Evaluating the Coverage of VerbNet

Weston Feely
Language Technologies Institute
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213-3891
WFeely@cs.cmu.edu

Claire Bonial, Martha Palmer
Department of Linguistics,
University of Colorado at Boulder
Hellems 290, 295 UCB
Boulder, CO 80309-0295
{CBonial, MPalmer}@colorado.edu

Abstract

This research presents a comparison of the syntactic behavior of verbs represented in an online verb lexicon, VerbNet, and the actual behavior of the verbs in the SemLink corpus. To complete this comparison, each verbal instance of the SemLink corpus is reformulated into a syntactic frame, e.g. Noun Phrase – Verb – Noun Phrase, and compared to syntactic frames listed in VerbNet. Through this effort, the coverage and accuracy of VerbNet is extended with the addition of new syntactic frames and thematic roles such that VerbNet is a more complete reflection of language in use.

1 Introduction

VerbNet (VN) (Kipper et al., 2008) is an online verb lexicon that provides valuable information on the relational semantics of approximately 6200 English verbs. VN is an extension of Levin's (1993) classification, in which verbs are organized according to their compatibility with certain syntactic, or "diathesis," alternations. For example, the verb *break* can be used transitively (*Tony broke the window*) or intransitively (*The window broke*). This represents one diathesis alternation, and other verbs that share the ability to alternate between these two syntactic realizations could be classified with *break*. Although the primary basis of Levin's classification is syntactic, the verbs of a given class do share semantic regularities as well. Levin hypothesized that this stems from the fact that the syntactic behavior of a verb is largely determined by its meaning; thus, there is a fundamental assumption that syntactic behavior is a reflection of semantics.

VN has extended Levin's work and the lexicon has proved to be a valuable resource for various NLP applications, such as automatic semantic role labeling (Swier & Stevenson, 2004), semantic inferencing (Zaenen, 2008), and automatic verb classification (Joanis et al., 2007). However, the utility of VN relies heavily on its coverage and accurate representation of the behavior of English verbs. Although VN is theoretically motivated, the coverage and accuracy of the lexicon has not been comprehensively investigated, with the exception of examinations of VN's representation of certain syntactic constructions (Bonial et al., 2011c; Bonial et al., 2012). This work compares the representations of syntactic behavior found in VN to actual syntactic behaviors found in the SemLink corpus. There are two primary purposes of this comparison: 1) coverage: to what extent does VN capture all syntactic realizations of a given verb? 2) accuracy: to what extent is VN's syntactic representation an accurate reflection of realization possibilities and probabilities? The findings herein can be used to improve both coverage and accuracy, thereby improving the utility of VN overall.

2 Background

In order to evaluate VN's coverage, syntactic and semantic information available in the verb lexicon VN and the annotated corpus SemLink were compared. These two resources are described in the next sections.

2.1 VerbNet Background

Class membership in VN is based on a verb's compatibility with certain syntactic frames and alternations. For example, all of the verbs in the Spray class have the ability to alternate the Theme or Destination as a noun phrase (NP) object or as a prepositional phrase (PP): *Jessica loaded the boxes into the wagon*, or *Jessica loaded the wagon with boxes*. VN's structure is somewhat hierarchical, comprised of superordinate and subordinate levels within each verb class. In the top level of each class, syntactic frames that are compatible with all verbs in the class are listed. In the lower levels, or "sub-classes," additional syntactic frames may be listed that are restricted to a limited number of members. In each class and sub-class, an effort is made to list all syntactic frames in which the verbs of that class can be grammatically realized. Each syntactic frame is detailed with the expected syntactic phrase type of each argument, thematic roles of arguments, and a semantic representation. For example:

Frame NP V NP PP.Destination
Example Jessica loaded boxes into the wagon.
Syntax Agent V Theme Destination
Semantics Motion(during(E), Theme)
Not(Prep-into (start(E), Theme, Destination))
Prep-into (end(E), Theme, Destination)
Cause(Agent, E)

2.2 SemLink Background

The SemLink corpus (Palmer, 2009; Loper et al., 2007) consists of 112,917 instances of the Wall Street Journal, each annotated with its corresponding VN class. Each instance is further annotated with PropBank (Palmer et al., 2005) arguments, which are numbered arguments that correspond to verb-specific roles. For example, these are the potential roles to be assigned for the verb *load*:

Roleset ID: load.01, *cause to be burdened*,
VN class: 9.7-2
Roles:
Arg0: loader, agent (VN role: 9.7-2-agent)
Arg1: beast of burden (VN role: 9.7-2-destination)
Arg2: cargo (VN role: 9.7-2-theme)
Arg3: instrument

Note that each verb sense, or "roleset," is mapped to its corresponding VN class, and each of the PropBank roles are mapped to VN thematic roles where possible. This roleset also demonstrates a sort of mismatch between PropBank and VN's treatment of *load*: PropBank treats the instrument as a numbered argument, whereas VN doesn't list an instrument as a semantic role for this verb.

Within the SemLink corpus, these mappings are made explicit such that with each instance, both PropBank and VN thematic roles are given for each argument. SemLink also contains mappings between PropBank rolesets, VN classes and FrameNet (Fillmore et al., 2002) frames, as well as corresponding mappings between PropBank arguments, VN thematic roles and FrameNet frame elements. Thus, SemLink is a resource created with the intent of allowing for interoperability amongst these resources.

2.3 Investigating VerbNet Using SemLink

The motivation for this project is to compare the set of syntactic frames listed in each VN class to the set of syntactic frames that actually occur in usage in the class's corresponding SemLink entries. Such a comparison is challenging because VN is a largely theoretical verb lexicon, which is still strongly rooted in Levin's original classification. SemLink, on the other hand, is an annotated corpus of real language in use, which often shows far more syntactic variability than assumed by theoretical linguistics. Thus, a comparison of VN with SemLink could provide a greater range of syntactic frames for most VN classes, simply because unexpected syntactic frames present themselves in the SemLink data.

This additional syntactic variation in the SemLink data should facilitate the primary goal of this project, which is to increase the coverage of VN's syntactic and semantic information. This is accomplished by using the empirically-derived information in the SemLink data to validate the class organization of VN by demonstrating which of VN's syntactic frames are present in the SemLink corpus for a given class, and which syntactic frames are present in the corpus that are not listed among the options for a given VN class. The additional syntactic frames detected can increase the coverage of each verb class's syntactic information, by augmenting each class's previous set of syntactic frames with empirically derived alternatives.

Additionally, the SemLink data will provide frequency information for syntactic frames, so that each syntactic frame in a VN class can be listed with how often it occurs in corpus data. This is especially important, because our empirical validation of the class organization of VN can be extended to: which syntactic frames are highly frequent in SemLink and present in a given VN class; which frames are highly frequent but missing from a given class; which frames are infrequent and present in a given class; and which frames are infrequent but missing from a given class.

3 Methods

The SemLink data for this project includes 70,270 SemLink instances, which are all the instances of the total 112,917 with a currently valid VN class assignment. Each of the SemLink instances included in the project data was processed for the necessary information to compare it to VN frames. This included the extraction of each SemLink instance's VN class assignment, the instance's PropBank roleset assignment, the syntactic frame from the Treebank parse, and the VN semantic roles for each constituent in the frame. After gathering this information from SemLink, frequencies were calculated for each syntactic frame type given its VN class assignment. The syntactic frames from SemLink were created using a Penn Treebank application-programming interface that automatically retrieved the syntactic constituents immediately dominating the part-of-speech tag for each of the words that were marked as arguments to the main verb in the SemLink instances. The rest of the information taken from SemLink was extracted directly from the SemLink Wall Street Journal annotations, using regular expressions.

The VN data for this project includes the frames (e.g. NP V NP) and corresponding semantic role argument structures (e.g. Agent V Theme) for all VN classes. These frames and argument structures were taken directly from the VN XML class files using regular expressions, with some small modifications to each frame. In order to facilitate matching with the SemLink frames, the constituents in each of VN's flat syntactic frames were stripped of additional tags, such as: redundant thematic roles (e.g. PP.Location; all roles are listed again in a separate line, e.g. Agent V Theme Location), syntactic alternation tags (e.g. NP-Dative), and other tags extraneous to the purpose at hand.

3.1 Frame Creation Method

The syntactic frames extracted from SemLink for this project were formed based on the linear order of syntactic constituents, as retrieved from the linear order of thematic role annotations in SemLink. In the case of arguments of the verb that were syntactically null elements, the last element in a movement chain was taken to form the frame, unless the null element was a passive or reduced relative clause marker, in which case the constituent one level above the trace constituent was taken. As an example, consider the following question: *Whom did she see?* In the Penn Treebank treatment of this sentence, there would be an object trace after *see* with an index indicating that the object trace corresponds to the question word *whom: Whom-1 did she see *T*-1?* The arguments identified for *see* would use the trace as the object position, resulting in the following frame: NP V NP, as opposed to the position of the realized constituents: NP NP V. In order to avoid interpreting passives as verb-initial frames, the passive and reduced relative constructions are treated differently and identified as such. Passives are currently excluded from this study as discussed in greater detail in Section 4.1.

3.2 Matching Conditions

After extracting the data from SemLink and VN, the data from each SemLink instance was matched against the set of [frame, argument structure] pairs in the corresponding VN class. This matching process was done using regular expressions in a three-step process.

First, the frame from the SemLink instance was checked against each of the frames in its corresponding VN class. If there was a match, the instance was counted as having matched a VN frame, and if the [VN class, frame] pair for this SemLink instance had not previously been matched, it was added to a list of frame types that matched VN. For example, consider the following SemLink instance, shown with its PropBank arguments and VN thematic role labels:

1. *The explosion of junk bonds and takeovers has…loaded corporations…with huge amounts of debt.*
Load, PropBank load.01, VN class Spray-9.7-2:
[The explosion of junk bonds and takeovers]$_{\text{ARG0, AGENT}}$ has…loaded$_{\text{RELATION}}$ [corporations…]$_{\text{ARG1, DESTINATION}}$ [with huge amounts of debt…]$_{\text{ARG2, THEME}}$.

This SemLink instance would be assigned the frame NP V NP PP, which matches a frame listed in its associated VN class:

Frame NP V NP.Destination PP.Theme
Example Jessica loaded the wagon with boxes.
Syntax Agent V Destination {with} Theme

Thus, this instance would be considered a frame match to VN.

Second, if the frame from the SemLink instance did not match any of the frames in the corresponding VN class, then the argument structure for the instance was checked against each of the argument structures in the corresponding VN class. If there was a match, the instance was counted as having matched VN, and if the [VN class, frame] pair for the SemLink instance had not previously been matched, it was added to a different back-off list of frame types that matched VN. The following instance is an example of this type of match:

2. *It doesn't mean unanimous…*
Mean, PropBank mean.01, VN class Conjecture-29.5:
$It_{ARG0,\ AGENT}$ $does[n't]_{NEGATIVE}$ $mean_{RELATION}$ $unanimous_{ARG1,\ THEME}$…

This frame syntactically is of type NP V ADJP, and VN only represents Themes realized as NPs. Thus, this frame was matched via arguments (Agent V Theme) rather than syntactic frames. It was quite common for a SemLink instance to include an unexpected constituent type such as the ADJP here, and it is this constituent information that can be used to expand the constituent types for frames in VN, discussed in Section 5. This particular instance also brings to light a problematic aspect of the SemLink corpus and the interoperability between VN and PropBank: PropBank has much more coarse-grained rolesets or senses than those found in the VN classes. Thus, this roleset, which would include instances of the sense of intentional "meaning" found in the Conjecture class, also includes this sense of unintentional "meaning". As a result, "It" above is treated as an Agent, although the status as an Agent is questionable.

Third, if the frame and its argument structure from the SemLink instance did not match any of the frames in the corresponding VN class, it was added to a final list of frame types that did not match VN. Consider the following unmatched examples of the relation *remain,* which belongs to the VN class Exist-47.1:

3. *Like just about everything else, that remains to be seen.*
[Like just about everything else,$]_{ADVERBIAL}$ $that_{ARG1,\ THEME}$ $remains_{RELATION}$ [to be seen$]_{ARG3}$ – NP V S

4. *The crowd remained good-natured, even bemused.*
[The crowd$]_{ARG1}$ $remained_{RELATION}$ [good-natured, even bemused$]_{ARG3}$ – NP V ADJP

These examples demonstrate a potential gap in VN's representation of verbs like *remain* in the Exist class. While the PropBank argument structure includes an Arg3 role that corresponds to "attribute" arguments for more abstract usages of *remain*, the VN class contains only the roles Theme and Location, and did not include frames with sentential complements or adjective phrases that could capture these attributes. This suggests one way that VN can be improved based on this empirical investigation of verbal behavior: the addition of an attribute argument to the Exist class for abstract usages.

The end result of this matching process was three counters and three lists. The counters are the portion of the total SemLink instances that 1) matched a VN frame, 2) did not match a frame but did match a VN argument structure, or 3) did not match VN at all. These token counters were converted into token percentages in Table 1 in Section 4 below. The lists contain frame types for each matching condition: frame types that were in VN, frame types that had argument structures that were in VN, and frame types that were not in VN. These type lists were converted into type percentages in Table 3 in Section 4.

This matching process was repeated for three frequency subdivisions of the SemLink frame types: high frequency, middle frequency, and low frequency. These frequency categories were defined as the top 30%, middle 40%, and bottom 30% of the SemLink frame types for each VN class, ranked by frequency. For this second matching process using frequency information, the SemLink frames that matched VN by frame and by argument structure were combined into one category of frame types that matched VN. The SemLink frames that did not match VN by frame or argument structure were left in a separate category of frame types that did not match VN. In the same manner as the first matching process, the end result was a set of counters for the frame tokens that matched VN, and a set of lists for the frame types that matched VN, subdivided by these frequency categories. The percentages of the SemLink frame tokens for each of these frequency subdivisions are in Table 2 of Section 4, and the percentages of the SemLink frame types for each of these frequency subdivisions are in Table 4 of Section 4.

3.3 Loose Match

For the particular instances in SemLink that contained WH-movement or topicalization, looser matching criteria were used: there was a successful argument structure match when the set of argument roles matched any set of argument roles in the corresponding VN class (ordering was not considered). This was done because transformations like topicalization and WH-movement allow variable movement of syntactic constituents along the syntactic parse, so this separate matching condition that disregards the linear order of argument roles was needed. Because these transformations are possible for all verbs, they are not the type of distinctive syntactic alternations that VN lists. For example:

5. *"It's been a steadily improving relationship," says Mr. Carpenter.*
Say, PropBank say.01, VN class Say-37.7:
["It's been a steadily improving relationship"]-1 says$_{RELATION}$ [**Trace**-1]$_{ARG1,\ TOPIC}$ [Mr. Carpenter]$_{ARG0,\ AGENT}$

This instance was recognized as a syntactic frame of the type V S NP, which VN does not include in the Say class. Since the frame did not match, the instance was tested for an argument match: V Topic Agent. However, this argument structure is also not represented in VN for the Say class. Nonetheless, the loose match condition recognizes the topicalization transformation, and with instances containing such movement, allows for a match based on sets of arguments. Because the roles of Agent and Topic are present in the class and the transformation was recognized, this instance was considered a match.

3.4 Semantic Role Updates

After the frame retrieval process, it was also necessary to update the set of semantic roles in each SemLink instance. This is due to the fact that the SemLink Wall Street Journal annotations are currently outdated, and awaiting an update in the near future. However, at the time of this writing the SemLink data used for this project was created using an old set of VN roles that are not current with the 3.2 version of VN (for a description of the recent VN semantic role updates, see Bonial et al., 2011a, Bonial et al., 2011b). Therefore, before the frame matching process could begin, the semantic roles in the argument structures retrieved from SemLink had to be updated using a type-to-type mapping of old VN roles to new VN roles. This update was done automatically.

4 Findings

The results of the matching process are discussed in the following sections.

4.1 Passives

Passive sentences in the Wall Street Journal section of SemLink were removed from the matching process to be considered separately, since previous attempts to include passives in the matching process created the largest source of error for the project. This is due to the fact that VN does not include passive versions of its frames in the frame listing for each verb class. This omission is purposeful, because common syntactic transformations like passivization and WH-movement are not considered to be syntactic alternations distinctive of verb classes, following Levin's original verb classification. Passives made up 26.7% of the original data set of 70,270 instances, and after removing them a set of 51,534 frame tokens remained to be considered for the matching process. Passive frames were included in a separate list of frames, potentially to be used for future augmentation of VN.

4.2 Matches

SemLink tokens that...	% of total SemLink frame tokens (51534)
Matched a VN Frame	51.23%
Matched a VN Argument Structure	24.30%
Did not match corresponding VN class	24.46%

Table 1: Results of Matching Process for SemLink Frame Tokens

If we focus on tokens, we see that the majority of frame tokens in SemLink match frames in VN. However, this needs to be qualified because the matches are highly skewed towards the high frequency frame token matches. This is shown in the following table.

Match/No match grouping	Frequency	% of total SemLink frame tokens
Matched VN	High Frequency (top 30%)	54.49%
	Middle Frequency (middle 40%)	20.63%
	Low Frequency (bottom 30%)	0.41%
Did not match VN	High Frequency (top 30%)	17.67%
	Middle Frequency (middle 40%)	5.47%
	Low Frequency (bottom 30%)	1.32%

Table 2: Results of Matching Process for SemLink Frame Tokens, Divided by Frequency

This demonstrates that the most frequent frame tokens make up the majority of the frame token matches. This is because a small number of highly frequent frame types bias the token matches towards the high frequency match category. For example, 34% of all frame tokens are NP V NP, the most frequent frame type. Therefore, it is important to also consider the SemLink frame type matches, which are available in the following tables.

SemLink frame types that...	% of total SemLink frame types (3721)
Matched a VN Frame	12.92%
Matched a VN Argument Structure	20.29%
Did not match corresponding VN class	66.78%

Table 3: Results of Matching Process for SemLink Frame Types

Match/No match grouping	Frequency	% of total SemLink frame types
Matched VN	High Frequency (top 30%)	18.57%
	Middle Frequency (middle 40%)	9.78%
	Low Frequency (bottom 30%)	4.86%
Did not match VN	High Frequency (top 30%)	19.99%
	Middle Frequency (middle 40%)	29.16%
	Low Frequency (bottom 30%)	17.63%

Table 4: Results of Matching Process for SemLink Frame Types, Divided by Frequency

When considering frame types, it is clear that the majority of unique syntactic frame types in SemLink do not match VN. Among the frame types that did match VN, the majority of these were high frequency, although the highest frequency frame types in each class only match VN frames of the class 18.57% of the time. This indicates that a wider set of constituents is needed in VN syntactic frames and possibly a wider range of semantic roles in several VN classes in order to account for abstract usages that will better match SemLink data.

5 Discussion

This research demonstrated that while the majority of frame tokens in SemLink match frames in VN, the frames listed in VN need a wider set of constituents because the prototypical constituents for a particular role (e.g. NP-Agent) are not always reflective of the prototypical syntactic realizations in SemLink. In this way, both coverage and accuracy of VN frames could be improved simply by expanding the constituent types that can make up a given frame. To address this issue, a supplementary resource has been created that lists all constituent types found in SemLink that match a particular frame type. For example, this frame exists in the Remove class:

Frame NP V NP
Example Doug removed the smudges
Syntax Agent V Theme

The drawback of this frame is that it assumes that the Agent and Theme roles will be realized as NPs for all verbs in the class in all cases. This investigation of SemLink shows that the Agent V Theme frame can truly be realized with each of the following orderings of constituents:

S_V_NP
NP_V_SBAR
NP_V_NP

The first two possibilities are likely not canonical usages, but in order for VN to fully capture verbal behavior, the resource should reflect both theoretically expected usage and actual usage. The mapping resource created through this research will, however, greatly increase the coverage of VN by including all possible constituent types. Additionally, this resource will help to facilitate interoperability between VN and corpus resources by allowing the information in VN to be more easily compared and applied to that of parsed corpora.

5.1 Assessment of Coverage

Overall, VN currently describes the prototypical syntactic and semantic behavior of many English verbs, but its coverage of a large text corpus like SemLink is fairly low. This is demonstrated by the figures in Table 3, which show that only 12.92% of the frame types in

SemLink are covered by VN's syntactic frames. An additional 20.29% of the frame types in SemLink can be covered using VN's thematic role labels, but this still leaves 66.78% of the syntactic frame types in SemLink unmatched to VN. This is a strong indication that there is a great amount of variability in the syntactic frame types that occur in real usage, which is not currently covered by VN.

When considering the impact of these results, it is important to remember that the organization of VN is based upon Levin's framework and hypothesis that semantic similarity underlies syntactic similarity. Accordingly, VN has focused on representing what can be thought of as typical, distinguishing frames and diathesis alternations of the verbs in a given class. The fact that these verbs participate in other syntactic behaviors not included in the classification is neither surprising nor does it necessarily undermine Levin's hypothesis, given that her classification was not originally intended to give a full enumeration of all behaviors, rather only distinctive behaviors. For the purposes of improving VN as a resource for NLP, the importance of coverage has become clear and is therefore the focus of this research. However, the focus of this research could easily be shifted to an examination of the frequency with which verbs participate in key diathesis alternations, and therefore an examination of Levin's hypothesis.

5.2 Increasing Coverage & Accuracy

Analysis of the SemLink instances that did not match VN frames revealed several classes that could be improved by the addition of a frame or thematic role, or both. In addition to the examples (3 & 4) of *remain* and its associated Exist class, which would require an additional Attribute role based on this study (discussed in Section 3.2), we found that a variety of other verbs and classes were characterized by roles and syntactic behaviors common to SemLink but not represented in VN. Unlike the examples of *remain,* some of these verbs represent new senses that may require entirely new classes. Consider these typical SemLink examples of the verb *add*, which take the following PropBank roleset:

Roleset id: add.03 , *achieve or gain*
Arg1: *Logical subject, patient, thing rising/gaining*
Arg2: *EXT, amount risen*
Arg4: *end point*

6. …*Nippon Mining added 15 to 960.*

...[Nippon Mining] $_{\text{ARG1}}$ added [15]$_{\text{ARG2}}$ [to 960]$_{\text{ARG4}}$

7. *Meanwhile, the broad-based Financial Times 100-share index added 30.4 points to end at 2142.6.*
[Meanwhile]$_{\text{ARGM-TEMPORAL}}$ [the broad-based Financial Times 100-share index]$_{\text{ARG1}}$ added$_{\text{RELATION}}$ [30.4 points]$_{\text{ARG2}}$ [to end at 2142.6]$_{\text{ARG4}}$

The verb *add* falls into several VN classes, Mix, Multiply and Say, of which the Multiply class is the closest fit. However, the Multiply class contains only the roles Agent, Theme, and Co-Theme, with frames such as:

Frame NP V NP PP
Example I multiplied x by y.
Syntax Agent V Theme {by} Co-Theme

This class does not reflect the realizations of the type seen in SemLink, which are particular to the financial domain. Thus, this study has revealed a gap in VN's coverage where the addition of a new (sub)class would be necessary to cover this sense of *add.*

The following table gives other examples of verbs, classes and actions required to increase the coverage of VN based on this study.

Verb	VN Class	Recommended Action
consent	Settle-89	Add NP V S frame: *Triton and Mr. Chase consented* ***to finding****…*
gain, rise, increase, climb	Calibratable-cos-45.6-1	Add Source/Result roles for beginning and final states: *Sales…rose 3%* ***to $29.3 million from $28.4 million****.*
get	-	Add class for *cause to do/be* sense: *We* ***can get that brought down to parity****…*
seek	Hunt-35.1	Add NP V S frame: *Cuba may seek* ***to postpone some sugar shipments****.*
stay, remain	Exist-47.1	Add Attribute role and frame NP V ADJP: *Oil prices stay* ***stable***
struggle	-	Add (sub)class for *try* sense: *The Sunday evening show* ***struggles to stay afloat***

5.3 Surprising Factors

One important factor revealed in the results of the frame matching process is the large number of frame mismatches that were the result of the frame creation process itself. In the case of null elements, the frame creation method described in 3.1 was largely based on anaphora, rather than cataphora. Examples such as the one below, which include cataphoric co-reference, caused the creation of erroneous frames:

8. **Null subject* to further load the stakes, Mr. Lane dreamed up a highly improbable romance…*
Load, PropBank Load.01, VN class Spray-9.7-2:
[**Null subject**]-1 to further$_{\text{ARGM-EXTENT}}$ load$_{\text{RELATION}}$ [the stakes]$_{\text{ARG1, DESTINATION}}$, [Mr. Lane]-1$_{\text{ARG0, AGENT}}$ dreamed up a highly improbable romance …

The frame retrieved from this example was V_NP_NP, with the argument structure V Destination Agent. Neither of these matched the expected syntactic frame, as shown in the VN entry below.

Frame NP V NP.Destination
Example Jessica sprayed the wall.
Syntax Agent V Destination

This mismatch occurred because the argument to the verb was considered to be the realized constituent "Mr. Lane," rather than its previous null subject index. The algorithm for the frame matching process was designed to prefer realized subjects over null subjects, which in many cases was quite successful. However, examples such as these show that sometimes null elements are preferable when forming syntactic frames from a parse, in cases of cataphora. This is an area of improvement that needs to be considered when updating the frame matching process for future work.

6 Conclusion

This comparison of syntactic behavior in SemLink and the syntactic facts represented in VN has allowed for an expansion of the coverage and accuracy of VN. Although the frame matching method described herein requires further refinement, this method has provided data that can be used to compare VN with real language use. This will be of great value to VN as a lexical resource, since many verb classes can be improved by the insights gained from examining the frame mismatches from this project. The supplementary resource described in Section 5 will expedite such a task because it can be used to directly compare the syntactic frames available in SemLink for a particular verb's argument structure with the syntactic frames already available to a VN class. However, this resource is still limited by the erroneous frames generated during the matching process, such as in the cataphora example in Section 5.3. Further revisions to the method of forming syntactic frames from a given parse could better reflect these types of usage.

7 Future Work

As stated in the sections above, the frame matching process described in this paper is still in need of some refinement, to handle all the syntactic variations that occur in SemLink. In particular, the passive syntactic frames will need to be added back into the frame matching process, after further consideration on how to handle such frames. It may be necessary to add passives to the loose matching condition that was applied to cases of topicalization and WH-movement. In addition, the frame retrieval process needs to be revised to account for cataphoric co-reference with a null subject, and other cases of null elements that cause problematic syntactic frames to be generated. Finally, the forthcoming new version of SemLink will be updated with the latest set of VN thematic roles and expanded, which should prove helpful when re-implementing the frame matching process described in this paper.

Once the frame matching process has been further refined, a more in-depth analysis of the impact of these findings will be undertaken. Specifically, while this research has focused on adding syntactic frames to VN in order to increase coverage, future research should focus on the extent to which verbs participate in the key diathesis alternations represented in both VN and Levin's classes. A focus on this question would allow for valuable discoveries in the validity of Levin's hypothesis that syntactic behavior stems from semantics.

The syntactic frame data generated by this project will also be useful for future work in automatic verb clustering. The syntactic frames alone may prove to be a great feature for predicting verb classification, and such an automatically structured classification could be usefully compared to the VN classification to further evaluate it. Perhaps most importantly, the results of this research should increase the value of VN as a NLP resource. The addition of new syntactic constituent types and thematic roles to VN classes based on the SemLink syntactic frames and argument structures should allow for VN to more accurately and comprehensively reflect English verbal behavior, which makes VN more practical for a range of NLP tasks.

Acknowledgments

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, DARPA/IPTO funding under the GALE program, DARPA/CMO Contract No. HR0011-06-C-0022, VN Supplement, and funding under the BOLT and Machine Reading programs. HR0011-11-C-0145 (BOLT) FA8750-09-C-0179 (M.R.) 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.

References

Bonial, C., Brown, S.W., Corvey, W., Palmer, M., Petukhova, V., and Bunt, H. 2011a. An Exploratory Comparison of Thematic Roles in VN and LIRICS. *Proceedings of the Sixth Joint ISO - ACL SIGSEM Workshop on Interoperable Semantic Annotation (ISA-6)*, 39-44.

Bonial, C., Corvey, W., Palmer, M., Petukhova, V., & Bunt, H. 2011b. A hierarchical unification of lirics and VN semantic roles. In *Proceedings of the ICSC Workshop on Semantic Annotation for Computational Linguistic Resources (SACL-ICSC 2011).*

Bonial, C., Brown, S.W. Hwang, J.D., Parisien, C., Palmer, M., and Stevenson, S. 2011c. Incorporating Coercive Constructions into a Verb Lexicon. *Proceedings of the ACL 2011 Workshop on Relational Models of Semantics,* Portland, Oregon, June 2011.

Bonial, C., Feely, W., Hwang, J.D., and Palmer, M. 2012. Empirically Validating VN Using SemLink. In *Proceedings of the ICSC Workshop on Semantic Annotation for Computational Linguistic Resources* (SACL-ICSC 2012), May, 2012.

Fillmore, C. J., Johnson, C. R., & Petruck, M. R. L. (2002). Background to Framenet. *International Journal of Lexicography,* 16(3):235–250.

Joanis, E., Stevenson, S. & James, D. (2007). A general feature space for automatic verb classification. *Natural Language Engineering,* 14(3):337–367.

Kipper, K. Korhonen, A., Ryant, N. and Palmer, M. 2008. A large-scale classification of English verbs. *Language Resources and Evaluation Journal*, 42, pp. 21-40.

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

Loper, E., Yi, S. & Palmer, M. (2007). Combining lexical resources: Mapping between PropBank and VN. In *Proceedings of the Seventh International Workshop on Computational Semantics (IWCS-7),* Tilburg.

Palmer, M. (2009). SemLink: Linking PropBank, VN and FrameNet. In *Proceedings of the Generative Lexicon Conference, GenLex-09,* Pisa, Italy.

Palmer, M., Gildea, D., and Kingsbury, P. 2005. The Proposition Bank: An annotated Corpus of Semantic Roles. *Computational Linguistics*, 31:1, pp. 71-105.

Swier, R.S. and Stevenson, S. 2004. Unsupervised semantic role labeling. *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, 95-102.

Zaenen, A., Condoravdi, C., and Bobrow, D. G. (2008). The encoding of lexical implications in VN. In *Proceedings of LREC 2008,* Morocco, May.