

CLAR: A Cross-Lingual Argument Regularizer for Semantic Role Labeling

Ishan Jindal^a, Yunyao Li^a, Siddhartha Brahma^{b*}, and Huaiyu Zhu^a

^aIBM Research, Almaden Research Center, CA 95120

^bGoogle Research, Berlin, Germany

ishan.jindal@ibm.com, {yunyaoli, huaiyu}@us.ibm.com,
sidbrahma@gmail.com

Abstract

Semantic role labeling (SRL) identifies predicate-argument structure(s) in a given sentence. Although different languages have different argument annotations, polyglot training, the idea of training one model on multiple languages, has previously been shown to outperform monolingual baselines, especially for low resource languages. In fact, even a simple combination of data has been shown to be effective with polyglot training by representing the distant vocabularies in a shared representation space. Meanwhile, despite the dissimilarity in argument annotations between languages, certain argument labels do share common semantic meaning across languages (e.g. *adjuncts* have more or less similar semantic meaning across languages). To leverage such similarity in annotation space across languages, we propose a method called *Cross-Lingual Argument Regularizer (CLAR)*. CLAR identifies such linguistic annotation similarity across languages and exploits this information to map the target language arguments using a transformation of the space on which source language arguments lie. By doing so, our experimental results show that CLAR consistently improves SRL performance on multiple languages over monolingual and polyglot baselines for low resource languages.

1 Introduction

Semantic Role Labeling (SRL) is the task of labeling each predicate and its corresponding arguments in a given sentence. SRL provides a more stable meaning representation across syntactically different sentences and has been seen to help a wide range of NLP applications such as question answering (Maqsd et al., 2014; Yih et al., 2016) and machine translation (Shi et al., 2016).

^b Work done while at IBM Research

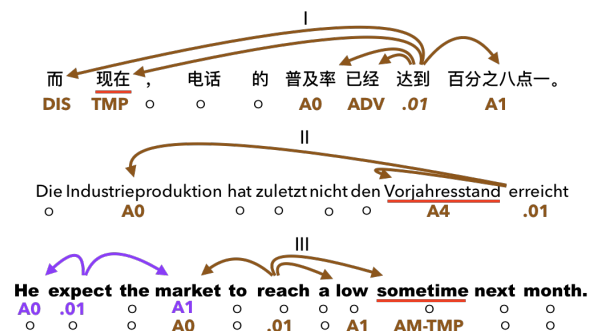


Figure 1: Example of predicate-argument structure from the CoNLL 2009 training data for I) Chinese, II) German, and III) English.

Recent end-to-end deep neural networks for SRL, though performing well for languages with large training data (Marcheggiani et al., 2017; Tan et al., 2018; He et al., 2018), are much less effective for low resources languages due to very limited annotated data for these languages. Methods such as *polyglot training* (Mulcaire et al., 2018) seek to make these models perform better on low resource languages by combining supervision from multiple languages. The key idea in polyglot training is to combine the training data from multiple languages by using multilingual word embeddings from a shared space and a common encoder model (e.g. an LSTM). The argument sets for the languages are kept separate by using different classification layers. The arguments sets are kept separate because the semantic label spaces are usually language-specific (Mulcaire et al., 2018).

However, despite the dissimilarity in argument annotations between languages, certain argument labels do share common semantic meaning across languages. Fig. 1 shows three different sentences from Chinese, German, and English, respectively, with defined predicate-argument structures. Although the predicates are essentially the same, their arguments are labeled differently across languages in the training data. For instance, all sentences

contain words representing the same underlying semantic meaning that is temporal but with different argument labels (TMP in Chinese, A4 in German, AM-TMP in English).

We hypothesize that we can improve the SRL performance of low resource languages during cross-lingual transfer by identifying such arguments with similar semantic meaning across languages and representing them close to each other in the feature space. This requires: (1) Detecting the correspondence between the labels in different languages; and (2) Representing arguments with similar semantic meaning in the feature space for better SRL performance.

We propose a method called *Cross-Lingual Argument Regularizer* (CLAR) with a two-step process:

Step 1: Pair Matching: Detecting a number of label pairs between the source and target languages during polyglot training. We call these arguments *common arguments*. Given the multilingual embedding already used in polyglot training, CLAR does not require additional cross-lingual alignments on parallel data.

Step 2: Regularization: Given the common arguments identified, find a transformation to bring the paired arguments close together. This transformation is learned and used in the polyglot training process so that the knowledge on the labels in the source language can be better transferred to knowledge in the corresponding labels in the target language.

We evaluate CLAR on the SRL portion of the CoNLL 2009 dataset (Hajjivc et al., 2009)¹ and compare its performance against baseline and polyglot training methods. The main contributions of this work are:

- We propose CLAR, a simple yet effective method for better cross-lingual transfer by detecting similar semantic role arguments between languages without requiring cross-lingual alignments or parallel data, and by learning a transformation for paired labels via regularization during SRL model training.
- We conduct comprehensive empirical studies and demonstrate the effectiveness of CLAR over both monolingual and polyglot baselines.

¹We do not evaluate CLAR on Japanese data due to licensing issues.

- We perform the ablation study and detailed analysis to understand why CLAR leads to better cross-lingual transfer and how its performance differs with different levels of correspondence among arguments.

The rest of the paper is organized as follows: Sec. 2 describes the base model. Sec. 3 describes CLAR. Sec. 4 demonstrate its efficacy with extensive empirical evaluation. Sec. 5 reviews the existing literature. Sec. 6 makes concluding remarks.

2 Base Model

The SRL task consists of four subtasks: 1) *predicate identification* (e.g., *reach*); 2) *sense disambiguation* of the identified predicate (e.g., *reach.01*); 3) *argument identification* for each predicate (e.g., *market*) and 4) *role classification* of the identified arguments (e.g., A0). Following Li et al. (2018) and Mulcaire et al. (2018), we focus on argument labeling and predicate sense disambiguation, both sequence tagging problems.

Model Architecture As shown in Fig. 2, our model architecture consists of four main modules: (1) *sentence encoder* takes the raw tokens sequentially and outputs a fixed sentence representation; (2) *role labeler* takes the sentence encoder output and identify and predicts roles of the tokens; (3) *predicate sense disambiguator* takes the sentence encoder output and predict the sense for each predicate; and (4) *CLAR regularizer* first detects the common arguments and then learns a manifold on which the arguments of the target languages lie. We now describe each of the modules in more details.

2.1 Sentence Encoder

Word Representation Knowing the predicate position has previously been shown to improve the argument labeling task (Li et al., 2018) and since the predicate position is marked in the CoNLL 2009 dataset, we use this information and obtain the predicate-specific word representations for each word in the sentence. In addition to predicate-specific flag w_i^f , we represent each word w_i in the sentence as a concatenation of several word features including randomly initialized word embeddings w_i^r , pre-trained word embeddings w_i^p , randomly initialized lemma embeddings w_i^l and randomly initialized POS tags embeddings w_i^s . Finally, each word is represented as $w_i = [w_i^r, w_i^p, w_i^l, w_i^s, w_i^f]$. Since we combine

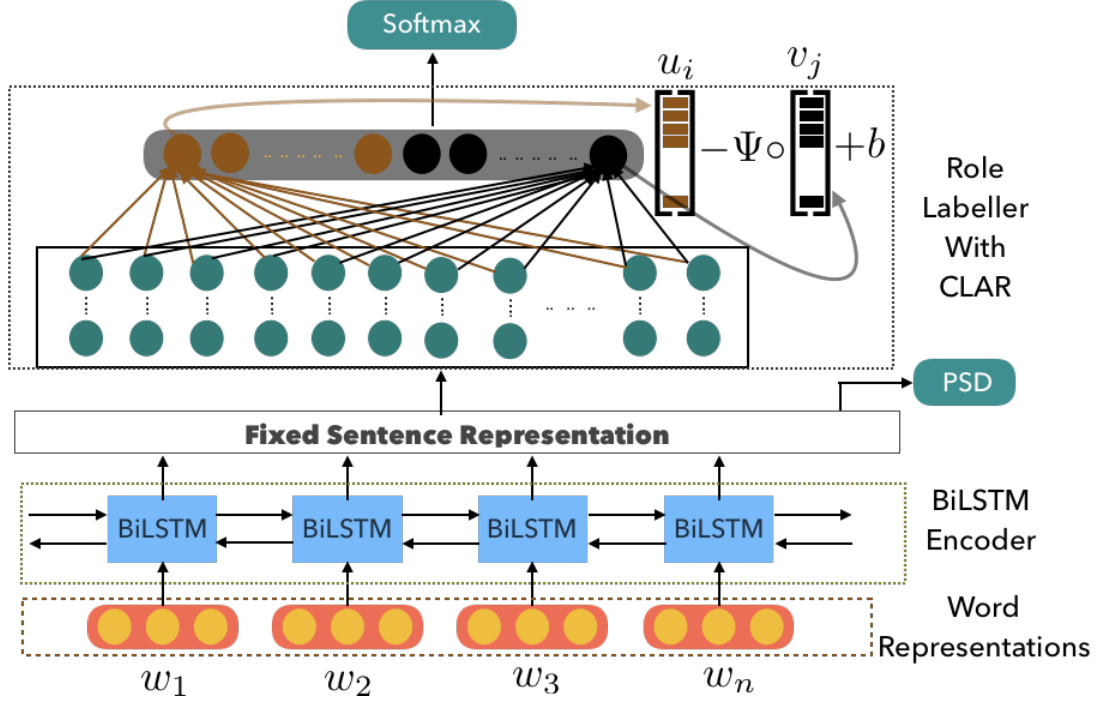


Figure 2: A multitask framework for predicate sense disambiguation and argument classification with CLAR argument regularization

the resources from a pair of languages similar to polyglot training (Mulcaire et al., 2018), we use the language-specific pre-trained word embeddings for w_i^p and train the SRL model on the source and target language simultaneously.

BiLSTM Encoder To model the sequential input we use Bi-directional Long Short Term Memory neural networks (Hochreiter and Schmidhuber, 1997), which take in concatenated word representation for each word in the j -th sentence $x_j = (w_{j1}, w_{j2}, \dots, w_{jn})$ and process them sequentially from both directions to obtain the contextual representations.

2.2 Semantic Role Labeler

Our role labeler consists of Multi-Layer Perceptron (MLP) layers with highway connections (Srivastava et al., 2015). It takes the contextualized word representations from the sentence encoder as an input and outputs a probability distribution over the set of argument labels for each word in the sentence. Given a sentence, we maximize the likelihood of labels for each word by minimizing

$$\mathcal{L}_{\text{Base}} = -\frac{1}{N} \sum_{i=1}^N p(y' = y_i | \mathbf{w}_i; \theta), \quad (1)$$

where y_i is the argument label, w_i represents the input token, θ represents the model parameters, and N denotes the total number of samples.

3 The CLAR Algorithm

The underlying motivation for polyglot training Mulcaire et al. (2018) is that arguments from different languages often help enhance each other. It is reasonable to assume that if corresponding arguments from source and target languages are located closer in the feature space, their mutual enhancements can be strengthened. The possibility for doing so is based on the following observation.

In neural network models that generate labels, the last layer is usually a softmax layer of the form

$$\mathbf{y}_i = \frac{\exp(\mathbf{H}\mathbf{a}_i)}{\sum \exp(\mathbf{H}\mathbf{a}_i)} \quad (2)$$

where $\mathbf{y}_i \in \mathbb{R}^k$, its k components corresponding to the k output argument labels. Given $\mathbf{a}_i \in \mathbb{R}^m$ as a representation of the input token i calculated by previous layers, the rows \mathbf{h}_k of the weights \mathbf{H} are responsible for distinguishing the different argument labels k from each other. During the simple polyglot training, the k argument labels consist of k_s for the source language and k_t for the target language. Splitting these \mathbf{h}_i s into two sets, \mathbf{u}_i for the source language and \mathbf{v}_i for the target language, we

observe that for arguments labels, the Euclidean distance between \mathbf{u}_i and \mathbf{v}_j are often small if the i and j are corresponding argument labels. These can be brought even closer together by an affine transform (linear transform and translation).

We therefore propose the following approach (CLAR) consisting of two steps:

Step 1: Pair Matching: Detect the best pairing of the arguments between a pair of languages.

Step 2: Regularization: Find a transformation that brings the feature vectors corresponding to the paired argument labels close to each other.

These two steps are described in detail below.

Pair Matching: The goal of this step is to identify matching label pairs in the two languages. We start with the simple polyglot training (Mulcaire et al., 2018) for the first few epochs without CLAR and collect the last layer weights for all the target and source language arguments.

Given the k_s vectors \mathbf{u}_i and k_t vectors \mathbf{v}_j , solve this constraint optimization problem

$$\text{minimize}_T \sum_i^{k_s} \sum_j^{k_t} T_{ij} \|\mathbf{u}_i - \mathbf{v}_j\|_2^2$$

subject to

$$\begin{aligned} \sum_i T_{ij} &\leq 1, \quad j = 1, \dots, k_t \\ \sum_j T_{ij} &\leq 1, \quad i = 1, \dots, k_s \\ \sum_{i,j} T_{ij} &\geq \min(k_t, k_s), \quad j = 1, \dots, k_t; \\ &\quad i = 1, \dots, k_s \\ T_{ij} &\in \{0, 1\}, \quad \forall i, j. \end{aligned} \quad (3)$$

Intuitively, this requires finding pairings between i and j such that the total squared distance between paired vectors $(\mathbf{u}_i, \mathbf{v}_j)$ is minimized, subject to the constraint that each source argument matches at most one target argument and vice versa, and that at least $K = \min(k_t, k_s)$ argument pairs are identified. This identifies K semantically similar argument pairs in source and target languages, represented in the binary matrix T , where $T_{ij} = 1$ means that argument i in source language and argument j in target language are paired together. Later on (Sec. 4.5) we will show that in certain situations

it makes sense to relax the ‘‘at most one’’ constraint and allow many-to-one or one-to-many matching.

This is an *Integer Linear Programming* problem, for which many excellent solvers exist. We use GLPK solver from CVXOPT².

We observe that the frequency distribution of the argument labels is quite skewed in the training dataset: a few labels (e.g., A0, A1) have much larger number of training examples than other labels. Experiments show that low-frequency labels cause noisy pair matching that degrades the output quality. Therefore, we consider only labels that have more than 1% of the total number occurrences in the respective language training data. Typically, 40 – 50% of the total labels in each language match this criterion. The k_s and k_t in the general algorithm are replaced by \hat{k}_s and \hat{k}_t for the number of arguments satisfying this criterion in the source and the target language, respectively.

Regularization: The goal of this step is to learn an affine transform to bring the target vectors closest to the corresponding source vectors. This step is performed iteratively during the overall training process.

Given the K pairs $(\mathbf{u}_i, \mathbf{v}_i)$ detected in the previous step, the objective of the overall optimization objective function is amended as follows

$$\mathcal{L}_{\text{CLAR}} = \mathcal{L}_{\text{Base}} + \lambda \sum_{i=1}^K \|\mathbf{u}_i - \Psi \mathbf{v}_i + b\|_2^2, \quad (4)$$

where $\Psi \mathbf{v}_i + b$ is the affine transform to bring \mathbf{v}_i close to \mathbf{u}_i , and λ controls the strength of the amendments by the paired labels. The transformation Ψ, b is learned iteratively by minimizing (4) during SRL model training.

4 Experiments

4.1 Dataset

We evaluate CLAR on CoNLL 2009 Shared Task dataset (Hajjivc et al., 2009) with English (EN) as the source language and five different languages, namely German (DE), Spanish (ES), Chinese (ZH), Czech (CS) and Catalan (CA), as target languages. The dataset includes no correspondence defined between the argument labels across languages. For instance, the argument label set in English contains

²<http://cvxopt.org/index.html>

		EN	CA	CS	DE	ES	ZH	avg
Zhao et al. (2009)		86.20	80.30	85.20	76.00	83.00	77.70	-
Roth and Lapata (2016)		87.70	-	-	80.10	80.20	79.40	-
Marcheggiani et al. (2017)		87.70	-	86.00	-	80.30	81.20	-
Cai et al. (2018)		89.60	-	-	-	-	84.30	-
Kasai et al. (2019)		90.20	-	-	-	83.00	-	-
Mulcaire et al. (2018)	Monolingual	86.54	77.31	84.87	66.71	75.98	81.26	77.22
	Polyglot	-	79.08	84.82	69.97	76.45	81.50	78.36
Base SRL + MUSE Embedding	Monolingual	86.47	78.92	89.78	68.73	78.09	81.34	79.37
	Polyglot	-	79.05	89.70	71.16	78.22	81.42	79.78
	CLAR	-	79.26	89.77	72.50	78.83	81.85	80.44
Base SRL + BERT Embedding	Monolingual	88.14	80.50	90.78	74.39	80.98	84.71	82.27
	Polyglot	-	81.87	90.67	74.45	81.88	84.79	82.73
	CLAR	-	82.18	90.81	75.33	82.13	85.04	83.09

Table 1: Semantic F1 scores (including sense) on CoNLL 2009 Shared task languages. The best reported performance on English and Spanish from (Kasai et al., 2019), Chinese from (Cai et al., 2018), German from (Roth and Lapata, 2016), Catalan from (Zhao et al., 2009) and Czech from (Marcheggiani et al., 2017). Underline shows the best performance among all methods.

(A0, A1, ...) while the argument label set in Spanish contains (Arg0-agt, Arg0-pat, ...). Further details on dataset is available in Appendix A.

EN	+CA	+CS	+DE	+ES	+ZH
86.47	87.12	86.70	87.09	86.68	86.90

Table 2: CLAR Semantic F1 scores (including sense) on EN test set for each language pair.

4.2 Setup

We compare CLAR with several *Monolingual* and *Polyglot* methods. For monolingual baselines, we train separate SRL models for each language. For Polyglot and CLAR methods, we train the SRL model on a pair of language. We use pre-trained multilingual embeddings to allow the multilingual sharing between languages. We use Multilingual Unsupervised and Supervised Embeddings (MUSE) (Conneau et al., 2017) for all the languages except Chinese³, where we use fastText aligned word embeddings (Joulin et al., 2018). We also use the pre-trained BERT multilingual cased embeddings (Devlin et al., 2019) in place of MUSE pre-trained embeddings to observe the effect of better multilingual embeddings. Details on model hyperparameters are presented in Appendix B. For all the experiments we fix the base model architecture. For the Polyglot training, we implement the *simple polyglot sharing* setup proposed by Mulcaire et al. (2018). Along with the reported results

³MUSE does not provide aligned vectors for the Chinese language

in Mulcaire et al. (2018) we also report the polyglot results with our model architecture keeping the same word representation to avoid any ambiguity between Polyglot and CLAR comparison.

4.3 Results

Comparison Against Polyglot and Monolingual Training: Table 1 summarizes the performance of CLAR and all baselines for SRL. As can be seen, for both MUSE and BERT embeddings, CLAR results in better SRL models than those obtained via monolingual and polyglot training for all target languages. The improvement is particularly noticeable for the languages with much fewer (< 1/3) training samples than those of EN (e.g. DE and ES). This result confirms that CLAR can effectively transfers knowledge from a high resource language (EN) to other languages with less resource.

Note that for CS, neither CLAR nor polyglot training shows performance gain over the baseline. CLAR outperforms the polyglot baseline but remains on par with the monolingual baseline. We present further investigation on this in Section 4.5.

Comparison Against SoTA: With the powerful BERT multilingual embeddings, CLAR surpasses the best previously reported results on 3 out of 6 languages (Table 1). In fact, its average performance surpasses that any previous-reported single system. The strong performance of CLAR confirms its great promise for cross-lingual transfer.

Cross-Lingual Transfer from Target to Source Language: Interestingly, cross-lingual transfer

Target	Source	Monolingual			Polyglot			CLAR		
		P	R	F1	P	R	F1	P	R	F1
CA	+EN	78.47	75.44	76.92	77.59	76.68	77.13	78.35	76.54	77.44
CS	+EN	80.36	76.00	78.12	80.32	75.69	77.93	79.91	76.50	78.17
DE	+EN	69.64	64.43	66.94	71.66	69.96	70.80	73.10	71.54	72.31
ES	+EN	78.22	75.63	76.90	78.37	75.83	77.07	79.77	76.22	77.95
ZH	+EN	78.27	75.07	76.64	79.04	74.50	76.68	79.36	75.43	77.34

Table 3: CLAR performance (argument classification only) on CoNLL 2009 Shared task languages and comparison with polyglot and monolingual methods.

by CLAR also helps improving the performance of languages with abundant training data. As illustrated in Table 2, transferring knowledge using CLAR from other languages to EN leads to small but consistent improvements for EN.

CLAR Performance on Arguments Alone:

Since CLAR mainly affects role labeling, we conduct further analysis of its performance on argument classification alone (i.e. predicate sense disambiguation is not evaluated). The results are summarized in Table 3 for Base SRL + MUSE embedding. One can observe that for all target languages, CLAR registers small but noticeable improvements (0.24% to 1.51%) for argument classification in comparison to both monolingual and polyglot methods. The consistent improvements confirm the effectiveness of CLAR in enabling better cross-lingual transfer.

4.4 What does CLAR do?

The results of our comparison studies clearly demonstrate that CLAR outperforms both baseline and polyglot training methods. In this subsection we first explain the intuition behind CLAR and then investigate how it regularizes the arguments.

Intuition: During Polyglot training we examine the last layer weights of the base SRL model and hypothesize that there exists a mapping between source and target language argument. To evaluate this hypothesis, we plot the weights of the output layer using SVD by keeping the two directions corresponding to top two largest eigenvalues learned by Polyglot (Row I) training in Fig. 3.

We draw a line between the arguments that are paired by Equation (3). As can be seen, the euclidean distance between some of the paired arguments is similar. For instance, the euclidean distance between the arguments A_1 and $ZH-A_1$ is similar to that between A_2 and $ZH-A_2$ in Fig. 3b. This

pattern emerges from the training data for most of the target languages. Further, we observe that the euclidean distances among the common arguments for the source and target languages are also similar. For example, in Fig. 3b, the euclidean distance between the source (EN) arguments A_1 and A_2 is similar to that between the target language arguments $ZH-A_1$ and $ZH-A_2$. This observation holds true for most of the arguments across the target languages (Fig. 3a - 3c).

The above observations confirm that there exists similar arguments in source and target languages. The arguments in target language lie on a manifold that is similar in structure, with some translation and/or rotation, to the manifold on which the source language argument lies.

Argument Matching and Regularization:

Therefore, we first match the arguments with similar meanings in the target and the source language. We observe that almost all the matched argument pairs have similar meaning: some are syntactically visible (e.g. $ES-argM-adv$ in ES and $AM-ADV$ in EN), whereas others are semantically similar (e.g. $ES-argM-fin$ and $AM-PNC$ having the same meaning *purpose*). After obtaining the matched argument pairs, we regularize the output layer weights of the matched target arguments by forcing them to live on a matched source arguments manifold in (4). A list of matched arguments for various language pairs is provided in Appendix C.

We plot the CLAR learned weight vectors in Fig. 3 (Row II). We can observe the uniformity in lines (in terms of length), which are drawn between paired target to source language arguments. Further, to quantify the length of these lines, we plot the euclidean distance matrix among the matched source language arguments. Among the target language arguments, we compute the correlation coefficient between the euclidean distance for EN-

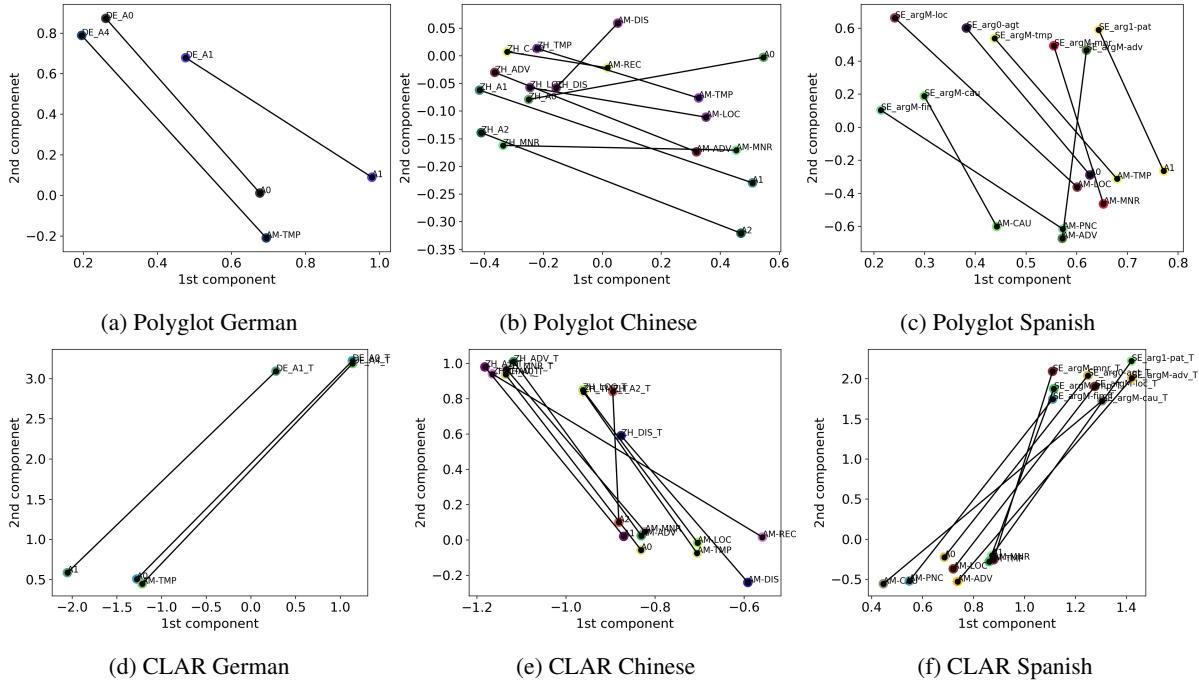


Figure 3: A low dimensional representation of output layer weights of the matched arguments in source and target language as determined by polyglot learned weight vectors in row I and by (3) in row II.

DE, EN-ZH, and EN-ES to be 0.9984, 0.9531, and 0.9352 respectively. The fact that all these coefficients are close to one indicates that CLAR is indeed able to detect a manifold for the target language arguments similar to the one for the source language arguments. Our experimental results (Table 3) demonstrate that allowing the paired target language arguments to lie on the detected manifold improves the argument classification performance.

4.5 Ablation and Analysis

Effect of K We also observe the impact of K on the argument classification performance in Table 4. We find that regularizing all the arguments obtained from (3), while performing better than polyglot, is not a great choice overall. We suspect that considering all the paired arguments adds noise in the system. This is likely because some of the arguments in the target languages are language-specific and might be matched with an argument in the source language which has no close correspondence, for example, the Chinese argument $ZH-C-C-A0$ has no direct corresponding argument in English.

Additionally, in some languages, arguments are labeled at a very granular level, and multiple arguments in these languages may correspond to a single argument in the source language. For example, multiple arguments in Czech frequently map to only one corresponding argument

Target	0	2	$K/2$	K
CA	77.13	77.13	77.44	77.20
CS	77.93	78.12	78.17	77.45
DE	70.80	72.08	72.31	71.20
ES	77.07	77.23	77.95	77.12
ZH	76.68	76.87	77.34	77.02

Table 4: Effect of K on argument classification performance ($K = 0$ represents Polyglot training)

in English.

Languages with Similar Linguistic Annotations:

To further study the effectiveness of CLAR, we analyze the cross-lingual transfer between the languages known to have similar linguistic annotations. We expect to observe better cross-lingual transfer between such language pairs. Specifically, we examine Spanish (ES) and Catalan (CA) from the same AnCora corpus (Taulé et al., 2008). We consider ES as the source language because it has more training samples than CA.

In Table 6 we show the paired arguments detected by CLAR along with the euclidean distance between them. It can be seen that the euclidean distance for all paired arguments are close to 1, confirming that CLAR can effectively match semantically similar arguments across languages.

The experimental results are summarized in Table 5. As expected, CLAR surpasses all prior re-

Training	Method	P	R	F1
CA	Baseline	78.47	75.44	76.92
+ES	Polyglot	79.10	75.90	77.47
	CLAR	78.72	77.91	78.31
+EN	Polyglot	77.59	76.68	77.13
	CLAR	78.35	76.54	77.44

Table 5: Catalan argument classification performance with Spanish as source language

Target	Source	Pair distance
CA-argM-tmp	ES-argM-tmp	0.9302
CA-argM-cau	ES-argM-cau	0.9523
CA-argM-atr	ES-argM-atr	0.9542
CA-arg2-ben	ES-arg2-ben	0.9608
CA-argM-fin	ES-argM-fin	0.9657
CA-arg1-null	ES-arg1-null	0.9672
CA-argM-mnr	ES-argM-mnr	0.9709
CA-argM-loc	ES-argM-loc	0.9790
CA-argM-adv	ES-argM-adv	0.9810
CA-arg0-cau	ES-arg0-cau	0.9839

Table 6: Paired arguments in the source (ES) and the target language (CA)

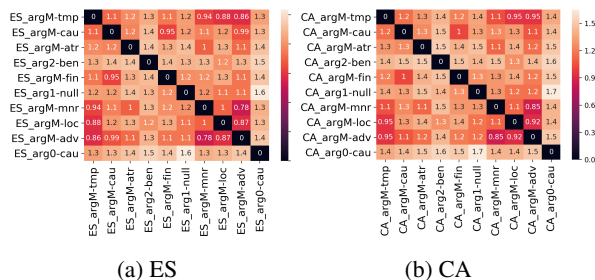


Figure 4: Euclidean distance between last layer weights for ES-CA cross-lingual transfer.

sults on CA. With the semantically similar language ES, the SRL performance on CA is better than the monolingual and polyglot training methods. Further, we observe a 0.87 point absolute gain in F1 score when the cross-lingual transfer occurred from a similar linguistic annotated language (ES) than a less similar language (EN), despite of much smaller training data size ($\leq 30\%$ of EN). This observation strengthens our hypothesis that by representing the semantically similar arguments across languages on similar manifolds improves the SRL performance.

To visualize the space on which the common source and target language argument lies, we plot the heatmap of the euclidean distance between the last layer weights of the learned model in Fig. 4. We plot the separate heatmaps among the paired arguments for each language, the source

language (in Fig. 4a) and the target language (in Fig. 4b). We observe these two heatmaps look very identical in distribution (a very high correlation coefficient 0.9996 and a low Frobenius norm square of the difference 1.793). This means that CLAR transforms the weight vectors of the corresponding target language arguments in such a way that the transformed weight vectors lie on a manifold, which is similar to another manifold on which source language argument weights lie but translated and/or rotated. The aforementioned is evident from Table 6 where we report the distance between these argument pairs.

Why is Czech an Exception? Though Czech (CS) has the most training samples in the CoNLL 2009 dataset, the cross-lingual transfer to and from CS is not very significant, as apparent both from Table 3 and previous work by Mulcaire et al. (2018). We observe that the arguments in CS are labeled at a significant finer granularity than those of other languages. For example, for temporal arguments alone, the argument set in Czech contains 9 different labels at the finest granularity. In contrast, each of the other languages has only one single label for temporal arguments. Since CLAR performs one-to-one mapping to and from the source language, we suspect that CLAR encounters challenges in choosing one among many fine grained arguments to map to a coarse argument in English. While it is possible to extend CLAR with many-to-one mapping, based on our preliminary study (Appendix D), it may introduce additional noise. We plan to explore this direction in the future.

5 Related Work

Models for SRL largely fall into two categories: *syntax-agnostic* and *syntax-aware*. For a long time, syntax was considered a prerequisite for better SRL performance (Punyakanok et al., 2008; Gildea and Jurafsky, 2002). In the absence of syntactic information, these methods struggle to capture the discriminatory features and thus perform poorly.

Recently, end-to-end deep neural models have been shown to extract useful discriminatory features even without syntactic information (Zhou and Xu, 2015; Marcheggiani et al., 2017; Tan et al., 2018; He et al., 2018) and achieve state-of-the-art performance. However, some works (Roth and Lapata, 2016; He et al., 2017; Strubell et al., 2018) argue that given a high-quality syntax parser, it is

possible to further improve the SRL performance. Along this line, (Marcheggiani and Titov, 2017) proposed a SRL model based on graph convolutional networks which incorporates syntactic information from a parser (Kiperwasser and Goldberg, 2016). Further, (Li et al., 2018) proposes a more general framework to integrate syntax into SRL tasks. All these methods have been shown to perform well on rich resource languages.

Several recent attempts have been made to transfer knowledge from rich source languages to low resource languages for SRL tasks (Mulcaire et al., 2018, 2019) such that the knowledge transfer helps the model to learn better feature representations for low resource languages. To some extent, in other NLP tasks such as named identity recognition (Xie et al., 2018), and syntactic dependency parsing (Ammar et al., 2016) this knowledge transfer seems to be helping low resource languages. Our experimental results further strengthen this claim and confirm that languages share knowledge at the semantic level as well.

An alternative line of work transfers cross-lingual knowledge to generate semantic labels for low resource languages by exploiting the monolingual SRL model and Multilingual parallel data (Akbik et al., 2016; Akbik and Li, 2016) with an assumption that the sentences in parallel corpora are semantically equivalent. Similarly, (Pražák and Konopík, 2017) converts the monolingual dependency tree to a universal dependency tree for cross-lingual transfer. Though these methods do not require the knowledge of semantic roles in the target language, they require the availability of massive parallel corpora. On the other hand, CLAR is able to detect the similarity among arguments between the language pairs even in the presence of less data.

6 Conclusion

We introduce CLAR, a Cross-Lingual Argument Regularizer. It explores linguistic annotation similarity across languages and exploits this obtained information during SRL model training to map the target language arguments as the deformation of a space on which source language arguments lie. We confirm the effectiveness of CLAR for SRL on CoNLL 2009 dataset over monolingual and polyglot methods, without prior knowledge of cross-lingual alignments or parallel data. This paper demonstrates the promise of understanding and exploiting linguistic annotation similarity across

languages during polyglot training. We plan to explore other ways of identifying and leveraging linguistic annotation similarity across languages.

Acknowledgments

The authors would like to thank Ranit Aharonov, Mo Yu, and Tyler Baldwin for their comments on an early draft of this work. We also thank our anonymous reviewers for their constructive comments and feedback.

References

- Alan Akbik, Vishwajeet Kumar, and Yunyao Li. 2016. Towards semi-automatic generation of proposition banks for low-resource languages. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 993–998.
- Alan Akbik and Yunyao Li. 2016. Polyglot: Multilingual semantic role labeling with unified labels. In *Proceedings of ACL-2016 System Demonstrations*, pages 1–6.
- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah A Smith. 2016. Many languages, one parser. *Transactions of the Association for Computational Linguistics*, 4:431–444.
- Jiaxun Cai, Shexia He, Zuchao Li, and Hai Zhao. 2018. A full end-to-end semantic role labeler, syntactic-agnostic over syntactic-aware? In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2753–2765.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational linguistics*, 28(3):245–288.
- Jan Hajič, Massimiliano Ciaramita, Richard Johanson, Daisuke Kawahara, Maria Antònia Martí, Lluís Màrquez, Adam Meyers, Joakim Nivre, Sebastian Padó, Jan vStřepánek, et al. 2009. The conll-2009 shared task: Syntactic and semantic dependencies in multiple languages. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–18. Association for Computational Linguistics.

- Jan Hajic, Jarmila Panevová, Zdenka Urevšová, Alevtina Bémová, Veronika Kolárová, and Petr Pajas. 2003. Pdt-vallex: Creating a large-coverage valency lexicon for treebank annotation. In *Proceedings of the second workshop on treebanks and linguistic theories*, volume 9, pages 57–68.
- Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018. Jointly predicting predicates and arguments in neural semantic role labeling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 364–369.
- Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. Deep semantic role labeling: What works and what’s next. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 473–483.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. 2018. Loss in translation: Learning bilingual word mapping with a retrieval criterion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Jungo Kasai, Dan Friedman, Robert Frank, Dragomir Radev, and Owen Rambow. 2019. Syntax-aware neural semantic role labeling with supertags. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 701–709.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. Simple and accurate dependency parsing using bidirectional lstm feature representations. *Transactions of the Association for Computational Linguistics*, 4:313–327.
- Zuchao Li, Shexia He, Jiayun Cai, Zhuosheng Zhang, Hai Zhao, Gongshen Liu, Linlin Li, and Luo Si. 2018. A unified syntax-aware framework for semantic role labeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2401–2411.
- Umar Maqsood, Sebastian Arnold, Michael Hülfenhaus, and Alan Akbik. 2014. Nerdle: Topic-specific question answering using wikia seeds. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: System Demonstrations*, pages 81–85.
- 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.
- 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 1506–1515.
- Phoebe Mulcaire, Jungo Kasai, and Noah A Smith. 2019. Polyglot contextual representations improve crosslingual transfer. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3912–3918.
- Phoebe Mulcaire, Swabha Swayamdipta, and Noah A Smith. 2018. Polyglot semantic role labeling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 667–672.
- Ondrej Prazák and Miloslav Konopík. 2017. Cross-lingual srl based upon universal dependencies. In *RANLP*, pages 592–600.
- Vasin Punyakanok, Dan Roth, and Wen-tau Yih. 2008. The importance of syntactic parsing and inference in semantic role labeling. *Computational Linguistics*, 34(2):257–287.
- 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.
- Chen Shi, Shujie Liu, Shuo Ren, Shi Feng, Mu Li, Ming Zhou, Xu Sun, and Houfeng Wang. 2016. Knowledge-based semantic embedding for machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2245–2254.
- Rupesh K Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Training very deep networks. In *Advances in neural information processing systems*, pages 2377–2385.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. Linguistically-informed self-attention for semantic role labeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5027–5038.
- Zhixing Tan, Mingxuan Wang, Jun Xie, Yidong Chen, and Xiaodong Shi. 2018. Deep semantic role labeling with self-attention. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

Mariona Taulé, M Antònia Martí, and Marta Recasens. 2008. Ancora: Multilevel annotated corpora for catalan and spanish. In *LREC 2008*.

Jiateng Xie, Zhilin Yang, Graham Neubig, Noah A Smith, and Jaime Carbonell. 2018. Neural cross-lingual named entity recognition with minimal resources. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 369–379.

Wen-tau Yih, Matthew Richardson, Chris Meek, Ming-Wei Chang, and Jina Suh. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–206.

Hai Zhao, Wenliang Chen, Jun’ichi Kazama, Kiyotaka Uchimoto, and Kentaro Torisawa. 2009. Multilingual dependency learning: Exploiting rich features for tagging syntactic and semantic dependencies. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning: Shared Task*, pages 61–66. 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.

A Dataset Description

Table 7 describes the training data statistics for each language. In the dataset, for every language, all sentences are marked with predicate-argument structures. Across the languages the argument label set is different.

B Hyperparameters

In our experiments, we randomly initialize the word and lemma embedding of dimension 100 each, the pos embedding of dimension 32, and the flag embedding of dimension 16. We use the same model parameters as mentioned in (Li et al., 2018): a 4-layer BiLSTM with 512 dimensional hidden units and 0.1 dropout rate for the sentence encoder. Our role labeler has 5 MLP highway layers with ReLU activations. We train the model with Adam optimizer (Kingma and Ba, 2014) and minimize the final categorical cross-entropy objective. We train each model for 20 epochs and use early stopping with patience 5 on target language development set. For all the experiments, we repeat with 3 different initialization and report the average F1 score along with precision and recall.

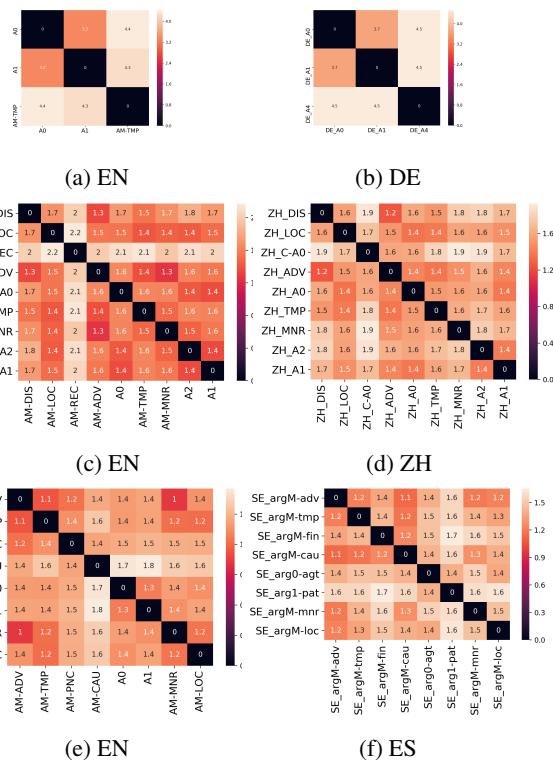


Figure 5: Euclidean distance between last layer weights for matched arguments. Row I: EN-DE, Row II: EN-ZH, Row III: EN-ES, column I: Source language, column II: Target language

C Paired Arguments

We present the list of matched arguments for source-target language pairs in Table 8. We observe that almost all the argument pairs have similar meaning: some are syntactically visible (e.g. ES-argM-adv in ES and AM-ADV in EN), whereas others are semantically similar (e.g. ES-argM-fin and AM-PNC having the same meaning purpose).

We also plot the the euclidean distance matrix among the matched source language arguments and among the target language arguments. In Fig. 5 we show the distance matrix for various language pairs. We compute the correlation coefficient between these matrices and All these coefficients are close to 1 which show that CLAR is indeed able to detect a manifold for the target language arguments similar to the one for the source language arguments.

D CLAR Extension to Many-to-one Mapping

We suspect that CLAR gets a difficulty in choosing one among many fine grained arguments to map to a coarse argument in source language. Here we perform the preliminary investigation on the many to one extension of CLAR. Since CS have

Dataset	Word	POS	Lemma	Arg Labels	Pred Labels	# Predicate	# Arguments	train/valid/test/ood
CA	31,079	15	22,388	39	14	37,431	84,367	13K/1.7K/1.8K/-
CS	75,572	15	35,310	62	116	414,237	365,255	38K/5.2K/4.2K/1.1K
DE	67,548	57	48,217	10	28	17,400	34,276	36K/1.6K/1.7K/707
EN	30,479	49	23,727	53	21	179,014	393,699	39K/1.3K/2.4K/425
ES	37,908	15	24,157	43	13	43,824	99,054	14K/1.6K/1.7K/-
ZH	40,351	38	40,351	37	10	102,813	231,869	22K/1.7K/2.5K/-

Table 7: Train data statistics for each language. Languages are coded with ISO 639-1 codes.

EN-ES		EN-ZH		EN-DE	
ES	EN	ZH	EN	DE	EN
ES-argM-adv	AM-ADV	ZH-DIS	AM-DIS	DE-A0	A0
ES-argM-tmp	AM-TMP	ZH-LOC	AM-LOC	DE-A4	AM-TMP
ES-argM-fin	AM-PNC	ZH-C-A0	AM-REC	DE-A1	A1
ES-argM-cau	AM-CAU	ZH-ADV	AM-ADV		
ES-argL-null	AM-REC	ZH-A0	A0		
ES-arg2-ext	C-AM-DIR	ZH-TMP	AM-TMP		
ES-arg0-agt	A0	ZH-MNR	AM-MNR		
ES-arg1-pat	A1	ZH-A2	A2		
ES-argM-mnr	AM-MNR	ZH-A1	A1		
ES-argM-loc	AM-LOC				

Table 8: Paired arguments in the source and the target language detected by pair matching algorithm during CLAR training.

fine grained labels and is good candidate to analyze many to one mapping, we allow many-to-one argument mapping from Czech to English by relaxing a constraint in the final optimization function and updating only this constraint

$$\sum_j T_{ij} \leq M, i = 1, \dots, \hat{k}_s, \quad (5)$$

while keeping all the other constraints intact. This modification allows at most M arguments in CS to pair with only one argument in EN. Now, following the training procedure, we observe that CLAR is able to efficiently capture many-to-one mappings with minimum noise. In Table 9, we present the argument pairs matched by CLAR. Interestingly, CLAR detects most of the argument pairs correctly, for example, {TWHEN, THL, THO} in CS are mapped to AM-TMP in EN, as expected. However, there are a few pairs that are wrongly mapped, for instance, DIR3 in CS is mapped to A2 in EN. We find that the detection of these noisy pairs is difficult to avoid as the *Prague Dependency Treebank 2.0*. (Hajic et al., 2003) (source of CS dataset) itself points the borderline cases associated with each argument label in CS. For example, ACOMP in CS has borderline

CS	EN	CS	EN
PAT	A1	MAT	A3
ACT	A0	BEN	A3
APP	A2	ACMP	AM-ADV
ADDR	A2	CAUS	AM-ADV
DIR3	A2	COND	AM-ADV
TWHEN	AM-TMP	COMPL	AM-DIS
THL	AM-TMP	CPHR	C-A1
THO	AM-TMP	EFF	AM-PNC
MANN	AM-MNR	AIM	AM-PNC
REG	AM-MNR	EXT	AM-EXT
MEANS	AM-MNR	DPHR	AM-DIR
LOC	AM-LOC	CRIT	R-AM-TMP
RSTR	AM-LOC	TTILL	R-AM-TMP
ID	AM-LOC	TSIN	R-AM-TMP
COMPL2	AM-LOC		
ORIG	AM-LOC		

Table 9: Paired arguments in the source (EN) and the target language (CS)

CLAR Mapping	P	R	F1
one-one	79.91	76.50	78.17
many-one	79.72	76.05	77.84
many-one (combined)	82.57	75.40	78.82

Table 10: Czech argument classification performance with many to one argument mapping.

cases with both COND and CAUS, therefore, they are mapped together to a single argument in EN.

Although CLAR with many-to-one mapping is able to match multiple target language argument labels to a single source language argument label, it actually leads to performance drop as compared to one-to-one mapping (Table 10). This drop in performance is likely because while learning many-to-one mappings, CLAR loses its discriminatory power among those multiple arguments which are mapped to a single label. To validate this phenomenon, at test time, we combine all the argument labels mapped to a single label both for the target and the prediction set; that is, we combine {TWHEN, THL, THO} and propose a new label (say TWHEN) and observe 1ppt \uparrow in F_1 on these combined labels. However, how to effectively leverage CLAR with many-to-one mapping for SRL model training remains an open question and requires further exploration in the future.