A WordNet Detour to FrameNet

Aljoscha Burchardt, Katrin Erk, and Anette Frank*

Saarland University and DFKI* Saarbrücken

{albu,erk,frank}@coli.uni-sb.de

Typ des Beitrags/Type of the paper Lecture

A WordNet Detour to FrameNet

Aljoscha Burchardt, Katrin Erk, and Anette Frank

In this paper, we present a rule-based system for the assignment of FrameNet frames by way of a "*detour* via WordNet". The system can be used to overcome sparse-data problems of statistical systems trained on current FrameNet data. We devise a weighting scheme to select the best frame(s) out of a set of candidate frames, and present first figures of evaluation.

In diesem Aufsatz beschreiben wir einen regelbasierten Ansatz zur Überbrückung fehlender Lexikoneinträge in FrameNet. WordNet Synsets dienen als „*Umweg* zu FrameNet". Auf der Basis von WordNet generieren wir für ein gegebenes Wort eine Menge von Kandidatenframes. Wir entwerfen ein Gewichtungsschema zur Auswahl des/der besten Frames und zeigen erste Evaluationsergebnisse.

1. Motivation

Recently, there has been a growing interest in more in-depth semantic analysis for practical NLP tasks, in particular as a basis for open-domain information access. Large-scale lexical semantic resources, such as WordNets (Fellbaum, 1998) have been developed and put to use for approximate semantic modeling in many applications. The FrameNet (Baker et al., 1998) and PropBank (Kingsbury et al., 2002) projects are developing lexical semantic resources that focus on the modeling of predicate-argument structure. The Berkeley FrameNet database groups words and expressions (lexical units, LUs for short) into semantic classes (frames) and lists semantic roles for each frame. This type of lexical semantic information is particularly useful for information access tasks, like Information Extraction (IE), or Question Answering (QA).

We are currently investigating the use of FrameNet frames for building partial text meaning representations (see Burchardt et al., 2005), to be used in applications like IE and QA as e.g. addressed by the Recognising Textual Entailment (RTE) Challenge[1]. Semantic representations building on frames provide normalisations over surface realisations (e.g. active/passive, verb/nominalization) and thus a sensible granularity for these applications.

[1] http://www.pascal-network.org/Challenges/RTE/.

Two major tasks in the (automatic) annotation of texts with frames are the *frame assignment problem***,** i.e. the identification of the proper frame for a given lexical unit, and the *semantic role assignment problem*, i.e., the assignment of the frame's semantic roles to major sentence constituents. In this paper, we are mainly concerned with the frame assignment problem.

As a base system for frame assignment we are using a system that treats this task as (supervised) word sense disambiguation, using the FrameNet collection of annotated sentences as training data (Erk 2005). But as FrameNet is still a growing resource, the annotation of contiguous text is confronted with two problems. The first is a problem of coverage: For example, for the 574 sentences of the RTE development corpus, with an average of 16.24 words/sentence, our WSD system yields an assignment of 2.7 frames and 3.6 frame elements per sentence. The second problem concerns lacking senses in the current FrameNet resource and is caused by the fact that FrameNet is being constructed one frame at a time, rather than one lemma at a time. While a lack of coverage leads to missing frame assignments, lacking senses result in wrong assignments.

In this paper we develop an approach to overcome these sparse data problems, by using WordNet as a "*detour* to FrameNet". We use WordNet synsets as an interface layer to propose LU-frame pairs that are missing in the FrameNet database. We employ an open-source WordNet-based WSD system of (Banerjee and Pedersen, 2003) to annotate lexical units in unseen texts with their contextually determined WordNet synset. Frame assignment can then proceed by using not just a single word, but also its synonyms and hypernyms. We devise a weighting scheme to select the best frame(s) out of a set of candidate frames, and present first figures of evaluation on the basis of held-out LU-frame associations. While the research reported here is devoted to English, the methods we propose will naturally carry over to German data, resources and linguistics analysis tools.

The paper is structured as follows. Section 2 provides background information on FrameNet and the context of our work. Section 3 describes the suite of systems we use for NL processing using FrameNet resources. Section 4 presents our rule-based frame-assignment system that assigns frames via WordNet. We describe the basic method and present the algorithm with the chosen weighting scheme for frame selection. We evaluate the system in two experiments. Section 5 concludes with an outlook on future work.

2. Background: FrameNet and SALSA

FrameNet (Baker et al, 1998) is based on Fillmore's Frame Semantics (Fillmore, 1976). Frame Semantics models the lexical meaning of predicates in terms of *frames*. A frame describes a conceptual structure or prototypical situation together with a set of semantic roles, or *frame elements (FEs),* that are involved in the situation. FrameNet currently contains about 600 frames of general conceptual classes.[2] As an example, consider the frame STATEMENT. This frame is evoked by *lexical units (LUs)* like *disclose.v, speak.v, suggest.v, comment.n,* etc., as illustrated in (1)–(4). The frame STATEMENT defines the core semantic roles SPEAKER, TOPIC, MESSAGE, and MEDIUM, and a set of *peripheral* (i.e., non-core) roles, e.g., ADDRESSEE and MANNER. Additionally, FrameNet defines *extra-thematic roles*, such as TIME in (4), which are not frame-specific.

(1) [Judy Rumbold $_{\text{SPEAKER}}$] *suggested* [in the Guardian $_{\text{MEDIUM}}$] [that one of the reasons for the maleness of your work was that you both had domineering mothers $_{\text{MESSAGE}}$].

(2) “[He $_{\text{SPEAKER}}$] *speaks* [highly $_{\text{MANNER}}$] [of you $_{\text{TOPIC}}$],” she said.

(3) “Did [Dominic $_{\text{SPEAKER}}$] ever *make* any *comments* [regarding Toby $_{\text{TOPIC}}$] [to you $_{\text{ADDRESSEE}}$]?”

(4) [In January 1990 $_{\text{TIME}}$] it was *disclosed* that [in 1976 a meltdown at Greifswald had only narrowly been averted $_{\text{MESSAGE}}$].

SALSA The research described below is conducted in the SALSA project (Saarbrücken Lexical Semantics Annotation and Acquisition Project). The aim of SALSA is to create a large lexical semantics resource for German based on Frame Semantics and to develop methods for the automated annotation and semantic analysis of corpora using frame semantic representations. SALSA is manually annotating the German TIGER corpus (Brants et al, 2002) with frames, following the definitions of frames and roles in the Berkeley FrameNet database (Erk et al, 2003b). SALSA further investigates probabilistic models for automatic frame annotation on the basis of manually annotated corpora (Baldewein et al, 2004) and explores the use of frame semantic annotations for dynamic semantic analysis in NLP tasks (Burchardt et al, 2005).

---

[2] For example: AWARENESS, COMMERCIAL_TRANSACTION, THEFT, MOTION, etc.; examples in this Section are from FrameNet: http://www.icsi.berkeley.edu/~framenet/.

3. Towards using Frames for NLP

3.1 Current Architecture

In our current work, we combine learning techniques for automated frame and role assignment with deep grammatical LFG parsing, to develop methods for open-domain frame-based information access and reasoning.

We employ the deep syntactic representations provided by large-scale LFG grammars (Butt et al, 2002)[3] as a syntactic basis for frame-based meaning assignment. In (Frank and Erk, 2004), we have designed a modular syntax-semantics interface to project frame semantic representations from the f-structure output of LFG parsing. We have built interfaces to the statistical frame and role assignment systems of (Erk, 2005) and (Baldewein et al, 2004) that propose disambiguated frame assignments for a given text. For further refinement of the frame semantic representations, we defined semantics construction rules for modifiers and named entities that realise extra-thematic semantic roles (e.g. TIME, LOCATION, PURPOSE, etc.). As an additional knowledge source, we have integrated the SUMO/MILO ontology (Niles and Pease, 2001), by way of the WSD system of (Banerjee and Pedersen, 2003), and the mapping from WordNet synsets to SUMO/MILO classes.

We are currently using this architecture as a base system for the Recognising Textual Entailment Challenge (RTE). In this task a system is presented with two text snippets, a so-called Text-Hypothesis pair, and has to decide whether the meaning of the Hypothesis is entailed by (can be inferred from) the Text. Examples (5,6), taken from the RTE corpus, may serve as an illustration.

(5) T: *Ostriches put their heads into the sand to avoid the wind.*
H: *Ostriches bury their heads in the sand.*

(6) T: *A Cuban American who is accused of espionage pleads innocent.*
H: *American accused of espionage.*

Figure 3, in the Appendix, displays the frame semantic annotation for the most probable LFG analysis for (5:H). The f-structure is enriched with a frame se-

[3] We are using the German LFG grammar developed at the IMS, University of Stuttgart, and the English LFG grammar developed at Parc (Riezler et al., 2002).

mantic projection (s::), according to the proposed assignments of the frame and role assignment systems (Erk 2005; Baldewein et al., 2004). SUMO/MILO concept classes are displayed as an additional projection (in the feature ONT). Figure 1 below displays the same structure, converted to the TIGER-SALSA format and visualised in the SALSA annotation tool (Erk et al, 2003a).

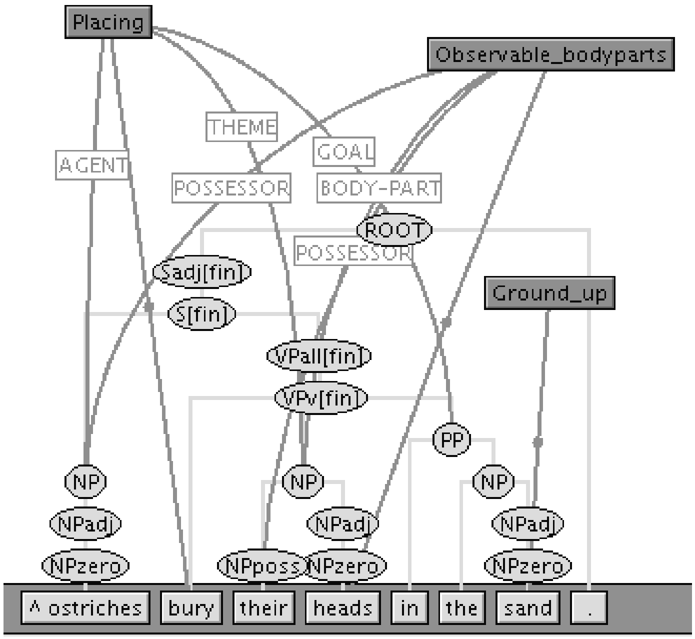

Figure 1: Frame and role assignments for (5:H) viewed in the SALSA annotation tool.

The verb *bury* is assigned the frame PLACING, with the roles AGENT, THEME and GOAL assigned to its SUBJ (*ostriches*), OBJ (*heads*), and OBLique (*in the sand*) arguments.[4] The noun *heads* evokes the frame OBSERVABLE_BODYPARTS, with roles POSSESSOR and BODY-PART assigned to SUBJ (*ostriches*) and OBJ SPEC POSS (*their*), and OBJ (*heads*), respectively.

The example further illustrates the coverage problem mentioned above: the frame OBSERVABLE_BODYPARTS is the only frame that could be assigned by the statistics-based frame assignment system. This gap in coverage is compensated by the system that assigns frames via WordNet, which we describe in Section 4. This system proposes all the frames as shown in Figure 1, including the frame GROUND_UP, which, however, is wrongly assigned.

---

[4] Functional edges are not made visible in Figure 1, but see Figure 3.

3.1 Frame Assignment as Word Sense Disambiguation

The task of frame assignment can be conceived as word sense disambiguation: A target expression may be listed as a lexical unit of several frames. Each of these frames can be seen as a sense of the target expression. For example, the verb *skim.v* is listed as a lexical unit for four frames: (All examples are from the FrameNet corpus. Some are abbreviated.)

- **READING:** ***Skimming*** *a chapter for its main idea may be done over coffee.*
- **REMOVING:** *Remove the vanilla pod,* ***skim*** *the jam, and let it cool.*
- **SCRUTINY:** *She* ***skimmed*** *through the newspaper clippings.*
- **SELF_MOTION:** *We* ***skimmed*** *across the surface of that sodding lake whilst all around us gathered the dark hosts of hell.*

The frames READING, REMOVING and SELF_MOTION constitute clearly distinguished senses of *skim.v*. READING and SCRUTINY are hard to distinguish as far as *skim.v* is concerned, even though in general they describe different situations, with different semantic roles and different LUs: READING includes *devore.v*, *peruse.v* while SCRUTINY has LUs like *analyze.v, search.v, survey.v*.

The frame assignment system of (Erk, 2005) views frame assignment as a supervised word sense disambiguation task, using the FrameNet annotated examples as its training corpus. The system uses the same rich set of features that (Florian et al, 2002) propose:

- A bag-of-words context including lemma and part of speech information for each word; the window size for the context is one sentence (since the FrameNet corpus does not provide contiguous annotated sentences).
- Word bigrams and trigrams centered on the target word.
- Head words and prepositions of complements and adjuncts of the target.

In addition to the features introduced by (Florian et al, 2003) we use features that describe the subcat frame of the target, both in terms of individual grammatical functions found for the target and as an n-gram combining all grammatical functions. For example, for the sentence *She skimmed through the newspaper clippings* there would be three subcat features: *subj, ppobj-through, ppobj-through+subj*. Grammatical functions are determined heuristically.

All feature types are weighted with a fixed, empirically determined weight. In addition, context word features are weighted by their distance from the target.

The system has a simple architecture: a Naïve Bayes classifier, which assigns each instance x the sense s that maximizes $P(s \mid x) = \frac{P(x \mid s)}{P(x)} P(s)$, with the simplifying (false) assumption that all features of the instance are independent: $P(x \mid s) = \prod_{\text{f in x}} P(f \mid s)$. Probabilities of features are determined using (smoothed) maximum likelihood estimation on the weighted features.

While the FrameNet corpus has been used successfully as a training corpus for frame-semantic role assignment, it is somewhat problematic as a training corpus for frame assignment. FrameNet is a growing resource, constructed frame by frame, which means that while each frame lists all LUs that introduce it, many LUs are still lacking some of their frames. For example, there is no frame for the sense of *treat.v* in *He treated them with patience*. Also, some frames have no annotated examples, and hence cannot be learned in a supervised learning setting, among them important frames like the Possession frame for *have.v*. To illustrate the problem, only 10.7% of the current 8000 LUs are ambiguous at all, and the baseline for the WSD task (assign each LU its most frequent frame) is already at 93% f-score. This is a problem of the FrameNet corpus, not of the FrameNet approach as such, as experiments with a snapshot of the German SALSA corpus confirm (Erk, 2005).

4. A WordNet Detour to FrameNet

4.1 More Coverage

As outlined above, applications using the current FrameNet release are facing coverage problems. In many cases, an appropriate frame exists for a given *target word,* but the word is not yet listed as an LU for that frame, or else the word may be listed, but there are not enough annotated sample sentences as training data for statistical frame assignment.

In the following we describe a system that addresses these problems. The system is rule-based and uses WordNet to generalize over a given target word in order to compensate for the missing LU and to assign the appropriate frame. For each target word, we consider a set of related words (*wordnet relatives*) and collect all *candidate frames* that are evoked by these words. To select, among the candidates, the frame to be assigned, we use a weighting function which is described in detail below.

This approach requires that we know the correct WordNet synset of the target word. To this end, we use an existing WordNet-based WSD system of (Banerjee and Pedersen, 2003) that disambiguates a word in its sentential context by computing the relatedness between the word and its context words. The underlying assumption is that semantically related words often co-occur. Relatedness is measured on the basis of WordNet, including overlap between the words that occurs in the glosses.

As an example of how our WordNet-based frame assignment system works, consider again example (5). For the target word *bury,* Pedersen's system returns the target synset 'bury#v#2'. Although *bury* is not yet listed in FrameNet, our system correctly assigns the frame PLACING because several WordNet relatives of this target are listed as LUs for this frame (e.g. *lay, put, place,…*). The second best candidate frame is ATTACK (which lists the WordNet relatives *set* and *lay*), but its weight is less than one third of the weight of PLACING. Our system therefore selects the latter.

4.2 Assigning Frames via WordNet

For a given input synset, the central algorithm of our system proceeds in three steps. First, we compute a **set of wordnet relatives** of the target word.[5] It contains all synonyms and hypernyms of the input word plus the antonyms of these words. The inclusion of antonyms is effective because antonyms are usually defined within the same FrameNet frame.

Second, we compute all **candidate frames evoked by the wordnet relatives**. These frames are in the first place those frames that list the respective words as LU. For words that are not listed as LU for any frame, we check whether they match any frame name, to exploit the fact that sometimes target words that correspond to a frame name are not yet listed as LUs, e.g. the noun *researcher* is

[5] We use a Perl package available at http://search.cpan.org/dist/WordNet-QueryData/.

not listed as LU in the frame RESEARCHER. As success criteria for matching, we currently require more than 50% overlap between the word and the frame name.

Finally, in order to select the best frame, **all candidate frames are weighted**. The overall weight for each frame is the sum of the weights for each wordnet relative that evokes this frame. These weights are computed according to the following formula:

$$\frac{\text{similarity(wordnet\_relative, target\_word)} * \text{BoostFactor}}{\text{Spreading\_factor(wordnet\_relative)}}$$

This formula takes into account three factors: (i) the similarity between the target word and the wordnet relative. We currently take the square of the WordNet path distance as similarity measure here.[6] (ii) A boost factor that rewards words that are listed as this frame's LU as opposed to those that only match the frame name. (iii) The spreading factor of the wordnet relative. This is the number of frames evoked by that word. For example, *go* is listed as LU for three frames (MOTION, COMPATIBILITY, NAME_BEARING) and thus has spreading factor 3. The more frames a synset evokes, the less discriminative and thus informative it is for our purpose. The complete algorithm is given in Figure 2.

As an example that highlights the algorithm's matching functionality and the interplay of the factors, take the target synset 'researcher#n#1'. The algorithm assigns the frame RESEARCH to this synset although the noun *researcher* is not yet listed as LU of this frame. But both *researcher* and its synonym *research worker* match the frame name. As both are in the target word's synset, the weight of this frame is relatively high. The second best candidate frame is PEOPLE_BY_VOCATION, evoked by the LU *scientist,* which is a hypernym of the target word. Although frames evoked by LU lookup are preferred, here a frame evoked by name similarity wins because of the two strong matches.

As an example for the effectiveness of using antonyms, consider the target word *new,* as in *There is a new U. S. law that bans e-mail.* While FrameNet does not yet list *new* as LU of the frame AGE, our system can assign the frame AGE, because its antonym *old* is listed for this frame.

[6] We use a Perl package (http://search.cpan.org/dist/WordNet-Similarity/) that allows for experimentation with a number of measures known from literature.

1. Search_words = {w | w $\in$ Input_synset} $\cup$ {w | $\exists$ Synset: hypernym(Synset, Input_synset) $\wedge$ w $\in$ Synset}

   Search_words = Search_words $\cup$ {w' | $\exists$ w $\in$ Search_words $\wedge$ antonym(w',w)}

2. forall F in Frames:
   ```
           Evoked_by_LU(F) = { }
           Evoked_by_Match(F) = { }
     end
     forall F in Frames, forall W in Search_words:
           if W is a LU of F then
                 Evoked_by_LU(F) = Evoked_by_LU(F) ∪ W
                 Spreading_factor(W) += 1
           elsif W matches name of F then
                 Evoked_by_Match(F) = Evoked_by_Match(F) ∪ W
                 Spreading_factor(W) += 1
           end
     end, end
   ```

3. forall synset,synset': similarity(synst,synset') = WordNet_Path_Distance(synset,synset')^2

   forall F in Frames:

$$\text{Weight(F)} = \sum_{Synset \in Evoked\_by\_LU(F)} \frac{similarity(Synset, Input\_synset) * BoostFactor}{Spreading\_factor(Synset)} + \sum_{Synset \in Evoked\_by\_Match(F)} \frac{similarity(Synset, Input\_synset)}{Spreading\_factor(Synset)}$$

   end

Figure 2: Algorithm for assigning frame(s) to a given synset.

4.3 Evaluation

For evaluation of the system we need suitable frame-annotated test data. The only data available as gold standard is the FrameNet corpus, yet this is not immediately suited for our evaluation context, because all words that are frame annotated there are already listed as LUs in FrameNet. Thus, our system would not have to use the detour via WordNet at all. In order to effectively evaluate our system on this data, we have devised a *detour-only* version. In this mode, we prevent the target word itself from being looked up as an LU, thus forcing the algorithm to take the detour via synonyms, antonyms, and hypernyms. This is equivalent to deleting the respective target word from any frame's LUs.

Table 1 shows the system's performance on the FrameNet corpus.[7] The first row shows in how many cases no, one, or more than one frame was assigned, thus measuring coverage. The second row indicates whether the gold standard frame assigned in the FrameNet corpus was equal to (contained in) the assigned frame(s), which is a weak measure of precision.

	frames assigned per synset		
	none	1	> 1
Total instances	13%	71%	16%
Gold standard frame equal to (contained in) assigned frame(s)	-	38%	7%

Table 1: Frame assignment of detour-only system (FrameNet corpus).

The detour-only system has high coverage: In 87% of the cases, it assigns one or more frames. So, if a target word is unknown to FrameNet, we have a 87% chance to assign some frame(s), in constrast to a statistical system, which could not assign any frame in this case. For the complete data set, precision in terms of containment of the gold standard frame is at 39%. If we only consider cases where some frame was assigned, the precision is 45%. For the non-matching cases, it is still to be determined how close the assignments are to the gold standard. Since there is no formal measure such as frame distance, we have computed a table that compares gold standard and assigned frames.

Gold standard frame	Frames assigned by system	Number of instances
MANUFACTURING	INVENTION	19
	INTENTIONALLY_CREATE	12
	BUILDING	11
	CAUSE_TO_START	5
	GETTING	4
	TRANSFORMATION	1

Table 2: Frames assigned by detour-only system and gold standard from FrameNet corpus.

Table 2 shows a small excerpt of this table for the gold standard frame MANUFACTURING. Most of the frames that the system assigns are either from the same

[7] By the time of writing we have evaluated on 80.000 frame annotated word instances (60.000 verb, 20.000 noun, 20.000 adj./adv.).

domain as or compatible with the gold standard frame. Inspection of more examples reveals that the frames we assign are in many cases semantically closely related, often only differing in aspect, perspective or specificity (e.g. CHOOSING vs. DECIDING, AMOUNTING_TO vs. ADDING_UP, TRAVEL vs. MOTION).

In addition to the detour-only condition, we also tested the full system on the FrameNet corpus. Here, coverage raises to 96%. For 83% the gold standard frame is contained in the set of assigned frames (for 67% the gold standard frame is unambiguously assigned). Inspection of the frames that do not match the gold standard shows that these are acceptable in the vast majority of cases.

In a second experiment we tested the two versions of the system – the full system and the detour-only version – on unseen text. We ran both versions on 560 sentences of the first development set of the RTE data (3.800 noun, verb, and adjective instances). Table 3 gives an overview of the distribution of the distance between the respective target synset and the synset(s) that finally triggered the frame assignment. We distinguish three cases: (i) both synsets are the same, (ii) the frame evoking synset is a direct hypernym of the target synset, and (iii) the frame evoking synset is a transitive hypernym of the target synset.

	Same synset (distance 0)	Hypernym synset (distance 1)	Transitiv hypernym (distance >1)
Full system	54%	18%	27%
Detour-only system	33%	31%	35%

Table 3:Distribution of distance between frame evoking and target synset (RTE data).

The table shows that even in the detour-only system, in almost two thirds of the cases the distance is 0 or 1. The only effect of switching to detour-only is that we get a 21% drop of cases with distance 0, yet two thirds of the cases only move to distance 1. In both cases, the frame assignment is triggered by a synonym, direct hypernym or antonym of the target word and one can expect the assigned frame to be of appropriate specificity. For the cases with distance greater 1 we also computed the average WordNet path distance to the target, which is 3 for both system modes. The frames we assign here are usually more general. For example, for the target 'life#n#6', which is the synset for the lifespan of e.g. a battery, we get the frame QUANTITY. It is evoked by 'measure#n#3' which is a hypernym of the target with distance 3.

5. Conclusion and Outlook

We have presented a rule-based system for mediating frame assignment by a "detour via WordNet". It uses WordNet synsets and relations to generalize over target words, in order to assign frames to words that are not yet listed in the FrameNet database. This method is suited to overcome the sparse-data problem caused by the current, incomplete FrameNet data set. In our current NLP architecture, the system is used in cases where statistical frame assignment fails.

A preliminary evaluation of the system yields promising results. Our method exploits the fact that sense discrimination in WordNet is in general more fine-grained than FrameNet senses, so that we can account for missing FrameNet entries. However, in certain cases, close wordnet relatives map to distinct FrameNet senses. Moreover, the method is dependent on the assignment of the correct synset by an independent WordNet-based WSD system. Thus, errors in the initial assignment of synsets immediately affect the quality of frame assignment. Our current evaluation is still restricted, in that we are lacking appropriate measures of distance to assess the quality of frame assignments that do not match the gold standard. Also, up to now, we have employed a weighting scheme that combines a limited number of factors: similarity (path distance), spreading factor, and assignment by LU-lookup vs. name matching. In future work, we will experiment with additional factors, threshold settings, and machine learning techniques for estimation of weights for the individual factors.

As frames are general conceptual classes and thus to a large extent language-independent, the method described will naturally carry over to other languages that dispose of WordNets. Since the Pedersen system is rule-based, it should be possible to apply the algorithm to German on the basis of GermaNet (Hamp and Feldweg, 97). Our system could thus directly be applied for frame assignment on the basis of the German FrameNet data created in the SALSA project.

References

Baker, C. F., Fillmore, C. J., & Lowe, J. B. (1998). The Berkeley FrameNet project. In *Proceedings of COLING-ACL*, Montreal, Canada.

Baldewein, U., Erk, K., Pado, S., & Prescher, D. (2004). Semantic Role Labelling with Similarity-Based Generalization Using EM-based Clustering. In *Proceedings of Senseval'04*, Barcelona, Spain.

Banerjee, S. & Pedersen, T. (2003). Extended Gloss Overlaps as a Measure of Semantic Relatedness. In *Proceedings of IJCAI 2003,* Acapulco, Mexico.

Brants, S., Dipper, S., Hansen, S., Lezius, W. & Smith, G. (2002). The TIGER Treebank. In *Proceedings of the Workshop on Treebanks and Linguistic Theories,* Sozopol.

Burchardt, A., Frank, A., & Pinkal M. (2005). Building Text Meaning Representations from Contextually Related Frames – A Case Study. In *Proceedings of IWCS-6*, Tilburg.

Butt, M., Dyvik, H., King, T. H., Masuichi, H., & Rohrer, C. (2002). The Parallel Grammar Project. *Proceedings of the Workshop on Grammar Engineering and Evaluation, COLING 2002*, Taipei, Taiwan.

Erk, K. (2005). Frame assignment as word sense disambiguation. In *Proceedings of IWCS-6*, Tilburg University, Tilburg, The Netherlands.

Erk. K., Kowalski, A., & Pado, S. (2003a). The SALSA Annotation Tool. In *Proceedings of the Workshop on Prospects and Advances in the Syntax/Semantics Interface,* Nancy.

Erk. K., Kowalski, A., Pado, S. & Pinkal, M. (2003b). Towards a Resource for Lexical Semantics: A Large German Corpus with Extensive Semantic Annotation. In *Proceedings of the ACL 2003,* Sapporo, Japan.

Fellbaum, C., (Ed.) (1998). WordNet: An Electronic Lexical Database. MIT Press.

Fillmore, C. J. (1976). Frame semantics and the nature of language. In *Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech*, volume 280, pages 20--32.

Florian, R., Cucerzan, S., Schafer, C. & Yarowsky, D. (2002). Combining classifiers for word sense disambiguation. In: *Journal of Natural Language Engineering*, 8(4): 327-431.

Frank, A. and Erk, K. (2004). Towards an LFG Syntax–Semantics Interface for Frame Semantics Annotation. In Gelbukh, A., (Ed.), *Computational Linguistics and Intelligent Text Processing*, LNCS, Heidelberg: Springer.

Hamp, B. and Feldweg, H (1997): GermaNet - a Lexical-Semantic Net for German. In *Proceedings of ACL workshop Automatic Information Extraction and Building of Lexical Semantic Resources for NLP Applications*. Madrid, 1997.

Kingsbury, P., Palmer, M., & Marcus, M. (2002). Adding semantic annotation to the Penn TreeBank. In *Proceedings of the HLT Conference*, San Diego.

Niles, I. and Pease, A. (2001). Toward a standard upper ontology. In *Proceedings of FOIS-2001,* Ogunquit, Maine. Association for Computing Machinery.

Riezler, S., King, T.H., Kaplan, R.M., Crouch, D., Maxwell, J.T. & Johnson, M. (2002). Parsing the Wall Street Journal using a Lexical-Functional Grammar and Discriminative Estimation Techniques. In *Proceedings of ACL'02*, Philadelphia, PA.

Appendix

Figure 3: F-structure with integrated/separate frame-semantic projection and ont. classes for Ex. (5:H).