Building Text Meaning Representations from Contextually Related Frames
– A Case Study –

Aljoscha Burchardt*, Anette Frank†*, and Manfred Pinkal*
Dept. of Computational Linguistics, Saarland University*
DFKI Saarbrücken†
{albu,frank,pinkal}@coli.uni-sb.de

Abstract

In this paper, we study frame semantic annotations as text meaning representations. We show how partially connected frame structures can be interlinked on the basis of FrameNet's frame relations and contextual relations from deep parsing to arrive at a shallow semantic representation that can be used in practical NLP tasks.

1 Introduction

With the recent success of broad-coverage statistical parsing systems, syntactic analysis is widely used in practical NLP applications, such as Question Answering, or Summarisation. At the same time, it is commonly recognised that more in-depth semantic analysis is needed for further achievements in the field of open-domain NLP-based information access.

In theoretical and computational semantics, truth-conditional semantic formalisms have been studied in-depth over the last decades. Still, we have not, as of today, seen robust and broad-coverage semantic analysis systems that provide deep semantic representations according to any major computational semantics framework.[1] Therefore, it is still an open question to what extent complex semantic representations as provided by these frameworks can be profitably used in broad-coverage NLP applications.

At the same time, large-scale lexical semantic resources such as WordNets (Fellbaum, 1998) have been developed and put to use for approximate semantic modeling in many applications. Recently, the FrameNet (Baker et al., 1998) and PropBank (Kingsbury et al., 2002) projects are developing lexical semantic resources that model predicate-argument structure.

[1] A notable exception is very recent work of (Bos et al., 2004), who compute semantic representations on the basis of a broad-coverage statistical CCG parser. However, due to the lack of proper evaluation standards, the results have still to be taken with care.

In this paper we build on a computational architecture that combines deep syntactic analysis with FrameNet's frame semantics (Frank and Erk, 2004).[2] We investigate frame annotations *in context* as a partial text meaning representation to be used in practical NLP tasks, such as Information Extraction or Question Answering.

FrameNet provides structured predicate-argument meaning and semantic classification, focusing on open class categories (verbs, nouns, adjectives). Frame semantic annotations of contiguous texts are therefore necessarily partial. Due to the missing constructional "glue" in semantics composition, argument and variable binding cannot be defined in a strictly compositional way, and we obtain partially connected graphs of frame structures.[3] A challenge in using frame semantic annotations as a partial text meaning representation structure is to produce more densely connected structures of frames by inducing co-reference relations between frames and frame roles.

We present a case study where we investigate different types of relations between frames when assigned to contiguous portions of text – contextual relations from deep parsing and lexico-semantic frame relations encoded in FrameNet – and show how specific patterns of such relations support the inference of co-referential relations between frames. We discuss possibilities of using learning techniques to induce such co-referential links between frames and sketch our current architecture for interfacing deep syntactic processing with frame semantics and frame-based reasoning.

The paper is structured as follows. Section 2 introduces frame semantics, in particular *frame relations*. Section 3 presents an outline of our investigations into frame-based meaning representation. We discuss how to connect frame annotations to obtain an interlinked (yet partial) semantic representation. We then present a worked-out example that illustrates how specific configurations of lexico-semantic and contextual relations can license the induction of co-referential links between frames. We argue that this process can be generalised and automated. In Section 4 we present the computational architecture we currently use for frame annotation of contiguous text, and an interface to a state-of-the-art reasoning architecture. In Section 5 we summarise our results and outline the next steps towards an architecture including variable-depth semantics construction and frame-based reasoning.

[2]The architecture builds on LFG-based processing (see Butt et al., 2002).

[3]Another aspect is that frame semantics provides relatively coarse-grained meaning descriptions: for example, predicates are not marked for polarity or factivity. That is, *like* and *dislike* are lexical units of the same frame without further meaning distinction; the same holds for predicates like *claim* and *confess*. Recent developments, within FrameNet, to include semantic types are a first step to address this point.

2 Frame Semantics

2.1 FrameNet

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* that are involved in the situation. FrameNet currently contains about 550 frames of general conceptual classes.[4] For our investigation, we concentrate on the domain of *criminal process*, which is particularly well worked out. As an example, consider the frame VERDICT with the semantic roles CASE, CHARGES, DEFENDANT, FINDING and JUDGE. This frame is evoked by words like *convict.v, find.v, verdict.n*, as in example (1). FrameNet further defines *extrathematic roles*, such as LOCATION in (1), which are not frame-specific.

(1) [Baragiola]$_{\text{DEFENDANT}}$ had previously been *convicted* [of murder]$_{\text{CHARGES}}$ [in Italy]$_{\text{LOCATION}}$, but had escaped in 1980 and obtained Swiss citizenship.

Examples (2) and (3) illustrate more linguistic variations of this frame with different instantiated roles and frame evoking elements.

(2) [The jury]$_{\text{JUDGE}}$ *convicted* [him]$_{\text{DEFENDANT}}$ [on the counts of theft]$_{\text{CHARGES}}$.

(3) On Thursday [a jury]$_{\text{JUDGE}}$ *found* [the youth]$_{\text{DEFENDANT}}$ [guilty of wounding Mr Lay]$_{\text{FINDING}}$.

2.2 Frame Relations – "FrameNet as a Net"

FrameNet defines a number of different types of relations between frames that provide more internal structure to the lexical database (Fillmore et al., 2004). The relevant relations for our purposes are the *Inheritance* and the *Subframe* relation. If a frame F_1 inherits from some frame F_2, then all roles of F_2 are also available at F_1 (modulo renaming). For example, the frame ARREST inherits the roles AGENT and PATIENT from the frame INTENTIONALLY_AFFECT (renamed into AUTHORITIES and SUSPECT).

The Subframe relation is used to model abstract 'scenario frames', such as CRIMINAL PROCESS or EMPLOYMENT. Scenario frames represent complex events with subframe relations holding between the scenario frame

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

and frames that describe (temporally ordered) sub-events. For example, the frame CRIMINAL_PROCESS has the subframes ARRAIGNMENT, ARREST, SENTENCING, and TRIAL. Subframes usually inherit roles from their super frame, e.g. CHARGE and DEFENDANT of ARRAIGNMENT inherit from the respective roles of CRIMINAL_PROCESS. The subframe relation will turn out particularly effective for establishing co-reference in frame-annotated texts.

3 Building Text Meaning Representations from Contextually Related Frames

3.1 Frame Semantics for Partial Text Meaning Representation

In this paper we study frame semantics as a framework for partial text meaning representation. By applying frames to contiguous portions of text – due to the lack of constructional "glue" – we obtain partially connected lexico-semantic predicate-argument structures in a network of frame-to-frame relations. In order to construct a more densely connected frame-based text meaning representation, we need to infer additional links between frames and frame elements. For this we can exploit the contextual relations between frames and frame elements as given by deep parsing: structural embedding or adjacency relations between neighbouring frames.

When trying to induce contextually linked frames, we have to distinguish two levels: the level of frame *instances*, where we can infer co-reference of events or role fillers, and the level of *types*, where we can infer intrinsic relations between frames and roles.

At the instance level, we can establish co-referential links between e.g. a filled role of one frame instance with an unfilled role of another frame instance provided we find sufficient supporting evidence. Two roles can be linked, for example, if – at the type level – the respective frames stand in a subframe relation with inheritance of roles and, in addition, the frame instances are contextually related in appropriate ways, e.g. by functional-syntactic, or semantic role embedding, or else by way of a discourse relation.

At the type level, we can induce relations between frames or roles on the basis of e.g. recurrent anaphoric linking patterns observed in texts. The induction of meaning relations at the type level is more involved and requires use of annotated corpora and learning techniques.

In both cases, the induction of co-reference relations between frames can only be heuristic, given that we build on a partial conceptual structure, not a fully specified truth-conditional semantic representation.

3.2 Frames in Context – A Case Study

In this section we present a case study that establishes systematic patterns of lexical-semantic and contextual relations that support the induction of co-referential relations between frames and roles. As an example we chose a short news wire text (4)[5] that pertains to the "scenario frame" CRIMINAL_PROCESS introduced in Section 2.2.

(4) In the first trial in the world in connection with the terrorist attacks of 11 September 2001, the Higher Regional Court of Hamburg has passed down the maximum sentence. Mounir al Motassadeq will spend 15 years in prison. The 28-year-old Moroccan was found guilty as an accessory to murder in more than 3000 cases.

Table 1 lists all frames and roles that are relevant for the example. Target predicates, such as *trial*, evoke the corresponding frame; the frame-specific semantic roles correspond to local constituents, which are displayed in the right column, e.g. the role TRIAL.CASE corresponds to the constituent *terrorist attacks*. Roles that cannot be associated with local constituents are left unfilled (e.g. ATTACK.VICTIM). Frame element fillers and co-references between frame elements that can be induced on the basis of frame relations, contextual relations or bridging inferences are displayed in brackets. For example, *Higher Regional Court* that fills the role SENTENCING.COURT can be induced as filler of the role TRIAL.COURT.

FrameNet relations. The frames evoked in the example pertain to the following frame relations: Both SENTENCING and TRIAL are subframes of CRIMINAL_PROCESS. VERDICT is again a subframe of TRIAL. Additionally, we assume[6] that ASSISTANCE inherits from INTENTIONALLY_ACT.

Contextual relations. The example features different types of contextual relations between frames and roles: functional syntactic embedding, frame semantic embedding, surface order or discourse relations, and co-reference. For example, SENTENCING and TRIAL are syntactically related by functional (adjunct) embedding; the ATTACK frame is embedded within the CASE role of TRIAL; the sentence projecting PRISON follows, and stands in a discourse relation (ELABORATION) to the sentence projecting SENTENCING. Finally, the referents corresponding to the roles PRISON.INMATES and VERDICT.DEFENDANT can be recognised as co-referent.

[5] `http://www.germnews.de/archive/dn/2003/02/19.html`

[6] This information is not contained in the current FrameNet release.

Target	Frame	Frame element	Filler (given vs. (induced))	
trial	TRIAL	CASE	*terrorist attacks*	(1)
		CHARGE	*(accessory to murder)*	(2)
		COURT	*(Higher Regional Court)*	(3)
		DEFENDANT ...	*(28-year-old Moroccan)*	(4)
attacks	ATTACK	ASSAILANT	*terrorist*	(5)
		VICTIM ...		(6)
		TIME (exth.)	*11 September 2001*	(7)
sentence	SENTENCING	CONVICT	*(Mounir al Motassadeq)*	(8)
		COURT	*Higher Regional Court*	(9)
		TYPE ...	*maximum sentence*	(10)
prison	PRISON	INMATES ...	*Mounir al Motassadeq*	(11)
		DURATION (exth.)	*15 years*	(12)
found guilty	VERDICT	CASE	*(terrorist attacks)*	(13)
		CHARGE	*accessory to murder*	(14)
		DEFENDANT	*28-year-old Moroccan*	(15)
		FINDING ...	*guilty*	(16)
accessory	ASSISTANCE	CO-AGENT		(17)
		FOCAL_ENTITY	*murder*	(18)
		HELPER ...	*(28-year-old Moroccan)*	(19)
murder	KILLING	KILLER		(20)
		VICTIM ...	*m.t. 3000 cases*	(21)

Table 1: Frame Annotations with Given/Inferred Frame Element Linkings

Inferred relations. Based on these lexico-semantic and contextual relations, we can infer further semantic relations between roles and frames, such as co-referential binding of unfilled roles. Figure 1 schematically illustrates the interaction of the central frame relations, contextual relations, and inferred relations that we identified in (4).

Closer study of the inferred relations reveals a number of underlying patterns of justifications, which we will exemplify in turn: In the majority of cases, we can infer role bindings on the basis of (a variety of) patterns of lexical semantic and contextual relations between frames and roles. In some cases, further lexical semantic knowledge is required, which is not yet encoded in FrameNet, such as 'semantic control', or causative relations between frames. We will finally discuss an example which motivates that additional semantic information, such as referential and temporal properties, needs to be considered for inducing role bindings.

Figure 2 illustrates an example of the first type, where we induce role

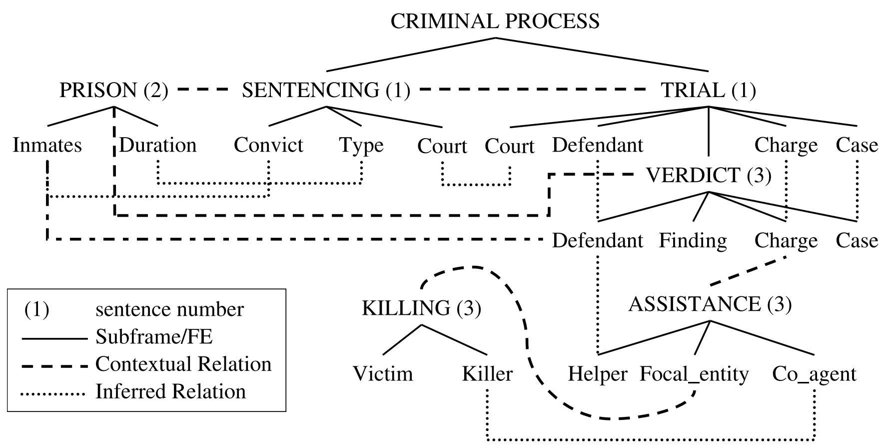

Figure 1: FrameNet Relations, Contextual Relations, and Inferred Relations.

identification of the role fillers of TRIAL.COURT (r_1)[7] and SENTENCING.COURT (r_2) (see (3) and (9) in Table 1). SENTENCING (F_2) and TRIAL (F_1) are subframes of CRIMINAL_PROCESS (F_0) and both role types (R_1,R_2) inherit from CRIMINAL_PROCESS.COURT (R_0). These inheritance and subframe relations are displayed by dashed lines (left). In addition, the frame instances (f_1,f_2) are in a functional (adjunct) embedding relation. This contextual relation is displayed by dotted lines (middle). On this basis, we assume that both frame instances are subframes of *one* CRIMINAL_PROCESS scenario instance (f_0). This leads to the linking of the roles (r_1) and (r_2) (right). Other examples of role identifications that follow this pattern are (1)-(13), (2)-(14), (4)-(15) in Table 1, which are based on the subframe relation between TRIAL and VERDICT.

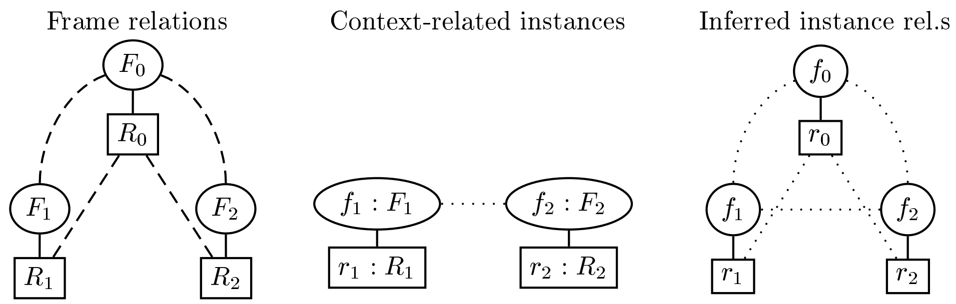

Figure 2: Inferring Instance Relations

[7]Frame (and role) types are printed in upper case, instances in lower case. $f_1 : F_1$ means that f_1 is an instance of frame F_1.

Figure 3 illustrates an example where role identification is induced on the type level, on the basis of a contextual co-reference relation. The frames PRISON (F_1) and VERDICT (F_2) are unrelated in FrameNet (left). In the text, the referents of the roles PRISON.INMATES (r_1) and VERDICT.DEFENDANT (r_2) are marked co-referent by means of a definite description (middle). In this case, we induce role identification at the type level by assumption of an 'anonymous' frame-to-frame relation that can be further specified, e.g. as a causation relation or a subframe relation within some scenario.

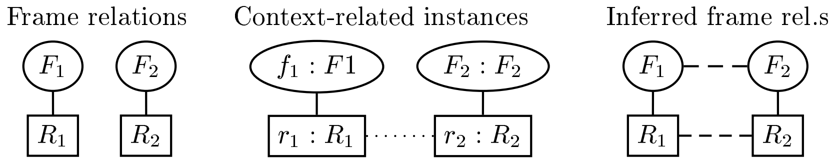

Figure 3: Inducing Frame Relations

In some cases, the (diverse patterns of) frame and contextual relations are not sufficient to induce role-identification. Here, we found that further lexical semantic information is required, in particular what we call *semantic control*, as a kind of meaning postulate: for some frames it is part of their inherent lexical meaning that a given role is co-referent with the agent/patient role of an embedded frame. For example, the defendant in a verdict is (found to be) the actor in the event that constitutes the charge of the verdict. This is represented in Figure 4 (left).

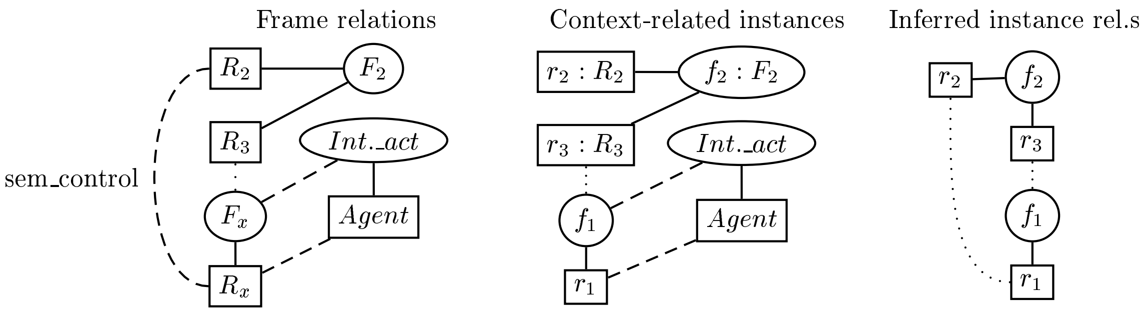

Figure 4: Inferring Instance Relations (by Semantic Control)

VERDICT (F_2) features semantic control, in that VERDICT.DEFENDANT (R_2) is marked identical to the agent of some frame F_x embedded within its CHARGE role (R_3) (dashed line). Agenthood is formally represented by inheritance from INTENTIONALLY_ACT.AGENT.

In the example (middle), VERDICT.CHARGE (r_3) embeds ASSISTANCE

(f_1). Furthermore, ASSISTANCE.HELPER (r_1) inherits from INTENTIONALLY_ACT.AGENT. We can thus conclude that the filler of VERDICT.DEFENDANT (r_2) is identical to the ASSISTANCE.HELPER (r_1) (right) ((15)-(19), Table 1). Other examples that involve semantic control are (17)-(20) and (8)-(11) (the latter assuming a causative relation between SENTENCING and PRISON[8]).

Finally, examples like (5) show that we need to enrich frame semantic representations with selected deeper semantic information to control the induction of role identification. We need to model referential properties, such as the introduction of new discourse referents (*a new trial*), and event modification by locational or temporal adjuncts. The former will be crucial to define 'blocking' factors for role identification rules, the latter will provide deeper semantic characterisations of contextual relations between frames, such as temporal sequence. This calls for a variable-depth semantics construction architecture that allows targeted refinement of the semantic representation.

(5) Mounir El Motassadeq (born April 3, 1974) is a Moroccan. In February 2003 he was convicted [...]. As of April 2004 he is the only person to have been convicted in direct relation to the September 11, 2001 attacks. The verdict and sentence were set aside on appeal [...]. A new trial is expected in mid-2004. (From Wikipedia)

3.3 Acquisition of role-linking patterns

We have identified various patterns of lexico-semantic and contextual relations that support the induction of co-reference relations between frames and roles: FrameNet's frame relations proved essential for linking contextually related (neighbouring) frame instances. Different types of contextual relations could be observed to support role identification: syntactic and semantic embedding, anaphoricity, connectedness by discourse relations or surface linearisation, as well as referential and temporal semantic properties.

More data needs to be investigated to determine the weight of the individual factors. In particular, we need to model referential properties of nouns and verbs in order to define 'blocking factors' for role identification. In future work, we will apply statistical methods for acquiring role-linking patterns from analysed (annotated) text samples of a restricted domain, like CRIMINAL PROCESS.[9] The aim is to learn weighted role-linking patterns that can be formalised as probabilistic inference rules.

[8] A *causation* relation is already defined in FrameNet but not yet broadly annotated.

[9] For experiments along these lines, see (Liakata and Pulman, 2004).

We have provided an abstract definition of semantic control in terms of the agent role marked by inheritance from the perspectivising frame INTENTIONALLY_ACT (the frame INTENTIONALLY_AFFECT additionally provides a patient role). This will facilitate the acquisition of lexical semantic control relations, yet it relies on the full specification of such inheritance relations in the FrameNet data (for the chosen domain).

Based on inferred or given role-linkings and subframe relations, we could also learn more involved patterns of 'bridging' inferences between frames. In (1) (repeated as (6)), given the contextually related (subject of VPs) roles SENTENCING.CONVICT and ESCAPE.ESCAPEE, and given the learned role-linking of SENTENCING.CONVICT and PRISON.INMATES, we can infer an instance of the PRISON frame, with PRISON.INMATES referentially bound to the ESCAPE.ESCAPEE.

(6) [Baragiola]$_{\text{CONVICT/ESCAPEE}}$ had previously been *convicted* of murder in Italy, but had *escaped* in 1980 and obtained Swiss citizenship.

4 Towards Automation

For automated processing of frame-based text meaning representations, we build on a computational syntax-semantics interface for frame assignment with interfaces to a frame-based reasoning architecture.

4.1 An LFG-based Syntax-Semantics Interface

We employ deep syntactic representations provided by large-scale LFG grammars (Butt et al., 2002) as a syntactic basis for frame-based meaning assignment. In Frank and Semecky (2004) we have built a modular syntax-semantics interface where frame semantic representations are projected from the f-structure output of LFG parsing. This architecture yields partially connected frame structures in the projected frame semantics layer.

We have built interfaces to a system for statistical frame and role assignment (see Baldewein et al., 2004) that provides disambiguated frame assignments for a given text. In addition, we have defined interfaces to incorporate co-reference information provided by external anaphora and co-reference resolution systems into the projected frame representations.

For further refinement of the frame semantic representations, we defined semantics construction rules for modifiers that realise extrathematic roles. In similar ways, we will introduce partial representations to model referential and temporal properties of nouns and verbs, respectively.

4.2 Logical Representation and Reasoning

The FrameNet data does not immediately lend itself for use in automated reasoning, as it does not yet come with a formal interpretation. In joint work in Baumgartner and Burchardt (2004), we have transferred the FrameNet frames and selected frame relations into normal *logic programs* to be interpreted under the *stable model semantics.* The paper gives arguments for choosing this framework instead of Description Logics which is currently proposed e.g. in the context of the Semantic Web.

As an additional knowledge source, we have integrated the SUMO/MILO ontology (Niles and Pease, 2001), using an existing Word Sense Disambiguation system and mapping from WordNet to SUMO/MILO classes. Disambiguation on the basis of WordNet also allowed us to access FrameNet by way of a 'detour' via WordNet synsets and relations. Thus, we can hypothesise frame projections for predicates that are not yet included in FrameNet, improving the coverage of our system.[10]

5 Conclusion and Outlook

We presented a case study that investigates frame semantic annotations of contiguous texts as shallow forms of text meaning representation. We established patterns of combination of lexico-semantic and contextual relations that can be used to enrich partially connected frame structures by heuristic inference of co-referential relations. In future work we will investigate the automated acquisition of role-linking patterns from annotated texts.

FrameNet's 'scenario' frames turned out particularly effective for establishing role-linking relations. The linking patterns are not scenario specific, and can thus be regarded as domain-independent methods for frame-based Information Extraction (similar to template filling and merging), where scenario frames serve as linguistically motivated 'domain models'.

With the choice of frame semantic structures as building blocks for a text meaning representation, we deliberately opted for partiality. We aim at an architecture for robust semantic processing with incremental depth of semantic analysis. Starting from robust frame semantic processing for coarse-grained information access, we want to allow for incremental enrichment of the semantic representations to handle special tasks that require more fine-grained and truth-conditional semantic information, such as e.g. answer validation in QA.

[10] A demonstration of this functionality (with manual WSD) can be found at `http://www.coli.uni-sb.de/~albu/cgi-bin/string2frames.cgi`.

References

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

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

Baumgartner, P. and Burchardt, A. (2004). Logic Programming Infrastructure for Inferences on FrameNet. In Alferes, J. J. and Leite, J. A., editors, *Logics in Artificial Intelligence, 9th European Conference, JELIA 2004*, Lecture Notes in Computer Science. Springer Verlag, Berlin, Heidelberg, New York.

Bos, J., Clark, S., Steedman, M., Curran, J., and Hockenmaier, J. (2004). Wide-Coverage Semantic Representations from a CCG Parser. In *Proceedings of Coling 2004*, Geneva, Switzerland.

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

Fellbaum, C., editor (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.

Fillmore, C. J., Baker, C. F., and Sato, H. (2004). FrameNet as a "Net". In *Proceedings of LREC 2004*, pages 61–65, Lisbon, Portugal.

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

Frank, A. and Semecky, J. (2004). Corpus-based Induction of an LFG Syntax-Semantics Interface for Frame Semantic Processing. In Hansen-Schirra, S., Oepen, S., and Uszkoreit, H., editors, *Proceedings of LINC 2004*, Geneva, Switzerland.

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

Liakata, M. and Pulman, S. (2004). Learning Theories from Texts. In *Proceedings of COLING 2004*, pages 183–190, Geneva, Switzerland.

Niles, I. and Pease, A. (2001). Toward a standard upper ontology. In *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)*, Ogunquit, Maine. Association for Computing Machinery.