover, we are focusing on the sentence representation of the *cause-displacement-cx* type CMCs (i.e. the results from the DISPLACE classifiers from Chapter 4).

First, we will briefly discuss our task of identifying VerbNet classes that only appear in CMC as cases of coercion and the challenges involved in Section 6.2. Rest of the chapter will present our results from the CMC sentence representation investigation.

6.2 On Coercion

Much of our effort has gone into the evaluation and categorization of VerbNet classes according to their conventionality and coerciveness with respect to CMC usages. In Bonial et al. (2011b), we examined each VerbNet class to determine its compatibility with caused motion usages. The results of this study was re-evaluated in detail under the new CMC guidelines. This effort resulted in an exhaustive list of all the VerbNet classes, identifying the caused motion construction types as coerced or conventional in each of these classes. The results of this re-evaluation is included in Appendix E.

The initial goal for the identification of coercive VerbNet classes was to mark up the relevant classes as being coercive or conventional, however, as we discussed, this task proved to be difficult. We made a final decision to clean up the VerbNet representations for the classes that already represent motion or paths of motion, as presented in Chapter 5 and consider the classes that have been already evaluated as carrying the semantics of motion and a path of motion as conventional CMCs. After all, if the semantic predicates in a class are already able to handle the caused motion usages, it would make sense to focus instead on the sentence representation for those classes that cannot handle them.

Throughout our work with the coercive usages of CMCs, one consistent issue we faced was achieving a clear consensus with respect to which VerbNet classes are coercive for caused motion usages. A consensus on whether or not a VerbNet class makes conventional use of CMCs is difficult because the decision comes down to an individual's acceptability judgments, which are dependent on frequency of usage and the individual's dialect or idiolect. For example, should the *Push-12* class, which includes verbs like *heave*, *yank*, or *shove* be considered a conventional use of a CMC? This aspect of semantics is underspecified in Levin's original classification: these verbs are grouped together for denoting an exertion of force. Thus, the decision would depend on whether or not the individual making the acceptability judgment decides that the path of motion is considered conventional for the class. We could argue that CMCs are conventional for the class because these verbs are frequently found in a caused motion context. We could counter this by saying that the path of motion is not semantically obligatory (Palmer, 1985).[1]

Here are a couple more examples. Should these instances be cases of coercion?

(94) a. Cynthia **carried** (*Carry-11.4*) her keys across the room.

b. Cynthia **shed** (*Substance_emission-43.4*) some light on the problem.

We normally don't question if an accompanied motion class like *Carry-11.4* has CMC as a conventional frame because we generally assume that it does. Part of the reason we assume motion is because the class is built with the semantics of motion in mind. The point of the class is to indicate an event in which two entities move – one by its own volition and the other by virtue of being carried. Strictly speaking, however, the semantics of the verb *carry* denotes only the holding event; the motion reading comes from the path preposition phrase in the sentence. Levin does make a note of this in the original classification: in the case of directed motion, the motion has to be made explicit through the prepositional phrase (Levin, 1993, pg. 136). Conversely, the *Substance_emission-43.4* class, which includes verbs like *bubble*, *foam*, *leak*, and *radiate* is a class

[1] According to Culicover and Jackendoff (2005), semantically optional meaning can be negated, while semantically obligatory meaning cannot. That is, since we can negate the path of motion in *"I pushed the pumpkin but it didn't move"*, but not for verbs like *throw "*I threw the pumpkin but it didn't move"*, the path of motion must not be implicit in the meaning of the *push* sentence.

that we would expect, after a cursory inspection, would not include CMC as a conventional usage. However, sentence 94b is highly conventionalized and the sentence no longer reads as a case of coercion.

Because the decision of coercion or conventionality is often dependent on the semantics of the individual verbs in the class rather than the semantics of the class as a whole, in the end the decision comes down a tug-of-war between what the behavior of the majority of the verbs would dictate and what makes sense for the prototypical semantics of the class as a whole. For the evaluations of the classes in Appendix E, we made specific efforts to abide by the prototypical semantics of the class rather than the idiosyncrasies of individual verbs.

One other challenging aspect of the coercive-conventional distinction is that the VerbNet classes can be both coercive and conventional depending on what CMC type we are talking about. Consider the following two sentences:

(95) a. Paul {bashed/smacked/beat} the vase into pieces.

b. Paul {bashed/smacked/beat} the ball into a small hole.

Both examples are from the *Hit-18.1* class. This class currently lists sentence 95a in its frames, a *cause-transform-cx* type of CMC, indicating resulting state. This sentence reads as a highly conventionalized use of the verb *bash* and many of the other verbs in this class. The *cause-displace-cx* or the change of location counterpart exemplified in 95b is not as conventionalized. In fact, the sentence feels more like a coerced CMC usage of the verb *bash* rather than a conventional one. This may indicate a constructional distinction between these two types of CMCs, suggesting that the question of whether or not a VerbNet class is coercive for CMC is too general, and the evaluation may have to be determined on the finer grained types of CMCs.

Future studies on the topic of coercion with respect to the VerbNet classes would require inquiries into further methods. For example, instead of relying on the judgments of a small group of linguists, the ratings for coercion and acceptability could easily be crowdsourced based on automatically generated sentences. While such a method may not overcome disagreements, tendencies and patterns observed in the annotator choices may be helpful in making sense of the status of VerbNet classes with regards to CMC usage.

6.3 Sentence Representation

6.3.1 Data

Out of the 436 DISPLACE instances represented in the WSJ, WEB, and BN test data from Chapter 4, a total of 311 instances were correctly classified as DISPLACE by the classifiers trained on the WSJ data.[2] VerbNet class information was gathered for each of the 311 instances. Of these 311 instances, 153 instances were labeled with the gold annotated VerbNet class labels provided by the SEMLINK project (Palmer, 2009), and 158 instances were classified by an automatic VerbNet classifier trained on SEMLINK's VerbNet annotations.[3] The automatic classifier was able to assign VerbNet labels to 150 out of the 158 instances. Thus, this provides us with 303 instances with VerbNet labels and 8 unclassifiable instances.

VerbNet class labels for the 303 instances were cross referenced with the VerbNet classes identified as being associated with the *cause-displacement-cx* type in Section 5.4.1.1. These VerbNet classes are marked with a "Y" in column four in Table D.1 in Appendix D. Out of the 303 instances, 207 instances were marked as belonging to the VerbNet classes that allow the *cause-displacement-cx* type of CMC in their class definitions. In other words, the 207 instances would be receiving their representation directly from their respective VerbNet classes.

The remaining 96 instances were manually analyzed for the appropriateness of their VerbNet labels. These instances did not match with any VerbNet classes currently associated with the *cause-displacement-cx* CMC type. In this analysis, we found 56 problematic instances: 49 instances had incorrect VerbNet label and they were reassigned to one of the classes associated with the *cause-displacement-cx* type (see examples 96a and 96b); 4 instances had incorrect VerbNet label and

[2] For the sentence representation, we use the results from the full feature classifiers trained on the WSJ training set from Section 4.4.4, in particular, the Downsample 2 system.

[3] The VerbNet classifier is an unpublished work by James Gung using semantically generalized features based on works by Kawahara and Palmer (2014) and Brown et al. (2011).

were reassigned to classes that do not include the semantics of CMC in their class definitions (see example 96c), and 3 instances were identified as CMC annotation mistakes (see example 96d).

(96) a. Environmentalist are trying to **clean** the oil from the scene birds that have been contaminated.

b. But the agency hasn't **yanked** psyllium off store shelves.

c. Their machine tools are even **bolted** to the shop floor.

d. Adoption agency insists on **introducing** the adopting parents to the birth mother.

The verb *clean* from example 96a was marked by the automatic VerbNet classifier as an instance of the *Preparing-26.3* class (this class includes verbs like *bake*, *brew*, and *prepare*). It was reassigned to the more semantically compatible *Clear-10.3* class (e.g. *"Strong winds cleared the sky"*). The verb in *yank* was manually assigned by the SEMLINK project as an instance of the *Hold-15.1* (e.g. *"She held his arm"*). This verb also belongs to the *Push-12* class with the sentence *"Nora yanked the button loose"* as one of the class examples. It is unclear why the original SEMLINK annotators chose to annotate this sentence as *Hold-15.1*, but it felt appropriate to reassign this instance to the *Push-12* class. Out of the 49 instances reanalyzed as instances of VerbNet classes including CMC usage, 30 instances were errors on the part of the automatic labels. The rest were SEMLINK's manually annotated labels that better fit with one of VerbNet's CMC classes.

The verb *bolt* in example 96c was marked by the automatic VerbNet classifier as an instance of the *Gobble-39.3* class (e.g. *"Cynthia ate/gobbled/bolted(down) the pizza"*). It was reassigned to *Tape-22.4*, which currently does not include CMC as one of its frames. The sentence 96d is an example of one of the 3 CMC label annotation error.

Thus, gathering together the remaining 40 instances with correct VerbNet labels (i.e., 96 minus mislabeled 56 = 40), the 4 instances with corrected VerbNet labels, and the 8 instances the automatic classifier could not classify, we have a total of 53 candidates for the sentence representation experiment.

6.3.2 Approach and Method

For the production of sentence representations, we use a simple rule-based implementation system that pulls from the semantic predicates of the DISPLACE category of CMC and the semantic predicates of the VerbNet class to which the verb is associated. For the recovery of the semantic arguments of the verb for populating semantic predicates, we use the techniques described in Section 4.3.4.2.

Producing sentence representations for caused motion instances with semantic membership in a class not known for its caused motion usage, requires the amalgamation of the semantic predicates representative of their respective class and the semantic predicates representative of the caused motion usage[4] . Because VerbNet semantic predicates are available at the frame level of description as opposed to the top class level, we pull the semantic predicates of a class from one of the following three sources:

- For Frame [**NP V (PP)**] we pull from the intransitive frames. For instances like *"Cynthia blinked the snow off her eyelashes"*, the semantic representation we want to pull for the verb *blink* is available through the intransitive frame in the *Hiccup-40.1.1* class (e.g. *"Cynthia blinked"*).
- For Frame [**NP V NP**] we pull from the transitive frames. For instances like *"Cynthia intimidated him into submission."*, the semantic representation we want to pull for the verb *tease* is available through the transitive frame in the *Amuse-31.1* class (e.g. *"Cynthia intimidated John"*).
- When possible, for Frame [**NP V NP PP**] we pull from the intransitive and simple transitive usages, as the semantic predicates in these frames will be descriptive of the semantics

[4] For a complete listing of semantic predicates and semantic roles associated with each CMC type, see Appendix F.

of the verb. If neither frame exists, we pull the semantic predicates relevant to the [NP V NP PP] frame. In such cases, the PPs do not represent paths. They are obliques modifying a transitive sentence. For example, *[It]*-NP *freed [him]*-NP *[of guilt]*-PP from the *Free-80* class is an instance of this frame.

For the *cause-displacement-cx* semantics, we join the following basic representation with the extracted VerbNet predicates:

cause(CAUSE,E)
motion(during(E), UNDERGOER)
path_rel(start(E),UNDERGOER,SOURCE,*ch_of_loc*,prep)
path_rel(during(E),UNDERGOER,PATH,*ch_of_loc*,prep)
path_rel(end(E),UNDERGOER,GOAL,*ch_of_loc*,prep)

If the semantic predicates originating from VerbNet do not already specify the `cause()` predicate, it is added to the final sentence representation. The `path_rel()` predicates are added as necessary for the SOURCE, PATH, and GOAL roles as applicable to the sentence. Finally, the `motion()` predicate is added to the representation. Consider the following sentence:

(97) Being a true Cajun, we love to **suck** (*Chew-39.2*) out the remaining mixture and flavor from the outside of the shell.

The verb *suck* is labeled as an instance of *Chew-39.2*. The *Chew-39.2* class specifies the roles AGENT and PATIENT and includes an [NP V NP] frame that covers sentences such as *"Cynthia nibbled the carrot"*, where Cynthia is the *Agent* and the carrot is the PATIENT of the event. The `take_in()` predicate, which represents the ingesting event, is true for the duration of the event and takes the two roles as its arguments:

The *Chew-39.2* class in VerbNet provides:
take_in(during(E),Agent,Patient)

In order to represent the path over which the sucked item travels, we pull from the *cause-*

displacement-cx predicates. Thus, the final sentence representation receives the `cause()` and `path_rel()` predicates from the construction. Since the source argument (i.e. *"from the outside of the shell"*) is specified, we need the `path_rel()` predicate for the start of the event; and since the goal argument is not specified we can say that the UNDERGOER is no longer located at the source location by the end of the event. Also, note that LocType constant *ch_of_loc* (see Section 5.3.1) is specified in both of the `path_rel` predicates.

The *cause-displacement-cx* construction provides:
cause(Actor,E) motion(during(E),Undergoer) path_rel(start(E),Undergoer,Source,ch_of_loc,prep)

Finally, there is a discrepancy between the AGENT role from the VerbNet class and ACTOR from the construction. The same is true for the PATIENT and UNDERGOER roles. This discrepancy is reconciled via the semantic role hierarchy: we recognize that the ACTOR role is a superrole to the AGENT role and that the UNDERGOER role is a superrole to the PATIENT role.

During the semantic role instantiation process, *"we"* is recognized as the ACTOR/AGENT of the event, the noun phrase *"the remaining mixture and flavor"* is recognized as playing the UNDERGOER/PATIENT role and the *"the outside of the shell"* is recognized as referring to the SOURCE information. Finally the preposition *from* is recognized as the preposition heading the SOURCE argument.

6.3.3 Results

In the previous section, we used example 97 to describe the process of sentence representation retrieval. Example 98 shows the result of the automatically generated representation of the example 97. The representation is instantiated.

(98)

*"Being a true Cajun, we love to **suck** out the remaining mixture and flavor from the outside of the shell."*
verb: *suck (out)*, VerbNet Class: CHEW-39.2
take_in(during(E)"we","the remaining mixture and flavor") cause("we",E) motion(during(E),"the remaining mixture and flavor") path_rel(start(E),"the remaining mixture and flavor", "the outside of the shell",ch_of_loc,"from") not(path_rel(end(E),"the remaining mixture and flavor", "the outside of the shell",ch_of_loc,"from"))

As discussed, the `take_in()` predicate comes from the *Chew-39.2* class, and the rest of the predicates are from the *cause-displacement-cx* construction. The predicate from the VerbNet class captures the fact that a certain item was ingested by the AGENT argument. The constructional predicates provide the source location of the the ingested item and the information that it is no longer at the source location at the end of the event.

Example 99 shows another acceptable representation; this time, for the verb *spew* from the *Substance_emission-43.4* class.

(99)

*"The photos showed active geysers on Io **spewing*** sulfurous material *190 miles into its atmosphere."*
verb: *spew*, VerbNet Class: SUBSTANCE_EMISSION-43.4
emit(during(E),?Source,"sulfurous material") cause("active geysers on Io",E) motion(during(E),"sulfurous material") not(path_rel(start(E),"sulfurous material","its atmosphere",ch_of_loc,"into")) path_rel(end(E),"sulfurous material","its atmosphere",ch_of_loc,"into")

Much like in the earlier example, the `emit()` predicate comes from the *Substance_emission-43.4* class definitions, with the rest of the predicates coming from the constructional definitions. Through the joint representation, we not only get the information that the *"sulfurous material"* is emitted from an unspecified SOURCE location during the course of the event, but we also get the reading that its motion is caused by the *"active gayzers on Io"*, and that by the end of the event it ends up in the atmosphere.

With the constructional representation, even if we lack the VerbNet class label, so long as we have identified the instance as a case of caused motion, we can give it an appropriate semantic representation. Example 100 shows one of the instances that the VerbNet classifier was not able to classify.

(100)

"Foreign exchange dealers said that *the currency market has begun to* ***distance*** *itself from the volatile stock exchange [...]"*
verb: *distance*, VerbNet Class: UNCLASSIFIABLE
motion(during(E),"itself") cause("the currency market",E) path_rel(start(E),"itself","the volatile stock exchange [...]",ch_of_loc,"from") not(path_rel(end(E),"itself","the volatile stock exchange [...]",ch_of_loc,"from"))

Even though we do not have the specific semantics of the verb *distance* in the representation, we are able to represent that the currency market is the cause of its own displacement from a certain source location/state. At the start of the event it is at the *"the volatile stock exchange"*, but no longer there by the end of the event.

Even the best of the representations in this batch are not without their problems. For example, one representation that is perhaps relevant for sentence 98, that is not represented by the predicates, is the directional particle *out* in *suck out*. Directional particles have been left out from the caused motion discussion, but in the future it may require further investigation. In example 99, relevant information to the path of motion that is not represented is the extent to which the sulfurous material travels in the atmosphere, namely, the *"190 miles"* specification. Since in this thesis we are not exploring the path information contained in the nominal phrases, focusing only on the prepositional phrases, we miss this type of information that could be used to elaborate further on the motion of the undergoer.

Nevertheless, what we get out of the joint VerbNet and CMC representation is the additional semantics of path of motion that may not be available from only the semantics the verb and its class . Furthermore, so long as we can identify a sentence as an instance of caused motion, we can use constructional definitions to assign the appropriate semantic predicate representations relating to motion and the changes of location/state, even if the verb is not found in VerbNet.

In the following subsections we discuss some of the issues relevant to the semantic representation of caused motion instances, detailing some of the challenges and issues faced, and identifying some of the future work that may be relevant to addressing these issues.

6.3.3.1 On LocType Constant

For the purposes of this experiment, we have kept the LocType constant fixed at *ch_of_loc*. However, this may be problematic for some of the representations. We have already seen an exemplar of such a case in example 100. The currency market moving away from the stock exchange is more akin to a change of state rather than a change of location. Examples 101 and 102 are further examples of both metaphoric or abstract motion: in 101, a freeze is placed on the utility firms and in 102, retirement is a state one can enter.

(101)

"The crisis began three years ago when California ***imposed*** *a price freeze on its private utility firms."*
verb: *freeze*, VerbNet Class: ENFORCE-63
enforce(during(E),"California","a price freeze") cause("California",E) motion(during(E),"a price freeze") not(path_rel(start(E),"a price freeze","its private utility firms",ch_of_loc,"on")) path_rel(end(E),"a price freeze","its private utility firms",ch_of_loc,"on")

(102)

"Mr. Guzman Cabrera put in more than 40 years at Pemex before being ***pushed*** *into retirement by La Quina."*
verb: *push*, VerbNet Class: FORCE-59
force(during(E),"La Quina",?Undergoer,"retirement") cause("La Quina",E) motion(during(E),?Undergoer) not(path_rel(start(E),?Undergoer,"retirement",ch_of_loc,"into")) path_rel(end(E),?Undergoer,"retirement",ch_of_loc,"into")

A better approach would be to appropriately assign the *ch_of_state* LocType constants for instances like this. In the long run, we would need an automatic system that can separate the abstract usages from the concrete usages. This would require further investigation into the pilot study on concreteness annotation discussed in Section 3.5. For the process of caused motion predicate implementation VerbNet, there will be a need to consistently assign the proper LocType constants to the frames we update in VerbNet.

Is the `motion()` predicate appropriate for representations in 101 and 102? As discussed in Chapter 5, the `motion()` predicate's use here is consistent with its current usage in VerbNet. Most of the current VerbNet classes do not make the distinction between the physical and abstract usages such as those seen in examples 101 and 102. Even the directed motion classes such as *Escape-51.1*, which includes member verbs like *come*, *go*, *escape* or *arrive*, do not make a distinction between the spatial motion from one physical location to another (e.g. *The convicts escaped the prison*) and the abstract or metaphorical motion from one state to another (e.g. *He never escaped from his fate*). Thus, rather than the `motion()` predicate, we need to focus on the proper assignment of the LocType constant that describes the nature of path of the `path_rel()` predicate.

This, however, brings up an issue related to the one in Section 3.3.3. As discussed in Section 5.3.1, the `motion()` predicate serves to represent the motion event of the *cause-displace-cx* and *cause-change-scale-cx* type CMCs. The other two CMC types, namely, *cause-transform-cx* and *cause-change-possession-cx* retain distinct predicates to represent the motion event. As discussed in Section 3.3.3, further investigation into the *cause-displacement-cx* type of CMCs would be needed to determine if we can identify further semantic categories that should take a different motion event predicate rather than the `motion()` predicate. For example, classes like *Enforce-63*, which includes verbs like *control*, *enforce*, and *impose* and is exemplified in 101, has no associated concrete usages like the *Force-59* class does in 102. Whether or not the `motion()` predicate makes sense with such classes is a topic of further discussion.

6.3.3.2 Semantic Role Conflicts

When the representations are joined, there are instances where the semantic roles from CMC do not line up with those of the VerbNet classes. We have two examples to illustrate the issue of semantic role conflict, the second one being a little more complicated than the first.

In example 103, we have the semantic predicates from the *Bump-18.4* class joining with the caused motion representation. The first four predicates come from the VerbNet class and the rest originate from the construction.

(103)

*"When a man **rammed** a bomb-laden car into a convoy of vehicles driven by Afghan employees of U.S. Protection and Investigations [...]"*
verb: *ram*, VerbNet Class: BUMP-18.4
manner(during(E),directedmotion,"a bomb-laden car") not(contact(during(E),"a bomb-laden car",**"a convoy of vehicles [...]"**)) manner(end(E),forceful,"a bomb-laden car") contact(end(E),"a bomb-laden car",**"a convoy of vehicles [...]"**) cause("a man",E) motion(during(E),"a bomb-laden car") not(path_rel(start(E),"a bomb-laden car",**"a convoy of vehicles [...]"**,ch_of_loc,"into")) path_rel(end(E),"a bomb-laden car",**"a convoy of vehicles [...]"**,ch_of_loc,"into")
manner(during(E),directedmotion,Undergoer) not(contact(during(E),Undergoer,**Location**)) manner(end(E),forceful,Undergoer) contact(end(E),Undergoer,**Location**) cause(Actor,E) motion(during(E),Undergoer) not(path_rel(start(E),Undergoer,**Goal**,ch_of_loc,prep)) path_rel(end(E),Undergoer,**Goal**,ch_of_loc,prep)

For sentences like *"The grocery cart hit against the wall"*, *Bump-18.4* has chosen to represent the prepositional phrase *"against the wall"* with the role of LOCATION rather than interpreting it as an instance of GOAL. The joining of the representations would require the recognition of such differences and would need to map the arguments to the appropriate roles on both the VerbNet and the construction side of the predicates. In reference to this particular *Bump-18.4* class, this may not be a conflict once VerbNet receives the necessary updates for VerbNet classes that include the semantics of path of motion. Other classes within the 18 class series like *Hit-18.1* and *Hit-18.2* have been slated to be evaluated for explicit inclusion of motion constructions (see Table D.7). Along with the other 18 class series, it is likely that during the second stage of the predicate implemention, *Bump-18.4* will be re-evaluated for the semantics of path of motion with the LOCATION role re-conceived as GOAL or DESTINATION roles.

Example 105 shows an example of a slightly more complicated case of semantic role mismatch. Unlike the previous example where we have two semantic roles referring to the same argument, here we have a case of similarly named semantic roles referring to different participants in the event. Consider the following sentence and the semantic roles that describe it.

(104) [It]-CAUSE freed [him]-SOURCE [of guilt]-THEME.

In the *Free-80* class the inconvenience or difficulty from which one is freed of is considered the THEME of the event and the entity that is made free of the inconvenience is assigned the SOURCE role. If we rephrase this as a caused-motion event, the inconvenience that is undergoing the motion (usually the UNDERGOER role) takes the SOURCE role and the location from which the the inconvenience leaves (usually the SOURCE information) is the THEME of the event for the *Free-80* class.

In the example below, a convicted murderer being released from the prison is likened to the prison being made free of the convicted murderer by the VerbNet side of the representation. The joining of the representation thus requires for the SOURCE role from the construction to be mapped to the THEME of the class, and the UNDERGOER role from the construction to be mapped to the SOURCE role in the class.

(105)

"A convicted murderer is ***released*** *from prison before serving his full sentence and then a nightmare becomes reality."*
verb: *release*, VerbNet Class: FREE-80
cause(?Actor,E) not(free(start(E),"prison","A convicted murderer")) free(end(E),"prison","A convicted murderer") path_rel(start(E),"A convicted murderer","prison",ch_of_loc, from) not(path_rel(end(E),"A convicted murderer","prison",ch_of_loc, from))
cause(?Actor,E) not(free(start(E),Theme,Source)) free(end(E),Theme,Source) path_rel(start(E),Undergoer,Source,ch_of_loc,prep) not(path_rel(end(E),Undergoer,Source,ch_of_loc,prep))
Role Mapping: Free-80.Theme = cause-displace-cx.Source Free-80.Source = cause-displace-cx.Undergoer

Sentences 103 and 105 exemplify issues of semantic role conflict that stems from a differing interpretations of the thematic relationship the arguments have with respect to the verb. In particular, what we see in example 105, is a result of the coercion: a verb from the *Free-80* class receives a set of participant roles from the constructional. Because semantics of caused motion differ from the semantics expressed in the class, the participant roles as used by the two events are at odds with each other. Consequently, we must resolve the differences in the roles by mapping the constructional roles to the appropriate counterparts in the class. While we do not expect semantic roles to conflict for all coercive usages, we do not expect this type of semantic role conflict to be an uncommon phenomenon. The trick of producing the right representation will come down to making the appropriate role mappings for the construction and the VerbNet classes.

6.4 Final Considerations and Future Work

In this chapter, we presented our preliminary work into the representation of coerced instances of CMCs and demonstrate that such a representation is capable of delivering semantic information that would not be available given the semantics of the verb alone. In fact, a CMC representation is possible for those instances whose verb has not been identified with a VerbNet label: so long as we can identify a sentence as a CMC, the relevant semantic information will be available via the constructional layer of description.

In terms of future work, we expect that once the VerbNet classes have been updated with the cleaner and more explicit predicate representations for path of motion, that make a distinction between a static location and a position in a path of motion, we will be in a better position to carry out a more definitive work on sentence representation. Moreover, once the VerbNet updates are complete, we can revisit our efforts to categorize VerbNet classes as coercive or conventional so that we can reassess our decision to fully rely on the existence of CMC frames in individual VerbNet classes when deciding whether or not a VerbNet class represents a case of CMC coercion.

.

Chapter 7

Conclusion

7.1 Summary

The work presented in this thesis is a step towards the larger goals of facilitating semantic representation for atypical usages of verbs. We have drawn from the theories of construction grammar that proposes that semantics can be interpreted at the syntactic level. We adapt the theories and definitions proposed by construction grammar to identify and represent caused motion constructions in corpus data. Through these efforts we seek to make the interpretation of sentences of caused motion possible even when the verb used in the sentence is atypical for the CMC in which it appears.

In the **first stage** of this thesis, we carried out a detailed analysis and annotation of CMCs. We defined a typological classification of a CMC's characteristics, and used it to create annotation guidelines for the corpus annotation of CMCs. In the **second stage** of this thesis, we utilized the annotated data from the first stage to train a model for the automatic classification of CMC and to identify the DISPLACE categories of CMCs. In the **third stage** of this thesis, we formulated new VerbNet predicates for the semantic representation of CMCs that will improve the overall representation of verbs that denote the semantics of path of motion. In the **fourth stage** of this thesis, we utilized the results from the automatic classification and the newly established VerbNet predicates for CMCs to give semantic representation to coercive instances of CMCs, which would not have a proper semantic representation from the VerbNet class alone.

7.2 Contributions

The major contributions of this thesis are as follows:

- **Semantic Type Classification of CMCs** This thesis provides definitions and semantic type classification of CMCs based on the corpus data. The definitions expand on the work done by Goldberg (1995), further detailing and refining the characteristics that define a CMC. The outcome of this effort is the semantic typing of CMCs and the development of the categorical types of CMCs as they occur in corpus data.

- **Annotated CMC Corpus** This thesis provides 20,809 instances of annotated corpus data. The annotated instances are marked with one of the CMC semantic type labels or as an instance of non-CMC.

- **Models for CMC Automatic Classification** This thesis demonstrates that CMCs can be reliably identified in the corpus data. The semantic and syntactic features used in the studies also show that they are scalable across the different corpora.

- **New VerbNet Representation for CMCs** This thesis includes new VerbNet predicates for the semantic representations of CMCs. The predicates are able to give consistent and explicit representation of the semantics of CMCs.

- **Improved VerbNet Representation for Path of Motion** The semantic predicates for CMC have been formulated so that they can be applicable to all VerbNet classes that include paths of motion in their semantics. Updates to VerbNet will ensure a consistent and systematic representation of the path of motion.

- **Preliminary Results for Sentence Representation** This thesis demonstrates that CMC representation can help give the proper semantic representation to sentences even when the verb in the sentence does not include the semantics of CMC.

The minor contributions of this thesis are as follows:

- **Concreteness Annotation** This thesis provides 2250 instances of annotated corpus data with concreteness labels. This small corpus is suitable for further study and evaluation.

- **Coercion Labelling** This thesis provides a coercion evaluation for all classes in VerbNet. This evaluation is independent of whether or not the VerbNet class already has a CMC frame. While this list was not integrated into the studies in this thesis, it could serve as a basis of further investigation into CMC's interactions with the VerbNet classes.

Aside from these specific contributions, we believe that the overall contribution of this work is the establishment of the processes involved in identifying and representing constructions in an empirical setting. This work assesses the necessary steps to define and annotate constructions in a corpus setting, train classifiers for constructions, and represent the semantics of constructions through VerbNet predicates. In this study, we provide a framework for implementing similar studies with other constructions and give a general expectations of the challenges that may be involved in such a process. While we have focused on the identification and representation of caused motion constructions, a similar corpus-driven study can be conducted for other constructions whose sentence representations would not be possible with the semantics of the verb alone.

7.3 Future Work

On the Definition and Semantic Typing of CMCs We only recognize a few of the metaphorical distinctions in our current analysis. In particular, the *cause-displace-cx* category is still semantically coarse-grained, allowing for both the highly metaphorical and highly concrete. In the future, we would like to extend our analysis over the *cause-displace-cx* category. To do this we will extend the study of the concreteness ratings and investigate further approaches for refining the category.

On the Automatic Classification of CMCs Future studies in automatic classification include the further investigation of features to see if we can raise the classifier performance beyond what this study presents. Moreover, in this study, the syntactic features are obtained from dependency parses that are automatically generated from gold phrase parses. In the future, experiments should be run on automatically parsed trees. Furthermore, we will extend our studies to the detection of not only the DISPLACE label, but also the other CMC types such as *cause-transform-cx*, *cause-change-possession-cx* and *cause-change-scale-cx*.

On VerbNet Predicate Representations of CMCs In conjunction with the future work on the further analysis of the *cause-displace-cx* category, we will be investigating the usage of the `motion()` predicate to determine if there is a need for more semantically specific predicates. Moreover, future work on the semantic predicates will include an assessment of the current VerbNet predicates to investigate if there is a way to introduce a predicate hierarchy that will allow semantically similar predicates to be grouped together. The outcome of such an effort will lessen the complexities associated with any future additions or updates to VerbNet's semantics.

On the Sentence Representation of CMCs The current study was an exploratory investigation into representing coercive CMC instances using the semantics from the CMC definitions and the available VerbNet class definitions. Once the implementation of the semantic predicates in VerbNet (as applicable to current classes) is complete, we can carry out further experiments on the sentence representation of CMCs.

Finally, in the future, we expect to extend this study beyond caused motion constructions to other constructional types. Other potential construction types we would like to examine in the near future are adjectival resultatives (e.g. *Mary hammered the metal flat*), which will be related to the *cause-transform-cx* category of CMC from this study; ditransitive constructions (e.g. *John lent me a bicycle*), which will be related to the *cause-change-possession-cx* category of CMC; and conative constructions (e.g. *Brian wiped at the counter with a damp rag*).

Bibliography

Agirre, Eneko and Philip Edmonds. 2006. Introduction. In Agirre and Edmonds, eds., Word Sense Disambiguation: Algorithms and Applications. Springer.

Bauer, Danie and Owen Rambow. 2011. Increasing coverage of syntactic subcategorization patterns in framenet using verbnet. In Fifth IEEE International Conference on Semantic Computing (ICSC 2011). Stanford, CA.

Bikel, Daniel M., Richard Schwartz, and Ralph Weischedel. 1999. An algorithm that learns what's in a name. Machine Learning: Special Issue on NL Learning 34.

Boas, Hans C. 2003. A Constructional Approach to Resultatives.. CSLI Publications.

Bonial, Claire, Susan Windisch Brown, William Corvey, Martha Palmer, Volha Petukhova, and Harry Bunt. 2011a. An exploratory comparison of thematic roles in VerbNet and LIRICS. In The Sixth Joint ISO-ACL SIGSEM Workshop on Interoperable Semantic Annotation.

Bonial, Claire, Susan Windisch Brown, Jena D. Hwang, Christopher Parisien, Martha Palmer, and Suzanne Stevenson. 2011b. Incorporating coercive constructions into a verb lexicon. In Proceedings of the ACL 2011 Workshop on Relational Models of Semantics. Portland, Oregon.

Brown, Susan, Dmitriy Dligach, and Martha Palmer. 2011. Verbnet class assignment as a wsd task. In Proceedings of the Ninth International Conference on Computational Semantics, pages 85–94. Oxford, UK: Association for Computational Linguistics.

Carreras, Xavier and Lluis Marquez. 2005. Introduction to the conll-2005 shared task: Semantic role labeling. In Proceedings of CoNLL-2005.

Chang, Chih-Chung and Chih-Jen Lin. 2001. LIBSVM: A library for support vector machines. ACM Transactions on Intelligent Systems and Technology 2(3):27:1–27:27.

Choi, Jinho. 2012. Optimization of Natural Language Processing Components for Robustness and Scalability. Ph.D. dissertation, University of Colorado at Boulder, Boulder, Colorado.

Croft, William. 1991. Syntactic Categories and Grammatical Relations. University of Chicago Press.

Culicover, Peter W. and Ray Jackendoff. 2005. Simpler Syntax. Oxford University Press.

Das, Dipanjan and Noah A. Smith. 2009. Semi-supervised frame-semantic parsing for unknown predicates. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies. Portland, Oregon.

Dorr, Bonnie. 1997. Large-scale dictionary construction for foreign language tutoring and interlingual machine translation. Machine Learning 12(4):271–322.

Dougherty, Ray C. 1974. The syntax and semantics of ”each other” constructions. Foundations of Language 12(1):1–47.

Dowty, David. 1979. Word Meaning and Montague Grammar: The Semantics of Verbs and Times in Generative Semantics and in Montague's PTQ. Springer Science Business Media.

Dowty, David. 1989. On the semantic content of the notion 'thematic role'. Properties, Types and Meaning 2:69–130.

Dowty, David. 1991. Thematic proto-roles and argument selection. Language 67(3):547–619.

Duffield, Cecily Jill, Jena D. Hwang, Susan Windisch Brown, Dmitriy Dligach, Sarah E. Vieweg, Jenny Davis, and Martha Palmer. 2007. Criteria for the manual grouping of verb senses. In Proceedings of the Linguistic Annotation Workshop held in conjunction with ACL-2007. Prague, the Czech Republic.

Fellbaum, Cristiane, Joachim Grabowski, and Shari Landes. 1998. Performance and confidence in a semantic annotation task. In C. Fellbaum, ed., WordNet: An Electronic Database. The MIT Press.

Fillmore, Charles J. 1966. Toward a modern theory of case. Tech. rep., Ohio State University.

Fillmore, Charles J. 1986. Pragmatically controlled zero anaphora. Berkeley Linguistics Society 12:95–107.

Fillmore, Charles J. 1988. Regularity and idiomaticity in grammatical constructions: the case of ”let alone”. Language 64(3):501–538.

Fillmore, Charles J. and Miriam R. L. Petruck. 2003. Framenet glossary. International Journal of Lexicography 16(3):235–250.

Fried, Mirjam. http://www.constructiongrammar.org/. What is construction grammar?

Gerber, Matthew and Joyce Y. Chai. 2010. Beyond nombank: A study of implicit arguments for nominal predicates. In 48th Annual Meeting of the Association for Computational Linguistics.

Gerber, Matthew and Joyce Y. Chai. 2013. Semantic role labeling of implicit arguments for nominal predicates. Computational Linguistics .

Gildea, Daniel and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. Computational Linguistics 28(2):245–288.

Giuglea, Ana-Maria and Alessandro Moschitti. 2006. Shallow semantic parsing based on framenet, verbnet, and propbank. In Proceedings of the 17th European Conference on Artificial Intelligence. Riva del Garda, Italy.

Gleitman, Lila. 1990. The structural sources of verb meanings. Language Acqusition 1:3–55.

Goldberg, Adele E. 1995. Constructions: A Construction Grammar Approach to Argument Structure. University Of Chicago Press.

Goldberg, Adele E. 2005. Constructions at Work. New York: Oxford University Press.

Goldberg, Adele E. and Ray Jackendoff. 2004. The english resultative as a family of constructions. Language 80(3):532–568.

Gruber, Jeffrey S. 1962. Studies in Lexical Relations. Ph.D. dissertation, Massachusetts Institute of Technology.

Guildea, Dan and Martha Palmer. 2002. The necessity of parsing for predicate argument recognition. In 40th Meeting of the Association for Computational Linguistics. Philadelphia, PA.

Habash, Nizar, Bonnie Dorr, and David Traum. 2003. Hybrid natural language generation from lexical conceptual structures. Machine Translation 18:81–127.

Hensman, Svetlana and John Dunnion. 2004. Automatically building conceptual graphs using verbnet and wordnet. In Proceedings of the 2004 international symposium on Information and communication technologies (ISICT '04).

Hovy, Edward H., Mitch Marcus, Martha Palmer, Sameer Pradhan, Lance Ramshaw, and Ralph M. Weischedel. 2006. Ontonotes: The 90In Proceedings of HLT/NAACL-2006. New York, NY.

Hripcsak, George and Adam S Rothschild. 2005. Agreement, the f-measure, and reliability in information retrieval. Journal of American Medical Informatics Association 12(3):296–298.

Hwang, Jena D., Rodney D. Nielsen, and Martha Palmer. 2010. Towards a domain independent semantics: Enhancing semantic representation with construction grammar. In Proceedings of the NAACL HLT Workshop on Extracting and Using Constructions in Computational Linguistics, pages 1–8. Los Angeles, California: Association for Computational Linguistics.

Hwang, Jena D., Martha Palmer, and Annie Zaenen. 2013. Representing paths of motion in representing paths of motion in VerbNet. In T. H. King and V. de Paiva, eds., From Quirky Case to Representing Space. CSLI Online Publications.

Jackendoff, Ray. 1983. Semantics and Cognition. The MIT Press.

Jackendoff, Ray. 1990. Semantic Structures. Cambridge, Massachusetts: The MIT Press.

Kawahara, Daisuke and Martha Palmer. 2014. Single classifier approach for verb sense disambiguation based on generalized features. In Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14). Reykjavik, Iceland.

Kay, Paul and Charles J. Fillmore. 1999. Grammatical constructions and linguistic generalizations: the what's x doing y? construction. Language .

Kingsbury, Paul and Martha Palmer. 2003. Propbank: The next level of treebank. In Proceedings of Treebanks and Lexical Theories. Växjö, Sweden.

Kipper, Karin, Anna Korhonen, Neville Ryant, and Martha Palmer. 2006. A large-scale extension of VerbNet with novel verb classes. In Proceedings of EURALEX. Turin, Italy.

Kipper, Karin, Anna Korhonen, Neville Ryant, and Martha Palmer. 2008. A large-scale classification of English verbs. Language Resources and Evaluation Journal 42(1):21–40.

Korhonen, Anna. 2002. Semantically motivated subcategorization acquisition. In Proceedings of the ACL Workshop on Unsupervised Lexical Acquisition, pages 51–58. Philadelphia, PA.

Kubat, Miroslav and Stan Matwin. 1997. Addressing the curse of imbalanced training sets: One-sided selection. Proceedings of the Fourteenth International Conference on Machine Learning .

Lakoff, George. 1987. Women, Fire, and Dangerous Things: What Categories Reveal About the Mind. Chicago: University of Chicago Press.

Lakoff, George. 1993. The contemporary theory of metaphor. In A. Ortony, ed., Metaphor and Thought. Cambridge University Press, 2nd edn.

Lakoff, George, Jane Espenson, Adele E. Goldberg, and Alan Schwartz. 1991. Master metaphor list.

Lakoff, George and Mark Johnson. 1980. Metaphors We Live By. Chicago: University of Chicago Press.

Langacker, Ronald W. 1987. Foundations of Cognitive Grammar: Theoretical Prerequisites. Stanford, CA: Stanford University Press. Vol 1, 1987(Hardcover), 1999(Paperback).

Levin, Beth. 1993. English Verb Classes and Alternations: A Preliminary Investigation. University of Chicago Press.

Levin, Beth and Malka Rappaport-Hovav. 1988. Building verb meanings. In The Projection of Arguments: Lexical and Compositional Factors. CSLI Publications.

Levin, Beth and Malka Rappaport-Hovav. 1995. Unaccusativity: At the Syntax-Lexical Semantics Interface. Cambridge, MA: MIT Press.

Levin, Beth and Malka Rappaport-Hovav. 2005. Argument Realization. New York: Cambridge University Press.

Marcus, Mitchell, Grace Kim, Mary Ann Marcinkiewicz, Robert Macintyre, Ann Bies, Mark Ferguson, Karen Katz, and Britta Schasberger. 1994. The penn treebank: Annotating predicate argument structure. In ARPA Human Language Technology Workshop, pages 114–119.

Marcus, Mitchell P., Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of english: the penn treebank. Computational Linguistics 19.

Merlo, Paola and Suzanne Stevenson. 2002. Structure and frequency of lexical semantic classes: Consequences for learning. In Workshop on Quantitative Investigations in Theoretical Linguistics (QITL). Osnabruek, Germany.

Merlo, Paola and Lonneke van der Plas. 2009. Abstraction and generalisation in semantic role labels: PropBank, VerbNet or both? In Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP, pages 288–296. Suntec, Singapore.

Michaelis, Laura A. 2004. Type shifting in construction grammar: An integrated approach to aspectual coercion. Cognitive Linguistics 15:1–67.

Michaelis, Laura A. 2005. Entity and event coercion in a symbolic theory of syntax. In Construction Grammar(s): Cognitive Grounding and Theoretical Extensions, vol. 3 of Constructional Approaches to Language, pages 45–87. Amsterdam: Benjamins.

Michaelis, Laura A. 2012. Making a case of construction grammar. In H. Boas and I. Sag, eds., Sign-Based Construction Grammar, pages 31–69. CSLI Publications.

Miller, George A. 1995. Wordnet: a lexical database for english. Communications of the ACM 38(11):39–41.

Ng, Hwee Tou and Hian Beng Lee. 1996. Integrating multiple knowledge sources to disambiguate word sense: An exemplar-based approach. In Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics, pages 40–47. Santa Cruz, California.

Palmer, Martha. 1985. Driving Semantics for Limited Domains. Ph.D. dissertation, University of Edinburgh.

Palmer, Martha. 2009. Semlink: Linking PropBank, VerbNet and FrameNet. In Proceedings of the Generative Lexicon Conference (GenLex-09). Pisa, Italy.

Palmer, Martha, Deborah Dahl, Rebecca Passonneau, Lynette Hirschman, Marcia Linebarger, and John Dowding. 1986. Recovering implicit information. In Proceedings of the Sixth Annual Meeting of the Association for Computational Linguistics. New York.

Palmer, Martha, Dan Guildea, and Paul Kingsbury. 2005. The Proposition Bank: An annotated corpus of semantic roles. Computational Linguistics 31(1):71–105.

Pinker, Steven. 1989. Learnability and Cognition. Cambridge, Massachusetts. London, England.: The MIT Press.

Pustejovsky, James. 1993. Type coercion and lexical selection. In J. Pustejovsky, ed., Semantics and the Lexicon. Dordrecht, The Netherlands: Kluwer Academic Publishers.

Pustejovsky, James and Elisabetta Jezek. 2008. Semantic coercion in language: Beyond distributional analysis. Italian Journal of Linguistics .

Rappaport-Hovav, Malka and Beth Levin. 2008. The English dative alternation: The case for verb sensitivity. Journal of Linguistics 44(01):129–167.

Reddy, Michael. 1979. The conduit metaphor: a case of frame conflict in our language about language. In A. Ortony, ed., Metaphor and Thought. Cambridge: Cambridge University Press.

Sag, Ivan A., Hans C. Boas, and Paul Kay. 2012. Introducing sign-based construction grammar. In H. C. Boas and I. A. Sag, eds., Sign-Based Construction Grammar. Stanford, CA: CSLI Publications.

Schneider, Nathan. 2014. Lexical Semantic Analysis in Natural Language Text. Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA.

Shi, Lei and Rada Mihalcea. 2005. Putting pieces together: Combining framenet, verbnet and wordnet for robust semantic parsing. In Computational Linguistics and Intelligent Text Processing; Sixth International Conference (CICLing 2005), pages 100–111.

Slobin, Dan. 2005. Linguistic representations of motion events: What is signifier and what is signified? In . W. H. C. Maeder, O. Fischer, ed., Iconicity Inside Out: Iconicity in Language and Literature 4. Amsterdam/Philadelphia: John Benjamins.

Srikumar, Vivek. 2013. Semantics of Role Labeling. Ph.D. dissertation, University of Illinois at Urbana-Champaign, Urbana, IL.

Subba, Rajen and Barbara Di Eugenio. 2009. An effective discourse parser that uses rich linguistic information. In Proceedings of HLT/NAACL-2009, pages 566–574.

Swier, Robert S. and Suzanne Stevenson. 2004. Unsupervised semantic role labelling. In Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing (EMNLP 2004), pages 95–102. Barcelona, Spain.

Talmy, Leonard. 1976. Semantic causative types. In , Syntax and Semantics 6: The Grammar of Causative Constructions, pages 43–116. New York: Academic Press.

Talmy, Leonard. 1991. Path to realization: A typology of event conflation. In Proceedings of the Seventeenth Annual Meeting of the Berkeley Linguistics Society: General Session and Parasession on The Grammar of Event Structure, pages 480–519.

Talmy, Leonard. 2000. Toward a Cognitive Semantics, vol. 1. The MIT Press.

Taylor, Ann. 1996. English Penn Treebank Guidelines.

Taylor, Ann, Mitchell Marcus, and Beatrice Santorini. 2003. The penn treebank: An overview. In A. Abeillé, ed., Building and Using Parsed Corpora, vol. 20 of Text, Speech and Language Technology. Springer.

Weischedel, Ralph, Sameer Pradhan, Lance Ramshaw, Jeff Kaufman, Michelle Franchini, Mohammed El-Bachouti, Nianwen Xue, Martha Palmer, Jena D. Hwang, Claire Bonial, Jinho Choi, Aous Mansouri, Maha Foster, Abdel aati Hawwary, Mitchell Marcus, Ann Taylor, Eduard Hovy, Robert Belvin, and Ann Houston. 2012. OntoNotes Release 4.99.

Weiss, Gary M. and Foster Provost. 2001. The effect of class distribution on classifier learning. Technical report, Rutgers University.

Yeh, Alexander. 2000. More accurate tests for the statistical significance of result differences. In Proceedings of the 18th Conference on Computational Linguistics, vol. 2.

Zaenen, Annie, Danny Bobrow, Cleo Condoravdi, and Raphael Hoffmann. 2010. Supporting rule-based representations with corpus-derived lexical information. In Proceedings of HLT/NAACL-2010 Workshop on Machine Reading..

Zaenen, Annie, Cleo Condoravdi, and Danny Bobrow. 2008. The encoding of lexical implications in VerbNet predicates of change of locations. In Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08). Marrakech, Morocco: European Language Resources Association (ELRA).

Appendix A

Finer Grained CMC Categories Discussion

Earlier in the study, we made an attempt to subdivide this category by a finer set of semantic divisions. Some of the categories we considered include (but are not limited to):

- **Incorporation:** This category was considered for instances where an object is made a member of a pre-existing group.
  - ∗ A capital-gains tax cut would be *included* in the final bill.
  - ∗ It no longer [...] *lump* the bikes into the same category as motorcycle.
- **Seduce or Force:** This category was considered for instances where a sentient entity is caused to do something, unwillingly or unwittingly.
  - ∗ The ”great disorder” that reigned at the agency **led** her into temptation.
  - ∗ They [...] **coax** investors into shifting some of their hoard into the stock market.
- **Spread:** This label was considered for sentences that invoke a image schema where the trajector is a mass that spreads out from a landmark (c.f. Lakoff's (1987) study on the preposition *out*).
  - ∗ The average household will **spread** 19 accounts over a dozen financial institutions.
  - ∗ They **distributed** the revisions [...] to U.S. attorneys around the country.

There were four major problems with such categorizations. First, there was an overlap between the semantics we were trying to disambiguate and the semantics projected by the lexical items. This was the case with the *Seduce or Force* category, and, to a lesser extent, the *Incorporation* category. This semantic information is already available through VerbNet, and there was no pressing need to make these distinctions at a constructional level. Second, in certain other cases the distinctions cut across the metaphorical boundaries we had created. While distinctions on image schema are an interesting idea to pursue, being faithful to the semantics would mean we would have to reach into other categories such as the *cause-change-possession-cx* category (e.g. the sentence *They distributed revisions to US attorney* is currently classified under *cause-change-possession-cx*). Thirdly, the task would quickly become intractable for the purpose of this thesis. There was neither sufficient time nor resources to devote to the proper subclassification of this category.[1] Last is perhaps the most compelling reason: we realized that the type of subcategorizations we were attempting were simply not capturing other metaphorical extensions available in this category. In the spirit of the other categories, it would be fitting for any new categories made to be metaphorically coherent.

[1] There really is enough material for this analysis to constitute a large chunk of a separate and independent thesis or paper.

Appendix B

Part of Speech Tags in Treebank

Following are tags Part of Speech tags used in various treebanks for English (Marcus et al., 1993; Weischedel et al., 2012).

Tag	Description	Tag	Description
CC	Coordinating conjunction	PRP$	Possessive pronoun
CD	Cardinal number	RB	Adverb
DT	Determiner	RBR	Adverb, comparative
EX	Existential there	RBS	Adverb, superlative
FW	Foreign word	RP	Particle
IN	Preposition or subordinating conjunction	SYM	Symbol
JJ	Adjective	TO	to
JJR	Adjective, comparative	UH	Interjection
JJS	Adjective, superlative	VB	Verb, base form
LS	List item marker	VBD	Verb, past tense
MD	Modal	VBG	Verb, gerund or present participle
NN	Noun, singular or mass	VBN	Verb, past participle
NNS	Noun, plural	VBP	Verb, non-3rd person singular present
NNP	Proper noun, singular	VBZ	Verb, 3rd person singular present
NNPS	Proper noun, plural	WDT	Wh-determiner
PDT	Predeterminer	WP	Wh-pronoun
POS	Possessive ending	WP$	Possessive wh-pronoun
PRP	Personal pronoun	WRB	Wh-adverb

Table B.1: Constituent Treebank Tags

Appendix C

VerbNet Role Hierarchy

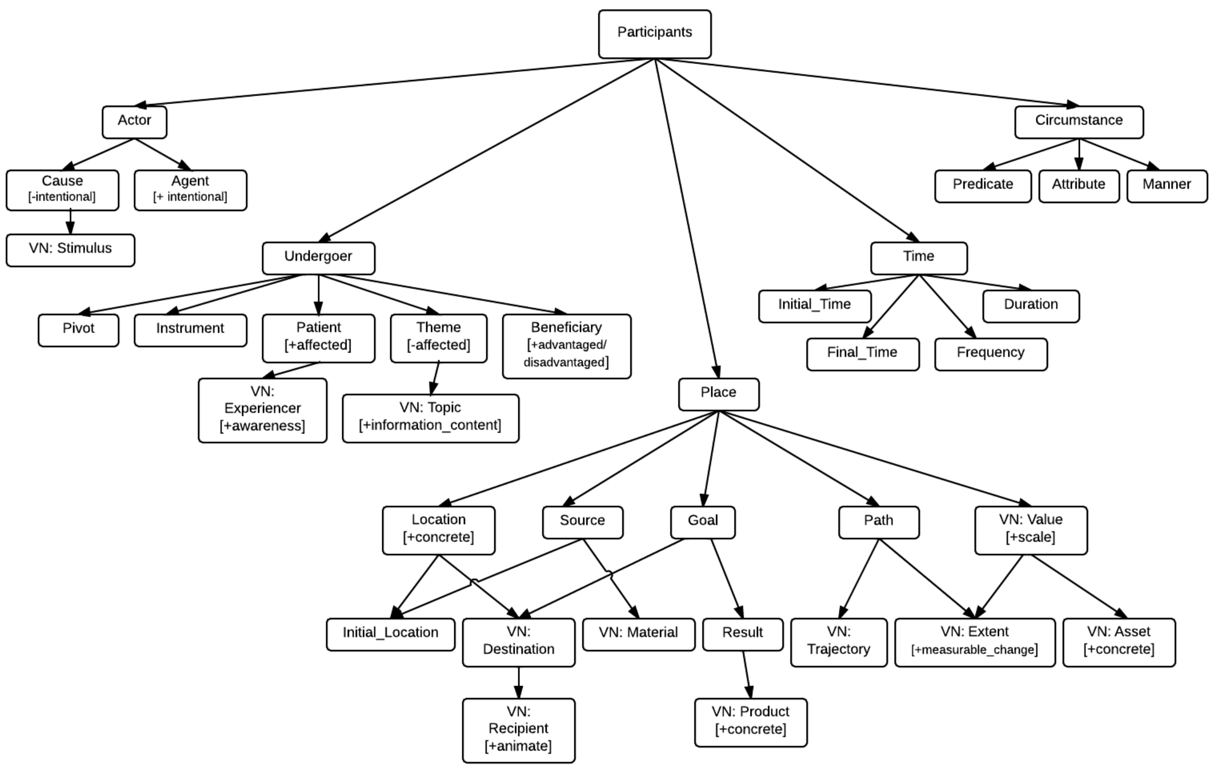

Appendix D

VerbNet Predicate Changes

+Motion Event; +Path of Motion; +/- *cause-displacement-cx*

VerbNet Class Name	**Class Number**	**`motion()` existing or new?**	**Allows cause-displace-cx?**	**Suggested Alternate?**
put-9.1	9.1	existing	Y	
put_spatial-9.2	9.2	existing	Y	
funnel-9.3	9.3	existing	Y	
put_direction-9.4	9.4	existing	Y	
pour-9.5	9.5	existing	Y	
coil-9.6	9.6	existing	Y	
spray-9.7	9.7	existing	Y	
fill-9.8	9.8	existing	Y	
butter-9.9	9.9	existing	Y	
pocket-9.10	9.10	existing	Y	
remove-10.1	10.1	new	Y	
fire-10.10	10.10	new		employment
resign-10.11	10.11	new		employment
banish-10.2	10.2	new	Y	
clear-10.3	10.3	new	Y	
wipe_manner-10.4.1	10.4.1	new	Y	
wipe_instr-10.4.2	10.4.2	new	Y	
pit-10.7	10.7	new	Y	
debone-10.8	10.8	new	Y	
mine-10.9	10.9	new	Y	

Table D.1: VerbNet classes expressing motion and path. (Part 1)

VerbNet Class Name	**Class Number**	`motion()` **existing or new?**	**Allows cause-displace-cx?**	**Suggested Alternate?**
send-11.1	11.1	existing	Y	
slide-11.2	11.2	existing	Y	
bring-11.3	11.3	existing	Y	
carry-11.4	11.4	existing	Y	
drive-11.5	11.5	existing	Y	
push-12	12	existing	Y	
hire-13.5.3	13.5.3	new		employment
concealment-16	16	new	Y	
throw-17.1	17.1	existing	Y	
pelt-17.2	17.2	existing	Y	
eat-39.1	39.1	new		remove frame
appear-48.1.1	48.1.1	new	Y	
disappearance-48.2	48.2	new	Y	
escape-51.1	51.1	existing	Y	
leave-51.2	51.2	existing	not possible	
roll-51.3.1	51.3.1	existing	Y	
run-51.3.2	51.3.2	existing	Y	
vehicle-51.4.1	51.4.1	existing	Y	
nonvehicle-51.4.2	51.4.2	existing	Y	
vehicle_path-51.4.3	51.4.3	existing	not possible	
waltz-51.5	51.5	existing	Y	
chase-51.6	51.6	existing	Y	
accompany-51.7	51.7	existing	Y	
reach-51.8	51.8	existing	not possible	
confine-92	92	new	Y	

Table D.2: VerbNet classes expressing motion and path. (Part 2)

+Motion Event; +Path of Motion; +/- *cause-change-possession-cx*

VerbNet Class Name	**Class Number**	**Information or possession?**	**`transfer()` existing or new?**	**Allows cause-change-poss-cx?**
steal-10.5	10.5	possession	new	Y
cheat-10.6	10.6	possession	new	Y
give-13.1	13.1	possession	existing	Y
contribute-13.2	13.2	possession	existing	Y
future_having-13.3	13.3	possession	new	not possible
fulfilling-13.4.1	13.4.1	possession	existing	Y
equip-13.4.2	13.4.2	possession	existing	not possible
get-13.5.1	13.5.1	possession	existing	Y
obtain-13.5.2	13.5.2	possession	existing	Y
invest-13.5.3	13.5.3	possession	existing	Y
exchange-13.6	13.6	possession	existing	not possible
berry-13.7	13.7	possession	existing	not possible
learn-14	14	information	new	Y
confess-37.10	37.1	information	new	Y
transfer_mesg-37.1.1	37.1.1	information	new	Y
inquire-37.1.2	37.1.2	information	new	not possible
interrogate-37.1.3	37.1.3	information	new	not possible
lecture-37.11	37.11	information	new	not possible
promise-37.13	37.13	information	new	not possible
tell-37.2	37.2	information	new	Y
manner_speaking-37.3	37.3	information	new	Y
instr_communication-37.4.1	37.4.1	information	new	Y
talk-37.5	37.5	information	new	Y
chit_chat-37.6	37.6	information	new	not possible
say-37.7	37.7	information	new	Y
complain-37.8	37.8	information	new	not possible
advise-37.9	37.9	information	new	not possible
nonverbal_expression-40.2	40.2	information	new	Y
wink-40.3.1	40.3.1	information	new	Y
crane-40.3.2	40.3.2	information	new	Y
curtsey-40.3.3	40.3.3	information	new	Y
pay-68	68	possession	existing	Y

Table D.3: VerbNet classes expressing transfer and path.

+Motion Event; +Path of Motion;
+/- *cause-change-scale-cx* or +/- *cause-transform-cx*

VerbNet Class Name	Class Number	event predicate for cause-transform-cx	event predicate for cause-ch-scale-cx
coloring-24	24	class specific	n/a
preparing-26.3	26.3	transform()	n/a
knead-26.5	26.5	transform()	n/a
turn-26.6.1	26.6.1	transform()	n/a
render-29.90	29.90	transform()	n/a
suffocate-40.7	40.7	class specific	n/a
poison-42.2	42.2	class specific	n/a
break-45.1	45.1	class specific	n/a
bend-45.2	45.2	class specific	n/a
cooking-45.3	45.3	class specific	n/a
other_cos-45.4	45.4	transform()	n/a
entity_specific_cos-45.5	45.5	transform()	change_value()
calibratable_cos-45.6	45.6	n/a	change_value()

Table D.4: VerbNet classes expressing transformation or scale change and path.

-Motion Event; -Path of Motion

VerbNet Class Name	Class Number	Remove `motion()`
keep-15.2	15.2	
animal_sounds-38	38	
light_emission-43.1	43.1	
sound_emission-43.2	43.2	
smell_emission-43.3	43.3	
substance_emission-43.4	43.4	
lodge-46	46	
entity_specific_modes_being-47.2	47.2	
modes_of_being_with_motion-47.3	47.3	Y
sound_existence-47.4	47.4	
swarm-47.5.1	47.5.1	Y
spatial_configuration-47.6	47.6	
assuming_position-50	50	Y
linger-53.1	53.1	
weekend-56	56	
own-100.1	100.1	

Table D.5: VerbNet classes expressing no motion.

Fictive Motion: -Motion Event; +Path of Motion

VerbNet Class Name	**Class Number**
continue-55.3	55.3
meander-47.7	47.7
terminus-47.9	47.9
image_impression-25.1	25.1
scribble-25.2	25.2
illustrate-25.3	25.3
transcribe-25.4	25.4

Table D.6: VerbNet classes expressing fictive motion.

VerbNet Class Name	**Class Series**	**new?**	**Motion Event Type(s)**	**Current Roles**		
				cause	**theme**	**path info**
hit-18.1	18		transform	Agent	Patient	Result
swat-18.2						
spank-18.3						
poke-19	19	x	motion	Agent	Instrument	Patient
cut-21.1	21		motion	Agent	Patient	Source
mix-22.1	22	x	motion	Agent	Patient	Co-Patient
shake-22.3		x				
tape-22.4						
separate-23.1	23	x	motion	Agent	Patient	Co-Patient
split-23.2		x				
build-26.1	26	x	transform	Agent	Material	Product
grow-26.2		x				
convert-26.6.2		x				
classify-29.10	29	x	motion	Agent	Theme	Goal
stimulus_subject-30.4	30		*none*			
initiate_comm.-37.4.2	37		*none*			
ferret-35.6	35	x	transfer	Agent	Theme	Source
feeding-39.7	39	x	transfer	Agent	Theme	Recipient
breathe-40.1.2	40		motion	Agent	Theme	Destination
remedy-45.7	45		*none*			
herd-47.5.2	47	x	motion	Agent	Theme	
force-59	59	x	motion	Agent	Patient	Result
enforce-63	63	x	motion	Agent	Theme	
deduce-97.2	97	x	transfer?	Agent	Theme	Source
involve-107	107	x	motion	Agent	Theme	Goal
become-109.1	109		*none*			

Table D.7: Remaining VerbNet classes for analysis.

Appendix E

VerbNet Classes and Coercion

The following is a list of 182 classes that have been judged as either conventional (V) or coerced (C) for the relevant caused motion usage. Each column represents a caused motion type. Empty cells mean that the construction is not applicable for the given VerbNet class. VerbNet classes that do not allow or are not coercible into any of the CMC types are not listed.

		cause-displace-cx	cause-transform-cx	cause-change-possess-cx	
				possesion	information
put-9.1	9.1	V			
put_spatial-9.2	9.2	V			
funnel-9.3	9.3	V			
put_direction-9.4	9.4	V			
pour-9.5	9.5	V			
coil-9.6	9.6	V			
spray-9.7	9.7	V			
fill-9.8	9.8	V			
butter-9.9	9.9	C			
pocket-9.10	9.10	V			
remove-10.1	10.1	V			
banish-10.2	10.2	V			
clear-10.3	10.3	V			
wipe_manner-10.4.1	10.4.1	V			
wipe_instr-10.4.2	10.4.2	V			
steal-10.5	10.5	V		V	
cheat-10.6	10.6	V		V	
pit-10.7	10.7	C			
debone-10.8	10.8	C			
mine-10.9	10.9	V			
fire-10.10	10.10	V			
resign-10.11	10.11	C			

		cause-displace-cx	cause-transform-cx	cause-change-possess-cx	
				possesion	information
send-11.1	11.1	V			
slide-11.2	11.2	V			
bring-11.3	11.3	V			
carry-11.4	11.4	V			
drive-11.5	11.5	V			
push-12	12	V			
give-13.1	13.1			V	
contribute-13.2	13.2			V	
fulfilling-13.4.1	13.4.1			V	
get-13.5.1	13.5.1			V	
obtain-13.5.2	13.5.2			V	
hire-13.5.3	13.5.3	C		V	
throw-17.1	17.1	V			
pelt-17.2	17.2	C	C		
hit-18.1	18.1	C	V		
swat-18.2	18.2	C	V		
spank-18.3	18.3	C	V		
bump-18.4	18.4	C			
poke-19	19	V			
cut-21.1	21.1	V	V		
carve-21.2	21.2	C	V		
mix-22.1	22.1	V			
amalgamate-22.2	22.2	V	V		
shake-22.3	22.3	V	V		
tape-22.4	22.4	V			
separate-23.1	23.1		C		
split-23.2	23.2		V		
disassemble-23.3	23.3		V		
coloring-24	24		C		
image_impression-25.1	25.1	V			
transcribe-25.4	25.4	V			
build-26.1	26.1		V		
grow-26.2	26.2		V		
preparing-26.3	26.3		C		
knead-26.5	26.5		V		
turn-26.6.1	26.6.1		V		
convert-26.6.2	26.6.2		V		
performance-26.7	26.7			V	
rehearse-26.8	26.8		C		

		cause-displace-cx	cause-transform-cx	cause-change-possess-cx	
				possesion	information
adjust-26.9	26.9	C			
declare-29.4	29.4				C
captain-29.8	29.8	C			
classify-29.10	29.10	V			
peer-30.3	30.3	C			
amuse-31.1	31.1	C			
admire-31.2	31.2	C			
marvel-31.3	31.3	C			
appeal-31.4	31.4	C			
want-32.1	32.1	C			
long-32.2	32.2	C			
judgement-33	33	C			
assessment-34.1	34.1	C			
hunt-35.1	35.1	C			
search-35.2	35.2	C			
stalk-35.3	35.3	C			
rummage-35.5	35.5	C			
ferret-35.6	35.6	V			
correspond-36.1	36.1	C			
marry-36.2	36.2	C			
meet-36.3	36.3	C			
battle-36.4	36.4	C			
transfer_mesg-37.1.1	37.1.1				V
tell-37.2	37.2				V
manner_speaking-37.3	37.3				V
instr_communication-37.4	37.4				V
talk-37.5	37.5	C			
chit_chat-37.6	37.6	C			
say-37.7	37.7				V
complain-37.8	37.8	C			V
advise-37.9	37.9	C			V
confess-37.10	37.10				V
overstate-37.12	37.12	C			C
animal_sounds-38	38	C			
eat-39.1	39.1	C			
chew-39.2	39.2	C			
gobble-39.3	39.3	C			
devour-39.4	39.4	C			
dine-39.5	39.5	C			
feeding-39.7	39.7			V	

		cause-displace-cx	cause-transform-cx	cause-change-possess-cx	
				possesion	information
hiccup-40.1.1	40.1.1	C		C	
breathe-40.1.2	40.1.2	V			C
exhale-40.1.3	40.1.3	V		C	
nonverbal_expression-40.2	40.2	C		C	
wink-40.3.1	40.3.1	C		C	
crane-40.3.2	40.3.2	C			V
curtsey-40.3.3	40.3.3			C	
snooze-40.4	40.4		C		
flinch-40.5	40.5	C		C	
body_internal_states-40.6	40.6	C			
suffocate-40.7	40.7		V		
pain-40.8.1	40.8.1	C			
tingle-40.8.2	40.8.2	C			
hurt-40.8.3	40.8.3	C	C		
groom-41.1.2	41.1.2		C		
floss-41.2.1	41.2.1	C	C		
braid-41.2.2	41.2.2	C			
simple_dressing-41.3.1	41.3.1		C		
dressing_well-41.3.2	41.3.2		C		
murder-42.1	42.1		C		
poison-42.2	42.2		C		
subjugate-42.3	42.3	C	C		
light_emission-43.1	43.1	C			
sound_emission-43.2	43.2	C			
smell_emission-43.3	43.3	C			
substance_emission-43.4	43.4	C			
destroy-44	44		C		
break-45.1	45.1	C	V		
bend-45.2	45.2	C	V		
cooking-45.3	45.3		C		
other_cos-45.4	45.4		C		
entity_specific_cos-45.5	45.5		C		
calibratable_cos-45.6	45.6		C		
entity_specific_modes_being-47.2	47.2	C			
modes_of_being_with_motion-47.3	47.3	V			
sound_existence-47.4	47.4	C			
swarm-47.5.1	47.5.1	C			
herd-47.5.2	47.5.2	V			
bulge-47.5.3	47.5.3	C			
spatial_configuration-47.6	47.6	C			
meander-47.7	47.7	C			

		cause-displace-cx	cause-transform-cx	cause-change-possess-cx	
				possesion	information
appear-48.1.1	48.1.1	C			
reflexive_appearance-48.1.2	48.1.2	C		V	
disappearance-48.2	48.2	C			
body_internal_motion-49	49	C			
assuming_position-50	50	C			
escape-51.1	51.1	C			
roll-51.3.1	51.3.1	V			
run-51.3.2	51.3.2	V			
vehicle-51.4.1	51.4.1	V			
nonvehicle-51.4.2	51.4.2	V			
waltz-51.5	51.5	C			
chase-51.6	51.6	V			
accompany-51.7	51.7	V			
rush-53.2	53.2	C			
fit-54.3	54.3	V			
bill-54.5	54.5	C			
establish-55.5	55.5	C			
weather-57	57	C	C		
urge-58.1	58.1	C			
force-59	59	C			
enforce-63	63	V			
admit-65	65	V			
pay-68	68			V	
help-72	72	C			
cooperate-73	73	C			
neglect-75	75		C		
withdraw-82	82	V			
focus-87.1	87.1	C			
settle-89	89	C			
adopt-93	93	C			
risk-94	94	C			
deduce-97.2	97.2	V			
confront-98	98	C			
patent-101	101	C			
promote-102	102	V			
require-103	103	C			
involve-107	107	V			
multiply-108	108	C			

Appendix F

Caused Motion Construction Definitions

cause-displacement-cx

- **Semantic Roles:**
  ACTOR, UNDERGOER, SOURCE, PATH, GOAL

- **Semantic Predicates:**
  `cause(`ACTOR, E)
  `motion(during`(E), UNDERGOER)
  `path_rel(start`(E), UNDERGOER, SOURCE, *ch_of_loc*,prep)
  `path_rel(during`(E), UNDERGOER, PATH, *ch_of_loc*,prep)
  `path_rel(end`(E), UNDERGOER, GOAL, *ch_of_loc*,prep)

- **Member VN Classes:**

accompany-51.7	drive-11.5	put-9.1
appear-48.1.1	escape-51.1	remove-10.1
banish-10.2	fill-9.8	roll-51.3.1
bring-11.3	funnel-9.3	run-51.3.2
butter-9.9	mine-10.9	send-11.1
carry-11.4	nonvehicle-51.4.2	slide-11.2
chase-51.6	pelt-17.2	spray-9.7
clear-10.3	pit-10.7	throw-17.1
coil-9.6	pocket-9.10	vehicle-51.4.1
concealment-16	pour-9.5	waltz-51.5
confine-92	push-12	wipe instr-10.4.2
debone-10.8	put_direction-9.4	wipe_manner-10.4.1
disappearance-48.2	put spatial-9.2	

cause-change-scale-cx

- **Semantic Roles:**
  ACTOR, UNDERGOER, SOURCE, EXTENT, GOAL

- **Semantic Predicates:**
  `cause`(ACTOR, E)
  `change_value`(`during`(E), UNDERGOER)
  `path_rel`(`start`(E), UNDERGOER, SOURCE, *ch_of_scale*,prep)
  `path_rel`(`during`(E), UNDERGOER, EXTENT, *ch_of_scale*,prep)
  `path_rel`(`end`(E), UNDERGOER, GOAL, *ch_of_scale*,prep)

- **Member VN Classes:**

  entity_specific_cos-45.5
  calibratable_cos.45.6

cause-transform-cx

- **Semantic Roles:**
  ACTOR, UNDERGOER, SOURCE, GOAL

- **Semantic Predicates:**
  `cause`(ACTOR, E)
  `transform`(`during`(E), UNDERGOER)
  `path_rel`(`start`(E), UNDERGOER, SOURCE, *ch_of_state*,prep)
  `path_rel`(`end`(E), UNDERGOER, GOAL, *ch_of_state*,prep)

- **Member VN Classes:**

bend-45.2	other_cos-45.4
break-45.1	poison-42.2
coloring-24	preparing-26.3
cooking-45.3	render-29.90
entity_specific_cos-45.5	suffocate-40.7
knead-26.5	turn-26.6.1

cause-receive-cx sub construction of *cause-possess-cx*

- **Semantic Roles:**
  ACTOR, UNDERGOER, SOURCE, GOAL
- **Semantic Predicates:**
  `cause`(ACTOR, E)
  `transfer`(`during`(E), UNDERGOER)
  `path_rel`(`end`(E), UNDERGOER, GOAL, *ch_of_poss* or *ch_of_info*,prep)
  `equals`(ACTOR, SOURCE)
- **Member VN Classes:**

(possession)	*(information)*	
contribute-13.2	confess-37.10	say-37.7
fulfilling-13.4.1	crane-40.3.2	talk-37.5
get-13.5.1	curtsey-40.3.3	tell-37.2
give-13.1	instr_communication-37.4.1	transfer_mesg-37.1.1
invest-13.5.3	manner_speaking-37.3	wink-40.3.1
pay-68	nonverbal expression-40.2	

cause-have-cx sub construction of *cause-possess-cx*

- **Semantic Roles:**
  ACTOR, UNDERGOER, SOURCE, GOAL
- **Semantic Predicates:**
  `cause`(ACTOR, E)
  `transfer`(`during`(E), UNDERGOER)
  `path_rel`(`start`(E), UNDERGOER, SOURCE, *ch_of_poss* or *ch_of_info*,prep)
  `equals`(ACTOR, GOAL)
- **Member VN Classes:**

(possession)	*(information)*
cheat-10.6	learn-14
get-13.5.1	
obtain-13.5.2	
steal-10.5	