

Unsupervised Semantic Frame Induction using Triclustering

Dmitry Ustalov[†], Alexander Panchenko[‡], Andrei Kutuzov^{*},
Chris Biemann[‡], and Simone Paolo Ponzetto[†]

[†]University of Mannheim, Germany

{dmitry, simone}@informatik.uni-mannheim.de

^{*}University of Oslo, Norway

andreku@ifi.uio.no

[‡]University of Hamburg, Germany

{panchenko, biemann}@informatik.uni-hamburg.de

Abstract

We use dependency triples automatically extracted from a Web-scale corpus to perform unsupervised semantic frame induction. We cast the frame induction problem as a *triclustering* problem that is a generalization of clustering for *triadic* data. Our replicable benchmarks demonstrate that the proposed graph-based approach, *Triframes*, shows state-of-the-art results on this task on a FrameNet-derived dataset and performing on par with competitive methods on a verb class clustering task.

1 Introduction

Recent years have seen much work on Frame Semantics (Fillmore, 1982), enabled by the availability of a large set of frame definitions, as well as a manually annotated text corpus provided by the FrameNet project (Baker et al., 1998). FrameNet data enabled the development of wide-coverage frame parsers using supervised learning (Gildea and Jurafsky, 2002; Erk and Padó, 2006; Das et al., 2014, *inter alia*), as well as its application to a wide range of tasks, ranging from answer extraction in Question Answering (Shen and Lapata, 2007) and Textual Entailment (Burchardt et al., 2009; Ben Aharon et al., 2010).

However, frame-semantic resources are arguably expensive and time-consuming to build due to difficulties in defining the frames, their granularity and domain, as well as the complexity of the construction and annotation tasks requiring expertise in the underlying knowledge. Consequently, such resources exist only for a few languages (Boas, 2009) and even English is lacking domain-specific frame-based resources. Possible inroads are cross-lingual semantic annotation transfer (Padó and Lapata, 2009; Hartmann

FrameNet	Role	Lexical Units (LU)
<i>Perpetrator</i>	Subject	kidnapper, alien, militant
<i>FEE</i>	Verb	snatch, kidnap, abduct
<i>Victim</i>	Object	son, people, soldier, child

Table 1: Example of a LU tricluster corresponding to the “Kidnapping” frame from FrameNet.

et al., 2016) or linking FrameNet to other lexical-semantic or ontological resources (Narayanan et al., 2003; Tonelli and Pighin, 2009; Laparra and Rigau, 2010; Gurevych et al., 2012, *inter alia*). But while the arguably simpler task of PropBank-based Semantic Role Labeling has been successfully addressed by unsupervised approaches (Lang and Lapata, 2010; Titov and Klementiev, 2011), fully unsupervised frame-based semantic annotation exhibits far more challenges, starting with the preliminary step of automatically inducing a set of semantic frame definitions that would drive a subsequent text annotation. In this work, we aim at overcoming these issues by automatizing the process of FrameNet construction through unsupervised frame induction techniques.

Triclustering. In this work, we cast the frame induction problem as a *triclustering* task (Zhao and Zaki, 2005; Ignatov et al., 2015), namely a generalization of standard clustering and bi-clustering (Cheng and Church, 2000), aiming at simultaneously clustering objects along three dimensions (cf. Table 1). First, using triclustering allows to avoid sequential nature of frame induction approaches, e.g. (Kawahara et al., 2014), where two independent clusterings are needed. Second, benchmarking frame induction as triclustering against other methods on dependency triples allows to abstract away the evaluation of the frame induction algorithm from other factors, e.g., the input corpus or pre-processing steps, thus allowing a fair comparison of different induction models.

The **contributions of this paper** are three-fold: (1) we are the *first to apply* triclustering algorithms for unsupervised frame induction, (2) we propose a *new approach to triclustering*, achieving state-of-the-art performance on the frame induction task, (3) we propose a *new method for the evaluation* of frame induction enabling straightforward comparison of approaches. In this paper, we focus on the simplest setup with *subject-verb-object* (SVO) triples and two roles, but our evaluation framework can be extended to more roles.

In contrast to the recent approaches like the one by [Jauhar and Hovy \(2017\)](#), our approach induces semantic frames without any supervision, yet capturing only two core roles: the subject and the object of a frame triggered by verbal predicates. Note that it is not generally correct to expect that the SVO triples obtained by a dependency parser are necessarily the core arguments of a predicate. Such roles can be implicit, i.e., unexpressed in a given context ([Schenk and Chiarcos, 2016](#)). Keeping this limitation in mind, we assume that the triples obtained from a Web-scale corpus cover most core arguments sufficiently.

Related Work. *LDA-Frames* ([Materna, 2012, 2013](#)) is an approach to inducing semantic frames using LDA ([Blei et al., 2003](#)) for generating semantic frames and their respective frame-specific semantic roles at the same time. The authors evaluated their approach against the CPA corpus ([Hanks and Pustejovsky, 2005](#)). *ProFinder* ([Cheung et al., 2013](#)) is another generative approach that also models both frames and roles as latent topics. The evaluation was performed on the in-domain information extraction task MUC-4 ([Sundheim, 1992](#)) and on the text summarization task TAC-2010.¹ [Modi et al. \(2012\)](#) build on top of an unsupervised semantic role labeling model ([Titov and Klementiev, 2012](#)). The raw text of sentences from the FrameNet data is used for training. The FrameNet gold annotations are then used to evaluate the labeling of the obtained frames and roles, effectively clustering instances known during induction. [Kawahara et al. \(2014\)](#) harvest a huge collection of verbal predicates along with their argument instances and then apply the Chinese Restaurant Process clustering algorithm to group predicates with similar arguments. The approach was evaluated on the verb

¹<https://tac.nist.gov/2010/Summarization>

cluster dataset of [Korhonen et al. \(2003\)](#).

A major issue with unsupervised frame induction task is that these and some other related approaches, e.g., ([O’Connor, 2013](#)), were all evaluated in completely different incomparable settings, and used different input corpora. In this paper, we propose a methodology to resolve this issue.

2 The Triframes Algorithm

Our approach to frame induction relies on graph clustering. We focused on a simple setup using two roles and the SVO triples, arguing that it still can be useful, as frame roles are primarily expressed by subjects and objects, giving rise to semantic structures extracted in an unsupervised way with high coverage.

Input Data. As the input data, we use SVO triples extracted by a dependency parser. According to our statistics on the dependency-parsed FrameNet corpus of over 150 thousand sentences ([Bauer et al., 2012](#)), the SUBJ and OBJ relationships are the two most common shortest paths between frame evoking elements (*FEEs*) and their roles, accounting for 13.5 % of instances of a heavy-tail distribution of over 11 thousand different paths that occur three times or more in the FrameNet data. While this might seem a simplification that does not cover prepositional phrases and frames filling the roles of other frames in a nested fashion, we argue that the overall frame inventory can be induced on the basis of this restricted set of constructions, leaving other paths and more complex instances for further work.

The Method. Our method constructs embeddings for SVO triples to reduce the frame induction problem to a simpler graph clustering problem. Given the vocabulary V , a d -dimensional word embedding model $v \in V \rightarrow \vec{v} \in \mathbb{R}^d$, and a set of SVO triples $T \subseteq V^3$ extracted from a syntactically analyzed corpus, we construct the triple similarity graph \mathcal{G} . Clustering of \mathcal{G} yields sets of triples corresponding to the instances of the semantic frames, thereby clustering frame-evoking predicates and roles simultaneously.

We obtain dense representations of the triples T by concatenating the word vectors corresponding to the elements of each triple by transforming a triple $t = (s, p, o) \in T$ into the $(3d)$ -dimensional vector $\vec{t} = \vec{s} \oplus \vec{p} \oplus \vec{o}$. Subsequently, we use the triple embeddings to generate the undirected graph

Algorithm 1 *Triframes* frame induction

Input: an embedding model $v \in V \rightarrow \vec{v} \in \mathbb{R}^d$,
a set of SVO triples $T \subseteq V^3$,
the number of nearest neighbors $k \in \mathbb{N}$,
a graph clustering algorithm CLUSTER.

Output: a set of triframes F .

- 1: $S \leftarrow \{t \rightarrow \vec{t} \in \mathbb{R}^{3d} : t \in T\}$
 - 2: $E \leftarrow \{(t, t') \in T^2 : t' \in \text{NN}_k^S(\vec{t}), t \neq t'\}$
 - 3: $F \leftarrow \emptyset$
 - 4: **for all** $C \in \text{CLUSTER}(T, E)$ **do**
 - 5: $f_s \leftarrow \{s \in V : (s, v, o) \in C\}$
 - 6: $f_v \leftarrow \{v \in V : (s, v, o) \in C\}$
 - 7: $f_o \leftarrow \{o \in V : (s, v, o) \in C\}$
 - 8: $F \leftarrow F \cup \{(f_s, f_v, f_o)\}$
 - 9: **return** F
-

$\mathcal{G} = (T, E)$ by constructing the edge set $E \subseteq T^2$. For that, we compute $k \in \mathbb{N}$ nearest neighbors of each triple vector $\vec{t} \in \mathbb{R}^{3d}$ and establish cosine similarity-weighted edges between the corresponding triples.

Then, we assume that the triples representing similar contexts appear in similar roles, which is explicitly encoded by the concatenation of the corresponding vectors of the words constituting the triple. We use graph clustering of \mathcal{G} to retrieve communities of similar triples forming frame clusters; a clustering algorithm is a function $\text{CLUSTER} : (T, E) \rightarrow \mathbb{C}$ such that $T = \bigcup_{C \in \mathbb{C}} C$. Finally, for each cluster $C \in \mathbb{C}$, we aggregate the subjects, the verbs, and the objects of the contained triples into separate sets. As the result, each cluster is transformed into a *triframe*, which is a triple that is composed of the subjects $f_s \subseteq V$, the verbs $f_v \subseteq V$, and the objects $f_o \subseteq V$.

Our frame induction approach outputs a set of triframes F as presented in Algorithm 1. The hyper-parameters of the algorithm are the number of nearest neighbors for establishing edges (k) and the graph clustering algorithm CLUSTER. During the concatenation of the vectors for words forming triples, the $(|T| \times 3d)$ -dimensional vector space S is created. Thus, given the triple $t \in T$, we denote the k nearest neighbors extraction procedure of its concatenated embedding from S as $\text{NN}_k^S(\vec{t}) \subseteq T$. We used $k = 10$ nearest neighbors per triple.

To cluster the nearest neighbor graph of SVO triples \mathcal{G} , we use the WATSET *fuzzy graph clustering* algorithm (Ustalov et al., 2017). It treats the vertices T of the input graph \mathcal{G} as the SVO triples, induces their senses, and constructs an intermedi-

ate sense-aware representation that is clustered using the Chinese Whispers (CW) hard clustering algorithm (Biemann, 2006). We chose WATSET due to its performance on the related synset induction task, its fuzzy nature, and the ability to find the number of frames automatically.

3 Evaluation

Input Corpus. In our evaluation, we use triple frequencies from the DepCC dataset (Panchenko et al., 2018), which is a dependency-parsed version of the Common Crawl corpus, and the standard 300-dimensional word embeddings model trained on the Google News corpus (Mikolov et al., 2013). All evaluated algorithms are executed on the same set of triples, eliminating variations due to different corpora or pre-processing.

Datasets. We cast the complex multi-stage frame induction task as a straightforward triple clustering task. We constructed a gold standard set of triclusters, each corresponding to a FrameNet frame, similarly to the one illustrated in Table 1. To construct the evaluation dataset, we extracted frame annotations from the over 150 thousand sentences from the FrameNet 1.7 (Baker et al., 1998). Each sentence contains data about the frame, FEE, and its arguments, which were used to generate triples in the form $(\text{word}_i : \text{role}_1, \text{word}_j : \text{FEE}, \text{word}_k : \text{role}_2)$, where $\text{word}_{i/j/k}$ correspond to the roles and FEE in the sentence. We omitted roles expressed by multiple words as we use dependency parses, where one node represents a single word only.

For the sentences where more than two roles are present, all possible triples were generated. Sentences with less than two roles were omitted. Finally, for each frame, we selected only two roles, which are most frequently co-occurring in the FrameNet annotated texts. This has left us with about 100 thousand instances for the evaluation. For the evaluation purposes, we operate on the intersection of triples from DepCC and FrameNet. Experimenting on the full set of DepCC triples is only possible for several methods that scale well (WATSET, CW, k -means), but is prohibitively expensive for other methods (LDA-Frames, NOAC).

In addition to the frame induction evaluation, where subjects, objects, and verbs are evaluated together, we also used a dataset of polysemous verb classes introduced in (Korhonen et al., 2003) and employed by Kawahara et al. (2014). Statis-

Dataset	# instances	# unique	# clusters
FrameNet Triples	99,744	94,170	383
Poly. Verb Classes	246	110	62

Table 2: Statistics of the evaluation datasets.

tics of both datasets are summarized in Table 2. Note that the polysemous verb dataset is rather small, whereas the FrameNet triples set is fairly large, enabling reliable comparisons.

Evaluation Measures. Following the approach for verb class evaluation by Kawahara et al. (2014), we employ *normalized modified purity* (nmPU) and *normalized inverse purity* (niPU) as the clustering quality measures. Given the set of the obtained clusters K and the set of the gold clusters G , normalized modified purity quantifies the clustering precision as the average of the weighted overlap $\delta_{K_i}(K_i \cap G_j)$ between each cluster $K_i \in K$ and the gold cluster $G_j \in G$ that maximizes the overlap with K_i : $\text{nmPU} = \frac{1}{N} \sum_i^{K_i} \max_{1 \leq j \leq |G|} \delta_{K_i}(K_i \cap G_j)$, where the weighted overlap is the sum of the weights c_{iv} for each word v in i -th cluster: $\delta_{K_i}(K_i \cap G_j) = \sum_{v \in K_i \cap G_j} c_{iv}$. Note that nmPU counts all the singleton clusters as wrong. Similarly, normalized inverse purity (collocation) quantifies the clustering recall: $\text{niPU} = \frac{1}{N} \sum_{j=1}^{|G|} \max_{1 \leq i \leq |K|} \delta_{G_j}(K_i \cap G_j)$. nmPU and niPU are combined together as the harmonic mean to yield the overall clustering F-score (F_1), which we use to rank the approaches.

Our framework can be extended to evaluation of more than two roles by generating more roles per frame. Currently, given a set of gold triples generated from the FrameNet, each triple element has a role, e.g., “Victim”, “Predator”, and “FEE”. We use fuzzy clustering evaluation measure which operates not on triples, but instead on a set of tuples. Consider for instance a gold triple (Freddy: Predator, kidnap: FEE, kid: Victim). It will be converted to three pairs (Freddy, Predator), (kidnap, FEE), (kid, Victim). Each cluster in both K and G is transformed into a union of all constituent typed pairs. The quality measures are finally calculated between these two sets of tuples, K , and G . Note that one can easily pull in more than two core roles by adding to this gold standard set of tuples other roles of the frame, e.g., (forest, Location). In our experiments, we focused on two main roles as our contribution is related to the application of triclustering methods. However, if more advanced

methods of clustering are used, yielding clusters of arbitrary modality (n -clustering), one could also use our evaluation schema.

Baselines. We compare our method to several available state-of-the-art baselines applicable to our dataset of triples.

LDA-Frames by Materna (2012, 2013) is a frame induction method based on topic modeling. We ran 500 iterations of the model with the default parameters. *Higher-Order Skip-Gram (HOSG)* by Cotterell et al. (2017) generalizes the Skip-Gram model (Mikolov et al., 2013) by extending it from word-context co-occurrence matrices to tensors factorized with a polyadic decomposition. In our case, this tensor consisted of SVO triple counts. We trained three vector arrays (for subjects, verbs and objects) on the 108,073 SVO triples from the *FrameNet* corpus, using the implementation by the authors. Training was performed with 5 negative samples, 300-dimensional vectors, and 10 epochs. We constructed an embedding of a triple by concatenating embeddings for subjects, verbs, and objects, and clustered them using k -means with the number of clusters set to 10,000 (this value provided the best performance). *NOAC* (Egurnov et al., 2017) is an extension of the Object Attribute Condition (OAC) triclustering algorithm (Ignatov et al., 2015) to numerically weighted triples. This incremental algorithm searches for dense regions in triadic data. A minimum density of 0.25 led to the best results. In the *Triadic* baselines, independent word embeddings of subject, object, and verb are concatenated and then clustered using a *hard clustering algorithm*: k -means, spectral clustering, or CW.

We tested various hyper-parameters of each of these algorithms and report the best results overall per clustering algorithm. Two trivial baselines are *Singletons* that creates a single cluster per instance and *Whole* that creates one cluster for all elements.

4 Results

We perform two experiments to evaluate our approach: (1) a frame induction experiment on the FrameNet annotated corpus by Bauer et al. (2012); (2) the polysemous verb clustering experiment on the dataset by Korhonen et al. (2003). The first is based on the newly introduced frame induction evaluation schema (cf. Section 3). The second one evaluates the quality of verb clusters only on a standard dataset from prior work.

Method	Verb			Subject			Object			Frame		
	nmPU	niPU	F ₁	nmPU	niPU	F ₁	nmPU	niPU	F ₁	nmPU	niPU	F ₁
Triframes WATSET	42.84	88.35	57.70	54.22	81.40	65.09	53.04	83.25	64.80	55.19	60.81	57.87
HOSG (Cotterell et al., 2017)	44.41	68.43	53.86	52.84	74.53	61.83	54.73	74.05	62.94	55.74	50.45	52.96
NOAC (Egurnov et al., 2017)	20.73	88.38	33.58	57.00	80.11	66.61	57.32	81.13	67.18	44.01	63.21	51.89
Triadic Spectral	49.62	24.90	33.15	50.07	41.07	45.13	50.50	41.82	45.75	52.05	28.60	36.91
Triadic <i>k</i> -Means	63.87	23.16	33.99	63.15	38.20	47.60	63.98	37.43	47.23	63.64	24.11	34.97
LDA-Frames (Materna, 2013)	26.11	66.92	37.56	17.28	83.26	28.62	20.80	90.33	33.81	18.80	71.17	29.75
Triframes CW	7.75	6.48	7.06	3.70	14.07	5.86	51.91	76.92	61.99	21.67	26.50	23.84
Singletons	0.00	25.23	0.00	0.00	25.68	0.00	0.00	20.80	0.00	32.34	22.15	26.29
Whole	3.62	100.0	6.98	2.41	98.41	4.70	2.38	100.0	4.64	2.63	99.55	5.12

Table 3: Frame evaluation results on the triples from the FrameNet 1.7 corpus (Baker et al., 1998). The results are sorted by the descending order of the Frame F₁-score. Best results are boldfaced.

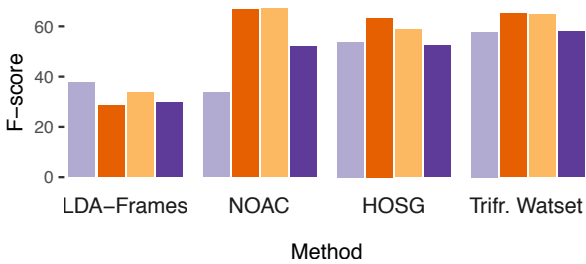


Figure 1: F₁-scores for verbs, subjects, objects, frames corresponding to Table 3.

Frame Induction Experiment. In Table 3 and Figure 1, the results of the experiment are presented. Triframes based on WATSET clustering outperformed the other methods on both Verb F₁ and overall Frame F₁. The *HOSG*-based clustering proved to be the most competitive baseline, yielding decent scores according to all four measures. The *NOAC* approach captured the frame grouping of slot fillers well but failed to establish good verb clusters. Note that *NOAC* and *HOSG* use only the graph of syntactic triples and do not rely on pre-trained word embeddings. This suggests a high complementarity of signals based on distributional similarity and global structure of the triple graph. Finally, the simpler *Triadic* baselines relying on hard clustering algorithms showed low performance, similar to that of *LDA-Frames*, justifying the more elaborate WATSET method.

While triples are intuitively less ambiguous than words, still some frequent and generic triples like (she, make, it) can act as hubs in the graph, making it difficult to split it into semantically plausible clusters. The poor results of the Chinese Whispers hard clustering algorithm illustrate this. Since the hubs are ambiguous, i.e., can belong to multiple clusters, the use of the WATSET fuzzy clustering algorithm that splits the hubs by disambiguating them leads to the best results (see Table 3).

Method	nmPU	niPU	F ₁
LDA-Frames	52.60	45.84	48.98
Triframes WATSET	40.05	62.09	48.69
NOAC	37.19	64.09	47.07
HOSG	38.22	43.76	40.80
Triadic Spectral	35.76	38.96	36.86
Triadic <i>k</i> -Means	52.22	27.43	35.96
Triframes CW	18.05	12.72	14.92
Whole	24.14	79.09	36.99
Singletons	0.00	27.21	0.00

Table 4: Evaluation results on the dataset of polysemous verb classes by Korhonen et al. (2003).

Verb Clustering Experiment. Table 4 presents results on the second dataset for the best models identified on the first dataset. The *LDA-Frames* yielded the best results with our approach performing comparably in terms of the F₁-score. We attribute the low performance of the Triframes method based on CW clustering to its hard partitioning output, whereas the evaluation dataset contains fuzzy clusters. Different rankings also suggest that frame induction cannot simply be treated as a verb clustering and requires a separate task.

5 Conclusion

In this paper, we presented the first application of *triclustering* for unsupervised *frame induction*. We designed a dataset based on the FrameNet and SVO triples to enable fair corpus-independent evaluations of frame induction algorithms. We tested several triclustering methods as the baselines and proposed a new graph-based triclustering algorithm that yields state-of-the-art results. A promising direction for future work is using the induced frames in applications, such as Information Extraction and Question Answering.

Additional illustrations and examples of extracted frames are available in the supplementary materials. The source code and the data are available online under a permissive license.²

²<https://github.com/uhh-lt/triframes>

Acknowledgments

We acknowledge the support of DFG under the “JOIN-T” and “ACQuA” projects and thank three anonymous reviewers for their helpful comments. Furthermore, we thank Dmitry Ignatov and Dmitry Gnatyshak for help in operating the NOAC method using the multimodal clustering toolbox. Besides, we are grateful to Ryan Cotterell and Adam Poliak for a discussion and an implementation of the HOSG method. Finally, we thank Bonaventura Coppolla for discussions and preliminary work on graph-based frame induction.

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 - Volume 1*, ACL '98, pages 86–90, Montreal, QC, Canada. Association for Computational Linguistics.
- Daniel Bauer, Hagen Fürstenau, and Owen Rambow. 2012. [The Dependency-Parsed FrameNet Corpus](#). In *Proceedings of the Eight International Conference on Language Resources and Evaluation*, LREC 2012, pages 3861–3867, Istanbul, Turkey. European Language Resources Association (ELRA).
- Roni Ben Aharon, Idan Szpektor, and Ido Dagan. 2010. [Generating Entailment Rules from FrameNet](#). In *Proceedings of the ACL 2010 Conference Short Papers*, pages 241–246, Uppsala, Sweden. Association for Computational Linguistics.
- Chris Biemann. 2006. [Chinese Whispers: An Efficient Graph Clustering Algorithm and Its Application to Natural Language Processing Problems](#). In *Proceedings of the First Workshop on Graph Based Methods for Natural Language Processing*, TextGraphs-1, pages 73–80, New York, NY, USA. Association for Computational Linguistics.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent Dirichlet Allocation](#). *Journal of Machine Learning Research*, 3:993–1022.
- Hans C. Boas. 2009. *Multilingual FrameNets in Computational Lexicography: Methods and Applications*. Trends in Linguistics. Studies and Monographs. Mouton de Gruyter.
- Aljoscha Burchardt, Marco Pennacchiotti, Stefan Thater, and Manfred Pinkal. 2009. [Assessing the impact of frame semantics on textual entailment](#). *Natural Language Engineering*, 15(4):527–550.
- Yizong Cheng and George M. Church. 2000. [Biclustering of Expression Data](#). In *Proceedings of the Eighth International Conference on Intelligent Systems for Molecular Biology*, pages 93–103. AAAI Press.
- Jackie C. K. Cheung, Hoifung Poon, and Lucy Vanderwende. 2013. [Probabilistic Frame Induction](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 837–846, Atlanta, GA, USA. Association for Computational Linguistics.
- Ryan Cotterell, Adam Poliak, Benjamin Van Durme, and Jason Eisner. 2017. [Explaining and Generalizing Skip-Gram through Exponential Family Principal Component Analysis](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 175–181, Valencia, Spain. Association for Computational Linguistics.
- Dipanjan Das, Desai Chen, André F. T. Martins, Nathan Schneider, and Noah A. Smith. 2014. [Frame-Semantic Parsing](#). *Computational Linguistics*, 40(1):9–56.
- Dmitry Egurnov, Dmitry Ignatov, and Engelbert M. Nguifo. 2017. [Mining Triclusters of Similar Values in Triadic Real-Valued Contexts](#). In *14th International Conference on Formal Concept Analysis - Supplementary Proceedings*, pages 31–47, Rennes, France.
- Katrin Erk and Sebastian Padó. 2006. [SHALMANESER — A Toolchain For Shallow Semantic Parsing](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation*, LREC 2006, pages 527–532, Genoa, Italy. European Language Resources Association (ELRA).
- Charles J. Fillmore. 1982. [Frame Semantics](#). In *Linguistics in the Morning Calm*, pages 111–137. Han-shin Publishing Co., Seoul, South Korea.
- Daniel Gildea and Martin Jurafsky. 2002. [Automatic Labeling of Semantic Roles](#). *Computational Linguistics*, 28(3):245–288.
- Iryna Gurevych, Judith Eckle-Kohler, Silvana Hartmann, Michael Matuschek, Christian M. Meyer, and Christian Wirth. 2012. [UBY – A Large-Scale Unified Lexical-Semantic Resource Based on LMF](#). In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, EACL '12, pages 580–590, Avignon, France. Association for Computational Linguistics.
- Patrick Hanks and James Pustejovsky. 2005. [A Pattern Dictionary for Natural Language Processing](#). *Revue Française de linguistique appliquée*, 10(2):63–82.
- Silvana Hartmann, Judith Eckle-Kohler, and Iryna Gurevych. 2016. [Generating Training Data for Semantic Role Labeling based on Label Transfer from Linked Lexical Resources](#). *Transactions of the Association for Computational Linguistics*, 4:197–213.

- Dmitry I. Ignatov, Dmitry V. Gnatyshak, Sergei O. Kuznetsov, and Boris G. Mirkin. 2015. **Triadic Formal Concept Analysis and triclustering: searching for optimal patterns**. *Machine Learning*, 101(1-3):271–302.
- Sujay Kumar Jauhar and Eduard Hovy. 2017. **Embedded Semantic Lexicon Induction with Joint Global and Local Optimization**. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 209–219, Vancouver, Canada. Association for Computational Linguistics.
- Daisuke Kawahara, Daniel W. Peterson, and Martha Palmer. 2014. **A Step-wise Usage-based Method for Inducing Polysemy-aware Verb Classes**. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics Volume 1: Long Papers*, ACL 2014, pages 1030–1040, Baltimore, MD, USA. Association for Computational Linguistics.
- Anna Korhonen, Yuval Krymolowski, and Zvika Marx. 2003. **Clustering Polysemic Subcategorization Frame Distributions Semantically**. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics - Volume 1*, ACL '03, pages 64–71, Sapporo, Japan. Association for Computational Linguistics.
- Joel Lang and Mirella Lapata. 2010. **Unsupervised Induction of Semantic Roles**. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 939–947, Los Angeles, CA, USA. Association for Computational Linguistics.
- Egoitz Laparra and German Rigau. 2010. **eXtended WordFrameNet**. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, LREC 2010, pages 1214–1219, Valletta, Malta. European Language Resources Association (ELRA).
- Jiří Materna. 2012. **LDA-Frames: An Unsupervised Approach to Generating Semantic Frames**. In *Computational Linguistics and Intelligent Text Processing, Proceedings, Part I*, CICLing 2012, pages 376–387, New Delhi, India. Springer Berlin Heidelberg.
- Jiří Materna. 2013. **Parameter Estimation for LDA-Frames**. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 482–486, Atlanta, GA, USA. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeffrey Dean. 2013. **Distributed Representations of Words and Phrases and their Compositionality**. In *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc., Harrahs and Harveys, NV, USA.
- Ashutosh Modi, Ivan Titov, and Alexandre Klementiev. 2012. **Unsupervised Induction of Frame-Semantic Representations**. In *Proceedings of the NAACL-HLT Workshop on the Induction of Linguistic Structure*, pages 1–7, Montréal, Canada. Association for Computational Linguistics.
- Srini Narayanan, Collin Baker, Charles Fillmore, and Miriam Petruck. 2003. **FrameNet Meets the Semantic Web: Lexical Semantics for the Web**. In *The Semantic Web - ISWC 2003: Second International Semantic Web Conference, Sanibel Island, FL, USA, October 20-23, 2003. Proceedings*, pages 771–787, Heidelberg, Germany. Springer Berlin Heidelberg.
- Brendan O’Connor. 2013. **Learning Frames from Text with an Unsupervised Latent Variable Model**. arXiv preprint arXiv:1307.7382.
- Sebastian Padó and Mirella Lapata. 2009. **Cross-lingual Annotation Projection of Semantic Roles**. *Journal of Artificial Intelligence Research*, 36(1):307–340.
- Alexander Panchenko, Eugen Ruppert, Stefano Faralli, Simone Paolo Ponzetto, and Chris Biemann. 2018. **Building a Web-Scale Dependency-Parsed Corpus from Common Crawl**. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation*, LREC 2018, pages 1816–1823, Miyazaki, Japan. European Language Resources Association (ELRA).
- Niko Schenk and Christian Chiarcos. 2016. **Unsupervised Learning of Prototypical Fillers for Implicit Semantic Role Labeling**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1473–1479, San Diego, CA, USA. Association for Computational Linguistics.
- 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)*, pages 12–21, Prague, Czech Republic. Association for Computational Linguistics.
- Beth M. Sundheim. 1992. **Overview of the Fourth Message Understanding Evaluation and Conference**. In *Proceedings of the 4th Conference on Message Understanding*, MUC4 '92, pages 3–21, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Ivan Titov and Alexandre Klementiev. 2011. **A Bayesian Model for Unsupervised Semantic Parsing**. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1445–1455, Portland, OR, USA. Association for Computational Linguistics.

- Ivan Titov and Alexandre Klementiev. 2012. **A Bayesian Approach to Unsupervised Semantic Role Induction**. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 12–22, Avignon, France. Association for Computational Linguistics.
- Sara Tonelli and Daniele Pighin. 2009. **New Features for FrameNet - WordNet Mapping**. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, pages 219–227, Boulder, CO, USA. Association for Computational Linguistics.
- Dmitry Ustalov, Alexander Panchenko, and Chris Biemann. 2017. **Watset: Automatic Induction of Synsets from a Graph of Synonyms**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2017, pages 1579–1590, Vancouver, Canada. Association for Computational Linguistics.
- Lizhuang Zhao and Mohammed J. Zaki. 2005. **TRICLUSTER: An Effective Algorithm for Mining Coherent Clusters in 3D Microarray Data**. In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, SIGMOD '05, pages 694–705, New York, NY, USA. ACM.