Building a computational lexicon and ontology with FrameNet

Myroslava O. Dzikovska[*], Mary D. Swift[†], James F. Allen[†]

[*] Human Communication Research Centre
University of Edinburgh, 2 Buccleuch Place, Edinburgh EH8 9LW, United Kingdom
m.dzikovska@ed.ac.uk
[†] Department of Computer Science
University of Rochester, Rochester, NY 14627-0226
{james,swift}@cs.rochester.edu

Abstract

This paper explores FrameNet as a resource for building a lexicon for deep syntactic and semantic parsing with a practical multiple-domain parser. The TRIPS parser is a wide-coverage parser which uses a domain-independent ontology to produce semantic interpretations in 5 different application domains. We show how semantic information from FrameNet can be useful for developing a domain-independent ontology. While we used FrameNet as a starting point for our ontology development, we were unable to use FrameNet directly because it does not have links between syntax and semantics, and is not designed to include selectional restrictions. We discuss changes that needed to be made to the FrameNet frame structure to convert it to our domain-independent LF Ontology, the additions we made to FrameNet lexicon, and the resulting differences between the systems.

1. Introduction

This paper explores FrameNet(Johnson and Fillmore, 2000) as a resource for building a lexicon for deep syntactic and semantic parsing with a practical multiple-domain parser. Semantic corpus annotation such as FrameNet is an important way to ensure reliability and ease of use of semantic representations. Achieving inter-annotator agreement results in semantic classes that can be reliably distinguished by humans, unlike, for example, WordNet synsets (Miller, 1995), which are often difficult to differentiate for human annotators. An open question, however, is whether the FrameNet classes and frame elements can be obtained and used automatically. There has been some work in this area, in particular, on learning FrameNet frame elements from corpora (Gildea and Jurafsky, 2002) and on using them in the SMARTKOM project (Chang et al., 2002). However, the extent to which FrameNet annotations will be usable in practical applications is still an open issue.

In this paper, we describe our experience in using FrameNet in the process of building a multi-domain conversational dialogue system. The TRIPS system is a dialogue assistant which has been applied to 5 different application domains. Our lexicon uses frame structures as a domain-independent semantic representation, and therefore FrameNet is an attractive source of semantic information. We used the FrameNet classes as a starting point for our ontology development.

We made our top-level ontology for parsing consistent with the FrameNet ontology, and this helped us by identifying the verb classes that can be reliably distinguished by human lexicon developers when defining entries in a computational lexicon. FrameNet also provides semantic roles, but it does not provide links between lexical entries and the frames, and it does not contain selectional restrictions. In creating those links, we changed the representation in order to simplify lexicon maintenance, making it easier to define syntax-semantics mappings and selectional restrictions in the lexicon and ontology.

We describe the needs of a wide coverage parser and grammar using the TRIPS parser as a realistic example in Section 3; we then discuss the changes that needed to be made in our domain-independent ontology from the FrameNet formalism (Section 4), and compare the resulting lexicons (Section 5). Our experience can be useful for the designers of other NLP systems, as well as guidance for further development of semantic annotation schemes which can be used in natural language understanding systems.

2. Background

Typically, a parsing and semantic interpretation system requires an ontology as a source of semantic types and a lexicon with the following information for every word:

- Syntactic features;
- Subcategorization frames;
- Semantic representation;
- For every subcategorization frame, the correspondence between syntactic and semantic structures.

A number of lexicon and ontology projects provide parts of the necessary information. Among the resources frequently used for natural language processing tasks are syntactic features and subcategorization frames in COMLEX (Macleod et al., 1994), word senses in WordNet (Miller, 1995) and EuroWordNet (Vossen, 1997), and semantic representations of world knowledge in CyC (Lenat, 1995). Of particular interest to our project is FrameNet, which provides semantic frame representations based on the analysis of corpus examples, and VerbNet (Kipper et al., 2000), which provides subcategorization frames and correspondence between those and verb semantics.

Though each of these lexicons and ontologies provides some of the requirements we listed above, there is no single resource which integrates all the information necessary for parsing. We found that FrameNet and VerbNet entries were the most useful for our purposes, as we discuss in more detail in the following sections. Integration of all the required information presents significant challenges, primarily in making sure that during parsing the correct semantic type can be chosen for the word, and correct semantic argument labels are assigned to all its arguments. We found that in a practical system simplifications may be necessary to achieve efficiency and accommodate the fact that the system cannot rely on the world knowledge available to humans annotating corpus examples.

3. The TRIPS parser

Before describing the use of FrameNet in the TRIPS ontology, we discuss in more detail the TRIPS parser and its representational requirements. The TRIPS parser is a chart parser which utilizes 3 main knowledge sources: a wide-coverage domain-independent grammar, a domain-independent lexicon, and a domain-independent ontology, as elaborated below.

Our wide-coverage domain-independent grammar has been developed and tested in 5 different spoken dialogue domains. It has been tested on human-human speech corpora (Swift et al., 2004), and provides good coverage of complex structures including gaps, relative clauses, complex noun phrases *etc*. The grammar rules build up a domain-independent logical form used for discourse processing, discussed below.

Our domain-independent lexicon provides word definitions for the grammar. Each word definition has to include the syntactic features, subcategorization frames and the linking between syntax and semantics to allow the parser to build the logical form. While our lexicon is not yet as large as the projects like WordNet, it offers wide coverage in our domains, which results in many ambiguous lexical entries. On average, there are 1.26 syntax-semantics patterns per word, and for verbs this figure is 1.60. The ambiguity in lexical entries necessitates the development of mechanisms for semantic disambiguation. In our project, we use domain-independent selectional restrictions expressed as feature sets as our primary disambiguation mechanism.[1]

Finally, our domain-independent ontology, which we call the *LF Ontology*, is the source of semantic types that provides the semantics for entries in the domain-independent lexicon.It includes the repository of all semantic types defined in the system, as well as selectional restrictions to help disambiguation. The relationship between the LF Ontology and FrameNet is discussed in the rest of the paper.

Using the domain-independent grammar and lexicon linked to the LF ontology, the TRIPS parser produces a domain-independent logical form. This is a flat unscoped neo-Davidsonian representation, using event arguments and semantic roles. It is similar to QLF (Alshawi et al., 1991) and Minimal Recursion Semantics (Copestake et al., 1995) in that it uses identifiers to link the (non-recursive) terms together. An example representation for *load the truck with oranges* is shown in Figure 1.

(SPEECHACT sa1 SA_REQUEST :content e123)
(F e123 LF::Filling*load :Agent pro1 :Theme v1 :Goal v2)
(IMPRO pro1 LF::Person :context-rel *YOU*)
(THE v1 (SET-OF LF::FOOD*orange))
(THE v2 LF::Vehicle*truck)

Figure 1: The LF representation of the sentence *load the oranges into the truck.*

The representation identifies the sense of the main verb *load* as an instance of concept LF::Filling, corresponding to the FrameNet frame *filling*. Moreover, it identifies *oranges* as a :Theme of the filling action, that is, the object being moved, and *truck* as a :Goal of the filling action. Since it is an imperative, the parser also infers an implicit pronoun as a subject of the sentence, corresponding to the :Agent role.

Unlike traditional QLF representations, which typically use n-place predicates, we use named arguments (which we call **semantic roles**) in our representations, as it is done in neo-Davidsonian representations and description logic. It makes it easier to provide uniform representations connected to different syntactic alternations (*e.g.*, the only difference between *the window broke* and *the hammer broke the window* is that the former does not have an instrument role filled in), and we hope to be able to use the role-based representations for some syntactic generalizations, as discussed in Section 6.

In this example, the role names defined for LF::Filling are exactly the same as those for the filling frame in FrameNet. This is not always the case, and the need to change the role structure for the LF ontology is discussed in Section 4.

In the rest of the section, we discuss the specific requirements the parsing system places on its lexicon and domain-independent ontology. These are the motivations for choosing FrameNet as an appropriate domain-independent ontology, but also for the changes needed for its use in a computational system.

3.1. Ontology design considerations

When providing the semantic information for parsing described in the previous section, the development of our system is influenced by two main goals: support for efficient wide-coverage parsing, and also fast lexicon acquisition. The first requirement means that the information provided in the lexicon should be sufficient to parse sentences encountered in the domains quickly. Therefore, we need to reduce the parser search ambiguity whenever possible while maintaining the wide coverage of the system. The second requirement means that new word definitions should be possible to define automatically, or, if defined by hand (as we are currently doing), the information necessary to define a lexical entry should be easy to obtain. Either the lexicon developer should be able to define a word from the examples of other similar words already defined in the lexicon, or, if no similar words were defined previously, the relevant information should be easy to obtain from online

[1] Another option would be statistical disambiguation, but it proves difficult for spoken dialogue domains, where corpus data are difficult and costly to collect. We have demonstrated that domain-independent selectional restrictions improve parsing speed and accuracy in our lexicon (Dzikovska, 2004).

resources. In particular, we would like to be able to obtain the semantic class of the word from FrameNet, and then find a way to link the syntactic structure with the frame elements.

Specifically, our decisions about the ontology were influenced by the following considerations:

- **The level of abstraction**. The semantic predicates used during interpretation must be specific enough to allow the system to draw reasonable inferences about the world. For example, using the same predicate Move to denote verbs such as *run*, *walk* and *drive* loses important distinctions between the meanings, such as speed and whether a vehicle is involved. At the same time, we want the semantic predicates to be such that the system has a reasonable chance of selecting the correct sense during the interpretation process. For example, WordNet lists 16 senses for the verb move, including "change location", "move as so to change position", "cause to move" and "change residence". Disambiguation between those senses is difficult even for human annotators, and extensive reasoning about context is necessary to select the correct sense is not feasible given the current state of the art for dialogue systems. FrameNet offers the appropriate level of abstraction for word senses, as discussed below.

- **The compositionality of meaning representations** In a domain-independent ontology, we would like the meanings of the complex phrases to be compositional, built from the meanings of their components. For example, consider a sentence *Submit a purchase order*. In a system that only knows about submitting purchase orders, this is an atomic action. Therefore, it can potentially be represented as a single concept in the system ontology, SubmitPurchaseOrder(p), where *p* is parameter which corresponds to the purchase order to submit. This representation may be the most efficient for domain reasoning, but if there are other things that can be submitted, such as proposals or application, this leads to a proliferation of concepts: SubmitProposal(p), SubmitApplication(a). This is not a desirable situation for parsing, because it results in additional ambiguity in constructions like *submit it*, which then become multiply ambiguous between interpretations with different possible meanings of *submit*.

- **Efficiency**. For a dialogue system, the speed of interpretation is crucial for effective operation, and we would like to use as much semantic information as possible during parsing to speed up and improve disambiguation.

- **Syntax-semantics mappings**. In order to use an ontology in a parsing system, we need to be able to link the syntactic structures to corresponding ontological representations. This needs to be specified in our lexicon; ideally, it should be available directly from a lexicon developed together with the ontology, otherwise, it needs to be acquired later, during construction of our parsing lexicon. The properties of the ontology, including the level of abstraction and compositionality, and also the arguments associated with each type, should facilitate syntax-semantics mapping. For example, if an ontology requires collecting phrases like *from Pittsford* and *to Avon* into a single PATH frame, then special handling for path adverbials has to be implemented in the grammar, adding to the complexity of the system. FrameNet has simple frame elements, which are easy to obtain during parsing. However, there are issues with disambiguating them, as discussed in Section 4.

In our analysis, the FrameNet frames offer the right level of abstraction for a computational system. The guideline we use in our lexicon is to consider two senses of a word different only if we can distinguish them automatically (i.e. based on subcategorization patterns and domain-independent selectional restrictions) in most circumstances. Because FrameNet was developed based on corpus examples, with frames which can be reliably distinguished by human annotators, the frame structures offer the right level of abstraction as word senses in a computational system. In addition, because the frames are expected to cover a large number of examples, they offer a good level of compositionality, representing generic situations with parameters to be filled in the roles.

3.2. Syntax-semantics templates and the LF Ontology

FrameNet is missing a crucial piece of information - syntax-semantics mappings, which are necessary to obtain our logical form representations. An example lexical entry in our lexicon is shown in Figure 2. It defines the verb *load* and 2 syntactic patterns. The pattern defined by AGENT-THEME-GOAL-TEMPL encodes the information that in a sentence *Load the oranges into the truck* the (implicit) subject will fill the :Agent role, the direct object is a noun phrase which will fill the :Theme role, and the prepositional complement is a prepositional phrase using the preposition *into*, and filling the :Goal role.

The syntax-semantic mappings have to be defined for all lexical entries. In defining them, we encounter issues with semantic role names similar to those we encountered when defining appropriate word senses. When a mapping between syntactic and semantic arguments is defined, the semantic arguments in the given frame must be defined on a level of abstraction appropriate to draw inferences about the world, but possible to disambiguate based on syntactic structure and selectional restrictions. We found that some FrameNet frame elements did not satisfy those criteria, which necessitated changes to the ontology structure discussed in Section 4.

The syntax-semantics templates are combined with selectional restrictions in our ontology to provide semantic disambiguation. Selectional restrictions are not part of the FrameNet database, we added them to our LF representation to provide the parser with the information necessary for disambiguation. For example, the LF ontology entry for LF::Filling is shown in Figure 3. It is a subtype of a more general LF::Motion frame (the addition of hierarchical structure to the LF Ontology is discussed in the next

```
(a) (load
     (wordfeats (morph (:forms (-vb))))
     (senses
       ((LF-Parent LF::Filling)
        (TEMPL AGENT-THEME-GOAL-TEMPL)
        (Example "Load the oranges into the truck"))
       ((LF-parent LF::Filling)
        (TEMPL AGENT-GOAL-THEME-TEMPL)
        (Example "Load the truck with oranges"))
       ) ))

(b) (AGENT-THEME-GOAL-TEMPL
     (SUBJ (NP) Agent)
     (DOBJ (NP) Theme)
     (COMP (PP (ptype into)) Goal optional)

(c) (AGENT-GOAL-THEME-TEMPL
     (SUBJ (NP) Agent)
     (DOBJ (NP) Goal)
     (COMP (PP (ptype with)) Theme)
```

Figure 2: Defining words in the lexicon (a) Lexicon definitions for the verb *load* in the LF::Filling sense; (b) The template used to define the syntactic pattern for *load the oranges into the truck* (c) The template used to define the syntactic pattern for *load the truck with oranges*

section). As such, it inherits a basic set of arguments, which are :Theme, :Source and :Goal.

```
(define-type LF::Motion
 :sem (Situation (Aspect Dynamic))
 :arguments
   (Theme (Phys-obj (Mobility Movable)))
   (Source (Phys-obj))
   (Goal (Phys-obj))

(define-type LF::Filling
 :parent LF::Motion
 :sem (Situation (Cause Agentive))
 :arguments
   (Agent (Phys-obj (Intentional +)))
   (Goal (Phys-obj (Container +))))
```

Figure 3: LF type definitions for LF::Motion and LF::Filling. In the lexicon, feature vectors from LF arguments are used to generate selectional restrictions based on mappings between subcategorization frames and LF arguments.

The LF definitions contain selectional restrictions on the arguments expressed in terms of semantic feature sets. Features encode basic meaning components used in semantic restrictions, such as form, origin and mobility for physical objects. For example, the :Theme argument is defined as *Phys-obj (Mobility Movable)* to reflect the fact that it has to be a mobile object, as opposed to generally fixed objects such as cities and mountains. LF::Filling places an additional restriction on its :Goal, requiring that it has to be a container.

The semantic feature set we utilize is a domain-independent feature set developed using EuroWordNet (Vossen, 1997) as a starting point, and extended by incorporating lexico-syntactic generalizations from other linguistic theories (Dzikovska et al., to appear). The set of features is limited to 3-10 per word. The small size of the feature set provides the lexicon developers with an easy to use framework in which to express semantic properties of words for selectional restrictions, because each word only needs to be classified along a small set of dimensions. However, the small feature set size limits the expressivity of the selectional restrictions, so not every possible restriction can be captured in it (see Section 4 for an example).

In our work on domain-independent lexicon development we found this approach a useful compromise. While it is small enough to keep lexicon development simple, it covers enough of the basic properties of words to significantly improve parsing speed and accuracy in two evaluation domains (Dzikovska, 2004). Selectional restrictions as feature sets offer further advantages in terms of efficient implementation and domain customization (Dzikovska et al., 2003). Therefore, in our lexicon we distinguish the word senses and semantic arguments which can be disambiguated based on syntactic structure and selectional restrictions expressible in terms of our feature set. This has a direct impact on the decision to simplify frame role structures discussed in the next section.

4. Adapting FrameNet for the TRIPS LF Ontology

We made two major changes to our ontology that diverged from FrameNet representation: we added a hierarchical structure and reduced the number of distinct frame elements (which we call roles). The FrameNet ontology is mostly flat, even though it contains many frames subsuming verbs that have identical argument structures. While FrameNet is designed to represent the hierarchies of frames, currently only about one-third of the frames in FrameNet inherit from other frames (Gildea, personal communication). In cases where frames included similar words but reflected finer meaning distinctions, we collected them under a common parent. For example, *Suasion1*, *Suasion2* and *Suasion3* include a group of verbs such as *encourage*, *convince*, *induce*, which have the same set of roles, but the difference in meaning comes from whether the addressee forms an intention to act. From the point of view of argument structure and selectional restrictions these frames are identical, so we collect them under a general parent and use the same set of selectional restrictions.

Table 1 shows the statistics about the number of LF types at different levels of our hierarchy. Level 0 types are types that do not inherit from anything, level 1 are types with 1 parent, and so on. The first 2 levels in our ontology were created artificially, because we needed special types for parsing: a unique root in the ontology, a type which unifies with nothing else ("-"), and another type which unifies with anything but "-". Thus, the contentful entries start at level 2, and we have 7 root entries that do not inherit from anything, 103 entries at depth 1. The majority of the types

Level	Frame Count
0	1
1	2
2	7
3	103
4	170
5	207
6	103
7	44
8	10
9	9

Table 1: The number of LF types at different levels of our LF hierarchy

we use are at depth 2 or 3 (170 and 207 respectively), but the hierarchy goes up to 6 levels deep, mostly in the parts of ontology where objects are classified.

In the process of developing our ontology, we had to add types to support problem solving and planning actions, which were absent in the version of FrameNet we utilised: FRAMENET II Release 1.0. For example, it did not have a classification for the word *need*, which occurs frequently in our dialogues, so we defined a new LF::Necessity frame in our lexicon.[2] Other words common in our task-oriented domains but not currently found in FrameNet are *cancel*, *revise*, *schedule*. Sometimes words were defined within FrameNet, but we needed to define additional senses because the FrameNet frame did not cover the common usage in our domain. For example, the word *change* is defined only as an instance of frame *Transformation*, where an entry is transformed into something else, like in *change the rabbit into a hat*. In one of our domains, a frequent usage is *Change the dial to VDC* (i.e., change the setting, but not the dial itself). So we created a new LF::Change-state frame to account for this sense. Similarly, the adjective *open* is defined as *Candidness* in FrameNet, corresponding to usages like *She was open with us about the party*, with synonyms such as *candid, forthright, etc*. In our domains, *open* has to do with physical accessibility, *The route is open*, or *there is an open door*. These senses are not suitable for the words grouped in the *Candidness* frame, thus we established the LF::Openness frame to account for them.

The hierarchical structure provides a level of generalisation in the ontology that makes it easier to include and maintain selectional restrictions. For that purpose, we also simplified the frame elements in our ontology. FrameNet utilises situation roles, so a *driving* situation involves a *driver* role, whereas the *communication* situation has a *communicator*. However, these roles may be seen as instances of a general *agent* role, which is an intentional being doing the action. A limited number of role names simplifies the inheritance in the LF Ontology by allowing us to define a general restriction (e.g., agents are intentional beings) high in the hierarchy tree.

For purposes of mapping between syntax and semantics, a smaller number of role names facilitates the definition of these mappings, because it creates opportunities for generalisation. For example, many motion verbs will use exactly the same set of syntax-semantics mappings, and not having the distinctions between "driver" and "self-mover" makes it easier to add new verbs by example.

More importantly, we found some frame elements too specific or too dependent on pragmatic information to be distinguishable during parsing. For example, the frame *closure* defines 2 separate frame elements: "Container-portal", for example *flap* in *Close the tent flap*, and "Containing-object", *coat* in *buttoned her coat*. Both can occur as direct objects of relevant verbs. Human annotators are able to distinguish those based on common sense knowledge. For parsing, however, selectional restrictions expressed with a limited set of semantic features are not specific enough to make this determination. Moreover, to our knowledge there is no reasoner able to make this distinction in a domain-independent manner. Therefore, we made the decision to define a more general :Theme role for our LF type LF::Closure, which covers both those semantic arguments. The relevant distinctions, if necessary, can be made by the domain specific reasoners using our customization mechanisms (Dzikovska et al., 2003).

The decision to use a reduced, more general set of roles has an advantage for fast acquisition of lexical entries. Many linguistic theories make syntactic generalisations based on semantic classes (see for example (Levin, 1993), (Jackendoff, 1990)). While we do not use such generalisations yet, we designed our ontology to facilitate those in the future, as discussed in Section 7. For example, the VerbNet lexicon defines the verb *close* with *agent*, *patient*[3] and *instrument* roles, and defines the corresponding subcategorization frames and syntactic variations. This generalization is only possible with more general role names, and we hope to use it in the future to speed up the development of syntax-semantics mappings.

5. Evaluation

In this section, we present statistics about our current lexicon, and how it compares with the FrameNet ontology.[4] Currently, our LF Ontology contains 656 LF types, corresponding to different concepts. The complete statistics for our lexicon is shown in Table 2. We have 2446 words total in our lexicon, 1999 of which are open class words - adjectives, nouns, verbs and adverbs, with 2248 different word senses. The system uses 37 semantic roles, considerably fewer than FrameNet, which has 554 frame elements.

We compared our lexicon with the FrameNet version 1.0. Table 3 shows the number of lexical items for each part of speech which were defined in both lexicons, in TRIPS lexicon only, and in FrameNet lexicon only.

It is interesting to note that while FrameNet is much larger in size than the TRIPS lexicon, there's a considerable number of lexical items, in all categories, which do

[2] *need*, and other words we cite in our examples, are also missing from the latest web version of FrameNet, FrameNet II release 1.1.

[3] which corresponds to our :Theme.

[4] The FrameNet statistics in this section are from FrameNet II Release 1.0 unless otherwise noted.

POS	Count	Senses	Synt. var.	Comment
ADJ	422	1.07	1.12	Adjectives
N	875	1.06	1.09	Nouns
ADV disc	36	1.08	1.11	Discourse adverbials
ADV	221	1.32	1.55	Adverbs (including adverbial prepositions)
V	490	1.29	1.71	Verbs
NAME	22	1.00	1.00	Names
PUNC	10	1.00	1.00	Punctuation signs
UTTWORD	121	1.01	1.01	Discourse words like OK, yes, yeah, etc.
OTHER	249	1.02	1.04	Other parts of speech for functional words, including ART, PREP, QUAN, CONJ, PRO, NUMBER
Total	2446	1.12	1.26	

Table 2: Lexicon statistics in our system

POS	Common	Trips only	FrameNet only
Adj	114	308	1072
N	285	582	2479
V	225	232	1774

Table 3: Lexicon statistics

not overlap between those lexicons. Part of the problem is that the comparison is with an older version of FrameNet II (release 1.0) and the current release (1.1) is much richer. However, manual inspection of the data and comparison with the release 1.1 data available on the Web still shows significant non-overlapping areas. For verbs, these include

- Verbs dealing with plans and goals: *achieve, accomplish, complete etc.*
- Verbs dealing with intentions and permissions: *need, authorize, assume, trust etc.*
- Verbs dealing with mutual understanding in a conversation: *recap, reformulate, misunderstand*
- Verbs with particles common in spoken language: *look for, back up, dig out etc.*

Verbs with particles do not appear to be consistently annotated in FrameNet, so the number of verbs listed as in TRIPS but not FrameNet may include some of those that in FrameNet are annotated as senses belonging to a verb ignoring the particle. When we excluded verbs with particles from the counting, the number of verbs defined in TRIPS but not FrameNet was 164, still a substantial difference. Moreover, when a particle is not included with the verb annotation, it poses a significant problem for a parser, because particles provide important syntactic clues during parsing and disambiguation, and loss of this information adds ambiguity to the process.

We did not analyze in detail the differences between nouns and adjectives, but based on several spot-checks, it appears that this is an area that has been developed in FrameNet II Release 1.1, which now defines many common adjectives and nouns such as colour names and common foods. The biggest differences appear to be in words that are essential for coverage in our domains, which are transportation and computer purchasing. Therefore, TRIPS defines the names for many physical objects such as *bus, dvd, cd-drive* which are not part of the FrameNet lexicon. This points to the issue we need to deal with in our future work. Our data suggest that the text corpora that are the basis of FrameNet are quite different from the task-oriented spoken dialogue corpora, and that's why there are a number of words important in our domains which are currently not included in the FrameNet database. If the LF types for those are added to our ontology, we need to address synchronization issues with further FrameNet updates.

5.1. Role structure evaluation

As discussed above, the names of semantic roles, much as the names of the frames themselves, have to be at the right level of abstraction in order to facilitate a connection with syntax. Therefore, during the development of the LF Ontology we needed to simplify the FrameNet role structure. The FrameNet version we evaluated contained 554 frame elements. We discussed in Section 4 the problems that this caused in efficiently acquiring lexical entries and in frame element disambiguation. In contrast, TRIPS has 37 roles used in subcategorization frames. This number is considerably easier to manage in defining syntax-semantics mappings, and for disambiguation.

The TRIPS role set, though developed independently, is similar in size and structure to the role set in another semantic lexicon, VerbNet (Kipper et al., 2000), which also aims to link syntactic and semantic structure. A detailed comparison can be found in (Dzikovska, 2004). In brief, VerbNet has 28 roles, 8 of which are the same as those used in the TRIPS LF lexicon. We did not conduct the formal evaluation of the consistency of the rest of the role set, but, generally speaking, the rest of the role sets intersect, but VerbNet makes finer distinctions in some cases (splitting :Theme into *theme* and *patient*). In addition, TRIPS contains semantic roles for classifying adjective, adverb and noun arguments, not covered by the VerbNet lexicon. We plan to resolve the differences and use VerbNet selectional restrictions and syntactic patterns to extend coverage of our verb lexicon as part of our future work.

In comparing the role sets it is also important to note that FrameNet intends to cover all parts of the sentence relevant to the event, be they verb arguments or adjuncts expressed by adverbs or even clauses. This results in some highly specific frame element names, such as “Abundant-entities”, “arguer” or “manifestation-of-bias”, each of which occurs only in a single frame. In our evaluation, 313 of frame elements appeared in one frame only. At the same time, the 6 most common frame elements, “Manner”, “Time”, “Degree”, “Place”, “Means” and “Purpose”, are handled as adverbial senses in the TRIPS lexicon, with the exception of a small number of verbs which subcategorize for them. For example, usually :Time-duration role is realized by an adverbial, as *for 5 minutes* in *She completed the task in 5 minutes*. But for 2 frames, LF::Take-time and LF::Leave-time explicitly subcategorize for it as a direct object, *e.g.*, It takes 5 minutes to complete. In the TRIPS lexicon there are 4 roles which appear with only 1 frame, and 2 of those are realized as adverbials in other constructs, so they are not unique labels for a given frame, but just exceptional cases of arguments typically handled by adverbials.

The large number of role names difficult to disambiguate for the parser is the main reason why we were unable to use FrameNet directly in our lexicon. The distinction between subcategorized arguments and adjuncts (generally coming from adverbials) is very important in parsing and semantic disambiguation, and that FrameNet does not mark it in its frame element structure makes it difficult to use directly in a practical NLP parser.

6. Future Work

Our work highlights both the usefulness of FrameNet as a basis for building a computational ontology and lexicon, and its limitations as a source representation for parsing. FrameNet provides word meanings which can be reliably distinguished by humans, which makes lexicon development easier, and frame representations are convenient for natural language processing because they are easy to obtain from linguistic structure and allow us to encode optional arguments. However, to facilitate connections to syntax and allow for possible syntactic generalisations, we needed to modify the information available in FrameNet by adding hierarchy and using a smaller set of role names.

In the future we plan to include syntactic generalizations based on syntactic alternations as done in VerbNet (Kipper et al., 2000). Currently there is no direct mapping between TRIPS and VerbNet classes. Our ontologies were developed independently, because the VerbNet database was unavailable at the time; additionally, VerbNet is not designed to cover other word classes, such as nouns and adjectives, and we developed our lexicon to provide semantic roles representations for all open-class words. As mentioned above, our analysis shows a significant overlap between our semantic roles and VerbNet roles. In our evaluation, we also noted a “core” set of roles, including “Agent”, “Cause”, “Source”, “Goal”, “Theme”, which, after the general frame elements typically implemented by adverbials we mentioned before, are the most frequent frame elements used in FrameNet. This raises issues of standardisation and developing a general set of roles suitable both for semantic analysis and for syntactic generalisations, and we are working on mappings between the TRIPS and VerbNet roles, and possibly between the TRIPS and FrameNet roles.

We also need to address the coordination between FrameNet and TRIPS ontologies. Our ontology is based on FrameNet, but it is not synchronised with the current FrameNet version, because of the changes and additional information necessary in our representations. Currently, instead of trying to synchronize our ontologies directly, we are working on a learning module which uses FrameNet and other resources to propose meanings of novel words as an aid to human lexicon developers.

7. Conclusions

In conclusion, this paper discusses FrameNet as a source of semantic information for a deep syntactic parser. Our wide coverage parser needs an ontology as a source of domain-independent word senses, and FrameNet provides a well-documented source of reliably distinguishable semantic classes. For use in our practical dialog system, however, we needed to streamline aspects of the FrameNet data for efficiency. There remain open questions, especially the extent to which such streamlining can be handled automatically as both systems develop in parallel, which need to be addressed in future work.

8. Acknowledgments

We thank Mark Core and Dan Gildea for useful comments and discussion for this paper. This material is based upon work supported by a grant from The Office of Naval Research under grants number N00014-01-1-10115 and N00014-03-1-0780 under subcontract to University of West Florida, from DARPA, under research grant number F30602-98-20133, DARPA under grant number NBCHD030010 via subcontract to SRI number 03-000223 and The National Science Foundation under grant number E1A-0080124. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of above named organizations.

9. References

Alshawi, Hiyan, David M. Carter, Björn Gambäck, Stephen G. Pulman, and Manny Rayner, 1991. Transfer through quasi logical form: A new approach to machine translation. Technical Report T91020, SICS, Stockholm, Sweden.

Chang, N., J. Feldman, R. Porzel, and K. Sanders, 2002. Scaling cognitive linguistics: Formalisms for language understanding. In *Proceedings of the First International Workshop on Scalable Natural Language Understanding*. Http://smartkom.dfki.de/Vortraege/icsi-ai2.pdf.

Copestake, Ann, Dan Flickinger, Rob Malouf, Susanne Riehemann, and Ivan Sag, 1995. Translation using Minimal Recursion Semantics. In *Proceedings of the 6th. International Conference on Theoretical and Methodological Issues in Machine Translation (TMI-95)*. Leuven, Belgium.

Dzikovska, Myroslava O., 2004. *A Practical Semantic Representation For Natural Language Parsing*. Ph.D. thesis, University of Rochester.

Dzikovska, Myroslava O., Mary D. Swift, and James F. Allen, 2003. Integrating linguistic and domain knowledge for spoken dialogue systems in multiple domains. In *Proceedings of IJCAI-03 Workshop on Knowledge and Reasoning in Practical Dialogue Systems*. Acapulco.

Dzikovska, Myroslava O., Mary D. Swift, and James F. Allen, to appear. Customizing meaning: building domain-specific semantic representations from a generic lexicon. In Harry Bunt (ed.), *Computing Meaning, Volume 3*, Studies in Linguistics and Philosophy. Kluwer Academic Publishers.

Gildea, Daniel and Daniel Jurafsky, 2002. Automatic labeling of semantic roles. *Computational Linguistics*, 28(3):245–288.

Jackendoff, R., 1990. *Semantic Structures*. The MIT Press.

Johnson, Christopher 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.

Kipper, Karin, 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)*. Menlo Park, CA: AAAI Press.

Lenat, Douglas B., 1995. CYC: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38.

Levin, Beth C., 1993. *English Verb Classes and Alternations: a Preliminary Investigation*. Chicago, IL: University of Chicago Press.

Macleod, Catherine, Ralph Grishman, and Adam Meyers, 1994. Creating a common syntactic dictionary of English. In *SNLR: International Workshop on Sharable Natural Language Resources*. Nara.

Miller, G., 1995. Wordnet: A lexical database for English. *Communications of the ACM*, 38(5).

Swift, Mary D., Myroslava O. Dzikovska, Joel R. Tetreault, and James F. Allen, 2004. Semi-automatic syntactic and semantic corpus annotation with a deep parser. In *Proceedings of LREC-2004*.

Vossen, Piek, 1997. Eurowordnet: a multilingual database for information retrieval. In *Proceedings of the Delos workshop on Cross-language Information Retrieval*.