

Evaluating Complement-Modifier Distinctions in a Semantically Annotated Corpus

Mark McConville, Myroslava O. Dzikovska

School of Informatics, University of Edinburgh
2 Buccleuch Place, Edinburgh, EH8 9LW, Scotland
Mark.McConville@ed.ac.uk, M.Dzikovska@ed.ac.uk

Abstract

We evaluate the extent to which the distinction between semantically core and non-core dependents as used in the FrameNet corpus corresponds to the traditional distinction between syntactic complements and modifiers of a verb, for the purposes of harvesting a wide-coverage verb lexicon from FrameNet for use in deep linguistic processing applications. We use the VerbNet verb database as our gold standard for making judgements about complement-hood, in conjunction with our own intuitions in cases where VerbNet is incomplete. We conclude that there is enough agreement between the two notions (0.85) to make practical the simple expedient of equating core PP dependents in FrameNet with PP complements in our lexicon. Doing so means that we lose around 13% of PP complements, whilst around 9% of the PP dependents left in the lexicon are not complements.

1 Introduction

The distinction between complements and modifiers is important for parsing and language interpretation, since it impacts upon both syntactic and semantic decisions. On the syntactic level, the distinction relates to questions such as whether a particular sentence is syntactically correct if a given dependent is not realised. On the semantic level, the complement-modifier distinction raises important issues such as whether a semantic representation is ‘complete’, with nothing needing to be inferred from context, and whether a particular preposition denotes an independent predicate, as opposed to being a contentless argument marker. The answers to these questions may significantly impact the quality of parsing and semantic interpretation. As noted by Meyers et al. (1996), incorrectly classifying a complement as a modifier may cause a syntactic parser to miss a parse, whilst incorrectly classifying a modifier as a complement may cause a parser to add a spurious parse. This would be a relevant distinction to make in a (deep) syntactic parser. In addition, in an accurate representation of predicate argument structure, syntactic phrase heads predicate of their complements, but modifiers predicate of the syntactic phrase heads. This distinction is therefore relevant for both deep parsers that combine syntax and semantics, and shallow semantic parsers that attempt to induce a predicate structure based on semantic role labelling.

Unfortunately, the precise boundary between complements and modifiers is notoriously difficult to define. Existing lexical semantic resources, for example PropBank (Palmer et al., 2005), FrameNet (Johnson and Fillmore, 2000), VerbNet (Kipper et al., 2000) and OntoNotes (Hovy et al., 2006), do include information related to the complement-modifier distinction, but each applies slightly different criteria, depending on whether the emphasis is syntactic or semantic. A number of recent projects have attempted to merge information from different resources (Kwon and Hovy, 2006; Crabbé et al., 2006) and use them in parsing (Shi and Mihalcea, 2005; McConville and Dzikovska, 2007). For such applications it is important to be able to

understand and evaluate to what extent the different approaches to making complement-modifier distinctions are compatible, and which approach is the most appropriate for a given application.

In this paper, we investigate whether the semantic criteria for distinguishing complements and modifiers used by the creators of the FrameNet corpus (i.e. ‘core’ versus ‘non-core’ semantic roles) correspond to syntactic intuitions, in particular to the (primarily) syntactic criteria used in the VerbNet lexicon, which only lists syntactic arguments of the verbs, based on whether they can participate in a number of syntactic alternations.

We show that while there is a reasonably good correlation (0.85 agreement) between the semantic ‘coreness’ of FrameNet verb dependents and the complements listed in VerbNet, these notions do not align perfectly, and discuss the implications for using FrameNet as a source of syntactic information for parsing. We argue that for deep parsers concerned with both syntactic and semantic representations, it may be beneficial to separate the semantic and syntactic aspects of complement/modifier distinction into separate features.

This paper proceeds as follows: **section 2** presents some necessary background, including an introduction to both FrameNet and VerbNet, as well as our lexical harvesting project; **section 3** discusses the methodology used in our investigation, in particular the way in which we combined use of VerbNet with our own linguistic intuitions in making a decision about complement-hood; **section 4** presents the results; and **section 5** discusses the implications for deep parsing.

2 Background

2.1 The FrameNet corpus

FrameNet (Johnson and Fillmore, 2000) is a corpus of 140,000 English sentences (mainly drawn from the BNC), each annotated with both syntactic and semantic information. Underlying the corpus is an ontology of 800 ‘frames’ (or semantic types), each of which is associated with a set

of ‘frame elements’ (or semantic roles). Take for example the following sentence from the corpus:

- (1) Overshadowed by Grigorovich, Kokonin nonetheless apparently eclipsed him in power in recent months.

In this example, the verb *eclipse* is associated with the *Surpassing* frame, which denotes situations where one entity is conceptualised as being superior to another in some way. This frame includes the following frame elements, among others:

- *Attribute* - a property that invokes a scale (e.g. ‘power’, ‘wealth’)
- *Item* - the entity located closest to the end of the scale (i.e. the ‘surpassor’)
- *Standard* - the entity located farthest from the end of the scale (i.e. ‘the surpassed’)
- *Time* - the time when the *Item* is higher on the scale

Other verbs which are listed in the FrameNet ontology as ‘evoking’ the *Surpassing* frame include *surpass*, *better*, *outdo* and *outshine*.

The FrameNet annotation process then runs as follows, assuming the example sentence in (1):

1. identify a target word for the annotation, for example the main verb *eclipsed*
2. identify the semantic frame which is evoked by the target word in the sentence - in this case the relevant frame is *Surpassing*
3. identify the sentential constituents which realise each frame element associated with the frame, i.e. *Overshadowed by Grigorovich, [Kokonin]_{Item} nonetheless apparently eclipsed [him]_{Standard} [in power]_{Attribute} [in recent months]_{Time}*.

Finally, some basic syntactic information about the target word and the constituents realising the various frame elements is also added:

- the part-of-speech of the target word (i.e. V, N, A or PREP)
- the syntactic category of each constituent realising a frame element (e.g. NP, PP, VP_{to}, S_{fin})
- the syntactic role, with respect to the target word, of each constituent realising a frame element (for example *Ext* (subject), *Obj* (object) or *Dep* (other dependent))

Thus, each sentence in the corpus can be seen to be annotated on at least three independent ‘layers’, as exemplified in Figure 1. The FrameNet corpus has proved to be a useful linguistic resource for a number of computational linguistics applications, for example semantic role labelling (Gildea and Jurafsky, 2002), information extraction (Surdeanu et al., 2003), and question answering (Kaisser and Webber, 2007).

2.2 Harvesting a verb lexicon from FrameNet

McConville and Dzikovska (2007) present a procedure for harvesting a wide-coverage verb lexicon, for use with a deep semantic parser, from the FrameNet corpus. The technique used was to read off lexical entries from annotated sentences, and subsequently filter out spurious subcategorisation frames. We took each sentence which had been annotated with respect to some target verb and converted it into a lexical entry whose subcategorisation frame contained the various annotated syntacto-semantic dependents. For example, the annotated sentence from Figure 1 was converted into the verb entry in Figure 2.¹

This simple approach to deriving a verb lexicon gave rise to a number of spurious subcategorisation frames, involving non-canonical verbal constructions and alternations (e.g. passives, imperatives, middles), which we then had to filter out. After the filtering process, the harvested lexicon had been reduced in size from 30,000 distinct verb/subcategorisation frame pairs to just 9,000. These were distributed across 2,600 verb senses, giving a ratio of 3.4 subcategorisation frames per verb sense. The process of filtering out spurious subcategorisation frames added value to the original FrameNet verb lexicon, making it more suitable for hooking up to a deep parser where non-canonical constructions are generally handled in the rule component (e.g. using lexical rules).

One of the main issues we encountered in filtering the FrameNet verb lexicon involved distinguishing between those *Dep* dependents which are true complements of the verb and those which are generic modifiers, and filtering out the latter. To return to the example sentence in (1), the prepositional phrase *in power* is a complement of the target verb *eclipse*, since:

- although it is optional, it defines an argument specific to this class of verb, hence its existence cannot be predicted from more general principles of grammar
- the preposition *in* is the only preposition which can be used to introduce the *Attribute* role of the verb *eclipse*

On the other hand, the prepositional phrase *in recent months* is a modifier of the target verb, since:

- it can be used with much the same meaning with almost all classes of verb, hence its existence (and optionality) can be predicted from more general principles of grammar
- the preposition *in* is *not* the only preposition which can be used to introduce the *Time* role — *Kokonin eclipsed Grigorovich {at the weekend, over three years, after a few days, on Saturday, ... }*

Thus, these two prepositional phrases need to be treated differently in parsing and construction of logical forms, as discussed in the introduction. In particular, we need to make sure that lexical subcategorisation frames in our harvested

¹Note that the most recent version of the FrameNet corpus includes these automatically-generated subcategorisation frames in lexical unit files.

	<i>Kokonin</i>	<i>apparently</i>	<i>eclipsed</i>	<i>him</i>	<i>in power</i>	<i>in recent months</i>
target			Surpassing			
frame element	Item			Standard	Attribute	Time
syntactic category	NP		V	NP	PP	PP
syntactic role	Ext			Obj	Dep	Dep

Figure 1: A FrameNet annotated sentence

ORTH	<i><eclipse></i>							
SYNCAT	V							
SEMTYPE	Surpassing							
ARGS	$\left[\begin{array}{l} \text{SYNROLE Ext} \\ \text{SYNCAT NP} \\ \text{SEMROLE Item} \end{array} \right]$		$\left[\begin{array}{l} \text{SYNROLE Obj} \\ \text{SYNCAT NP} \\ \text{SEMROLE Standard} \end{array} \right]$		$\left[\begin{array}{l} \text{SYNROLE Dep} \\ \text{SYNCAT PP} \\ \text{SEMROLE Attribute} \end{array} \right]$		$\left[\begin{array}{l} \text{SYNROLE Dep} \\ \text{SYNCAT PP} \\ \text{SEMROLE Time} \end{array} \right]$	

Figure 2: A harvested verb entry

verb lexicon contain only those dependents which are subjects or complements, with modifiers filtered out.

In order to do this simply and straightforwardly, we availed ourselves of one of the features built in to the FrameNet ontology — the ‘coreness’ feature on frame elements.

In the ontology, the frame elements associated with each frame have been partitioned into two main groups, *Core* and *non-Core*, a distinction which (according to the annotation guidelines) is meant to cover the ‘semantic spirit’ of the distinction between complements and modifiers. Thus, for example, obligatory complements are always *Core*, as are:

- those which, when omitted, receive a definite interpretation, e.g. the *Goal* argument of the verb *arrive* (cf. *John arrived*)
- those whose semantics cannot be predicted from their form, e.g. the *Intermediary* argument of the verb *relies* as in *John relies [on Mary]* — the preposition *on* does not encode the *Intermediary* role with any other class of verbs outwith the *Reliance* frame

Non-Core frame elements are themselves partitioned into two classes:

Peripheral frame elements which do not introduce additional or distinct events from the main reported event, i.e. *Time*, *Manner*, *Place*, *Degree*, etc.

Extra-thematic situate an event against a backdrop of another state-of-affairs, i.e. *Frequency*, *Containing_event*, *Beneficiary*, etc.

To go back to the example sentence in Figure 1 involving the verb *eclipse* from the *Surpassing* frame, the FrameNet ontology classes the *Item*, *Standard* and *Attribute* roles as *Core*, and the *Time* role as *Peripheral*. This appears to correspond exactly with linguistic intuitions about which dependents are complements and which are modifiers of the verb. For this rea-

son, when it came to filtering out spurious subcategorisation frames from the verb lexicon we had harvested from the FrameNet corpus, we decided on the simple expedient of deleting all and only those *Dep* dependents which evoke a *non-Core* frame element. This process resulted in the elimination of many spurious subcategorisation frames — the number of verb/subcategorisation frame pairs in the harvested lexicon was cut by 45%, from 16,000 down to 9,000. However, it was clear from the start that the correlation between the semantic ‘coreness’ and syntactic complementhood is far from perfect. For instance, it was not difficult to find examples where syntactic complements evoked semantically *non-Core* frame elements. A number of constituents in the FrameNet corpus have been marked as direct objects, despite invoking *non-Core* frame elements, as in:

- (2) $[John]_{\text{Agent}} \textit{ripped} [the\ top]_{\text{Subregion}} [from\ his\ packet\ of\ cigarettes]_{\text{Patient}}$

In this instance, the verb *rip* has been assigned by the annotators to the frame *Damaging*, where the *Subregion* frame element is marked as being *Peripheral*, based on examples like *John ripped his trousers [below the knee]*. In this particular case, the problem was probably caused by annotators not being careful enough when assigning verbs with different subcategorisation alternations to frames — it would have been better to have assigned the verb *rip* as used in (2) to the *Removing* frame, where the direct object invokes a *Core* frame element (i.e. *Theme*). Thus, the decision to retain all senses of the verb *rip* within the same frame has led to a situation where semantic and syntactic coreness have become dislocated.

The aim of the project reported here was to investigate the extent to which this kind of problem impacts upon the effectiveness of the expedient we chose to distinguish complements from modifiers in the verb lexicon we harvested from the FrameNet corpus. In other words, we wanted to ascertain to what extent the ‘coreness’ feature on frame elements in the FrameNet ontology corresponds with linguis-

tic intuitions as to which dependents are complements and which are modifiers.

2.3 The VerbNet verb lexicon

Unfortunately, judgments on which verb dependents are complements and which are adjuncts are notoriously difficult to make consistently. Although there are many cases which are clearly complements and many which are clearly modifiers, there are a large number of borderline cases where it is hard to make any kind of definite decision either way. As noted by the creators of the Penn TreeBank (Marcus et al., 1994): “After many attempts to find a reliable test to distinguish between arguments and adjuncts, we have abandoned structurally marking this difference”. The situation is muddled still further by the fact that certain of the suggested criteria found in the literature are plagued by issues of gradient grammaticality. For example, Meyers et al. (1996) states that pseudo-passivisation (e.g. *John is relied on by many people*) is a property of complement PPs but not modifier ones. However, Tseng (2006) points out that there is a continuum of acceptability for pseudo-passives, and many PPs which are clearly modifiers can be pseudo-passivised, e.g. *David always takes that seat in the corner because he hates being sat next to*.

In order to help us in deciding whether a given verb dependent is really a complement, we decided to make use of the VerbNet verb lexicon (Kipper et al., 2000) as a syntactic gold standard.

VerbNet is a lexicon of around 5,000 English verb senses, partitioned into 237 top-level classes. Each verb class specifies, among other things, a set of associated subcategorisation frames listing the arguments (i.e. subjects and complements) that are appropriate for all the verbs in the class. Take for example the VerbNet class *exceed-90*, which is the nearest equivalent to the *Surpassing* frame in FrameNet. This class encompasses verbs like *surpass*, *top* and *outstrip*, and specifies the following two subcategorisation frames:²

- NP:Theme1 V NP:Theme2
- NP:Theme1 V NP:Theme2 P NP:Attrib

In order to use VerbNet as our gold standard for making distinctions between complements and modifiers, we assumed that verb dependents which are listed in the relevant class are definitely complements. However, we were unable to assume straightforwardly that any verb dependents which are *not* so listed are modifiers, since VerbNet is an incomplete resource, not being strictly corpus-based. We thus report two sets of results: those where we assume that VerbNet is a literal gold standard (i.e. pretend that it is complete), and those where we allow ourselves to make use of other criteria in deciding whether a dependent which is not listed in VerbNet is a complement or modifier.

In addition to VerbNet, we considered ComLex (Grishman et al., 1994) and PropBank as possible sources of syntactic information. We decided that ComLex unsuitable because

²VerbNet subcategorisation frames are best thought of as flattened representations of LTAG elementary trees.

it is a purely syntactic resource and does not specify semantic roles for the complements listed in its subcategorisation frames, thus making it impossible to decide in many cases whether a given FrameNet dependent is listed or not. In addition, the PropBank proved to be of little use, since, as pointed out in Palmer et al. (2005), ‘We make no attempt to adhere to any linguistic distinction between arguments and adjuncts’.

3 Methodology

We took the lexicon we had harvested from the FrameNet corpus³, and extracted all and only those entries (incorporating an orthographic base form, a semantic type, and a subcategorisation frame) which specify at least one PP dependent. Along with the information about annotated dependents, each entry was also associated with the corpus sentence which it had been harvested from.

A total of 17,000 verb entries were extracted in this way. The next step was to select a random sample of these for manual checking. Unfortunately, this proved to be somewhat more problematic than just picking a random subset. The FrameNet project’s approach to annotation has proceeded on a ‘frame-by-frame’ basis rather than focusing on fully annotating running text, meaning that each frame in turn is fully annotated with respect to its lexical units, before moving on to the next identified frame. As a result, some frames contain many more lexical units than others, and thus are associated with many more annotated instances. Thus, of the 261 FrameNet frames implicated in our set of extracted verb entries, the most common 10% account for 65% of the total number of entries. In particular, just the one frame (*Self_motion*) accounts for a quarter of all the extracted entries.

We attempted to counteract this bias by limiting each frame to a maximum of *two* entries, and moreover restricting each verb in a frame to a maximum of *one* entry. We were thus left with 593 verb entries in our sample, involving a total of 432 subcategorised *Core* PP dependents and 204 non-*Core* ones.

The annotation task involved going through each PP dependent in turn and deciding whether or not VerbNet classifies it as a complement of the target verb. The relevant decision tree runs as follows:

1. If the verb is listed in VerbNet with the appropriate sense:
 - if the VerbNet class lists the relevant subcategorisation frame, including the PP dependent in question, then mark the dependent as a ‘complement’
 - if the VerbNet entry *does not* list the relevant subcategorisation frame, including the PP dependent in question:
 - if you think that a more complete VerbNet entry for the relevant verb would list the relevant subcategorisation frame, including the PP dependent

³More precisely, the lexicon after having removed all spurious frames involving non-canonical constructions like passives and imperatives, but retaining *all* dependents, whether *Core* or non-*Core*.

in question, then mark the dependent as a ‘complement’

- if not, then mark the dependent as a ‘non-complement’

2. If the verb is *not* listed in VerbNet with the appropriate sense:

- if you think that a more complete VerbNet which *did* contain an entry for the relevant verb would list the appropriate subcategorisation frame, including the PP dependent in question, then mark the dependent as a ‘complement’
- if not, then mark the dependent as a ‘non-complement’

Thus, we first of all determined whether the relevant sense of the target verb was included in VerbNet. If so, and if a matching subcategorisation frame including the dependent as a complement was listed, then it was deemed to be a complement PP.

In cases where either the relevant sense does not exist in VerbNet, or where the relevant verb sense does exist but a matching subcategorisation frame is not listed, things become a little more complicated, due to the incomplete nature of VerbNet as a resource.

Take for example, the example sentence in Figure 2. Although the verb *eclipse* does not appear in VerbNet, the synonymous verb *surpass* does (cf. *John surpassed/eclipsed Mary in raw talent*). By close inspection of the subcategorisation frames listed in the VerbNet class in which *surpass* appears (i.e. *exceed-90*), we judge that the *Attribute* dependent of the verb *eclipse* in Figure 2 is a VerbNet complement — the class contains the subcategorisation frame NP:Theme1 V NP:Theme2 P NP:Attrib, where Theme1, Theme2 and Attrib match *Item*, *Standard* and *Attribute* respectively.

In cases where the relevant verb sense *does* appear in VerbNet, but there is no matching subcategorisation frame including the PP in question we applied our own linguistic intuitions (using standardly assumed criteria for distinguishing complements from modifiers) to judge whether a more complete version of VerbNet *should* list the PP dependent as a complement.

4 Results

One quarter of the verb entries in our sample were annotated by both authors, in order to test for inter-annotator agreement. The results were as follows:

- Is the appropriate sense of the target verb listed in VerbNet? Agreement: 0.95, kappa: 0.90
- Assuming both annotators agree that the target verb appears in VerbNet, is the PP dependent listed as a complement? Agreement: 0.97, kappa: 0.93
- Assuming that both annotators agree that the target verb appears in VerbNet but the PP dependent is not listed, is it a complement of the target verb? Agreement: 0.80, kappa: 0.60

- Assuming that both annotators agree that the target verb is *not* listed in VerbNet, is the PP dependent a complement of the verb? Agreement: 0.94, kappa: 0.87

The results of our investigation into the relation between coreness and complement-hood are presented in Table 1.

The three ‘experiments’ listed are as follows:

- **Experiment 1** only takes into account the 433 dependents whose verb senses were adjudged to be listed in VerbNet, ignoring annotator judgements about whether an unlisted PP dependent is a complement or not (i.e. it assumes that VerbNet is a complete resource for the verbs it lists)
- **Experiment 2** also assumes the subset of 433 dependents whose verb senses were adjudged to be listed in VerbNet, but includes annotator judgements for unlisted PP dependents
- **Experiment 3** includes results for all 634 PP dependents in the sample, including those whose target verbs do not appear in VerbNet

For each experiment, we present both the total assignments of dependents to classes (the columns represent FrameNet’s *Core* versus *non-Core* distinction and the rows represent the judgements about syntactic complement-hood made by the annotators, in conjunction with VerbNet) as well as the interannotator agreement and Cohen’s kappa scores. Note that agreement increases significantly when we take into account annotator judgement in cases where VerbNet fails to list the relevant PP dependent. Note also that the amount of chance agreement depends on the relative proportion of *Core* and *non-Core* dependents annotated in the FrameNet corpus. It appears that the former have been annotated more completely, since over the corpus as a whole there are twice as many *Core* PPs listed than *non-Core* ones — we were able to find many PP modifiers which had been ignored by annotators.

In conclusion, it appears that around 13% of syntactic PP complements will be lost if we simply delete all *non-Core* dependents from our harvested lexicon, and 9% of the dependents retained will *not* be syntactic complements.

Finally, we manually examined all instances where a FrameNet *Core* PP dependent was judged *not* to be a complement of the relevant target verb. A significant proportion of these (around one third) appeared to involve some kind of bracketing mismatch between syntax and semantics, where the *Core* dependent annotated in FrameNet is *not* a syntactic dependent of the target verb. Take the following example:

- (3) She looked away quickly, and unfastened [the waistband]_{FASTENER} [of her uniform skirt]_{CONTAINING-OBJECT}.

Here the target verb is *unfasten*. The FrameNet annotation recognises both *the waistband* and *of her uniform skirt* as distinct dependents of the verb. However, on the syntactic level it is the entire noun phrase *the waistband of her uniform skirt* which is a syntactic dependent of the verb (the direct object).

	Experiment 1		Experiment 2		Experiment 3	
	Core	non-Core	Core	non-Core	Core	non-Core
complements	199	37	258	49	395	59
non-complements	82	115	23	103	37	145
agreement	0.73		0.83		0.85	
kappa	0.65		0.75		0.65	

Table 1: Results

5 Discussion

As can be seen from our analysis, while the distinction between *Core* and *non-Core* semantic roles in the FrameNet ontology is highly correlated with the kind of syntactic criteria for complement-hood used in verb lexicons like VerbNet, the two do not align perfectly. In other words, there are a significant number of cases where either a *non-Core* semantic dependent is realised by a verbal complement or a *Core* semantic dependent is realised by a verbal modifier. Although many of these mismatches can be put down to either annotation errors (generally failing to assign a particular use of some verb to the optimal frame) or questionable annotation policy (e.g. annotating as semantic dependents phrases which are definitely not syntactic dependents of the target verb), there remain a large number of such mismatches which are simply a result of the incompatibility of semantic and syntactic notions of ‘coreness’.

The fact that the complement-modifier distinction is so difficult to pin down has important implications for parsing and semantic interpretation. As pointed out in the introduction, a parser needs to have access to a complete, accurate list of what complements go with which verbs, if correct parses are not to be missed or spurious parses to be added. Consider the following contrasting examples:

- (4) (a) John relied on the map
 (b) John fell on the table
 (c) John slept on the table

The semantic representations corresponding to possible interpretations of those utterances are shown in Table 2. In (4a), ‘rely’ *requires* an *on*-PP complement (denoting the thing relied upon). In (4b), the *on*-PP is optional syntactically, and can be either a complement or an adjunct semantically, depending on whether it denotes the trajectory of the fall, or the location where John fell. In (4c), the *on*-PP is syntactically and semantically optional, and it is definitely not a complement, since the location of the event is in no way unique to the *sleep* predicate. Unless the parser has access to this kind of information, then it will not be able to judge, for example, that the second sentence has two distinct interpretations, whereas the other two only have one interpretation each.

In addition, note that there is an important contrast between the *on*-PP complements of *rely* and *fall*. In the former, the preposition *on* does not contribute any meaning to the sentence, other than to make clear what role its NP complement plays in the situation (i.e. it is basically a case-marker, cf. *John trusted the map*). The intrinsic meaning

John	relied	on the map
Protagonist		complement Intermediary
John	fell	on the table
Theme		complement Goal
Theme		modifier Place
John	slept	on the table
Sleeper		modifier Place

Table 2: Possible semantic representations using FrameNet roles for utterances in example (4), and their relation to the complement/modifier distinction

of the preposition *on*, involving some object being in contact with the top horizontal surface of some other object, plays no role in the predicate argument structure of this sentence, which can best be represented as something like *rely_on(John,map)*. Furthermore, *on* is the only preposition which can be used here to encode the relevant semantic role (apart from its stylistic variant *upon*).

In the second sentence, on the other hand, the preposition *does* contribute its intrinsic meaning to the predicate argument structure — as the place where the trajectory of the fall ends. Thus, the predicate argument structure of this sentence is more like *fall(John,to(on(table)))*, with the inference that, at the culmination of the event of falling, John is in fact ‘on the table’. Also, a number of other prepositions can be used here to realise the relevant semantic role, e.g. *off*, *under*, *over*, *through*, etc.

The contrast between argument-marker and predicative uses of prepositions is important both for applications connecting language to reasoning, and for applications using shallow semantic representations. For example, Kaiser and Webber (2007) describe the use of FrameNet in question answering where questions are paraphrased using verbs in the same frame. If we consider the paraphrases that apply in our cases, an appropriate paraphrase for (4a) would be *John trusted the map*, leaving out the preposition *on*, while an appropriate paraphrase for (4b) would be *John dropped on the table*, where the preposition is retained, since it is crucial for the meaning.

With this in mind, we decided to investigate how different parsers handle the distinction between predicative

and non-predicative uses of prepositions. We considered: (a) shallow semantic parsers that output FrameNet frames (Gildea and Jurafsky, 2002); (b) the LinGO English Resource Grammar (Copestake and Flickinger, 2000), a deep HPSG grammar that produces semantic representations using first-order logic,⁴; and (c) the TRIPS parser Allen et al. (2007), a unification based parser for dialogue systems. These systems represent several different ways of using logical forms: semantic parsers are used for question answering (Kaisser and Webber, 2007) and information retrieval (Surdeanu et al., 2003); the LinGO ERG has been used for translating spoken dialogue (Kay et al., 1994); and the TRIPS parser produces semantic representations that are easy to map into representations used by domain specific reasoners common in dialogue systems.

It turns out that these three parsers have all taken different approaches to the distinction between predicative and non-predicative prepositions. Shallow semantic parsers output frame representations that identify frame element names, but not their specific meanings. So for (4a), they will identify *on the map* as a `Intermediary`, and in (4b) treat *on the table* as `Goal`, leaving open the question whether it is *the table* alone, or the whole PP, that fills the appropriate slot. The LinGO ERG parser represents some prepositions as argument markers, but not always consistently, e.g. *on* in (4a) will be represented as a case marker, but *for* in *John left for Boston* will be represented as a regular preposition, as will *on* in (4b). In contrast, the TRIPS dialogue parser makes strictly semantic decisions, and marks all PPs where the preposition is not clearly an argument marker as complements (thus ensuring that the the preposition predicates are included in the logical form), but at the expense of being unable to rule out certain syntactically anomalous utterances, such as **John put it*.

These differences in approach are not surprising, given that the various syntactic and semantic criteria for identifying complements may not align, as we have shown in this paper, and therefore creating a lexicon that accurately identifies such distinctions is a difficult task. We are currently working on methods to better represent this distinction in our lexicon.

We propose that for parsing lexicons, especially for deep parsers, a possible solution is to replace a single distinction with several finer-grained features, addressing the key issues raised in the introduction: Is some dependent syntactically required to complete the utterance? Is some (optional) prepositional or adjectival phrase a possible dependent of a given verb? And does a particular preposition correspond to an independent predicate (regardless of whether the dependent can be classified as a complement)?

The method adopted by VerbNet of defining syntactic frames based on alternations is a good approach to defining syntactic complements. It also takes the first step towards identifying those prepositions that may contribute meaning towards the predicate vs. those which do not, by defining classes of equivalent prepositions (such as location and di-

rection) which may appear in the same PP.

However, since VerbNet is incomplete, we are currently developing a corpus-based approach to answer these questions. We are particularly interested in the answer to the third question, whether a preposition in a given PP dependent is an argument marker, or contributes meaning to the logical form. This information is, to our knowledge, not coded in any existing resources. Adding it would provide essential information for building semantic representations, and therefore make such representations more usable in interpretation tasks.

6 Conclusion

In this paper, we have attempted to evaluate the extent to which the distinction between semantically `Core` and non-`Core` dependents, as used in the FrameNet corpus, corresponds to the traditional distinction between syntactic complements and modifiers of a verb. We used the VerbNet verb database as our gold standard for making judgements about complement-hood, in conjunction with our own intuitions in cases where we considered VerbNet to be incomplete. We concluded that there is enough agreement between the two notions (0.85) to make practical the simple expedient of equating core PP dependents in FrameNet with PP complements in the wide-coverage verb lexicon we harvested from FrameNet. Doing so means that we lose around 13% of PP complements, whilst around 9% of the PP dependents left in the lexicon are not complements. We then discussed the implications of this result for deep parsing, suggesting that for parsers concerned with both syntactic and semantic representations, it may be beneficial to separate the semantic and syntactic aspects of complement/modifier distinction into separate features

Acknowledgements

The work reported here was supported by grants N000140510043 and N000140510048 from the Office of Naval Research.

7 References

- James Allen, Myroslava Dzikovska, Mehdi Manshadi, and Mary Swift. 2007. Deep linguistic processing for spoken dialogue systems. In *Proceedings of the ACL'07 Workshop on Deep Linguistic processing*, pages 49–56.
- Ann Copestake and Dan Flickinger. 2000. An open source grammar development environment and broad-coverage English grammar using HPSG. In *Proceedings of the 2nd International Conference on Language Resources and Evaluation*, Athens, Greece.
- Benoit Crabbé, Myroslava O. Dzikovska, William de Beaumont, and Mary D. Swift. 2006. Increasing coverage of a domain independent dialogue lexicon with VerbNet. In *Proceedings of the Third International Workshop on Scalable Natural Language Understanding (ScaNaLU 2006)*, New York City.
- Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational Linguistics*, 28(3):245–288.

⁴To be more precise, the LinGO ERG produces MRS representations that encode scope ambiguities. However, MRS representations resolve to fully instantiated logic formulas, which are our concern here

- Ralph Grishman, Catherine Macleod, and Adam Meyers. 1994. COMLEX syntax: Building a computational lexicon. In *Proceedings of COLING'94*.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. Ontonotes: The 90% solution. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA, June. Association for Computational Linguistics.
- Christopher Johnson and Charles J Fillmore. 2000. The FrameNet tagset for frame-semantic and syntactic coding of predicate-argument structure. In *Proceedings ANLP-NAACL 2000*, Seattle, WA.
- Michael Kaisser and Bonnie Webber. 2007. Question answering based on semantic roles. In *Proceedings of the ACL'07 Workshop on Deep Linguistic processing*.
- Martin Kay, Jean Mark Gawron, and Peter Norvig. 1994. *Verbmobil: A Translation System for Face-To-Face Dialog*. CSLI Press, Stanford, California.
- Karin Kipper, Hoa Trang Dang, and Martha Palmer. 2000. Class-based construction of a verb lexicon. In *Proceedings of the 7th Conference on Artificial Intelligence (AAAI-00) and of the 12th Conference on Innovative Applications of Artificial Intelligence (IAAI-00)*, pages 691–696, Menlo Park, CA, July 30– 3. AAAI Press.
- Namhee Kwon and Eduard H. Hovy. 2006. Integrating semantic frames from multiple sources. In Alexander F. Gelbukh, editor, *CICLing*, volume 3878 of *Lecture Notes in Computer Science*, pages 1–12. Springer.
- 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 *Proceedings of the ARPA Human Language Technology Workshop*.
- Mark McConville and Myroslava O. Dzikovska. 2007. Extracting a verb lexicon for deep parsing from FrameNet. In *Proceedings of the ACL-07 Workshop on Deep Linguistic Processing*, Prague, Czech Republic, June.
- Adam Meyers, Catherine Macleod, and Ralph Grishman. 1996. Standardization of the complement/adjunct distinction. In *Proceedings of the Seventh Euralex International Congress, Gothenburg, Sweden*.
- Martha Palmer, Paul Kingsbury, and Daniel Gildea. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.
- Lei Shi and Rada Mihalcea. 2005. Putting pieces together: Combining framenet, verbnet and wordnet for robust semantic parsing. In *Proceedings of the Sixth International Conference on Intelligent Text Processing and Computational Linguistics*, Mexico.
- Mihai Surdeanu, Sanda M. Harabagiu, John Williams, and Paul Aarseth. 2003. Using predicate-argument structures for information extraction. In *Proceedings of ACL'03*, pages 8–15.
- Jesse Tseng. 2006. English prepositional passives in HPSG. In *Proceedings of Formal Grammar'06*, pages 147–159.