

Unsupervised Learning of Prototypical Fillers for Implicit Semantic Role Labeling

Niko Schenk and Christian Chiarcos

Applied Computational Linguistics Lab

Goethe University Frankfurt am Main, Germany

{n.schenk, chiarcos}@em.uni-frankfurt.de

Abstract

Gold annotations for supervised implicit semantic role labeling are extremely sparse and costly. As a lightweight alternative, this paper describes an approach based on unsupervised parsing which can do without iSRL-specific training data: We induce prototypical roles from large amounts of explicit SRL annotations paired with their distributed word representations. An evaluation shows competitive performance with supervised methods on the SemEval 2010 data, and our method can easily be applied to predicates (or languages) for which no training annotations are available.

1 Introduction

Semantic role labeling (SRL) (Gildea and Jurafsky, 2002) has become a well-established and highly important NLP component which directly benefits various downstream applications, such as text summarization (Trandabăţ, 2011), recognizing textual entailment (Sammons et al., 2012) or QA systems (Shen and Lapata, 2007; Moreda et al., 2011). Its goal is to detect verbal or nominal predicates, together with their associated arguments and semantic roles, either by PropBank/NomBank (Palmer et al., 2005; Meyers et al., 2004) or FrameNet (Baker et al., 1998) analysis. In its traditional form, however, SRL is restricted to the *local syntactic context of the predicate* as in the following example from Ruppenhofer et al. (2010):

[_{GOAL/NI}In the centre of this room] there was an upright beam, [_{THEME}which] had been **placed** [_{TIME}at some period] as a support for

the old worm-eaten baulk of timber which spanned the roof.

In a FrameNet-style analysis of the sentence, the predicate *place* evokes the PLACING frame, with two frame elements (roles) overtly expressed (THEME and TIME) but with one role – GOAL – beyond the embedded relative clause and thus beyond the scope of the SRL parser. Such implicit roles, or *null instantiations* (NIs) (Fillmore, 1986; Ruppenhofer, 2005) are much harder to detect automatically, as they require to broaden the analysis to the surrounding discourse, commonly also to preceding (or following) sentences.

State-of-the-art approaches to implicit SRL (iSRL) are supervised and need a groundwork of hand-annotated training data – which is costly, extremely sparse, limited to only a handful of predicates, and requires careful feature engineering (Gerber and Chai, 2012; Silberer and Frank, 2012; Li et al., 2015). A first attempt has been made to combine the scarce resources available (Feizabadi and Padó, 2015), but given the great diversity of predicate-specific roles and enormous complexity of the task, the main issues remain (Chen et al., 2010).

A promising exploratory effort recently made by Gorinski et al. (2013) aims to overcome the annotation bottleneck by using distributional methods to infer evidence for elements filling null instantiated roles. The authors do not rely on gold annotations but instead learn distributional properties of fillers induced from a large corpus.

Our Contribution: We propose an extension of the distributional idea for unsupervised iSRL to loosen the need for annotated training data. Specifically, we

propose to induce predicate and role-specific **prototypical fillers** from large amounts of SRL annotated texts in order to resolve null instantiations as (semantically and syntactically) similar elements found in the context. Parts of our approach have been successfully applied in traditional SRL (Hermann et al., 2014), but not yet to implicit roles. Our work differs from Gorinski et al. (2013) in that we extend discrete context vectors to SRL-guided embeddings and experiment with a variety of different configurations. We intend *not* to set a new benchmark beating the current state-of-the-art for *supervised* iSRL, but rather provide a simple and alternative strategy which does not rely on manually annotated gold data. Still, we demonstrate that our method is highly competitive with supervised methods on one out of two standard evaluation sets and that it can easily be extended to other predicates for which no implicit gold annotations are available.

2 Method

2.1 Prototypical Fillers

We use large amounts of *explicit* SRL annotations to compute predicate-specific *protofillers* (prototypical fillers) for each frame element (role) individually:

$$\vec{v}^{\text{protofiller}} = \frac{1}{N} \sum_{i=0}^N E(w_i) \quad (1)$$

where N is the total number of tokens filling a particular role and $E(\cdot)$ is an embedding function which maps a word w_i to its distributed representation, i.e., a precomputed vector of d dimensions. Note that only those words contribute to the protofiller of a frame element which occur in this role.

2.2 Identifying Null Instantiations

Our approach generalizes over labeled filler instances of the frame (PLACING in the example) as found in corpus data, e.g., *placed on the middle picture, planted on the top of the church, hung over the river, laid on the table*, etc. We exploit their syntactic (here: prepositional) and semantic properties (inanimate, spacial NPs) in order to capture a composed meaning and thus to approximate the correct implicit role *in the centre of this room*. We measure similarity between a trained protofiller \vec{v}^p

and a candidate constituent \vec{v}^c by cosine similarity $\cos(\theta) = \frac{\vec{v}^p \cdot \vec{v}^c}{\|\vec{v}^p\| \|\vec{v}^c\|}$ and predict a candidate as null instantiation which maximizes the inner product with the protofiller. As candidate constituents for an implicit argument we initially consider all terminal and non-terminal nodes in a context window of the predicate, ruling out those categories which never occur as implicit arguments, which do contain the target predicate and/or which are already overt arguments. The result set comprises mainly nouns, verbs and PPs. Candidate constituents in our evaluation data are available from their respective (manual) syntax annotation, but could easily be extracted using automated phrase-structure parsers. The candidate *vectors* for arbitrary length n -grams are derived in the same way (by means of Equation 1).

2.3 Training Resources & Tools

In accordance with domain-specific evaluation data, we chose to learn protofillers on two distinct corpora: *The Corpus of Late Modern English Texts, CLMET* (Smet, 2005) (≈ 35 M tokens, 18th–20th century novels) and a subset of the English *Gigaword corpus* (Graff and Cieri, 2003) (≈ 500 M tokens of Newswire texts). We label the first one with SEMAFOR¹ (Das et al., 2014), a FrameNet-style semantic parser. We employ MATE² (Björkelund et al., 2009) to obtain a PropBank/NomBank analysis for each sentence in Gigaword.

	CLMET	Gigaword
# explicit roles	21.9M	264.0M
# predicate instances	9.5M	122.5M
# roles per predicate	2.3	2.2
# predicates per sentence	7.6	4.2

Table 1: Statistics on the number of explicit fillers used for training protofillers.

Table 1 highlights general statistics on the number of predicates collected from both corpora. Two observations are worth noting: While on average the number of explicitly realized roles/frame elements per predicate/frame in both data sets is similar, we find more predicate instances in CLMET than in Gigaword. This is due to the FrameNet lexicon and its more fine-grained modeling of lexical units, as

¹<http://www.cs.cmu.edu/~ark/SEMAFOR/>

²<https://code.google.com/p/mate-tools/>

opposed to PropBank. Also note that FrameNet currently specifies 9.7 frame elements per lexical frame³ which – despite the fact that this number also comprises non-core arguments – is much larger than what can explicitly be labeled by the SRL systems.

Regarding the distributional component, we experimented with a variety of distributed word representations: We chose out of the box vectors; Collobert et al. (2011), dependency-based word embeddings (Levy and Goldberg, 2014) and the pre-trained Google News vectors from *word2vec*⁴ (Mikolov et al., 2013). Using the same tool, we also trained custom embeddings (bag-of-words and skip-gram) with 50 dimensions on our two corpora.

3 Evaluation

In order to assess the usefulness of our approach, a quantitative evaluation has been conducted on two iSRL test sets which have become a de facto standard in this domain: a collection of fiction novels from the SemEval 2010 Shared Task with manual annotations of null instantiations (Ruppenhofer et al., 2010), and Gerber and Chai (2010)’s augmented NomBank data set. Table 2 shows some general statistics on the number of implicit roles and candidate phrases involved in our experiments. As to have a comparison with the supervised approaches referred to in this study, we also provide the size of the training data.

	SemEval	NomBank
# predicate instances		
in training set	1,370	816
in test set	1,703	437
# implicit arguments		
in training set	245	650
in test set	259	246
# of candidate phrases		
per predicate instance	27.6	52.2
proportion of single tokens	63.4%	47.9%
proportion of phrases	36.6%	52.1%
∅ length of candidate phrase (in tokens)	5.8	7.1

Table 2: Statistics on implicit arguments and candidate phrases from the test sections of the two evaluation sets.

³https://framenet.icsi.berkeley.edu/frndrupal/current_status, accessed March 2016.

⁴<https://code.google.com/p/word2vec/>

3.1 SemEval Data

In Table 4, we report the classification scores for the (*NI-only*) null instantiation *linking* task on the SemEval data, given the parsed candidate phrases and the gold information about the missing frame element.⁵ For space reasons, we only include the results of our best-performing configuration, obtained from protofillers trained on the late modern English texts and Collobert et al. (2011) embeddings (C&W) with the search space for candidate NIs limited to the current and previous sentence. As a reference, we compare our results to the two best models (M_1 and $M_{1'}$) by Silberer and Frank (2012), the vector-based resolver (VEC) by Gorinski et al. (2013) – which is most similar to ours – and, finally, their ensemble combination of four semantically informed resolvers by majority vote (4X).

	<i>P</i>	<i>R</i>	<i>F</i> ₁
Silberer and Frank (2012) M_1	30.8	25.1	27.7
Silberer and Frank (2012) $M_{1'}$	35.6	20.1	25.7
Gorinski et al. (2013) VEC	21.0	18.0	19.0
Gorinski et al. (2013) 4X	26.0	24.0	25.0
This paper: C&W embeddings	27.2	25.7	26.4

Table 4: NI linking performance on the SemEval test data.

The figures in Table 4 suggest that our approach clearly outperforms the vector-based method by Gorinski et al. (2013) and is best in terms of overall recognition rate (recall) among all systems. One potential reason for that might be that, in contrast to the VEC resolver, we do not compute mere context vectors but do rely on the valuable annotations obtained from explicit SRL structures. Also, we do not restrict our analysis to *head* words only, as we have seen that syntactic information from function words is crucial for the resolution of null instantiated roles, too. Moreover, our distributional protofiller method is highly competitive with state-of-the-art performance by Silberer and Frank (2012), yet does not yield better results in terms of *F*₁ score. Note however that, in contrast to their approach, ours is largely unsupervised and does neither rely on gold coreference chains, nor do we need to train on im-

⁵ This avoids error propagation from NI *detection* and allows us to directly compare our results to previous approaches on the same task. Note that Laparra and Rigau (2012) do only report their accuracies for the full pipeline.

	<i>B</i>	Gerber & Chai			Laparra & Rigau			Proto C&W			Proto W2Vcbow		
<i>predicates:</i>	<i>F</i> ₁	<i>P</i>	<i>R</i>	<i>F</i> ₁	<i>P</i>	<i>R</i>	<i>F</i> ₁	<i>P</i>	<i>R</i>	<i>F</i> ₁	<i>P</i>	<i>R</i>	<i>F</i> ₁
sale	36.2	47.2	41.7	44.2	41.2	39.4	40.3	61.0	29.6	39.8	60.8	26.8	37.2
price	15.4	36.0	32.6	34.2	53.3	53.3	53.3	14.7	25.8	18.7	21.8	36.6	27.3
investor	9.8	36.8	40.0	38.4	43.0	39.5	41.2	22.5	48.3	30.7	24.1	57.2	33.9
bid	32.3	23.8	19.2	21.3	52.9	51.0	52.0	30.4	31.5	30.9	40.0	41.5	40.7
plan	38.5	78.6	55.0	64.7	40.7	40.7	40.7	41.1	43.2	42.1	44.3	51.0	47.4
cost	34.8	61.1	64.7	62.9	56.1	50.2	53.0	32.5	19.1	24.0	49.9	29.3	36.9
loss	52.6	83.3	83.3	83.3	68.4	63.5	65.8	54.8	73.1	62.6	54.7	63.8	58.9
loan	18.2	42.9	33.3	37.5	25.0	20.0	22.2	33.9	49.0	40.1	33.2	44.2	37.9
investment	0.0	40.0	25.0	30.8	47.6	35.7	40.8	29.1	21.8	24.9	39.2	34.3	36.6
fund	0.0	14.3	16.7	15.4	66.7	33.3	44.4	100.0	33.3	50.0	75.0	25.0	37.5
Overall	26.5	44.5	40.4	42.3	47.9	43.8	45.8	30.2	35.2	32.5	33.5	39.2	36.1

Table 3: Classification scores for implicit argument labeling on the NomBank test section. Baseline *B* from Gerber & Chai (2010): uses previous occurrence of same predicate. Gerber & Chai (2010): supervised logistic regression classifier trained on implicit fillers. Laparra & Rigau (2013): algorithm based on coherence relationship between predicates and fillers. Our best-performing protofillers are obtained by Collobert et al. (2011) embeddings (**Proto C&W**) and custom trained vectors (**Proto W2Vcbow**) using Gigaword SRL annotations.

licit semantic roles in a supervised setting. An error analysis of our method reveals that it is particularly effective for NIs encountered in the *same sentence* as the target predicate (44.4% accuracy), which seems plausible given the contextual setup in which protofillers are derived.

3.2 NomBank Data

Compared to the SemEval data, Gerber and Chai (2010)’s augmented NomBank resource covers only ten nominal predicates, which allows us to nicely visualize the distributional profile based on their prototypical fillers. For each predicate, we simply concatenate all per role computed protofillers and apply multidimensional scaling to project the so obtained vectors onto two dimensions (cf. Figure 1).

We observe that the predicate grouping is now based on the prototypical fillers that they co-occur with: In the Wall Street Journal texts, *loss*, *loan* and *investment* are similar because their proto-agents (A0 fillers) who lose, lend and invest resp. are semantically shared (i.e. companies, banks). Similarly, *bid*, *cost* and *fund* are related in that the targets or commodities (A2) are all money-financed. Finally, the predicates *sale* and *plan* are to be expected as outliers as they are less homogeneous in their prototypical argument structure.

We have empirically evaluated our protofiller




Figure 1: Clustered projection of the ten nominal predicates from Gerber and Chai (2010) in protofiller space.

method also on this data set: Table 3 reports the classification scores for implicit argument resolution compared to the state-of-the-art (Laparra and Rigau, 2013). We restrict the search for implicit arguments to certain predicate-specific parts-of-speech, since some syntactic constituents (e.g., SBAR) never occur as implicit arguments. For choosing the final implicit arguments for each individual predicate instance, we follow the same deterministic strategy as described in Gerber and Chai (2010), which informally states that, if a certain role is not overtly expressed (within a chain of mentions of the same predicate in previous sentences), it is an implicit

candidate. POS lists and cosine similarity thresholds which trigger an actual prediction have been optimized on the development set. The context window for candidate NIs is optimal for the current and previous two sentences in our setting, which explains why the the number of candidate constituents is approximately twice as large for the NomBank predicates (cf. Table 2).

Our best-performing protofillers are again obtained by Collobert et al. (2011) embeddings substituting explicit SRL annotations in the Gigaword corpus, and with custom-trained embeddings using the continuous bag-of-words model. Overall, our results significantly exceed the highly informed baseline but cannot beat the state-of-the art on this test set. For some predicates, the protofillers seem to generalize better (higher recall), and in particular for the low-frequency predicates (*fund*), precision can be increased. Also, we found that the dependency-based word embeddings do perform slightly worse (not shown), compared to our optimal two configurations. This might be due to the fact that the inherent properties of dependency-based contexts mostly focus on relations between semantically valuable nouns, ignoring (“skipping”) functional words and categories.⁶ The same pertains to the pre-computed Google News vectors which come with a frequency cutoff excluding stop words, again a constraint which is harmful for the correct identification of implicit roles. Furthermore, skip-gram embeddings perform significantly worse than the embeddings derived by the continuous bag-of-words implementation (relative decrease in F_1 by more than 30%). Finally, we observed that inferring implicit roles for nominal predicates is much more challenging because our collected fillers exhibit a much greater variation. For example, the protoagents of *loan* can roughly be divided into two categories, institutions and countries. This in turn introduces noise and has a negative effect on the quality of the singleton protofillers which by vector average capture neither of the two groups perfectly. Promising alternatives could operate on (topic-like) protofiller **clusters** which we leave for future work.

⁶This is also nicely illustrated in Levy and Goldberg (2014).

4 Summary & Conclusion

We have described a lightweight approach for the resolution of implicit semantic roles which does not rely on manual gold annotations. For each predicate-specific role, our method generalizes over explicit SRL-guided annotations incorporating pre-trained word embeddings. This allows us to capture their idiosyncratic properties and use the so-inferred protofillers to find null instantiated roles by means of distributional similarity.

Our method has proven to be generally useful, in particular on the SemEval data, where it is competitive with supervised systems. Its greatest benefit stems from its simplicity and from that fact that it allows to induce null-instantiated roles for arbitrary predicates. As it is applicable even if no iSRL training data is available, it represents a promising technique to address iSRL data scarcity issues.

In our experiment, we employed PropBank/NomBank-style (i)SRL annotations, and our general design clearly benefits from using small-scale inventories of semantic roles. It should be noted though, that our approach is not restricted to any particular SRL tagset, but can be equally applied to other role inventories with similar degrees of consistence and size. Beyond SRL annotations in a strict sense, this might even extend to syntactic dependency annotations that are occasionally taken as a substitute for semantic roles proper. In particular, we see potential in combining our experiments with on-going efforts to cross-lingual projection, adaptation and harmonization of syntax annotations along the lines of Sukhareva and Chiarcos (2014, 2016) and related approaches based on frameworks such as the Universal Dependencies (Nivre, 2015, UD).⁷ If successful, an adaptation using grammatical relations rather than semantic roles represents a promising possibility to create iSRL annotation and iSRL annotation tools for other languages, as Universal Dependencies are becoming increasingly available for major and low-resourced languages and can be projected to others.

The protofillers involved in this study are available at: <http://acoli.cs.uni-frankfurt.de/resources>.

⁷<http://universaldependencies.github.io>

Acknowledgments

The authors would like to thank Joyce Chai and Hendrik De Smet for providing us access to their resources and corpora. We are grateful to Egoitz Laparra for sending us their evaluation script and also thank the anonymous reviewers for their valuable feedback and insightful comments.

References

- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics (ACL-COLIG 1998)*, pages 86–90, Montreal, Quebec, Canada.
- Anders Björkelund, Love Hafdel, and Pierre Nugues. 2009. Multilingual Semantic Role Labeling. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009): Shared Task*, pages 43–48, Boulder, Colorado, June.
- Desai Chen, Nathan Schneider, Dipanjan Das, and Noah A. Smith. 2010. SEMAFOR: Frame Argument Resolution with Log-linear Models. In *Proceedings of the 5th International Workshop on Semantic Evaluation (SemEval-2010)*, pages 264–267, Los Angeles.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural Language Processing (Almost) from Scratch. *Journal of Machine Learning Research*, 12:2493–2537.
- Dipanjan Das, Desai Chen, André F. T. Martins, Nathan Schneider, and Noah A. Smith. 2014. Frame-semantic Parsing. *Computational Linguistics*, 40(1):9–56.
- Parvin S. Feizabadi and Sebastian Padó. 2015. Combining Seemingly Incompatible Corpora for Implicit Semantic Role Labeling. In *Proceedings of the 4th Joint Conference on Lexical and Computational Semantics (*SEM 2015)*, pages 40–50, Denver, CO.
- Charles J. Fillmore. 1986. Pragmatically Controlled Zero Anaphora. In *Proceedings of Berkeley Linguistics Society*, pages 95–107, Berkeley, CA.
- Matthew Gerber and Joyce Chai. 2010. Beyond NomBank: A Study of Implicit Arguments for Nominal Predicates. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics (ACL-2010)*, pages 1583–1592, Uppsala, Sweden.
- Matthew Gerber and Joyce Chai. 2012. Semantic Role Labeling of Implicit Arguments for Nominal Predicates. *Computational Linguistics*, 38(4):755–798.
- Daniel Gildea and Daniel Jurafsky. 2002. Automatic Labeling of Semantic Roles. *Computational Linguistics*, 28(3):245–288.
- Philip Gorinski, Josef Ruppenhofer, and Caroline Sporleder. 2013. Towards Weakly Supervised Resolution of Null Instantiations. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)*, pages 119–130, Potsdam, Germany.
- David Graff and Christopher Cieri. 2003. English Gigaword. Linguistic Data Consortium, Philadelphia, LDC2003T05. Web Download.
- Karl Moritz Hermann, Dipanjan Das, Jason Weston, and Kuzman Ganchev. 2014. Semantic Frame Identification with Distributed Word Representations. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL-2014)*, Baltimore, Maryland.
- Egoitz Laparra and German Rigau. 2012. Exploiting Explicit Annotations and Semantic Types for Implicit Argument Resolution. In *Proceedings of the 6th International Conference on Semantic Computing (ICSC-2012)*, pages 75–78.
- Egoitz Laparra and German Rigau. 2013. ImpAr: A Deterministic Algorithm for Implicit Semantic Role Labeling. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL-2013)*, pages 1180–1189, Sofia, Bulgaria.
- Omer Levy and Yoav Goldberg. 2014. Dependency-Based Word Embeddings. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL-2014)*, pages 302–308, Baltimore, MD.
- Ru Li, Juan Wu, Zhiqiang Wang, and Qinghua Chai. 2015. Implicit Role Linking on Chinese Discourse: Exploiting Explicit Roles and Frame-to-Frame Relations. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2015)*, pages 1263–1271, Beijing, China.
- Adam Meyers, Ruth Reeves, Catherine Macleod, Rachel Szekely, Veronika Zielinska, Brian Young, and Ralph Grishman. 2004. The NomBank Project: An Interim Report. In *Proceedings of the HLT-NAACL 2004 Workshop on Frontiers in Corpus Annotation*, pages 24–31, Boston, Massachusetts.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *CoRR*, abs/1301.3781.
- Paloma Moreda, Hector Llorens, Estela Saquete, and Manuel Palomar. 2011. Combining Semantic Information in Question Answering Systems. *Journal of Information Processing and Management*, 47(6):870–885.

- Joakim Nivre. 2015. Towards a Universal Grammar for Natural Language Processing. In *Proceedings of the 16th International Conference on Computational Linguistics and Intelligent Text Processing (CICLing-2015)*, pages 3–16, Cairo, Egypt. LNCS 9041, Springer.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational Linguistics*, 31(1):71–106.
- Josef Ruppenhofer, Caroline Sporleder, Roser Morante, Collin Baker, and Martha Palmer. 2010. SemEval-2010 Task 10: Linking Events and Their Participants in Discourse. In *Proceedings of the 5th International Workshop on Semantic Evaluation (SemEval-2010)*, pages 45–50, Los Angeles, CA.
- Josef Ruppenhofer. 2005. Regularities in Null Instantiation. Ms, University of Colorado.
- Mark Sammons, V.G.Vinod Vydiswaran, and Dan Roth. 2012. Recognizing Textual Entailment. In Daniel M. Bikel and Imed Zitouni, editors, *Multilingual Natural Language Applications: From Theory to Practice*, chapter 6, pages 209–258. IBM Press.
- Dan Shen and Mirella Lapata. 2007. Using Semantic Roles to Improve Question Answering. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL 2007)*, pages 12–21, Prague, Czech Republic.
- Carina Silberer and Anette Frank. 2012. Casting Implicit Role Linking as an Anaphora Resolution Task. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics (*SEM-2012)*, pages 1–10, Montreal, Quebec, Canada.
- Hendrik De Smet. 2005. A Corpus of Late Modern English Texts. *International Computer Archive of Modern and Medieval English (ICAME)*, 29:69–82.
- Maria Sukhareva and Christian Chiarcos. 2014. Diachronic Proximity vs. Data Sparsity in Cross-lingual Parser Projection. A Case Study on Germanic. In *COLING-2014 Workshop on Applying NLP Tools to Similar Languages, Varieties and Dialects (VarDial-2014)*, Dublin, Ireland.
- Maria Sukhareva and Christian Chiarcos. 2016. Combining Ontologies and Neural Networks for Analyzing Historical Language Varieties. A Case Study in Middle Low German. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC-2016)*, Portorož, Slovenia.
- Diana Trandabăț. 2011. Using Semantic Roles to Improve Summaries. In *Proceedings of the 13th European Workshop on Natural Language Generation (ENLG-2011)*, pages 164–169, Nancy, France.