

Transferring Semantic Roles Using Translation and Syntactic Information

Maryam Aminian¹, Mohammad Sadegh Rasooli², Mona Diab¹

¹Department of Computer Science, The George Washington University, Washington

²Department of Computer Science, Columbia University, New York

¹{aminian, mtdiab}@gwu.edu

²rasooli@cs.columbia.edu

Abstract

Our paper addresses the problem of annotation projection for semantic role labeling for resource-poor languages using supervised annotations from a resource-rich language through parallel data. We propose a transfer method that employs information from source and target syntactic dependencies as well as word alignment density to improve the quality of an iterative bootstrapping method. Our experiments yield a 3.5 absolute labeled F-score improvement over a standard annotation projection method.

1 Introduction

Semantic role labeling (SRL) is the task of automatically labeling predicates and arguments of a sentence with shallow semantic labels characterizing “Who did What to Whom, How, When and Where?” (Palmer et al., 2010). These rich semantic representations are useful in many applications such as question answering (Shen and Lapata, 2007) and information extraction (Christensen et al., 2011), hence gaining a lot of attention in recent years (Zhou and Xu, 2015; Täckström et al., 2015; Roth and Lapata, 2016; Marcheggiani et al., 2017). Since the process of creating annotated resources needs significant manual effort, SRL resources are available for a relative small number of languages such as English (Palmer et al., 2005), German (Erk et al., 2003), Arabic (Zaghouni et al., 2010) and Hindi (Vaidya et al., 2011). However, most languages still lack SRL systems. There have been some efforts to use information from a resource-rich language to develop SRL systems for resource-poor languages. *Transfer* methods address this problem by transferring information from a resource-rich language

(e.g. English) to a resource-poor language.

Annotation projection is a popular transfer method that transfers supervised annotations from a *source* language to a *target* language through parallel data. Unfortunately this technique is not as straightforward as it seems, e.g. *translation shifts* lead to erroneous projections and accordingly affecting the performance of the SRL system trained on these projections. Translation shifts are typically a result of the differences in word order and the semantic divergences between the source and target languages. In addition to translation shifts, there are errors that occur in translations, automatic word alignments as well as automatic semantic roles, hence we observe a cascade of error effect.

In this paper, we introduce a new approach for a *dependency-based SRL* system based on annotation projection without any semantically annotated data for a target language. We primarily focus on improving the quality of annotation projection by using translation cues automatically discovered from word alignments. We show that exclusively relying on partially projected data does not yield good performance. We improve over the baseline by filtering irrelevant projections, iterative bootstrapping with relabeling, and weighting each projection instance differently with data-dependent cost-sensitive training.

In short, contributions of this paper can be summarized as follows; We introduce a weighting algorithm to improve annotation projection based on cues obtained from syntactic and translation information. In other words, instead of utilizing manually-defined rules to filter projections, we define and use a customized cost function to train over noisy projected instances. This newly defined cost function helps the system weight some projections over other instances. We then utilize this algorithm in a bootstrapping framework. Un-

like traditional bootstrapping, ours relabels every training instance (including labeled data) in every self-training round. Our final model on transferring from English to German yields a 3.5 absolute improvement labeled F-score over a standard annotation projection method.

2 Our Approach

We aim to develop a dependency-based SRL system which makes use of training instances projected from a source language (SLang) onto a target language (TLang) through parallel data. Our SRL system is formed as a pipeline of classifiers consisting of a predicate identification and disambiguation module, an argument identification module, and an argument classification module. In particular, we use our re-implementation of the greedy (local) model of Björkelund et al. (2009) except that we use an averaged perceptron algorithm (Freund and Schapire, 1999) as the learning algorithm.

2.1 Baseline Model

As our baseline, we apply automatic word alignment on parallel data and preserve the intersected alignments from the source-to-target and target-to-source directions. As our next step, we define a *projection density criteria* to filter some of the projected sentences. Given a target sentence from TLang with w words where f words have alignments ($f \leq w$), if the source sentence from SLang has p predicates for which p' of them are projected ($p' \leq p$), we define projection density as $(p' \times f)/(p \times w)$ and prune out sentences with a density value less than a certain threshold. The threshold value is empirically determined during tuning experiments performed on the development data. In this criteria, the denominator shows the maximum number of training instances that could be obtained by projection and the nominator shows the actual number of relevant instances that are used in our model. In addition to speeding up the training process, filtering sparse alignments helps remove sentence pairs with a significant translation shifts. Thereafter, a supervised model is trained directly on the projected data.

2.2 Model Improvements

As already mentioned, the quality of projected roles is highly dependent on different factors including translation shifts, errors in automatic word

alignments and the SLang supervised SRL system. In order to address these problems, we apply the following techniques to improve learning from partial and noisy projections, thereby enhancing the performance of our model:

- Bootstrapping to make use of unlabeled data;
- Determining the quality of a particular projected semantic dependency based on two factors: 1) source-target syntactic correspondence; and, 2) projection completeness degree. We utilize the above constraints in the form of a data-dependent cost-sensitive training objective. This way the classifier would be able to learn translation shifts and erroneous instances in the projected data, hence enhancing the overall performance of the system.

Bootstrapping Bootstrapping (or self-training) is a simple but very useful technique that makes use of unlabeled data. A traditional self-training method (McClosky et al., 2006) labels unlabeled data (in our case, fill in missing SRL decisions) and adds that data to the labeled data for further training. We report results for this setting in §3.1 as *fill-in*. Although fill-in method is shown to be very useful in previous work (Akbik et al., 2015), empirically, we find that it is better to *relabel all* training instances (including the already labeled data) instead of only labeling unlabeled raw data. Therefore, the classifier is empowered to discover outliers (resulting from erroneous projections) and change their labels during the training process. Figure 1 illustrates our algorithm. It starts with training on the labeled data and uses the trained model to label the unlabeled data and relabel the already labeled data. This process repeats for a certain number of epochs until the model converges, i.e., reaches its maximum performance.

Data-dependent cost-sensitive training In our baseline approach, we use the standard perceptron training. In other words, whenever the algorithm sees a training instance x_i with its corresponding label y_i , it updates the weight vector θ for iteration t based on the difference between the feature vector $\phi(x_i, y_i)$ of the gold label and the feature vector $\phi(x_i, y_i^*)$ of the predicted label y_i^* (Eq. 1).

$$\theta^t = \theta^{t-1} + \phi(x_i, y_i) - \phi(x_i, y_i^*) \quad (1)$$

In Eq. 1, the algorithm assumes that every data point x_i in the training data $\{x_1, \dots, x_n\}$ has the

Inputs: 1) Projected data $\mathcal{D} = \mathcal{D}^{\mathcal{L}} \cup \mathcal{D}^{\mathcal{U}}$ where $\mathcal{D}^{\mathcal{L}}$ and $\mathcal{D}^{\mathcal{U}}$ indicate labeled and unlabeled instances in the projected data; 2) Number of self-training iterations m .

Algorithm:

Train model θ^0 on $\mathcal{D}^{\mathcal{L}}$

for $i = 1$ to m **do**

$\mathcal{D}^{\mathcal{U}} \leftarrow$ Label data $\mathcal{D}^{\mathcal{U}}$ with model θ^{i-1} .

$\mathcal{D}^{\mathcal{L}} \leftarrow$ Relabel data $\mathcal{D}^{\mathcal{L}}$ with model θ^{i-1} .

Train model θ^i on $\mathcal{D}^{\mathcal{L}} \cup \mathcal{D}^{\mathcal{U}}$

Output: The model parameters θ^m .

Figure 1: The iterative bootstrapping algorithm for training SRL on partially projected data.

same importance and the cost of wrongly predicting the best label for each training instance is uniform. We believe this uniform update is problematic especially for the transfer task in which different projected instances have different qualities. To mitigate this issue, we propose a simple modification, we introduce a cost $\lambda_i \in [0, 1]$ for each training instance x_i . Therefore, Eq. 1 is modified as follows in Eq. 2.

$$\theta^t = \theta^{t-1} + \lambda_i (\phi(x_i, y_i) - \phi(x_i, y_i^*)) \quad (2)$$

In other words, the penalty of making a mistake by the classifier for each training instance depends on the importance of that instance defined by a certain cost. The main challenge is to define an effective cost function, especially in our framework where we don't have supervision. Accordingly, we experiment with the following cost definitions:

- **Projection completeness:** Our observation shows that the density of projection is a very important indicator of projection quality. We view it as a rough indicator of translation shifts: the more alignments from source to target, the less we have a chance of having translation shifts. As an example, consider the sentence pair extracted from English-German Europarl corpus: “*I sit here between a rock and a hard place*” and its German translation “*Zwei Herzen wohnen ach in meiner Brust*” which literally reads as “*Two hearts dwell in my chest*”. The only words that are aligned (based on the output of Giza++) are the English word “*between*” and the German word “*in*”. In fact, the German sentence is an idiomatic translation of the English sentence. Consequently predicate-argument structure of these sentences vary

tremendously; The word “*sit*” is predicate of the English sentence while “*wohnen (dwell)*” is the predicate of the German sentence.

We use the definition of *completeness* from Akbik et al. (2015) to define the sparsity cost (λ^{comp}): this definition deals with the proportion of a verb or direct dependents of verbs in a sentence that are labeled.

- **Source-target syntactic dependency match:** We observe that when the dependency label of a target word is different from its aligned source word, there is a higher chance of a projection mistake. However, given the high frequency of source-target dependency mismatches, it is harmful to prune those projections that have dependency mismatch; instead, we define a different cost if we see a training instance with a dependency mismatch. For an argument x_i that is projected from source argument s_{x_i} , we define the cost λ_i^{dep} according to the dependency of the source and target words $dep(x_i)$ and $dep(s_{x_i})$ as Eq. 3.

$$\lambda_i^{dep} = \begin{cases} 1 & \text{if } dep(x_i) = dep(s_{x_i}) \\ 0.5 & \text{otherwise} \end{cases} \quad (3)$$

As an example, consider Fig. 2 that demonstrates an English-German sentence pair from EuroParl “*I would urge you to endorse this*” with its German translation that literally reads as “*I ask for your approval*”. As we can see, there is a shift in translation of English clausal complement “*to endorse this*” into German equivalent “*um Zustimmung (your approval)*” which leads the difference in the syntactic structure of source and target sentences. Therefore, neither the predicate label of English verb “*endorse*” nor the argument “*A2*” should not be projected to the German noun “*Zustimmung*”. Dashed edges between sentences show intersected word alignments. Here, projecting semantic role of “*endorse*” (A2) to to the word “*Zustimmung*” through alignment will lead to the wrong semantic role for this word.

- **Completeness + syntactic match:** We employ the average of λ^{dep} and λ^{comp} values as defined above. This way, we simultaneously encode both the completeness and syntactic similarity information.

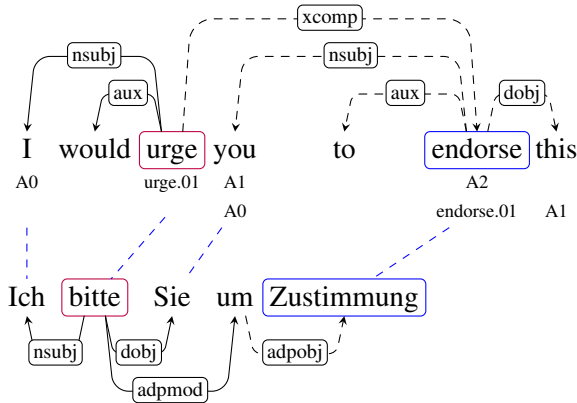


Figure 2: Example of English-German sentences from Europarl with dependency structure. Different dependencies are shown with dashed arcs. Predicate–argument structure of the English sentence is shown below each word.

3 Experiments

Data and Setting We use English as the source and German as the target language. In our setting, we assume to have supervised part-of-speech tagging and dependency parsing models for both the source (SLang) and target (TLang) languages. We use the Universal part-of-speech tagset of Petrov et al. (2011) and the Google Universal Treebank (McDonald et al., 2013). We ignore the projection of the *AM* roles to German since this particular role does not appear in the German dataset.

We use the standard data splits in the CoNLL shared task on SRL (Hajič et al., 2009) for evaluation. We replace the POS and dependency information with the predictions from the Yara parser (Rasooli and Tetreault, 2015) trained on the Google Universal Treebank.¹ We use the parallel Europarl corpus (Koehn, 2005) and Giza++ (Och and Ney, 2003) for extracting word alignments. Since predicate senses are projected from English to German, comparing projected senses with the gold German predicate sense is impossible. To address this, all evaluations are conducted using the Gold predicate sense.

After filtering projections with density criteria of §2.1, 29417 of the sentences are preserved. The number of preserved sentences after filtering sparse alignments is roughly one percent of the

¹Our ideal setting is to transfer to more languages but because of the semantic label inconsistency across CoNLL datasets, we find it impossible to evaluate our model on more languages. Future work should work on defining a reliable conversion scheme to unify the annotations in different datasets.

Model	Cost	Lab. F1
Baseline	×	60.3
Bootstrap–fill-in	×	61.6
Bootstrap–relabel	×	62.4
Bootstrap–relabel	comp.	63.0 _(+1.0)
Bootstrap–relabel	dep.	63.4 _(+1.8)
Bootstrap–relabel	comp.+dep.	63.8 _(+1.3)
Supervised	–	79.5

Table 1: Labeled F-score for different models in SRL transfer from English to German using gold predicates. Cost columns shows the use of cost-sensitive training using projection completeness (“comp.”), source-target dependency match (“dep.”) and both (“comp.+dep.”). The numbers in parenthesis show the absolute improvement over the Bootstrap–fill-in method.

original parallel data (29K sentences out of 2.2M sentences). Density threshold is set to 0.4 determined based on our tuning experiments on development data.

3.1 Results and Discussion

Table 1 shows the results of different models on the German evaluation data. As we can see in the table, bootstrapping outperforms the baseline. Interestingly, relabelling all training instances (Bootstrap–relabel) gives us 0.8 absolute improvement in F-score compared to when we just predict over instances without a projected label (Bootstrap–fill-in). Here, the fill-in approach would label only the German word “um” in Fig. 2 that does not have any projected label from the English side. While the relabeling method will overwrite all projected labels with less noisy predicted labels.

We additionally observe that the combination of the two cost functions improves the quality further. Overall, the best model yields 3.5 absolute improvement F-score over the baseline. As expected, none of the approaches improves over supervised performance. We further analyzed the effects of relabeling approach on identification and classification of non–root semantic dependencies. Figure 3 shows precision, recall and F–score of the two most frequent semantic dependencies (predicate pos + argument label): VERB+A0, VERB+A1 throughout relabeling iterations. As demonstrated in the graph, both precision and recall improve by cost-sensitive re-




Figure 3: Precision, recall and F-score of VERB+A0 and VERB+A1 during relabeling iterations on the German development data. Horizontal axis shows the number of iterations and vertical axis shows values of precision, recall and F-score.

labeling for VERB+A0. In fact, cost-sensitive training helps the system refine irrelevant projections at each iteration and assigns more weight on less noisy projections, hence enhancing precision. Our analysis on VERB+A0 instances shows that source-target dependency match percentage also increases during iterations leading to increase the recall. In other words, weighting projection instances based on dependency match helps classifier label some of the instances which were dismissed during projection, thereby will increase the recall. While similar improvement in precision is observed for VERB+A1, Figure 3 shows that the recall is almost descending by relabeling. Our analysis shows that unlike VERB+A0, percentage of source-target dependency match remains almost steady for VERB+A1. This means that cost-sensitive relabeling for this particular semantic dependency has not been very successful in labeling unlabeled data.

4 Related Work

There have been several studies on transferring SRL systems (Padó and Lapata, 2005, 2009; Mukund et al., 2010; van der Plas et al., 2011, 2014; Kozhevnikov and Titov, 2013; Akbik et al., 2015). Padó and Lapata (2005), as one of the earliest studies on annotation projection for SRL using parallel resources, apply different heuristics and techniques to improve the quality of their model by focusing on having better word and constituent alignments. van der Plas et al. (2011) improve an annotation projection model by jointly training a transfer system for parsing and SRL. They solely focus on fully projected annotations and train only on verbs. In this work, we train on all predicates

as well as exploit partial annotation. Kozhevnikov and Titov (2013) define shared feature representations between the source and target languages in annotation projection. The benefit of using shared representations is complementary to our work encouraging us to use it in future work.

Akbik et al. (2015) introduce an iterative self-training approach using different types of linguistic heuristics and alignment filters to improve the quality of projected roles. Unlike our work that does not use any external resources, Akbik et al. (2015) make use of bilingual dictionaries. Our work also leverages self-training but with a different approach: first of all, ours does not apply any heuristics to filter out projections. Second, it trains and relabels all projected instances, either labeled or unlabeled, at every epoch and does not gradually introduce new unlabeled data. Instead, we find it more useful to let the target language SRL system rule out noisy projections via relabeling.

5 Conclusion

We described a method to improve the performance of annotation projection in the dependency-based SRL task utilizing a data-dependent cost-sensitive training. Unlike previous studies that use manually-defined rules to filter projections, we benefit from information obtained from projection sparsity and syntactic similarity to weigh projections. We utilize a bootstrapping algorithm to train a SRL system over projections. We showed that we can get better results if we relabel the entire train data in each iteration as opposed to only labeling instances without projections.

For the future work, we consider experimenting with newly published Universal Proposition Bank (Wang et al., 2017) that provides a unified labeling scheme for all languages. Given the recent success in SRL systems with neural networks (Marcheggiani et al., 2017; Marcheggiani and Titov, 2017), we plan to use them for further improvement. We expect a similar trend by applying the same ideas in a neural SRL system.

Acknowledgments

This work has been partly funded by DARPA LORELEI Grant and generous support by Leidos Corp. for the 1st and 3rd authors. We would like to acknowledge the useful comments by three anonymous reviewers who helped in making this publication more concise and better presented.

References

- Alan Akbik, Iulia Chiticariu, Marina Danilevsky, Yunyao Li, Shivakumar Vaithyanathan, and Huaiyu Zhu. 2015. [Generating high quality proposition banks for multilingual semantic role labeling](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 397–407. Association for Computational Linguistics.
- Anders Björkelund, Love Hafdel, and Pierre Nugues. 2009. *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL 2009): Shared Task*, chapter Multilingual Semantic Role Labeling. Association for Computational Linguistics.
- Janara Christensen, Mausam, Stephen Soderland, and Oren Etzioni. 2011. [An analysis of open information extraction based on semantic role labeling](#). In *Proceedings of the Sixth International Conference on Knowledge Capture, K-CAP '11*, pages 113–120, New York, NY, USA. ACM.
- Katrin Erk, Andrea Kowalski, Sebastian Padó, and Manfred Pinkal. 2003. [Towards a resource for lexical semantics: A large german corpus with extensive semantic annotation](#). In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 537–544, Sapporo, Japan. Association for Computational Linguistics.
- Yoav Freund and Robert E. Schapire. 1999. [Large margin classification using the perceptron algorithm](#). *Machine Learning*, 37(3):277–296.
- Jan Hajič, Massimiliano Ciaramita, Richard Johanson, Daisuke Kawahara, Antònia Maria Martí, Lluís Màrquez, Adam Meyers, Joakim Nivre, Sebastian Padó, Jan Štěpánek, Pavel Straňák, Mihai Surdeanu, Nianwen Xue, and Yi Zhang. 2009. *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL 2009): Shared Task*, chapter The CoNLL-2009 Shared Task: Syntactic and Semantic Dependencies in Multiple Languages. Association for Computational Linguistics.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86.
- Mikhail Kozhevnikov and Ivan Titov. 2013. [Cross-lingual transfer of semantic role labeling models](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1190–1200. Association for Computational Linguistics.
- Diego Marcheggiani, Anton Frolov, and Ivan Titov. 2017. [A simple and accurate syntax-agnostic neural model for dependency-based semantic role labeling](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 411–420. Association for Computational Linguistics.
- Diego Marcheggiani and Ivan Titov. 2017. [Encoding sentences with graph convolutional networks for semantic role labeling](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1507–1516. Association for Computational Linguistics.
- David McClosky, Eugene Charniak, and Mark Johnson. 2006. [Effective self-training for parsing](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*.
- Ryan McDonald, Joakim Nivre, Yvonne Quirbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu Castelló, and Jungmee Lee. 2013. [Universal dependency annotation for multilingual parsing](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 92–97, Sofia, Bulgaria. Association for Computational Linguistics.
- Smruthi Mukund, Debanjan Ghosh, and Rohini Srihari. 2010. [Using cross-lingual projections to generate semantic role labeled annotated corpus for urdu - a resource poor language](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 797–805. Coling 2010 Organizing Committee.
- Franz Josef Och and Hermann Ney. 2003. [A systematic comparison of various statistical alignment models](#). *Computational Linguistics*, 29(1).
- Sebastian Padó and Mirella Lapata. 2005. [Cross-lingual bootstrapping of semantic lexicons: The case of framenet](#). In *Proceedings of the National Conference on Artificial Intelligence*, volume 20, pages 1087–1092.
- Sebastian Padó and Mirella Lapata. 2009. [Cross-lingual annotation projection for semantic roles](#). *Journal of Artificial Intelligence Research*, 36(1):307–340.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. [The proposition bank: An annotated corpus of semantic roles](#). *Computational Linguistics, Volume 31, Number 1, March 2005*.
- Martha Palmer, Daniel Gildea, and Nianwen Xue. 2010. [Semantic role labeling](#). *Synthesis Lectures on Human Language Technologies*, 3(1):1–103.
- Slav Petrov, Dipanjan Das, and Ryan McDonald. 2011. [A universal part-of-speech tagset](#). *arXiv preprint arXiv:1104.2086*.
- Lonneke van der Plas, Marianna Apidianaki, and Chenhua Chen. 2014. [Global methods for cross-lingual](#)

- semantic role and predicate labelling. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1279–1290. Dublin City University and Association for Computational Linguistics.
- Lonneke van der Plas, Paola Merlo, and James Henderson. 2011. [Scaling up automatic cross-lingual semantic role annotation](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 299–304. Association for Computational Linguistics.
- Mohammad Sadegh Rasooli and Joel Tetreault. 2015. Yara parser: A fast and accurate dependency parser. *arXiv preprint arXiv:1503.06733*.
- Michael Roth and Mirella Lapata. 2016. [Neural semantic role labeling with dependency path embeddings](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1192–1202. 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.
- Oscar Täckström, Kuzman Ganchev, and Dipanjan Das. 2015. [Efficient inference and structured learning for semantic role labeling](#). *Transactions of the Association of Computational Linguistics*, 3:29–41.
- Ashwini Vaidya, Jinho Choi, Martha Palmer, and Bhuvana Narasimhan. 2011. [Analysis of the hindi proposition bank using dependency structure](#). In *Proceedings of the 5th Linguistic Annotation Workshop*, pages 21–29, Portland, Oregon, USA. Association for Computational Linguistics.
- Chenguang Wang, Alan Akbik, laura chiticariu, Yunyao Li, Fei Xia, and Anbang Xu. 2017. [Crowd-in-the-loop: A hybrid approach for annotating semantic roles](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1923, Copenhagen, Denmark. Association for Computational Linguistics.
- Wajdi Zaghouani, Mona Diab, Aous Mansouri, Sameer Pradhan, and Martha Palmer. 2010. [The revised arabic propbank](#). In *Proceedings of the Fourth Linguistic Annotation Workshop*, pages 222–226, Uppsala, Sweden. Association for Computational Linguistics.
- Jie Zhou and Wei Xu. 2015. [End-to-end learning of semantic role labeling using recurrent neural networks](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on*
- Natural Language Processing (Volume 1: Long Papers)*, pages 1127–1137. Association for Computational Linguistics.