Criteria for Identifying and Annotating Caused Motion Constructions in Corpus Data

Jena D. Hwang[1], Annie Zaenen[2], Martha Palmer[1]

[1]Department of Linguistics, University of Colorado at Boulder
[2]Center for the Study of Language and Information, Stanford University
hwangd@colorado.edu, azaenen@stanford.edu, martha.palmer@colorado.edu

Abstract

While natural language processing performance has been improved through the recognition that there is a relationship between the semantics of the verb and the syntactic context in which the verb is realized, sentences where the verb does not conform to the expected syntax-semantic patterning behavior remain problematic. For example, in the sentence *"The crowed laughed the clown off the stage"*, a verb of non-verbal communication *laugh* is used in a caused motion construction and gains a motion entailment that is atypical given its inherent lexical semantics. This paper focuses on our efforts at defining the semantic types and varieties of caused motion constructions (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. The annotation will serve as training and test data for classifiers for CMCs, and the CMC definitions developed throughout this study will be used in extending VerbNet to handle representations of sentences in which a verb is used in a syntactic context that is atypical for its lexical semantics.

Keywords: VerbNet, caused motion constructions, semantic coercion, lexical semantics, sentence representation

1. Introduction

While natural language processing performance has been improved through the recognition that there is a relationship between the semantics of the verb and the syntactic context in which the verb is realized (Guildea and Palmer, 2002), sentences where the verb does not conform to the expected syntax-semantic patterning behavior remain problematic. Consider the following sentences:

1. The goalie kicked the ball into the field.
2. The crowd laughed the clown off the stage.
3. 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 1, the goalie causes the ball to go into the field, in 2, the crowd causes the clown to go off the stage, and in 3, 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) (Goldberg, 1995). The verb *laugh* of sentence 2 is normally considered an intransitive manner of speaking verb (e.g. *The crowd laughed 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 laughing*). The verb *tilt* is a verb of spatial configuration normally taking, as its object argument, the inclined item (e.g. *He tilted the bottle*). In 3, the verb is not only coerced into the caused motion reading, the coerced meaning is also abstract rather than physical (e.g. *He tilted the liquid into his mouth and swallowed*). Whether the motion is physical or abstract, the semantics parallel one other: 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.

Thus, if the semantic interpretation is strictly based on the expected semantics of the verb and its arguments, it fails to include the relevant information from the CMC. An accurate semantic role labelling for such sentences requires that NLP classifiers to accurately identify these coerced usages in data. Furthermore, once the CMCs identified and the semantic roles are properly assigned, the sentence would require an accurate semantic interpretation with appropriate representations that include the semantics of the CMCs. Making a semantic analysis available for both conventional and coerced caused motion instances would be useful in making inferences related to the states or locations pre- and post-event (Zaenen et al., 2008).

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 current effort is aimed at providing a larger set of high-quality annotated data for further training and testing of CMC classifiers. In this study, we develop detailed criteria for identifying CMCs that will aid in the production of consistent annotation with high inter-annotator agreement. In turn, successful annotation of the data will be used to establish whether or not the descriptive criteria are indeed useful in characterizing CMCs.

For semantic representation, we turn to the lexical resource VerbNet and the semantic predicates it provides for sentence representation. 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 1. Verb-

Net does not currently have effective ways of handling instances in which the verb is used in an atypical syntactic context, through which it gains new meaning. That is, because VerbNet classes are designed to focus on the protypical or conventional behavior of verbs, semantic coercion that often cross-cuts 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 to both the coercive and the conventional usages of verbs.

This paper, thus, 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. Moreover, we will also present an overview of the work involved in incorporating CMC definitions in VerbNet, outlining the changes necessary to make CMC representation possible.

2. Bigger Picture: VerbNet

VerbNet (Kipper et al., 2008) is a lexical resource that expands on Levin's verb classification (Levin, 1993). In accordance with Levin's work, VerbNet's classification of verbs is based on the hypothesis that verbs that are 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. For example, the verbs *loan* and *rent* are grouped together in GIVE-13.1, and the verbs *deposit* and *situate* are grouped into PUT-9.1.

Our effort at detailing semantic behavior of CMCs through annotation is part of a greater effort at equipping VerbNet with systematic ways for dealing with coercive usages of verbs. VerbNet is currently useful at providing an analysis for 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 (e.g. *laugh* gains a caused motion reading when used in *They laughed the clown off the stage*). 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 meaning of the member verbs. Consider the following three sentences:

4. John slouched himself into the chair.
5. The crowd laughed the clown off the stage.
6. Cynthia blinked snow off her eyelashes.

The verb *slouch* in example 4 is a posture verb appearing in the ASSUMING_POSITION-50 class along with verbs like *hunch, lean* and *slump*. VerbNet currently interprets the CMC usage of this verb by including it to 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[1]. Finally the verb *blink* is in the WINK-40.3.1 class, which includes other verbs of gesture involving a body part such as *clap*, *nod*, and *point*. Unlike the other two instances, this particular usage of CMC is not at this time addressed by VerbNet.

A better way, we believe, is introducing constructional definitions – VerbNet constructions, if you will – to interact with the current VerbNet classes to project the constructional meaning to the sentence where the inherent semantics of the verb does not include it. VerbNet should give a uniform semantic treatment for all sentences of caused motion regardless of the specific lexical meaning of the verb. Figure 1 is a visualization of an instance of CMC instantiated by a verb typical of motion.

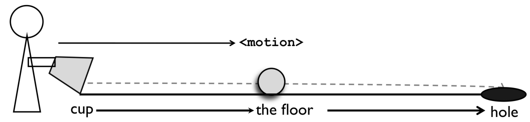

Figure 1: *"Cynthia rolled the ball out of the cup, across the floor, into a hole."*

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 are neither verbs of motion or verbs of locational change. Nonetheless, VerbNet should give a similar semantic representation when they appear in a caused motion context as in example 6. Figure 2 shows the visualization of the caused motion event and the desirable sentence representation for the sentence.

[1]The inclusion of the caused motion frame in the NONVERBAL_EXPRESSION-51.3.2 class comes directly from Levin (1993). This class been reanalyzed and the CMC frame has been removed. While this change serves to strengthen the semantics unity of the member verbs, much like the *blink* example in sentence 6, the sentence representation for 5 becomes unattainable.

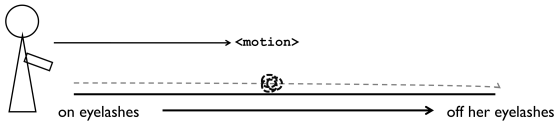


```
cause(Cynthia, E)
motion(E, the snow)
body_process(E, Cynthia)
rel_to_path(start(E), the snow, her eyelashes)
rel_to_path(end(E), the snow, not(her eyelashes))
```

Figure 2: *"Cynthia blinked snow off her eyelashes."*

One notable difference between the two representations is that the *blink* sentences that include the CMC carry an additional semantic component, namely, the act of blinking itself. As we have noted, the semantics of motion and path information would be projected from the semantics of the constructions. What the VerbNet HICCUP-40.1.1 class would provide is then the specialized semantics that is 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.
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 would have been produced by the classes by themselves, since the CMC frame is inherent in the member verbs (e.g. *She carelessly tossed her book on the table*). However, when paired with the HICCUP-40.1.1 class, which includes verbs such as *blink*, and *shiver*, it produces the correct analysis for sentences like *She blinked the snow off her eyelashes* or *The man shivered the cold off his bones*.

3. Applying Linguistic Theory

We take Goldberg's (1995) analysis of CMCs as the starting point for our definitions. As described above, a CMC requires three roles: the causal argument, the patient argument undergoing motion, and one or more path arguments; and it roughly means "the causal argument directly causes the patient to move along a path specified by the prepositional phrase". Goldberg's analysis includes both physical and abstract/metaphorical motion. Its prototypical caused motion instances include those occurring with verbs of motion (e.g. *throw*, *push*) and the coerced instances as seen in sentences 2 and 6.

3.1. Entailments of Motion

Goldberg recognizes that depending on the verb that instantiates the CMC, the construction can retain a different set of inferences when it comes to motion. Consider the following examples:

7. Harry locked Joe into the room.
8. Harry allowed Joe into the room.
9. Harry invited Joe into the house.
10. Harry helped Joe into the house.

Sentence 7, Harry prevents Joe's motion by introducing a barrier – presumably a door – in Joe's potential path of motion. The verb *allow* in sentence 8 has the opposite effect: Harry makes Joe's motion possible by removing the prohibition that may have kept Joe out of the room. Sentence 9's semantics is similar to that of 8, but entailment of motion is dependent on the condition that Joe accepts Harry's invitation. Finally, sentence 10 is likely the only example here where Joe's motion is unconditionally entailed: it is not possible to negate Joe's movement into the house (e.g. **Harry helped Joe into the house, but Joe never went into it*). Goldberg identifies these distinctions as four senses related to the 'central' sense CMC. The four examples above are examples of prevented motion, enabled motion, implied motion and assisted motion, respectively.
We take a slightly narrower view of what defines CMCs. Specifically, we focus on a subset of Goldberg's constructions for which the path of motion is strictly entailed. Since the annotation's goal is to be used for automatic identification of the CMCs, the annotated label should describe meaningful and coherent phenomena. That is, if the descriptive features of a construction are too varied, then automatic identification is less likely. If the label is semantically too general, no meaningful inferences can be made from its identification. This decision, effectively, selects for only one of the four related senses as a part of our CMC definition. By restricting CMCs to instances in which motion is strictly entailed, we can 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.

3.2. Expanding the Definition

In order to account for the idiosyncratic usages of the construction, Goldberg introduces a series of semantic constraints for CMCs (c.f. section 7.4 in Goldberg (1995)). However, Goldberg's central sense, we find, is too general to identify the members of the construction reliably in an annotation setting. Instances expressing clear-cut motion (e.g. *He kicked the ball into the bin* or coerced instances like *The crowd laughed him off the stage*) are easier for human annotators. However, instances that are not specifically discussed in the analysis can only be evaluated given general criteria from Goldberg's analysis, and the more abstract the meaning, the more difficult is the evaluation. Therefore, we further expand and elaborate on the definitions proposed by Goldberg (1995).
We find that there are some natural categories that emerge in the examination of data. Consider the following examples of CMCs:

11. Ron Brierley raised its stake in the company Friday to 15.02% from about 14.6% Thursday.
12. They plan to recycle them into fresh sealing clay.

The difficulty in sentence 11 hangs on the interpretation of the nature of path invoked in the sentence (i.e. where does the stake go?). Once the path of motion is conceived of as a linear scale (Lakoff, 1993), the motion is, then, the movement from one point to another point in that scale. In a similar way, the path expressed by the prepositional phrase seen in example 12 is a state rather than a literal location.

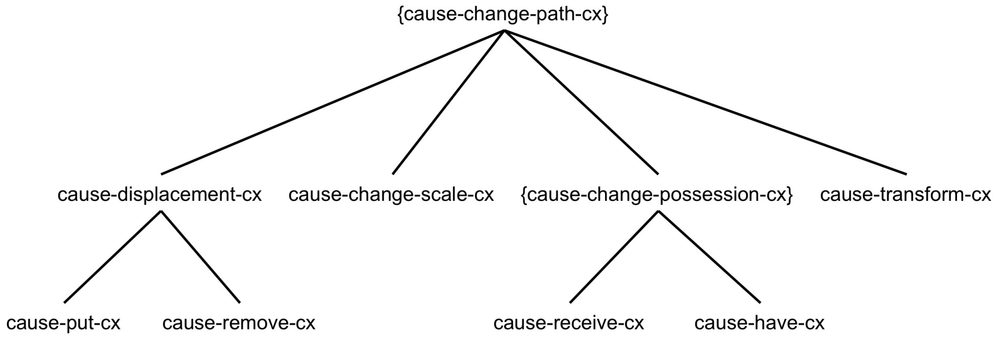

Figure 3: CMC Classification

The path, here, is conceived of as a change of state over which the patient is moved. In particular, the attribute of the patient undergoes a change of state, moving from a certain initial state to a final state (Goldberg and Jackendoff, 2004). These types of elaborations are necessary to provide the annotators with further information that will allow them to make decisions about sentences like this as they are encountered in the data.
The following example of CMC is an instance of change of possession:

13. Vector will sell WordPerfect to someone who [...] will capitalize the product.

Identifying this sentence as CM is based on the interpretation that while the thing possessed is not literal motion, change of ownership is well understood as a metaphorical extension of physical transfer (c.f. Jackendoff (1972; Rappaport-Hovav and Levin (2008)). This relation that is also recognized in Goldberg (1995), we make explict in our definitions. In the next section we discuss our CMC classification and the categories created in order to account for such idiosyncrasies in CMC usage.

4. CMC Classification

Through a series of iterative annotation steps we will discuss in Section 5., we developed a semantically fine-grained sub classification of CMCs based on corpus data. The classification is shown in Figure 3. 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 CMC categories details the particular characteristics and semantic types that they lend to consistent semantic inferences.

4.1. 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. Examples 1-6, and 10 seen so far in this paper are examples of this category. 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 path arguments. Following are examples in these categories:

14. (PUT) She put the book *[from the shelf]-SOURCE [on the table]-GOAL.
15. (REMOVE) She removed the book [from the shelf]-SOURCE *[onto the table]-GOAL.

4.2. 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 in which the patient argument moves along a path that is expressed as a linear scale to include sentences like the one seen in example 11. The movement can be in any direction. Following are further examples of this category:

16. I marked the price down to 2 dollars.
17. The heavy downpour raised the level of the lake to 1000 ft.

4.3. cause-change-possession-cx

The category *cause-change-possession-cx* (TRANSFER) includes caused motion instances of transfer of possession, where 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-have-cx* (GIVE), which includes instances that syntactically realize only the GOAL (or recipient) argument and *cause-receive-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:

18. (GIVE) Mary gave my coat [to the salvation army]-GOAL.
19. (RECEIVE) Mary purchased my coat [from the salvation army]-SOURCE.

In 18, Mary is both the causal entity and the source entity (i.e. the coat was in possession of Mary at the beginning of the event). Conversely in 19, Mary is the causal entity and the goal entity (i.e. the coat is in possession of Mary at the end of the event. Consequently their parent category TRANSFER that presumably allows for both GOAL and SOURCE as prepositional phrase is not instantiable (e.g. **Mary purchased my coat from salvation army to my mother*). Verbs such as *mail* or *send* that can appear in the presence of both arguments (e.g. *Mary mailed my coat from Denver to Mexico*) are treated as a case of DISPLACE category as the causal argument is one causing the change of location – the question of possession (i.e. who does the coat belong to?) is irrelevant to such sentences.

4.4. cause-transform-cx

The category of *cause-transform-cx* (TRANSFORM) are 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 12. The following are further examples of this category.

20. I broke the vase into little pieces.
21. This company renders crayon scribblings of toddlers into some kind of abstract art.

In example 20, 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 21, the company transforms – presumably by reinterpretation – 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 (20) or abstract (21) sentences.

4.5. Annotation Labels and Guidelines

The classification, thus, introduces a total of 7 CM labels for annotation: DISPLACE, PUT, REMOVE, SCALE, TRANSFORM, GIVE, and RECEIVE. In addition to these 7 labels, annotators were provided with a miscellaneous OTHER category to which they could classify sentences they judged as CMC but did not feel that they 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).

5. Data Annotation

The data for this study was pulled from the Wall Street Journal (WSJ) [2], Web Text (WEB), and the CNN Broadcasting News (BN) corpora of the OntoNotes project (Weischedel et al., 2011). 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.

For the first two passes of annotation, the sentences of WSJ (totalling 15184 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. 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 and an 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.

The third and fourth passes of the annotation were conducted under the newly established guidelines. For the third pass of the annotation, instances matching 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 CMC syntactic forms were selected from the BN corpus (totalling 3753 instances). For this final pass of the annotation, two new annotators unfamiliar with CMC were trained with the guidelines for the annotation of BN.

Table 1 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 4.5.). The “Overall IAA” reports the annotator agreement over the overall caused motion label (CMC vs. NON-CMC)[3].

	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 1: Inter-annotator agreement rates and f-scores. Significant increase ($p<0.05$) in f-score, when compared against pre-guideline annotation is labeled with an * in the last column.

The agreement numbers show 88.3% and 88.1% agreement for the expert annotators for pre-guidelines and post-guidelines 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 (77.2% in WSJ, 71.0% in WEB, 81.1% 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

[2] 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)).

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

annotators performed on the 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 chi-squared test at a significance level of $p<0.05$ (Yeh, 2000). They are reported in Table 1.

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 annotation before the guidelines. Additionally the annotators qualitatively reported that the guidelines help 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).

In summary, the inter-annotator rates in the different passes of annotation show a comparable agreement of over 83%. However, the F-scores suggest that there indeed was a significant difference in the annotation agreement when the agreement is calculated based on the positive labels only. The numbers suggest that the guidelines were useful for the expert annotators. Despite the fact that more labels for annotation normally means higher cognitive load, it is encouraging to find that the usefulness for guidelines in annotation outweighed the load to produce a higher performance[4]. 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.

6. Introducing CMCs to VerbNet

One significant prerequisite for introducing a constructional layer of meaning to VerbNet, as described in Section 2., is a thorough analysis of the verbs of motion and the verbs of locational change, especially those taking the path information that is represented in the classes.

6.1. Semantic Predicates for Path

In an earlier study, we carried out such an analysis of paths of motion, firstly, to identify the inconsistencies and limitations of VerbNet's current representation of path phrases, and secondly, to suggest a more explicit and semantically informed representation for paths of motion in VerbNet (Hwang et al., 2013). In it, we note the lack of consistency in the semantic predicates that represent sentences involving paths of motion. For example, we observe that VerbNet has general tendencies to represent the path of motion using one of the following four predicates: `location()`,`via()`, `direction()` and `Prep()`[5]. The problem lies in that there is no VerbNet wide semantic motivation behind the preferred use of one particular predicate over another – the preference seems generally class-specific.

Our efforts, now, are directed at an overhaul of VerbNet predicates to introduce a single predicate, namely, the `rel_to_path` predicate, to represent locations or states that are components in the path of motion described in a sentence. For example, path in DISPLACE type CMCs will be represented as following:

```
rel_to_path(start(E), UNDERGOER, SOURCE)
rel_to_path(during(E), UNDERGOER, PATH)
rel_to_path(end(E), UNDERGOER, GOAL)
```

The above `rel_to_path` predicates specify the location at which the UNDERGOER will be found relative to the path of motion. Thus, the path expressed in the sentence *Cynthia rolled the ball across the floor into a hole* will instantiate the following representation:

```
rel_to_path(during(E), the ball, the floor)
rel_to_path(end(E), the ball, a hole)
```

For TRANSFORM type CMCs, in which the path is conceived of as a change of state, the initial and final states will be specified:

```
rel_to_path(start(E),UNDERGOER,INITIAL_STATE)
rel_to_path(end(E), UNDERGOER, RESULT)
```

Thus, the resultant state *little pieces* in the sentence *I broke the vase into little pieces* will be the state to which *the vase* could be attributed at the end of the event:

```
rel_to_path(end(E), the vase, little pieces)
```

6.2. Semantic Roles

Additionally, through a separate effort, VerbNet's semantic roles have been re-evaluated, resulting in a detailed hierarchical definition of the current semantic roles (Bonial et al., 2011). We now leverage the newly introduced hierarchy for the linking of VerbNet classes and constructions. Consider the following two CMCs of DISPLACE type:

22. Jake forced [the cat]-PATIENT into the carrier.
23. Mary jumped [the lion]-THEME through the hoop.

The patient argument in a CMC can be a PATIENT or a THEME depending on the VerbNet class of the verb in the sentence. Since both PATIENT or THEME are possible semantic roles for the argument, a CMC would specify their supertype UNDERGOER as a role corresponding to the argument in motion.

In addition to this, introduction of CMCs to VerbNet will require the introduction of new roles. Take the following examples, for instance:

24. He rolled the ball [across the floor]-PATH.
25. Bill baked the crust from [a golden yellow]-INITIAL_STATE [to a crispy brown]-RESULT.

[4]The IAA over the individual labels for the post-guideline passes show 79.5% and 79.2% agreement rates for WEB and BN 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.

[5]Out of the four predicates shown, the `Prep()` predicate has been the focus of our attention. This predicate is not semantically meaningful; rather, it serves to specify a prepositional phrase.

The new roles include the PATH role to designate elements of path as seen in 24, which is distinguished from strict locative adjuncts (e.g. *in her garage* in the sentence *She threw a party in her garage*), and the INITIAL_STATE role to designate states at the beginning of the event. The role will be used in TRANSFORM type of CMC, in which the path is the movement from a state to another as exemplified in 25. INITIAL_STATE is, thus, the counterpart role to the semantic role RESULT, which already exists in VerbNet.

7. Conclusion

In this study, we have carried out a detailed analysis and annotation of CMCs, systematizing its defining characteristics into a typological classification and creating annotation guidelines that will improve corpus annotation of CMCs. Through these efforts we seek to make interpretation of sentences of caused motion possible even when the verb used in the sentence is atypical for the CMC in which it appears. The annotation thus provides data for training and testing of automatic CMC classifiers. Moreover, our classification of CMCs gives insight into the types and varieties of semantic inferences entailed by CMCs in corpus data. Such a classification of constructions, when used with verb specific semantics already available through resources such as VerbNet (Kipper et al., 2008) will aid in the proper semantic analysis of sentences and the drawing of inferences where caused motion is atypical for the verb.

Implementation of the CMC representation in VerbNet is currently underway. In the future, we plan to extend 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 semantics of the verb alone.

Additionally, we have observed during annotation that as the meaning of the sentence moves further away from the concrete usage, the judgement of CMC becomes increasingly difficult. In order to gain a better understanding of the effect abstractness has on the construction, we have manually annotated all of the CMC instances in the WSJ, as identified in the current annotation effort, with a concreteness rating. We have not included this part of the study in this paper, as we are still in the earlier stages of investigation. Our plan is to fold these abstractness ratings into the CMC definition as the study progresses.

Finally, on the front of automatic classification of CMCs, we are currently restricting the training of the classifier to instances labeled as DISPLACE CMCs. We plan on systematically moving on to other labels as we continue with our experiments. We will also be experimenting with the abstractness levels to determine if concreteness or abstractness of an event affects performance levels of the classifier.

8. Acknowledgements

We gratefully acknowledge the support of the National Science Foundation Grant NSF-IIS-1116782, A Bayesian Approach to Dynamic Lexical Resources for Flexible Language Processing and DARPA FA8750-09-C-0179 (via BBN) Machine Reading: Ontology Induction. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or DARPA.

9. References

Claire Bonial, Susan Windisch Brown, William Corvey, Martha Palmer, Volha Petukhova, and Harry Bunt. 2011. An exploratory comparison of thematic roles in VerbNet and LIRICS. In *The Sixth Joint ISO-ACL SIGSEM Workshop on Interoperable Semantic Annotation*.

Adele E. Goldberg and Ray Jackendoff. 2004. The english resultative as a family of constructions. *Language*, 80(3):532–568, September.

Adele E. Goldberg. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. University Of Chicago Press.

Dan Guildea and Martha Palmer. 2002. The necessity of parsing for predicate argument recognition. In *40th Meeting of the Association for Computational Linguistics*, Philadelphia, PA, July.

George Hripcsak and Adam S Rothschild. 2005. Agreement, the f-measure, and reliability in information retrieval. *Journal of American Medical Informatics Association*, 12(3):296–298.

Jena D. Hwang, 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, June. Association for Computational Linguistics.

Jena D. Hwang, Martha Palmer, and Annie Zaenen. 2013. Representing paths of motion in representing paths of motion in VerbNet. In Tracy Holloway King and Valeria de Paiva, editors, *From Quirky Case to Representing Space*. CSLI Online Publications.

Ray Jackendoff. 1972. *Semantic Interpretation in Generative Grammar.* The MIT Press, Cambridge, Massachusetts.

Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. 2008. A large-scale classification of English verbs. *Language Resources and Evaluation Journal*, 42(1):21–40.

George Lakoff. 1993. The contemporary theory of metaphor. In Andrew Ortony, editor, *Metaphor and Thought*. Cambridge University Press, 2nd edition.

Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press.

Mitchell Marcus, 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.

Malka Rappaport-Hovav and Beth Levin. 2008. The English dative alternation: The case for verb sensitivity. *Journal of Linguistics*, 44(01):129–167, March.

Alexander Yeh. 2000. More accurate tests for the statistical significance of result differences. In *Proceedings of the 18th Conference on Computational Linguistics*, volume 2.

Annie Zaenen, 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, May. European Language Resources Association (ELRA).