

Cross-Lingual Semantic Role Labeling with High-Quality Translated Training Corpus

Hao Fei¹ and Meishan Zhang^{2*} and Donghong Ji¹

1. Department of Key Laboratory of Aerospace Information Security and Trusted Computing, Ministry of Education, School of Cyber Science and Engineering, Wuhan University, China

2. School of New Media and Communication, Tianjin University, China

{hao.fe, dhji}@whu.edu.cn and mason.zms@gmail.com

Abstract

Many efforts of research are devoted to semantic role labeling (SRL) which is crucial for natural language understanding. Supervised approaches have achieved impressive performances when large-scale corpora are available for resource-rich languages such as English. While for the low-resource languages with no annotated SRL dataset, it is still challenging to obtain competitive performances. Cross-lingual SRL is one promising way to address the problem, which has achieved great advances with the help of model transferring and annotation projection. In this paper, we propose a novel alternative based on corpus translation, constructing high-quality training datasets for the target languages from the source gold-standard SRL annotations. Experimental results on Universal Proposition Bank show that the translation-based method is highly effective, and the automatic pseudo datasets can improve the target-language SRL performances significantly.

1 Introduction

Semantic role labeling (SRL), which aims to capture the high-level meaning of a sentence, such as *who did what to whom*, is an underlying task for facilitating a broad range of natural language processing (NLP) tasks (Shen and Lapata, 2007; Liu and Gildea, 2010; Genest and Lapalme, 2011; Gao and Vogel, 2011; Wang et al., 2015; Khan et al., 2015). Currently, the majority of research work on SRL is dedicated to the English language, due to the availability of large quantity of labeled data. With this regard, cross-lingual SRL, especially the one transferring the advantage of the source language with affluent amount of resources (e.g., English) to the target language where the labeled data is scarce or even not available, is of great importance

*Corresponding author.




Figure 1: Illustration of cross-lingual SRL methods. Our method is in the dotted blue box.

(Kozhevnikov and Titov, 2013; He et al., 2019; Aminian et al., 2019).

Previous work on cross-lingual SRL can generally be divided into two categories: model transferring and annotation projection. The former builds cross-lingual models on language-independent features such as cross-lingual word representations and universal POS tags which can be transferred into target languages directly (McDonald et al., 2013; Swayamdipta et al., 2016; Daza and Frank, 2019). The latter bases on a large-scale parallel corpus between the source and target languages where the source-side sentences are annotated with SRL tags automatically by a source SRL labeler, and then the source annotations are projected onto the target-side sentences in accordance of word alignments (Yarowsky et al., 2001; Hwa et al., 2005; van der Plas et al., 2011; Kozhevnikov and Titov, 2013; Pado and Lapata, 2014; Akbik et al., 2015). In addition, the annotation projection can be combined with model transferring naturally.

Particularly, the projected SRL tags in annotation projection could contain much noise because of the source-side automatic annotations. A straightforward solution is the translation-based approach, which has been demonstrated effective for cross-lingual dependency parsing (Täckström et al., 2012; Rasooli and Collins, 2015; Guo et al., 2016; Zhang et al., 2019). The key idea is to translate the gold-standard source training data into target

language side by translation directly, avoiding the problem of the low-quality source annotations. Fortunately, due to recent great advances of neural machine translation (NMT) (Bahdanau et al., 2015; Wu et al., 2016), this approach could have great potentials for cross-lingual transferring.

To this end, in this paper, we study the translation-based method for cross-lingual SRL. Figure 1 illustrates the differences between previous approaches. Sentences of the source language training corpus are translated into the target language, and then the source SRL annotations are projected into the target side, resulting in a set of high-quality target language SRL corpus, which is used to train the target SRL model. Further, we merge the gold-standard source corpus and the translated target together, which can be regarded as a combination of the translation-based method and the model transferring. Our baseline is a simple BiLSTM CRF model by using multilingual contextualized word representations (Peters et al., 2018; Devlin et al., 2019). For a better exploration of the blended corpus, we adopt a parameter generation network (PGN) to enhance the BiLSTM module, which can capture the language differences effectively (Platanios et al., 2018; Jia et al., 2019).

We conduct experiments based on Universal Proposition Bank corpus (v1.0) (Akbik et al., 2015; Akbik and Li, 2016) over seven languages. First, we verify the effectiveness of our method on the single-source SRL transferring, where the English language is adopted as the source language and the remaining are used as the target languages. Results show that the translation-based method is highly effective for cross-lingual SRL, and the performances are further improved when PGN-BiLSTM is used. Further, we conduct experiments on the multi-source SRL transferring, where for each target language all the remaining six languages are used as the source languages. The same tendencies as the single-source setting can be observed. We conduct detailed analysis work for both settings to understand our proposed method comprehensively.

In summary, we make the following two main contributions in this work:

- We present the first work of the translation-based approach for unsupervised cross-lingual SRL. We build a high-quality of pseudo training corpus for a target language, and then verify the effectiveness of the corpus under a range of settings.

- We take advantage of the multilingual contextualized word representations, and strengthen the multilingual model training with PGN-BiLSTM model.

All codes and datasets are released publicly available for the research purpose¹.

2 Related Work

There exists extensive work for cross-lingual transfer learning (van der Plas et al., 2011; Kozhevnikov and Titov, 2013; Pado and Lapata, 2014; Rasooli and Collins, 2015; Tiedemann and Agic, 2016; Zhao et al., 2018; Chen et al., 2018, 2019; Aminian et al., 2019). Model transferring and annotation projection are two mainstream categories for the goal. The first category aims to build a model based on the source language corpus, and then adapt it to the target languages (Yarowsky et al., 2001; Hwa et al., 2005; Tiedemann, 2015). The second category attempts to produce a set of automatic training instances for the target language by a source language model and a number of parallel sentences, and then train a target model on the dataset (Björkelund et al., 2009; McDonald et al., 2013; Lei et al., 2015; Swayamdipta et al., 2016; Mulcaire et al., 2018; Daza and Frank, 2019).

For cross-lingual SRL, annotation projection has received the most attention (Pado and Lapata, 2014). A range of strategies have been proposed to enhance the SRL performance of the target language, including improving the projection quality (Tiedemann, 2015), joint learning of syntax and semantics (Kozhevnikov and Titov, 2013), iterative bootstrapping to reduce the influence of noise target corpus (Akbik et al., 2015), and joint translation and SRL (Aminian et al., 2019).

Our work is mainly inspired by the recent work of treebank translation of cross-lingual dependency parsing (Tiedemann et al., 2014; Tiedemann, 2015; Rasooli and Collins, 2015; Guo et al., 2016; Tiedemann and Agic, 2016; Conneau et al., 2018; Zhang et al., 2019), which is referred to as the translation-based approaches. These approaches directly project the gold-standard annotation into the target side, alleviating the problem of erroneous source annotations in standard annotation projection. In addition, we combine the approach with model transferring, which has been concerned little for cross-lingual SRL. The model transferring

¹<https://github.com/scofield7419/XSRL-ACL> under Apache License 2.0.


benefits much from the recent advance of cross-lingual contextualized word representations (He et al., 2019).

The development of universal annotation schemes for a variety of NLP tasks can greatly facilitate cross-lingual SRL, including POS tagging (Petrov et al., 2012), dependency parsing (McDonald et al., 2013; Przepiórkowski and Patejuk, 2018), morphology (Sylak-Glassman et al., 2015) and SRL (Aminian et al., 2019). Our work makes use of the publicly available Universal Proposition Bank (UPB) (Akbik et al., 2015; Akbik and Li, 2016), which annotates the predicate and semantic roles following the frame and role schemes of the English Proposition Bank 3.0 (Kingsbury and Palmer, 2002; Palmer et al., 2005).


Supervised SRL models are also closely related to our work (He et al., 2017, 2018a; Xia et al., 2019). A great deal of work attempts for an end-to-end solution with sophisticated neural networks, detecting the predicates as well as the corresponding argument roles in one shot (He et al., 2017; Tan et al., 2018; Li et al., 2019). Also there exist a number of studies which aims for adapting various powerful features for the task (Strubell et al., 2018; Li et al., 2018). In this work, we exploit a multilingual PGN-BiLSTM model (Jia et al., 2019) with contextualized word representations (He et al., 2019), which can obtain state-of-the-art performance for cross-lingual SRL.

3 SRL Translation


We induce automatic target data from the gold-standard source data by full translation and then project the SRL predicates and arguments into their corresponding words by aligning, producing the final translated SRL corpus for the target language automatically. The method has been demonstrated effective for cross-lingual dependency parsing (Tiedemann et al., 2014; Tiedemann, 2015; Tiedemann and Agic, 2016; Zhang et al., 2019). Compared with annotation projection, we can ensure the annotation quality at the source side, thus higher quality target corpus is also expected. In addition, dependency-based SRL could benefit more by this method, as only predicate words and their arguments are required to be projected into the target side, while dependency parsing should concern all sentential words. The overall process is accomplished by two steps: translating and projecting.



(a) One-to-one projection.



(b) Predicate-argument collision. Only keep predicate.



(c) Argument-argument collision. Only keep the one with higher confidence.

Figure 2: Examples of SRL projection from English (upper) to German (lower), where the thick blue solid arrows indicate successful projections, and the thin red dotted arrows indicate invalid projections.

Translating. First, we use a state-of-the-art translation system to produce the target translations for the sentences of the source SRL data. Give a source sentence $e_1 \cdots e_n$, we translate it into $f_1 \cdots f_m$ of the target language. It is worth noting that the recent impressive advances in NMT (Bahdanau et al., 2015; Wu et al., 2016) facilitate our work greatly, which enables our method to have high-quality translations.

Projecting. Then we incrementally project the corresponding predicates or arguments of a source sentence $e_1 \cdots e_n$ to its target $f_1 \cdots f_m$. We adopt two kinds of information to assist the projection: (1) the alignment probabilities $a(f_j|e_i)$ from the source word e_i into f_j , which can be calculated by a word-alignment tool, and (2) the POS tag distributions $p(t_*|f_j)$ of the target sentential words, which can be derived from a supervised target language POS tagger, where $i \in [1, n]$, $j \in [1, m]$, and t_* denotes an arbitrary POS tag.

We focus on SRL-related words of the source sentence only, and perform the process gradually at the predicate level. For each predicate in a sentence, we collect the predicate word as well as its role words, and then project their role labels into the target sentence. Formally, for each of these words (i.e., e_i), we have the SRL role tag r_{e_i} as

well as its POS tag t_{e_i} , both of which have been already annotated in the UPB. First, we find its target word f_j with the highest alignment probability, regarding the word f_j as the corresponding projection carrying the semantic role r_{e_i} . Then we calculate the confidence score of this projection by the following formula:

$$\text{score}(e_i \rightarrow f_j, r_{e_i}) = a(f_j|e_i)p(t_{e_i}|f_j), \quad (1)$$

which is a joint probability of word alignment corresponding and POS tag consistency.

The one-one target-source alignment 2(a) is the ideal condition of the projection. However, there could be many-to-one cases for the given words, leading to semantic role conflicts at the target language words. For these cases, we take precedence for the predicate projections, and otherwise keep only the highest confidence projections. Figure 2(b) shows a predicate-argument conflict example, where the predicate projection is reserved, and Figure 2(c) shows an argument-argument conflict example where the projection with the higher confidence score is reserved.

Finally, we set a threshold value α to remove low confidence projections. If the confidence score of a predicate projection is below α , all the roles of this predicate are removed as well. For the argument projections whose confidence is below α , we remove the single arguments directly, with no influence on the other projections.

4 The SRL Model

In this work, we focus on dependency-based SRL, recognizing semantic roles for a given predicate (He et al., 2017). The task can be treated as a standard sequence labeling problem, and a simple multi-layer BiLSTM-CRF model is exploited here, which has archived state-of-the-art performance with contextualized word representations (He et al., 2018b; Xia et al., 2019; He et al., 2019). In particular, we adapt the model to better support multilingual inputs by using a PGN module on the BiLSTM (Hochreiter and Schmidhuber, 1997). Figure 3 shows the overall architecture.

4.1 Word Representation

Given an input sentence $s = w_1 \cdots w_n$ of a specific language \mathcal{L} and w_p (p denotes the position) is the predicate word, we use three sources of features to represent each word: (1) the word form, (2) the

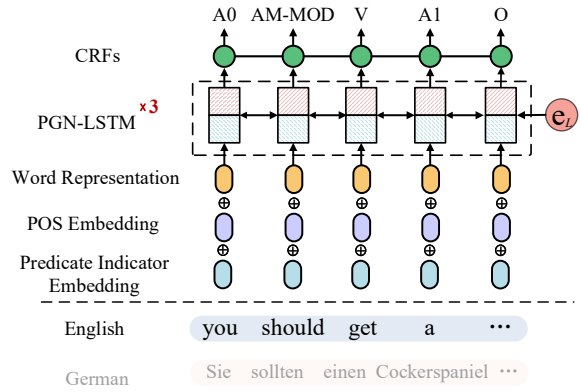


Figure 3: The overall architecture of the SRL model.

POS tag and (3) the predicate indicator:

$$\mathbf{x}_i = \mathbf{v}_{w_i} \oplus \mathbf{v}_{t_i} \oplus \mathbf{v}_{(i==p)}, \quad (2)$$

where $t_1 \cdots t_n$ is the universal POS tag sequence for the input sentence. For the POS tags and the predicate indicators, we use the embedding method to obtain their vectorial representations.

We compare three kinds of word form representations for cross-lingual SRL: (1) multilingual word embeddings, (2) multilingual ELMo representation (Peters et al., 2018), and (3) multilingual BERT representation (Devlin et al., 2019). Note that we use the averaged vector of the inner-word piece representations from BERT outputs as the full word representation.

4.2 Encoding Layer

We employ the PGN-BiLSTM (Platanios et al., 2018; Jia et al., 2019) to encode the input sequence $\mathbf{x}_1 \cdots \mathbf{x}_n$, which is first introduced for cross-domain transfer learning to capture domain difference. Here we use it for the multilingual setting aiming to model the language characteristics.

Compared with the vanilla BiLSTM module, PGN-BiLSTM dynamically selects the language-aware parameters for BiLSTM. Let \mathbf{V} be the flattened vector of all the parameters of a BiLSTM cell, the language-aware $\mathbf{V}_{\mathcal{L}}$ is produced by:

$$\mathbf{V}_{\mathcal{L}} = \mathbf{W}_{\text{PGN}} \times \mathbf{e}_{\mathcal{L}}, \quad (3)$$

where \mathbf{W}_{PGN} denotes the parameters of vanilla BiLSTM part in the PGN-BiLSTM, including the weights of the input, forget, output gates and the cell modules, and $\mathbf{e}_{\mathcal{L}}$ is the embedding representation of language \mathcal{L} . The mechanism of parameter generation of PGN-BiLSTM is illustrated in Figure 4. Following, we derive module parameters from




Figure 4: The mechanism of the PGN-BiLSTM.

V_L to compute the BiLSTM outputs. The overall process can be formalized as:

$$\begin{aligned} h_1 \cdots h_n &= \text{PGN-BiLSTM}(x_1 \cdots x_n, e_L) \\ &= \text{BiLSTM}_{V_L}(x_1 \cdots x_n) \end{aligned} \quad (4)$$

which differs from the vanilla BiLSTM in that e_L is one extra input to obtain BiLSTM parameters. Specifically, we adopt a 3-layer bidirectional PGN-LSTM as the encoder.

4.3 Output Layer

Given the encoder output $h_1 \cdots h_n$ for sentence $s = w_1 \cdots w_n$, we use CRFs (Lafferty et al., 2001) to compute the probability of each candidate output $y = y_1 \cdots y_n$:

$$\begin{aligned} o_i &= \mathbf{W}h_n, i \in [1, n] \\ p(y|s) &= \frac{\exp\{\sum_i(o_i + \mathbf{T}_{y_{i-1}, y_i})\}}{Z} \end{aligned} \quad (5)$$

where \mathbf{W} and \mathbf{T} are the parameters of CRFs, and Z is a normalization factor for probability calculation. The Viterbi algorithm is used to search for the highest-probability output SRL tag sequence.

5 Experiments

5.1 Universal Proposition Bank

Our experiments are based on the Universal Proposition Bank (UPB, v1.0)², which is built upon Universal Dependency Treebank (UDT, v1.4)³ and Proposition Bank (PB, v3.0)⁴. In UPB, consistent dependency-based universal SRL annotations are constructed across all languages. In particular, we assemble the English SRL dataset based on the English EWT subset from the UDT v1.4 and the English corpus in PB v3.0. Finally, we choose a total of seven languages as our datasets, including English (EN) and German (DE) of the IE.German

²<https://github.com/System-T/UniversalPropositions>

³<https://lindat.mff.cuni.cz/repository/xmlui/handle/11234/1-1827>

⁴<http://proppbank.github.io/>

Fam.	Lang.	Train	Dev	Test	Pred.	Arg.
IE.Ge	EN	10,907	1,631	1,633	41,359	100,170
	DE	14,118	799	977	23,256	58,319
IE.Ro	FR	14,554	1,596	298	26,934	44,007
	IT	12,837	489	489	26,576	56,893
	ES	28,492	3,206	1,995	81,318	177,871
	PT	7,494	938	936	19,782	41,449
Ura	FI	12,217	716	648	27,324	60,502

Table 1: Statistics of the UPB, where Fam. indicates the language family, IE.Ge refers to the Indo-European Germanic, IE.Ro refers to the Indo-European Romance, and Ura represents Uralic.

family, French (FR), Italian (IT), Spanish (ES)⁵ and Portuguese (PT) of the IE.Romance family, and Finnish (FI) of the Uralic family. Table 1 shows the data statistics in detail.

5.2 SRL Translation

We focus on unsupervised cross-lingual SRL, assuming that no gold-standard target-language SRL corpus is available. Our goal is to construct pseudo training datasets by corpus translation from the gold-standard source-language SRL datasets. The Google Translation System⁶ is adopted for sentence translation, and the fastAlign toolkit (Dyer et al., 2013) is used to obtain word alignments. In order to obtain accurate word alignment, we collect a set of parallel corpora to augment the training dataset of fastAlign.⁷ The universal POS tags of the translated sentences are produced by supervised monolingual POS taggers, which are trained on the corresponding UDT v1.4 datasets, respectively.⁸

5.3 Settings

Multi-lingual word representations. As mentioned in Section 4.1, we investigate three kinds of multilingual word representations: (1) Word Embedding (Emb): MUSE is exploited to align all monolingual fastText word embeddings into a universal space (Lample et al., 2018).⁹ (2) ELMo: A blended dataset¹⁰ of the seven languages is used to train multilingual ELMo (Mulcaire et al., 2019).

⁵We merge the Spanish and Spanish-AnCora as one.

⁶<https://translate.google.com/>, Oct. 1 2019

⁷<http://opus.nlpl.eu/>, Europarl v8.

⁸A simple BiLSTM-CRF POS tagging model with monolingual ELMo representations is used, which can achieve accuracies of 96.54%(EN), 97.15%(DE), 94.42%(FR), 97.21%(IT), 94.12%(ES), 95.86%(PT) and 92.16%(FI), respectively.

⁹<https://github.com/facebookresearch/MUSE>

¹⁰CoNLL2017 corpus: <https://lindat.mff.cuni.cz/repository/xmlui/handle/11234/1-1989>




Figure 5: Performances with the translated target data under varying projection threshold α .

(3) BERT: the official released multilingual BERT (base, cased version) is used directly (Devlin et al., 2019).¹¹

Hyperparameters. For SRL translation, there is only one hyperparameter, the projection confidence threshold α , for filtering low-quality translated SRL sentences. Figure 5 shows the performances in the preliminary experiments for each languages under different α . Accordingly, we set α universally for all languages to 0.4. For the neural SRL models, the dimension sizes of multilingual word embeddings, ELMo and BERT are 300, 1024 and 768, respectively. The POS tag, predicate-indicator and language ID embedding sizes are 100, 100 and 32, respectively. The hidden size of LSTM is set to 650. We exploit online training with a batch size of 50, and the model parameters are optimized by using the Adam algorithm with an initial rate of 0.0005. The training is performed over the whole training dataset without early-stopping for 80 iterations on bilingual transfer, and 300 iterations on multilingual transfer.

Baselines. In order to test the effectiveness of our *PGN* model, we compare it with several baselines as well. First, we denote our model by using the vanilla BiLSTM instead as *BASIC*, and in particular, this model is exploited for all monolingual training all through this work. Further, we adopt two much stronger baselines, the *MoE* model proposed by Guo et al. (2018) and the *MAN-MoE* model proposed by Chen et al. (2019), respectively. Both the two models are designed to train a model effectively based on corpora from multiple languages, similar to our *PGN* model.

Evaluation. We use the F1 score as the major metric to measure the model performance for each

¹¹<https://github.com/google-research/bert>

Model	DE	FR	IT	ES	PT	FI	Avg
SRC							
Emb	42.7	51.0	42.6	40.1	43.9	30.0	41.7
BERT	43.2	53.1	44.4	41.2	44.2	31.6	43.0
ELMo	46.8	54.6	43.0	42.1	46.1	33.9	44.4
TGT							
Emb	49.4	51.3	45.5	48.4	46.9	38.7	46.7
BERT	53.0	54.3	49.1	51.3	48.8	41.1	49.6
ELMo	54.6	55.3	49.7	53.6	49.8	43.9	51.1
SRC & TGT (ELMo)							
<i>BASIC</i>	59.2	61.7	55.1	58.3	53.7	47.6	55.8
<i>PGN</i>	65.0	64.8	58.7	62.5	56.0	54.5	60.3
<i>MoE</i>	63.2	63.3	56.7	60.3	55.0	50.6	58.2
<i>MAN-MoE</i>	64.3	65.3	57.1	62.8	55.2	52.3	59.4

Table 2: Results of cross-lingual transfer from English.

target language. Each model is trained five times and the averaged value is reported. We conduct significance tests by using the Dan Bikel’s randomized parsing evaluation comparator¹².

5.4 Cross-Lingual Transfer from English

We first conduct experiments on cross-lingual transfer from the English source to the rest of the other six target languages, respectively, which has been a typical setting for cross-lingual investigations (Wang et al., 2019). The results are summarized in Table 2. We list the F-scores by using only the source corpus (SRC), only the translated target corpus (TGT) and the mixture corpus of source and target (SRC & TGT), comparing the performances of different multilingual word representations as well as different multilingual SRL models.

Multilingual word representations. First, we evaluate the effectiveness of the three different multilingual word representations exploited. We compare their performances under two settings, by using SRC and TGT corpus, respectively. According to the results, we find that the multilingual contextualized word representations (i.e. BERT and ELMo) are better in both two settings, which is consistent with previous studies (Mulcaire et al., 2019; Schuster et al., 2019). Interestingly, the multilingual BERT performs worse than the ELMo, which can be explained by that the ELMo representation is pre-trained based on the corpus which involves in the focused seven languages. This indicates that the official released multilingual BERT can be further improved, since monolingual BERT has been demonstrated to produce better performances than

¹²<http://www.cis.upenn.edu/~dbikel/software.html#comparator>

ELMo (Tenney et al., 2019).

Translated target. Next, We consider taking the translated target as only the training data to examine the effectiveness of the pseudo datasets. As shown in Table 2, we find that the translated datasets can bring significantly better performances than the source baseline overall languages, resulting in an averaged F1 score increase of $51.1 - 44.4 = 6.7$. The results demonstrate that corpus translation is one effective way for cross-lingual SRL. The observation is in line with the previous work for cross-lingual dependency parsing (Tiedemann and Agic, 2016; Zhang et al., 2019). By direct gold-standard corpus translation, the produced pseudo training data can not only remain high-quality SRL annotations but also capture the language divergences effectively, which leads to better performance than the source baseline model.

Combining source and pseudo target. Further, we combine the pseudo translated target corpus with the source language corpus together to train the target SRL models. According to the numbers in Table 2, we see that further gains can be achieved for all languages, where the averaged improvement is $55.8-51.1=4.7$ (*BASIC* is used for a fair comparison). Note that since several source sentences are filtered during translation which might be the reason for the gains, we make a fairer comparison off-the-line by setting $\alpha=0$ (i.e., no sentence filtering). Similar gains can be achieved still. Considering that the translated sentences are semantically equal to their counterparts in the gold-standard source, the possible reasons could be two hands: (1) the translated sentences may be biased in linguistic expression due to the data-driven translation models, (2) the discarded conflicted annotations in corpus translation are important, which are complementary to our model.

Language-aware encoder. Finally, we investigate the effectiveness of PGN-BiLSTM module, which is exploited to capture language-specific information when the mixture corpus of both source and target datasets are used for training. As shown in Table 2, we can see that the language-aware encoder by *PGN* can boost the F1 scores significantly, achieving an averaged improvement by $60.3-55.8=4.5$. In addition, we report the results of *MoE* and *MAN-MoE*, respectively, which also exploit the language information. All the results demonstrate the usefulness of language-specific informa-

Model	EN	DE	FR	IT	ES	PT	FI	Avg
SRC								
Emb	50.3	49.2	52.4	44.9	46.7	51.0	36.4	47.3
BERT	51.8	50.6	54.0	45.3	51.3	51.8	38.1	49.0
ELMo	53.6	51.6	56.7	51.3	57.4	52.6	39.7	51.8
TGT								
Emb	56.5	51.6	55.2	47.1	50.0	53.2	40.4	50.6
BERT	59.8	55.5	57.0	52.6	54.3	56.6	44.0	54.3
ELMo	60.7	57.8	59.9	54.8	56.7	58.8	46.9	56.5
SRC & TGT (ELMo)								
<i>BASIC</i>	61.9	64.8	60.3	56.4	61.1	63.1	50.7	59.8
<i>PGN</i>	65.7	68.8	66.1	64.8	68.7	69.2	58.6	66.0
<i>MoE</i>	63.2	67.8	63.1	62.6	65.2	67.5	54.2	63.4
<i>MAN-MoE</i>	64.0	68.5	67.2	65.7	67.5	69.0	57.5	65.6

Table 3: Cross-lingual transfer with multiple sources.

Source	EN	DE	FR	IT	ES	PT	FI
EN		65.0	64.8	58.7	62.5	56.0	54.5
DE	63.2		63.9	60.4	65.8	53.4	50.5
FR	60.1	53.7		63.3	63.6	62.1	51.3
IT	60.2	58.9	65.3		65.1	58.6	48.6
ES	60.1	57.3	64.9	64.1		67.0	50.7
PT	57.3	58.6	65.1	63.5	67.8		40.9
FI	50.7	52.1	64.6	53.6	60.3	51.6	
ALL	65.7	68.8	66.1	64.8	68.7	69.2	58.6

Table 4: The results of bilingual transferring.

tion, and our *PGN* model is most effective.

5.5 Multi-Source Transfer

Further, we investigate the setting of multi-source transfer learning, where all other languages except a given target language are used as the source languages, aiming to study the effectiveness of our translation-based method comprehensively.

Overall performances. The results on multiple source SRL transferring are shown in Table 3. Generally, the results share similar tendencies with the single-source cross-lingual transfer from the source English, where the multilingual ELMo performs the best, the SRL models trained on the translated target datasets show better performances than those trained with the source datasets, and the mixture corpus with both source and target language datasets bring the best performances, which can be further improved by our final *PGN* model with language-aware encoders. We compare the *PGN* model with the *MoE* and *MAN-MoE* as well, showing slightly better performances, which indicates the effectiveness of the PGN-BiLSTM module. In addition, we can see that multi-source models outperform the single-source models in all cases, which is intuitive and consistent with previous studies (Lin et al., 2019).

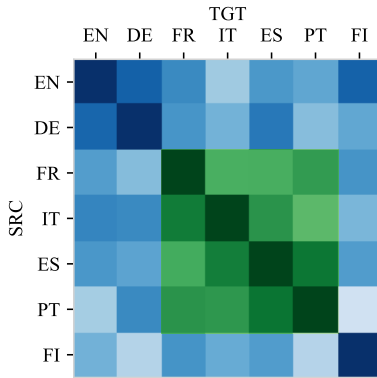


Figure 6: Similarity heatmap of language embeddings for source languages to target languages. Deeper color indicates higher similarity.

Fine-grained bilingual transfer. Following, we investigate the individual bilingual SRL transferring by examining the performance of each source-target language pair, aiming to uncover which language benefits a target most and trying to answer whether all source languages are useful for a target language. Table 4 shows the results, where the cross-lingual models are trained on the mixture corpus of the source and translated target datasets. First, we can see that the languages belonging to a single family can benefit each other greatly, bringing better performances than the other languages in the majority of cases (i.e., EN–DE, FR–IT–ES–PT). Second, the multi-source transfer as indicated by All is able to obtain better performances across all languages, which further demonstrates its advantages over the single-source transfer.

Further, we look into the *PGN* model in detail, aiming to understand their capabilities of modeling linguistic-specific information. We examine it by simply visualizing the language ID embeddings $e_{\mathcal{L}}$ of each source-target language pair, respectively, where their Euclidean distances are depicted. Intuitively, better performance can be achieved if the distance between the target and the source languages is closer. Figure 6 shows the heatmap matrix. We can see the overall tendency is highly similar to the results in Table 4, which is consistent with our intuition.

5.6 Analysis

Here we conduct detailed analysis to understand the gains from the translated target datasets. We select three representative languages for analysis, including German (DE), French (FR) and Finnish (FI), one language for each family, and compare

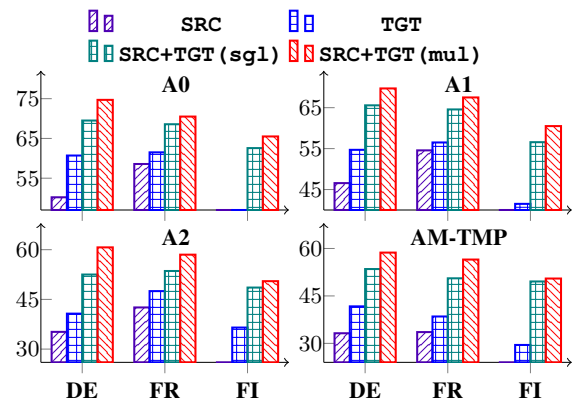


Figure 7: Performances on different argument label.

four models mainly, including three models (i.e., SRC, TGT and SRC & TGT with *PGN*) of the single-source transfer from English and the final *PGN* model of multi-source transfer.

Performances by the SRL roles. First, we investigate the cross-lingual SRL performances in terms of SRL Roles. We select four representative roles for comparison, including A0 (*Agent*), A1 (*Patient*), A2 (*Instrument, Benefactive, Attribute*) and AM-TMP (*Temporal*), and report their F1 scores. Figure 7 shows the results. As a whole, the role A0 achieves the best F1 scores across all languages and all models, A1 ranks the second, and A2 and AM-TMP are slightly worse. The tendency could be accounted for by the distribution of these labels, where A0 is the most frequent and A2 and AM-TMP have lower frequencies than A0 and A1. The second possible reason could be due to that the majority of the A0 and A1 words are notional words which could be more easily transferred by cross-lingual models.

In addition, we can see that the tendencies across different models for all three languages and all labels are identical, where multi-source transfer performs the best, single-source SRC+TGT ranks the second and our baseline model is the last. The observation is consistent with the overall tendency, demonstrating the stability and also further verifying the effectiveness of our proposed models.

Performances by the distances to the predicate.

Second, we study the SRL performances in terms of the distance to the predicate word. Intuitively, long-distance relations are more difficult, thus we expect that the SRL performance would decrease as the distance increases, as SRL actually detects the relationship between the role words and their

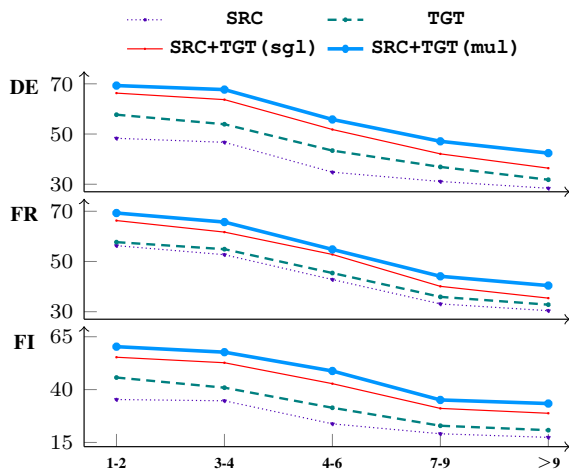


Figure 8: Performances by surface distance between predicates and arguments.

predicates. Figure 8 shows the F1 scores. First, for all the settings we can see that the SRL performance drops by longer distances, which confirms our intuition. In addition, the tendency between different models is the same as the overall results, demonstrating the effectiveness of our method.

6 Conclusion

We proposed a translation-based alternative for cross-lingual SRL. The key idea is to construct high-quality datasets for the target languages by corpus translation from the gold-standard SRL annotations of the source languages. In addition, we combined the gold-standard source SRL corpora and the pseudo translated target corpora together to enhance the cross-lingual SRL models. We investigated cross-lingual SRL models with different kinds of multilingual word representations. Further, we presented a PGN-BiLSTM encoder to better exploit the mixture corpora of different languages. Experimental results on the UPB v1.0 dataset show that the translation-based method is an effective method for cross-lingual SRL transferring. Significant improvements can be achieved by using the translated datasets for all selected languages, including both single-source and multi-source transfer. Experiment analysis is offered to understand the proposed method in depth.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No.61772378 and 61602160), the National Key Research and Development Program of China (No.2017YFC1200500),

the Research Foundation of Ministry of Education of China (No.18JZD015), and the Major Projects of the National Social Science Foundation of China (No.11&ZD189).

References

- Alan Akbik, Laura 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 ACL*, pages 397–407.
- Alan Akbik and Yunyao Li. 2016. Polyglot: Multilingual semantic role labeling with unified labels. In *Proceedings of ACL*, pages 1–6.
- Maryam Aminian, Mohammad Sadegh Rasooli, and Mona Diab. 2019. Cross-lingual transfer of semantic roles: From raw text to semantic roles. In *Proceedings of the IWCS*, pages 200–210.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the ICLR*.
- Anders Björkelund, Love Hafdell, and Pierre Nugues. 2009. Multilingual semantic role labeling. In *Proceedings of the CoNLL*, pages 43–48.
- Xilun Chen, Ahmed Hassan Awadallah, Hany Hassan, Wei Wang, and Claire Cardie. 2019. Multi-source cross-lingual model transfer: Learning what to share. In *Proceedings of the ACL*, pages 3098–3112.
- Xilun Chen, Yu Sun, Ben Athiwaratkun, Claire Cardie, and Kilian Q. Weinberger. 2018. Adversarial deep averaging networks for cross-lingual sentiment classification. *Transactions of the Association for Computational Linguistics*, 6:557–570.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. XNLI: Evaluating cross-lingual sentence representations. In *Proceedings of the EMNLP*, pages 2475–2485.
- Angel Daza and Anette Frank. 2019. Translate and label! an encoder-decoder approach for cross-lingual semantic role labeling. In *Proceedings of the EMNLP*, pages 603–615.
- 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 NAACL*, pages 4171–4186.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. A simple, fast, and effective reparameterization of IBM model 2. In *Proceedings of the NAACL*, pages 644–648.

- Qin Gao and Stephan Vogel. 2011. Corpus expansion for statistical machine translation with semantic role label substitution rules. In *Proceedings of the ACL*, pages 294–298.
- Pierre-Etienne Genest and Guy Lapalme. 2011. Framework for abstractive summarization using text-to-text generation. In *Proceedings of the Workshop on Monolingual Text-To-Text Generation*, pages 64–73.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2016. A representation learning framework for multi-source transfer parsing. In *Proceedings of the AAAI*, pages 2734–2740.
- Jiang Guo, Darsh Shah, and Regina Barzilay. 2018. Multi-source domain adaptation with mixture of experts. In *Proceedings of the EMNLP*, pages 4694–4703.
- Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018a. Jointly predicting predicates and arguments in neural semantic role labeling. In *Proceedings of the ACL*, 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 ACL*, pages 473–483.
- Shexia He, Zuchao Li, and Hai Zhao. 2019. Syntax-aware multilingual semantic role labeling. In *Proceedings of the EMNLP*, pages 5353–5362.
- Shexia He, Zuchao Li, Hai Zhao, and Hongxiao Bai. 2018b. Syntax for semantic role labeling, to be, or not to be. In *Proceedings of the ACL*, pages 2061–2071.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Rebecca Hwa, Philip Resnik, Amy Weinberg, Clara I. Cabezas, and Okan Kolak. 2005. Bootstrapping parsers via syntactic projection across parallel texts. *Natural Language Engineering*, 11(3):311–325.
- Chen Jia, Xiaobo Liang, and Yue Zhang. 2019. Cross-domain NER using cross-domain language modeling. In *Proceedings of the ACL*.
- Atif Khan, Naomie Salim, and Yogan Jaya Kumar. 2015. A framework for multi-document abstractive summarization based on semantic role labelling. *Appl. Soft Comput.*, 30:737–747.
- Paul R. Kingsbury and Martha Palmer. 2002. From treebank to propbank. In *Proceedings of the LREC, Spain*.
- Mikhail Kozhevnikov and Ivan Titov. 2013. Cross-lingual transfer of semantic role labeling models. In *Proceedings of the ACL*, pages 1190–1200.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the ICML*, pages 282–289.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *Proceedings of the ICLR*.
- Tao Lei, Yuan Zhang, Lluís Màrquez, Alessandro Moschitti, and Regina Barzilay. 2015. High-order low-rank tensors for semantic role labeling. In *Proceedings of the NAACL*, pages 1150–1160.
- Zuchao Li, Shexia He, Jiaxun 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 EMNLP*, pages 2401–2411.
- Zuchao Li, Shexia He, Hai Zhao, Yiqing Zhang, Zhuosheng Zhang, Xi Zhou, and Xiang Zhou. 2019. Dependency or span, end-to-end uniform semantic role labeling. In *Proceedings of the AAAI*, pages 6730–6737.
- Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, Antonios Anastasopoulos, Patrick Littell, and Graham Neubig. 2019. Choosing transfer languages for cross-lingual learning. In *Proceedings of the ACL*, pages 3125–3135.
- Ding Liu and Daniel Gildea. 2010. Semantic role features for machine translation. In *Proceedings of the COLING*, pages 716–724.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-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 ACL*, pages 92–97.
- Phoebe Mulcaire, Jungo Kasai, and Noah A. Smith. 2019. Low-resource parsing with crosslingual contextualized representations. *CoRR*, abs/1909.08744.
- Phoebe Mulcaire, Swabha Swayamdipta, and Noah A. Smith. 2018. Polyglot semantic role labeling. In *Proceedings of the ACL*, pages 667–672.
- Sebastian Pado and Mirella Lapata. 2014. Cross-lingual annotation projection for semantic roles. *Journal of Artificial Intelligence Research*, 36(36):307–340.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.

- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *in Proceedings of the NAACL*, pages 2227–2237.
- Slav Petrov, Dipanjan Das, and Ryan McDonald. 2012. A universal part-of-speech tagset. In *Proceedings of the LREC*, pages 2089–2096.
- Lonneke van der Plas, Paola Merlo, and James Henderson. 2011. Scaling up automatic cross-lingual semantic role annotation. In *Proceedings of the ACL*, pages 299–304.
- Emmanouil Antonios Platanios, Mrinmaya Sachan, Graham Neubig, and Tom Mitchell. 2018. Contextual parameter generation for universal neural machine translation. In *Proceedings of the EMNLP*, pages 425–435.
- Adam Przepiórkowski and Agnieszka Patejuk. 2018. Arguments and adjuncts in universal dependencies. In *Proceedings of the COLING*, pages 3837–3852.
- Mohammad Sadegh Rasooli and Michael Collins. 2015. Density-driven cross-lingual transfer of dependency parsers. In *Proceedings of the EMNLP*, pages 328–338.
- Tal Schuster, Ori Ram, Regina Barzilay, and Amir Globerson. 2019. Cross-lingual alignment of contextual word embeddings, with applications to zero-shot dependency parsing. In *Proceedings of the NAACL*, pages 1599–1613.
- Dan Shen and Mirella Lapata. 2007. Using semantic roles to improve question answering. In *Proceedings of the EMNLP*, pages 12–21.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. Linguistically-informed self-attention for semantic role labeling. In *Proceedings of the EMNLP*, pages 5027–5038.
- Swabha Swayamdipta, Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2016. Greedy, joint syntactic-semantic parsing with stack LSTMs. In *Proceedings of the CoNLL*, pages 187–197.
- John Sylak-Glassman, Christo Kirov, David Yarowsky, and Roger Que. 2015. A language-independent feature schema for inflectional morphology. In *Proceedings of the ACL*, pages 674–680.
- Oscar Täckström, Ryan T. McDonald, and Jakob Uszkoreit. 2012. Cross-lingual word clusters for direct transfer of linguistic structure. In *Proceedings of the NAACL*, pages 477–487.
- Zhixing Tan, Mingxuan Wang, Jun Xie, Yidong Chen, and Xiaodong Shi. 2018. Deep semantic role labeling with self-attention. In *Proceedings of the AAAI*, pages 4929–4936.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, and Ellie Pavlick. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. In *Proceedings of the ICLR*.
- Jörg Tiedemann. 2015. Improving the cross-lingual projection of syntactic dependencies. In *Proceedings of the Nordic Conference of Computational Linguistics*, pages 191–199.
- Jörg Tiedemann and Željko Agić. 2016. Synthetic treebanking for cross-lingual dependency parsing. *Journal of Artificial Intelligence Research*, 55:209–248.
- Jörg Tiedemann, Željko Agić, and Joakim Nivre. 2014. Treebank translation for cross-lingual parser induction. In *Proceedings of the CoNLL*, pages 130–140.
- Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester. 2015. Machine comprehension with syntax, frames, and semantics. In *Proceedings of the ACL*, pages 700–706.
- Yuxuan Wang, Wanxiang Che, Jiang Guo, Yijia Liu, and Ting Liu. 2019. Cross-lingual BERT transformation for zero-shot dependency parsing. In *Proceedings of the EMNLP*, pages 5720–5726.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144.
- Qingrong Xia, Zhenghua Li, Min Zhang, Meishan Zhang, Guohong Fu, Rui Wang, and Luo Si. 2019. Syntax-aware neural semantic role labeling. In *Proceedings of the AAAI*, pages 7305–7313.
- David Yarowsky, Grace Ngai, and Richard Wicentowski. 2001. Inducing multilingual text analysis tools via robust projection across aligned corpora. In *Proceedings of the International Conference on Human Language Technology Research*.
- Meishan Zhang, Yue Zhang, and Guohong Fu. 2019. Cross-lingual dependency parsing using code-mixed TreeBank. In *Proceedings of the EMNLP*, pages 996–1005.
- Han Zhao, Shanghang Zhang, Guanhang Wu, João Paulo Costeira, José M. F. Moura, and Geoffrey J. Gordon. 2018. Multiple source domain adaptation with adversarial learning. In *Proceedings of the ICLR*.

Target	SRC			TGT			SRC+TGT	SRC+TGT
	Emb	BERT	ELMo	Emb	BERT	ELMo	BASIC+ELMo	PGN+ELMo
Source: DE								
EN	47.32	51.62	52.82	55.04	59.20	60.48	61.05	63.21
FR	46.00	49.94	50.99	52.37	55.77	57.02	59.91	63.90
IT	40.90	43.68	45.06	48.01	51.62	52.91	57.94	60.38
ES	39.01	42.57	43.67	49.59	52.84	53.92	60.80	65.89
PT	38.25	41.73	43.07	41.44	45.76	46.94	49.14	53.40
FI	29.93	33.64	34.95	41.78	44.09	45.21	45.74	50.53
Source: FR								
EN	35.47	39.49	40.80	48.57	53.04	54.12	56.91	60.05
DE	40.01	43.86	45.24	41.33	45.16	46.54	50.53	53.69
IT	47.12	50.06	51.33	51.45	53.38	53.62	60.31	63.34
ES	40.46	44.01	45.09	50.36	53.77	54.79	58.61	63.62
PT	44.68	47.47	48.65	52.12	55.58	56.67	59.47	62.08
FI	26.44	30.92	32.07	40.71	43.97	45.05	48.76	51.31
Source: IT								
EN	37.07	39.49	40.96	47.10	51.26	52.40	54.05	60.13
DE	39.75	42.67	43.74	45.84	50.03	51.34	55.90	58.91
FR	47.39	50.08	51.28	54.45	57.29	58.78	60.03	65.30
ES	44.29	47.92	49.14	52.09	55.68	55.08	60.56	65.09
PT	42.18	46.54	47.60	49.38	53.85	54.96	57.02	58.65
FI	31.12	33.72	35.05	40.80	43.90	44.97	46.37	48.62
Source: ES								
EN	41.63	44.37	45.45	48.37	52.01	53.10	55.08	60.05
DE	36.32	39.65	40.73	44.21	47.90	49.37	51.11	57.27
FR	46.74	50.84	52.29	52.38	55.34	56.39	59.58	64.93
IT	41.39	45.42	46.82	50.10	53.01	54.01	58.83	64.09
PT	47.52	50.46	51.68	53.44	56.49	57.54	62.30	67.01
FI	29.46	32.19	33.33	39.27	42.95	44.07	47.91	50.72
Source: PT								
EN	34.83	38.09	39.27	43.16	46.09	47.50	53.12	57.30
DE	37.11	41.64	42.73	46.80	49.41	50.62	55.95	58.64
FR	42.05	46.28	47.61	49.07	52.78	54.15	58.64	65.12
IT	38.55	42.35	43.72	47.09	51.20	52.24	56.22	63.51
ES	39.58	44.01	45.02	46.57	50.61	52.06	60.84	67.81
FI	21.54	25.01	26.24	33.53	36.91	38.03	38.50	40.90
Source: FI								
EN	32.99	35.99	37.38	37.83	40.48	41.65	46.80	50.70
DE	30.98	34.59	35.70	40.75	44.41	45.84	50.68	52.12
FR	39.52	43.85	44.97	48.35	50.82	52.30	56.82	64.63
IT	33.82	36.39	37.66	41.81	46.01	47.07	52.61	53.65
ES	35.23	39.56	40.61	43.43	47.81	49.10	55.48	60.37
PT	27.93	31.90	33.30	33.84	38.21	39.43	47.75	51.61

Table 5: Results of fine-grained bilingual transfer.

A Bilingual Transfer by Each Source

In the paper, we investigate the individual bilingual SRL transferring of each source target language pair. We here list the detailed results of the bilingual transfer in Table 5.

B Extended SRL Analysis

We conduct detailed analysis on the detailed role labeling and the distance to the predicate, further

for the Italian, Spanish and Portuguese languages. Figures 9 and Figure 10 show the results.




Figure 9: Performances on different argument label.




Figure 10: Performances by surface distance between predicates and arguments.