

Cross-lingual Spoken Language Understanding with Regularized Representation Alignment

Zihan Liu, Genta Indra Winata, Peng Xu, Zhaojiang Lin, Pascale Fung

Center for Artificial Intelligence Research (CAiRE)

Department of Electronic and Computer Engineering

The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

zihan.liu@connect.ust.hk, pascale@ece.ust.hk

Abstract

Despite the promising results of current cross-lingual models for spoken language understanding systems, they still suffer from imperfect cross-lingual representation alignments between the source and target languages, which makes the performance sub-optimal. To cope with this issue, we propose a regularization approach to further align word-level and sentence-level representations across languages without any external resource. First, we regularize the representation of user utterances based on their corresponding labels. Second, we regularize the latent variable model (Liu et al., 2019a) by leveraging adversarial training to disentangle the latent variables. Experiments on the cross-lingual spoken language understanding task show that our model outperforms current state-of-the-art methods in both few-shot and zero-shot scenarios, and our model, trained on a few-shot setting with only 3% of the target language training data, achieves comparable performance to the supervised training with all the training data.¹

1 Introduction

Data-driven neural-based supervised training approaches have shown effectiveness in spoken language understanding (SLU) systems (Goo et al., 2018; Chen et al., 2019; Haihong et al., 2019). However, collecting large amounts of high-quality training data is not only expensive but also time-consuming, which makes these approaches not scalable to low-resource languages due to the scarcity of training data. Cross-lingual adaptation has naturally arisen to cope with this issue, which leverages the training data in rich-resource source languages and minimizes the requirement of training data in low-resource target languages.

¹The code is available in <https://github.com/zliucr/crosslingual-slu>.




Figure 1: Illustration of cross-lingual spoken language understanding systems, where English is the source language and Spanish is the target language.

In general, there are two challenges in cross-lingual adaptation. First, the imperfect alignment of word-level representations between the source and target language limits the adaptation performance. Second, even though we assume that the word-level alignment is perfect, the sentence-level alignment is still imperfect owing to grammatical and syntactical variances across languages. Therefore, we emphasize that cross-lingual methods should focus on the alignments of word-level and sentence-level representations, and increase the robustness for inherent imperfect alignments.

In this paper, we concentrate on the cross-lingual SLU task (as illustrated in Figure 1), and we consider both few-shot and zero-shot scenarios. To improve the quality of cross-lingual alignment, we first propose a **Label Regularization (LR)** method, which utilizes the slot label sequences to regularize the utterance representations. We hypothesize that if the slot label sequences of user utterances are close to each other, these user utterances should have similar meanings. Hence, we regularize the distance of utterance representations based on the corresponding representations of label sequences to further improve the cross-lingual alignments.

Then, we extend the latent variable model (LVM)

proposed by Liu et al. (2019a). The LVM generates a Gaussian distribution instead of a feature vector for each token, which improves the adaptation robustness. However, there are no additional constraints on generating distributions, making the latent variables easily entangled for different slot labels. To handle this issue, we leverage Adversarial training to regularize the LVM (ALVM). We train a linear layer to fit latent variables to a uniform distribution over slot types. Then, we optimize the latent variables to fool the trained linear layer to output the correct slot type (one hot vector). In this way, latent variables of different slot types are encouraged to disentangle from each other, leading to a better alignment of cross-lingual representations.

The contributions of our work are summarized as follows:

- We propose LR and ALVM to further improve the alignment of cross-lingual representations, which do not require any external resources.
- Our model outperforms the previous state-of-the-art model in both zero-shot and few-shot scenarios on the cross-lingual SLU task.
- Extensive analysis and visualizations are made to illustrate the effectiveness of our approaches.

2 Related Work

Cross-lingual Transfer Learning Cross-lingual transfer learning is able to circumvent the requirement of enormous training data by leveraging the learned knowledge in the source language and learning inter-connections between the source and the target language. Artetxe et al. (2017) and Conneau et al. (2018) conducted cross-lingual word embedding mapping with zero or very few supervision signals. Recently, pre-training cross-lingual language models on large amounts of monolingual or bilingual resources have been proved to be effective for the downstream tasks (e.g., natural language inference) (Conneau and Lample, 2019; Devlin et al., 2019; Pires et al., 2019; Huang et al., 2019). Additionally, many cross-lingual transfer algorithms have been proposed to solve specific cross-lingual tasks, for example, *named entity recognition* (Xie et al., 2018; Mayhew et al., 2017; Liu et al., 2020a), *part of speech tagging* (Kim et al., 2017; Zhang et al., 2016), *entity linking* (Zhang et al., 2013; Sil et al., 2018; Upadhyay et al., 2018b), personalized

conversations (Lin et al., 2020), and dialog systems (Upadhyay et al., 2018a; Chen et al., 2018).

Cross-lingual Task-oriented Dialog Systems

Deploying task-oriented dialogue systems in low-resource domains (Bapna et al., 2017; Wu et al., 2019; Liu et al., 2020b) or languages (Chen et al., 2018; Liu et al., 2019a,b), where the number of training samples is limited, is a challenging task. Mrkšić et al. (2017) expanded Wizard of Oz (WOZ) into multilingual WOZ by annotating two additional languages. Schuster et al. (2019) introduced a multilingual SLU dataset and proposed to leverage bilingual corpus and multilingual CoVe (Yu et al., 2018) to align the representations across languages. Chen et al. (2018) proposed a teacher-student framework based on a bilingual dictionary or bilingual corpus for building cross-lingual dialog state tracking. Instead of highly relying on extensive bilingual resources, Qin et al. (2020) introduced a data augmentation framework to generate multilingual code-switching data for cross-lingual tasks including the SLU task. Liu et al. (2019b) leveraged a mixed language training framework for cross-lingual task-oriented dialogue systems. And Liu et al. (2019a) proposed to refine the cross-lingual word embeddings by using very few word pairs, and introduced a latent variable model to improve the robustness of zero-shot cross-lingual SLU. Nevertheless, there still exists improvement space for the cross-lingual alignment. In this paper, we propose to further align the cross-lingual representations so as to boost the performance of cross-lingual SLU systems.

3 Methodology

Our model architecture and proposed methods are depicted in Figure 2, and combine label regularization (LR) and the adversarial latent variable model (ALVM) to conduct the intent detection and slot filling. In the few-shot setting, the input user utterances are in both the source and target languages, while in the zero-shot setting, the user utterances are only in the source language. Note that both the source and target languages contain only one language.

3.1 Label Regularization

3.1.1 Motivation

Intuitively, when the slot label sequences are similar, we expect the corresponding representations of user utterances across languages to be similar.

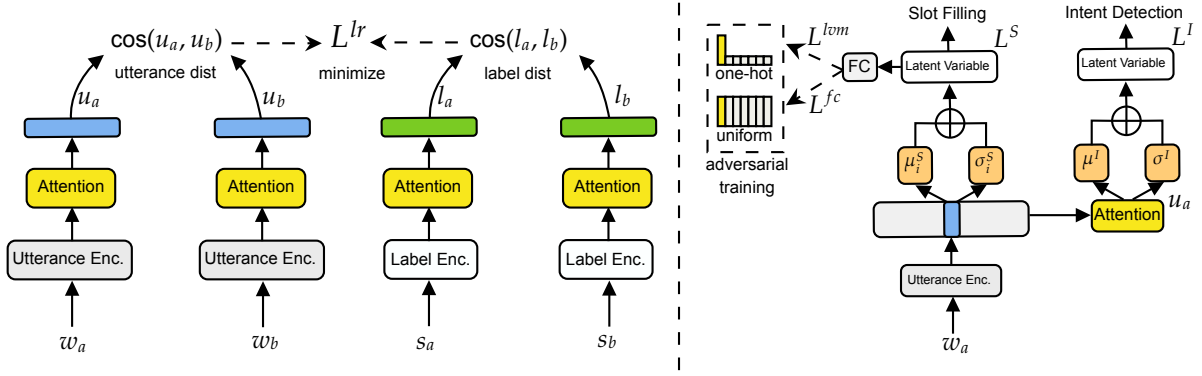


Figure 2: Left: Illustration of label regularization (LR). Right: The model architecture with adversarial latent variable model (ALVM), where FC consists of a linear layer and a softmax function.

For example, when the slot label sequences contain the weather slot and the location slot, the user utterances should be asking for the weather forecast somewhere. However, the representations of utterances across languages can not always meet these requirements because of the inherent imperfect alignments in word-level and sentence-level representations. Therefore, we propose to leverage existing slot label sequences in the training data to regularize the distance of utterance representations.

When a few training samples are available in the target language (i.e., few-shot setting), we regularize the distance of utterance representations between the source and target languages based on their slot labels. Given this regularization, the model explicitly learns to further align the sentence-level utterance representations across languages so as to satisfy the constraints. Additionally, it can also implicitly align the word-level BiLSTM hidden states across languages because sentence-level representations are produced based on them.

When zero training samples are available in the target language (i.e., zero-shot setting), we regularize the utterance representations in the source language. It can help better distinguish the utterance representations and cluster similar utterance representations based on the slot labels, which increases the generalization ability in the target language.

3.1.2 Implementation Details

Figure 2 (Left) illustrates an utterance encoder and a label encoder that generate the representations for utterances and labels, respectively.

We denote the user utterance as $\mathbf{w} = [w_1, w_2, \dots, w_n]$, where n is the length of the utterance. Similarly, we represent the slot label sequences as $\mathbf{s} = [s_1, s_2, \dots, s_n]$. We combine a

bidirectional LSTM (BiLSTM) (Hochreiter and Schmidhuber, 1997) and an attention layer (Felbo et al., 2017) to encode and produce the representations for user utterances and slot label sequences. The representation generation process is defined as follows:

$$[h_1^w, h_2^w, \dots, h_n^w] = \text{BiLSTM}_{\text{utter}}(\mathbf{E}(\mathbf{w})), \quad (1)$$

$$[h_1^s, h_2^s, \dots, h_n^s] = \text{BiLSTM}_{\text{label}}(\mathbf{E}(\mathbf{s})), \quad (2)$$

$$m_i^w = h_i^w v^w, \quad \alpha_i^w = \frac{\exp(m_i^w)}{\sum_{t=1}^n \exp(m_t^w)}, \quad (3)$$

$$m_i^s = h_i^s v^s, \quad \alpha_i^s = \frac{\exp(m_i^s)}{\sum_{t=1}^n \exp(m_t^s)}, \quad (4)$$

$$u = \sum_{i=1}^n \alpha_i^w h_i^w, \quad l = \sum_{i=1}^n \alpha_i^s h_i^s, \quad (5)$$

where the superscript w and s represents utterance and label, respectively, v is a trainable weight vector in the attention layer, α_i is the attention score for each token i , \mathbf{E} denotes the embedding layers for utterances and label sequences, and u and l denotes the representation of utterance \mathbf{w} and slot label \mathbf{s} , respectively.

In each iteration of the training phase, we randomly select two samples for the label regularization. As illustrated in Figure 2 (Left), we first calculate the cosine similarity of two utterance representations u_a and u_b , and the cosine similarity of two label representations l_a and l_b . Then, we minimize the distance of these two cosine similarities. The objective functions can be described as follows:

$$\cos(u_a, u_b) = \frac{u_a \cdot u_b}{\|u_a\| \|u_b\|}, \quad (6)$$

$$\cos(l_a, l_b) = \frac{l_a \cdot l_b}{\|l_a\| \|l_b\|}, \quad (7)$$

$$L^{lr} = \sum_{a,b} \text{MSE}(\cos(u_a, u_b), \cos(l_a, l_b)), \quad (8)$$

where the superscript lr denotes label regularization, and MSE represents mean square error. In the zero-shot setting, both samples u_a and u_b come from the source language. While in the few-shot setting, one sample comes from the source language and the other one comes from the target language.

Since the features of labels and utterances are in different vector spaces, we choose not to share the parameters of their encoders. During training, it is easy to produce expressive representations for user utterances due to the large training samples, but it is difficult for label sequences since the objective function L^{lr} is the only supervision. This supervision is weak at the beginning of the training since utterance representations are not sufficiently expressive, which leads to the label regularization approach not being stable and effective. To ensure the representations for slot label sequences are meaningful, we conduct pre-training for the label sequence encoder.

3.1.3 Label Sequence Encoder Pre-training

We leverage the large amount of source language training data to pre-train the label sequence encoder. Concretely, we use the model architecture illustrated in Figure 2 to train the SLU system in the source language, and at the same time, we optimize the label sequence encoder based on the objective function L^{lr} in Eq (8). The label sequence encoder learns to generate meaningful label sequence representations that differ based on their similarities since the extensive source language training samples ensure the high quality of the utterance encoder.

3.2 Adversarial Latent Variable Model

In this section, we first give an introduction to the latent variable model (LVM) (Liu et al., 2019a), and then we describe how we incorporate the adversarial training into the LVM.

3.2.1 Latent Variable Model

Point estimation in the cross-lingual adaptation is vulnerable due to the imperfect alignments across languages. Hence, as illustrated in Figure 2 (Right), the LVM generates a Gaussian distribution with mean μ and variance σ for both word-level and

sentence-level representations instead of a feature vector, which eventually improves the robustness of the model’s cross-lingual adaptation ability. The LVM can be formulated as

$$\begin{bmatrix} \mu_i^S \\ \log(\sigma_i^S)^2 \end{bmatrix} = \mathbf{W}_i^S h_i, \quad \begin{bmatrix} \mu^I \\ \log(\sigma^I)^2 \end{bmatrix} = \mathbf{W}_i^I u, \quad (9)$$

$$z_i^S \sim q_i^S(z|h_i), \quad z^I \sim q^I(z|u), \quad (10)$$

$$p_i^S(s_i|z_i^S) = \text{Softmax}(\mathbf{W}_p^S z_i^S), \quad (11)$$

$$p^I(I|z^I) = \text{Softmax}(\mathbf{W}_p^I z^I), \quad (12)$$

where W_i^S and W_i^I are trainable parameters to generate the mean and variance for word-level hidden states h_i and sentence-level representations r , respectively, from user utterances. $q_i^S \sim \mathcal{N}(\mu_i^S, (\sigma_i^S)^2 \mathbf{I})$ and $q^I \sim \mathcal{N}(\mu^I, (\sigma^I)^2 \mathbf{I})$ are the generated Gaussian distributions, which latent variables z_i^S and z^I are sampled from, and p_i^S and p^I is the predictions for the slot of the i^{th} token and the intent of the utterance, respectively.

During training, all the sampled points from the same generated distribution will be trained to predict the same slot label, which makes the adaptation more robust. In the inference time, the true mean μ_i^S and μ^I is used to replace z_i^S and z^I , respectively, to make the prediction deterministic.

3.2.2 Adversarial Training

Since there are no constraints enforced on the latent Gaussian distribution during training, the latent distributions of different slot types are likely to be close to each other. Hence, the distributions for the same slot type in different user utterances or languages might not be clustered well, which could hurt the cross-lingual alignment and prevent the model from distinguishing slot types when adapting to the target language.

To improve the cross-lingual alignment of latent variables, we propose to make the latent variables of different slot types more distinguishable by adding adversarial training to the LVM. As illustrated in Figure 2 (Right), we train a fully connected layer to fit latent variables into a uniform distribution over slot types. At the same time, the latent variables are regularized to fool the trained fully connected layer by predicting the correct slot type. In this way, the latent variables are trained to be more recognizable. In other words, the generated distributions for different slot types are more likely to repel each other, and for the same slot type are more likely to be close to each other, which

leads to a more robust cross-lingual adaptation. We denote the size of the whole training data as J and the length for data sample j as $|Y_j|$. Note that in the few-shot setting, J includes the number of data samples in the target language. The process of adversarial training can be described as follows:

$$p_{jk} = \mathcal{FC}(z_{jk}^S), \quad (13)$$

$$L^{fc} = \sum_{j=1}^J \sum_{k=1}^{|Y_j|} \text{MSE}(p_{jk}, \mathcal{U}), \quad (14)$$

$$L^{lvm} = \sum_{j=1}^J \sum_{k=1}^{|Y_j|} \text{MSE}(p_{jk}, y_{jk}^S), \quad (15)$$

where \mathcal{FC} consists of a linear layer and a Softmax function, and z_{jk}^S and p_{jk} is the latent variable and generated distribution, respectively, for the k^{th} token in the j^{th} utterance, MSE represents the mean square error, \mathcal{U} represents the uniform distribution, and y_{jk}^S represents the slot label. The slot label is a one-hot vector where the value for the correct slot type is one and zero otherwise. We optimize L^{fc} to train only \mathcal{FC} to fit a uniform distribution, and L^{lvm} is optimized to constrain the LVM to generate more distinguishable distributions for slot predictions. Different from the well-known adversarial training (Goodfellow et al., 2014) where the discriminator is to distinguish the classes, and the generator is to make the features not distinguishable, in our approach, the \mathcal{FC} layer, acting as the discriminator, is trained to generate uniform distribution, and the generator is regularized to make latent variables distinguishable by slot types.

3.3 Optimization

The objective functions for the slot filling and intent detection tasks are illustrated as follows:

$$L^S = \sum_{j=1}^J \sum_{k=1}^{|Y_j|} -\log(p_{jk}^S \cdot (y_{jk}^S)^\top), \quad (16)$$

$$L^I = \sum_{j=1}^J -\log(p_j^I \cdot (y_j^I)^\top), \quad (17)$$

where p_{jk}^S and y_{jk}^S is the prediction and label, respectively, for the slot of the k^{th} token in the j^{th} utterance, and p_j^I and y_j^I is the intent prediction and label, respectively, for the j^{th} utterance.

The optimization for our model is to minimize the following objective function:

$$L = L^S + L^I + L^{lr} + \alpha L^{fc} + \beta L^{lvm}, \quad (18)$$

# Utterance	English	Spanish	Thai
Train	30,521	3,617	2,156
Validation	4,181	1,983	1,235
Test	8,621	3,043	1,692

Table 1: Number of utterances for the multilingual SLU dataset. English is the source language, and Spanish and Thai are the target languages.

where α and β are hyper-parameters, L^{fc} only optimizes the parameters in \mathcal{FC} , and L^{lvm} optimizes all the model parameters excluding \mathcal{FC} .

4 Experiments

4.1 Dataset

We conduct our experiments on the multilingual spoken language understanding (SLU) dataset proposed by Schuster et al. (2019), which contains English, Spanish, and Thai across the weather, reminder, and alarm domains. The corpus includes 12 intent types and 11 slot types, and the data statistics are shown in Table 1.

4.2 Training Details

The utterance encoder is a 2-layer BiLSTM with a hidden size of 250 and dropout rate of 0.1, and the size of the mean and variance in the latent variable model is 150. The label encoder is a 1-layer BiLSTM with a hidden size of 150, and 100-dimensional embeddings for label types. We use the Adam optimizer with a learning rate of 0.001. We use *accuracy* to evaluate the performance of intent detection and BIO-based *f1-score* to evaluate the performance of slot filling. For the adversarial training, we realize that the latent variable model is not able to make slot types recognizable if the \mathcal{FC} is too strong. Hence, we decide to first learn a good initialization for \mathcal{FC} by setting both α and β parameters in Eq (18) as 1 in the first two training epochs, and then we gradually decrease the value of α . We use the refined cross-lingual word embeddings in Liu et al. (2019a)² to initialize the cross-lingual word embeddings in our models and let them not be trainable. We use the delexicalization (delex.) in Liu et al. (2019a), which replaces the tokens that represent numbers, time, and duration with special tokens. We use **36** training samples in Spanish and **21** training samples in Thai on the 1% few-shot setting, and **108** training samples in Spanish and **64** training samples in Thai on the 3%

²Available at <https://github.com/zliuocr/Crosslingual-NLU>

Model	Spanish				Thai			
	Intent Acc.		Slot F1		Intent Acc.		Slot F1	
Few-shot settings								
	1%-shot	3%-shot	1%-shot	3%-shot	1%-shot	3%-shot	1%-shot	3%-shot
BiLSTM-CRF	93.03	93.63	75.70	82.60	81.30	87.23	52.57	66.04
+ LR	93.08	95.04	77.04	84.09	84.04	89.20	57.40	67.45
BiLSTM-LVM	92.86	94.46	75.19	82.64	83.51	89.08	55.08	67.26
+ LR	93.79	95.16	76.96	83.54	86.33	90.80	59.02	70.26
+ ALVM	93.78	95.27	78.35	83.69	85.40	90.70	59.75	69.38
+ LR & ALVM	93.82	95.20	78.46	84.19	87.43	90.96	61.44	70.88
+ LR & ALVM & delex.	94.71	95.62	80.82	85.18	87.67	91.61	62.01	72.39
XL-SLU	92.70	94.96	77.67	82.22	84.04	89.59	55.57	67.56
M-BERT	92.77	95.56	80.15	84.50	83.87	89.19	58.18	67.88
Zero-shot settings								
XL-SLU		90.20		65.79		73.43		32.24
+ LR		91.51		71.55		74.86		32.86
+ ALVM		91.48		71.21		74.35		32.97
+ LR & ALVM		92.31		72.49		75.77		33.28
MLT		86.54		74.43		70.57		28.47
CoSDA-ML		94.80		80.40		76.80		37.3
M-BERT		74.91		67.55		42.97		10.68
Multi. CoVe		53.34		22.50		66.35		32.52
+ Auto-encoder		53.89		19.25		70.70		35.62
Translate Train		85.39		72.89		95.89		55.43
All-shot settings								
Target [†]		96.08		86.03		92.73		85.52
Source & Target [‡]		98.06		87.65		95.58		88.11

Table 2: Cross-lingual SLU results (averaged over three runs). [†]denotes supervised training on all the target language training samples. [‡]denotes supervised training on both the source and target language datasets. The bold numbers denote the best results in the few-shot or zero-shot settings. The underlined numbers represent that the results are comparable (distances are within 1%) to the all-shot experiment with all the target language training samples. The results of Multi. CoVe and Multi. CoVe + Auto-encoder are taken from Schuster et al. (2019), and the results of XL-SLU in the zero-shot settings are taken from Liu et al. (2019a).

few-shot setting. Our models are trained on GTX 1080 Ti. The number of parameters for our models is around 5 million.

4.3 Baselines

We compare our model to the following baselines.

BiLSTM-CRF This is the same cross-lingual SLU model structure as Schuster et al. (2019).

BiLSTM-LVM We replace the conditional random field (CRF) in BiLSTM-CRF with the LVM proposed in Liu et al. (2019a).

Multi. CoVe Multilingual CoVe (Yu et al., 2018) is a bidirectional machine translation system that tends to encode phrases with similar meanings into similar vector spaces across languages. Schuster et al. (2019) used it for the cross-lingual SLU task.

Multi. CoVe w/ auto-encoder Based on Multilingual CoVe, Schuster et al. (2019) added an auto-encoder objective so as to produce better-aligned

representations for semantically similar sentences across languages.

Multilingual BERT (M-BERT) It is a single language model pre-trained from monolingual corpora in 104 languages (Devlin et al., 2019), which is surprisingly good at cross-lingual model transfer.

Mixed Language Training (MLT) Liu et al. (2019b) utilized keyword pairs to generate mixed language sentences for training cross-lingual task-oriented dialogue systems, which achieves promising zero-shot transfer ability.

CoSDA-ML Qin et al. (2020) proposed a multilingual code-switching data augmentation framework to enhance the cross-lingual systems based on M-BERT (Devlin et al., 2019). It is a concurrent work of this paper.

XL-SLU It is a previous state-of-the-art model in the zero-shot cross-lingual SLU task, which combines Gaussian noise, cross-lingual embeddings

refinement, and the LVM (Liu et al., 2019a).

Translate Train Schuster et al. (2019) trained a supervised machine translation system to translate English data into the target language, and then trained the model on the translated dataset.

All-shot Settings We train the BiLSTM-CRF model (Lample et al., 2016) on all the target language training samples, and on both the source and target language training set.

5 Results & Discussion

5.1 Few-shot Setting

Quantitative Analysis The few-shot results are illustrated in Table 2, from which we can clearly see consistent improvements made by label regularization and adversarial training. For example, on the 1% few-shot setting, our model improves on BiLSTM-LVM in terms of accuracy/f1-score by 1.85%/1.16% in Spanish, and by 4.16%/6.93% in Thai. Our model also surpasses a strong baseline, M-BERT, while our model based on BiLSTM has many fewer parameters compared to M-BERT. For example, on the 1% few-shot setting, our model improves on M-BERT in terms of accuracy/f1-score by 3.80%/3.83% in Thai. Instead of generating a feature point like CRF, the LVM creates a more robust cross-lingual adaptation by generating a distribution for the intent or each token in the utterance. However, distributions generated by the LVM for the same slot type across languages might not be sufficiently close. Incorporating adversarial training into the LVM alleviates this problem by regularizing the latent variables and making them more distinguishable. This improves the performance in both intent detection (a sentence-level task) and slot filling (a word-level task) by 0.92%/3.16% in Spanish and by 1.89%/4.67% in Thai on the 1% few-shot setting. This proves that both sentence-level and word-level representations are better aligned across languages.

In addition, LR aims to further align the sentence-level representations of target language utterances into a semantically similar space of source language utterances. As a result, there are 0.93%/2.82% improvements in intent detection for Spanish/Thai on the 1% few-shot setting after we add LR to BiLSTM-LVM. Interestingly, the performance gains are not only on the intent detection but also on the slot filling, with an improvement of 1.77%/3.94% in Spanish/Thai. This is attributed to

Model	Thai	
	Intent	Slot
<i>few-shot on 5% target language training set</i>		
BiLSTM-CRF	90.05	72.11
+ LR	91.11	73.71
BiLSTM-LVM	91.02	73.11
+ LR	91.45	75.18
+ ALVM	91.08	74.67
+ LR & ALVM	91.58	75.87
+ LR & ALVM & delex.	92.51	77.03
XL-SLU	91.05	73.43
M-BERT	92.02	75.52

Table 3: Results of few-shot learning on 5% Thai training data, which are averaged over three runs. We make the training samples in Thai the same as the 3% Spanish training samples (108).

the fact that utterance representations are produced based on word-level representations from BiLSTM. Therefore, the alignment of word-level representations will be implicitly improved in this process. Furthermore, incorporating LR and ALVM further tackles the inherent difficulties for the cross-lingual adaptation and achieves the state-of-the-art few-shot performance. Notably, by only leveraging 3% of target language training samples, the results of our best model are on par with the supervised training on all the target language training data.

Adaptation ability to unrelated languages

From Table 2, we observe impressive improvements in Thai, an unrelated language to English, by utilizing our proposed approaches, especially when the number of target language training samples is small. For example, compared to the BiLSTM-LVM, our best model significantly improves the accuracy and f1-score by $\sim 4\%$ / $\sim 7\%$ in intent detection and slot filling in Thai in the few-shot setting on 1% data. Additionally, in the same setting, our model surpasses the strong baseline, M-BERT, in terms of accuracy and f1-score by $\sim 4\%$. This illustrates that our approaches provide strong adaptation robustness and are able to tackle the inherent adaptation difficulties to unrelated languages.

Comparison between Spanish and Thai

To make a fair comparison for the few-shot performance in Spanish and Thai, we increase the training size of Thai to the same as 3% Spanish training samples, as depicted in Table 3. We can see that there is still a performance gap between the Spanish and Thai (3.11% in the intent detection task and 8.15% in the slot filling task). This is because

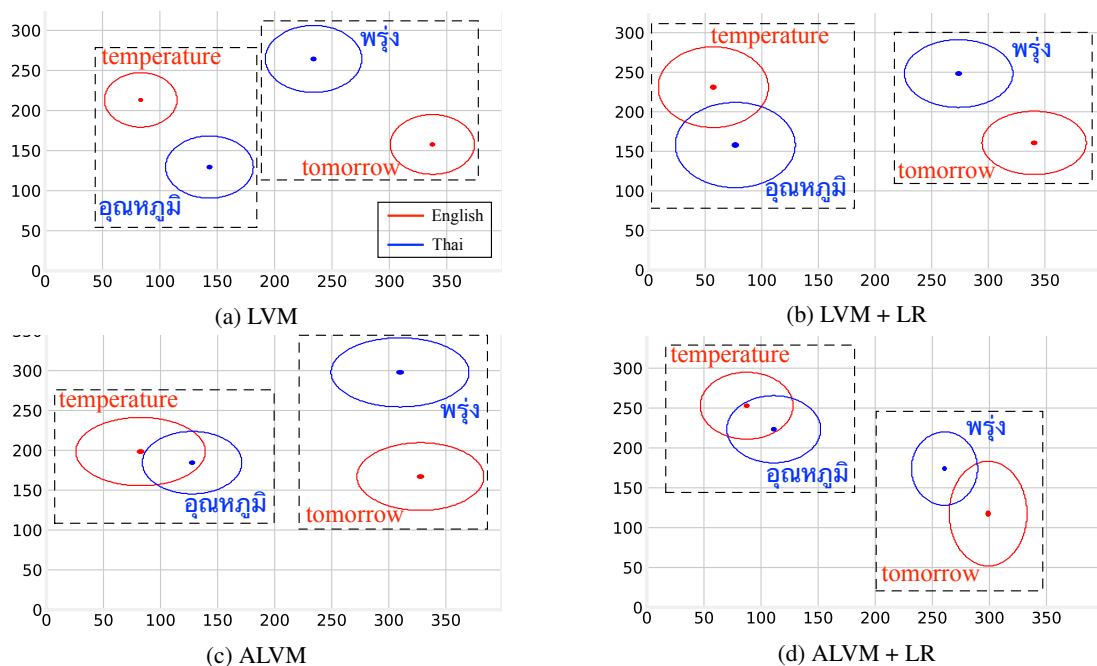


Figure 3: Visualization for latent variables of parallel word pairs in English and Thai over different models trained on 1% target language training set. We choose the word pairs “temperature- อุณหภูมิ” and “tomorrow- พรุ่ง” from the parallel sentences “what will be the temperature tomorrow” and “อุณหภูมิ จะ อยู่ ท เท่า ไหร่ พรุ่ง” in English and Thai, respectively. To draw the contour plot, we sample 3000 points from the distribution of latent variables for the selected words, use PCA to project those points into 2D and calculate the mean and variance for each word.

Model	Spanish		Thai	
	Intent	Slot	Intent	Slot
<i>few-shot on 1% target language training set</i>				
Our Model	93.82	78.46	87.43	62.44
w/o Pre-training	92.75	77.11	86.29	60.20
<i>few-shot on 3% target language training set</i>				
Our Model	95.20	84.19	90.97	70.88
w/o Pre-training	94.51	82.83	89.72	69.66
<i>zero-shot setting</i>				
Our Model	92.31	72.49	75.77	33.28
w/o Pre-training	91.02	71.72	75.18	32.69

Table 4: Results of the ablation study for the label sequence encoder pre-training (averaged over three runs). Our model refers to the one that combines LR, ALVM and delex. with BiLSTM-LVM.

Spanish is grammatically and syntactically closer to English than Thai, leading to a better quality of cross-lingual alignment.

Visualization of Latent Variables The effectiveness of the LR and ALVM can be clearly seen from Figure 3. The former approach decreases the distance of latent variables for words with similar semantic meanings in different languages. For the latter approach, to make the distributions for different slot types distinguishable, our model reg-

ularizes the latent variables of different slot types far from each other, and eventually it also improves the alignment of words with the same slot type. Incorporating both approaches further improves the word-level alignment across languages. It further proves the robustness of our proposed approaches when adapting from the source language (English) to the unrelated language (Thai).

5.2 Zero-shot Setting

From Table 2, we observe the remarkable improvements made by LR and ALVM on the state-of-the-art model XL-SLU in the zero-shot setting, and the slot filling performance of our best model in Spanish is on par with the strong baseline Translate Train, which leverages large amounts of bilingual resources. LR improves the adaptation robustness by making the word-level and sentence-level representations of similar utterances distinguishable. In addition, integrating adversarial training with the LVM further increases the robustness by disentangling the latent variables for different slot types. However, the performance boost for slot filling in Thai is limited. We conjecture that the inherent discrepancies in cross-lingual word embeddings and language structures for topologically different lan-

guages pairs make the word-level representations between them difficult to align in the zero-shot scenario. We notice that Multilingual CoVe with auto-encoder achieves slightly better performance than our model on the slot filling task in Thai. This is because this baseline leverages large amounts of monolingual and bilingual resources, which largely benefits the cross-lingual alignment between English and Thai. CoSDA-ML, a concurrent work of our model, utilizes additional augmented multilingual code-switching data, which significantly improves the zero-shot cross-lingual performance.

5.3 Effectiveness of Label Sequence Encoder Pre-training

Label sequence encoder pre-training helps the label encoder to generate more expressive representations for label sequences, which ensures the effectiveness of the label regularization approach. From Table 4, we can clearly observe the consistent performance gains made by pre-training in both few-shot and zero-shot scenarios.

6 Conclusion

Current cross-lingual SLU models still suffer from imperfect cross-lingual alignments between the source and target languages. In this paper, we propose label regularization (LR) and the adversarial latent variable model (ALVM) to regularize and further align the word-level and sentence-level representations across languages without utilizing any additional bilingual resources. Experiments on the cross-lingual SLU task illustrate that our model achieves a remarkable performance boost compared to the strong baselines in both zero-shot and few-shot scenarios, and our model has a robust adaptation ability to unrelated target languages in the few-shot scenario. In addition, visualization for latent variables further proves that our approaches are effective at improving the alignment of cross-lingual representations.

Acknowledgments

This work is partially funded by ITF/319/16FP and MRP/055/18 of the Innovation Technology Commission, the Hong Kong SAR Government.

References

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2017. Learning bilingual word embeddings with (almost)

no bilingual data. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462.

Ankur Bapna, Gokhan Tür, Dilek Hakkani-Tür, and Larry Heck. 2017. Towards zero-shot frame semantic parsing for domain scaling. *Proc. Interspeech 2017*, pages 2476–2480.

Qian Chen, Zhu Zhuo, and Wen Wang. 2019. Bert for joint intent classification and slot filling. *arXiv preprint arXiv:1902.10909*.

Wenhu Chen, Jianshu Chen, Yu Su, Xin Wang, Dong Yu, Xifeng Yan, and William Yang Wang. 2018. Xlnbt: A cross-lingual neural belief tracking framework. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 414–424.

Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems 32*, pages 7057–7067. Curran Associates, Inc.

Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *International Conference on Learning Representations (ICLR)*.

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.

Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1615–1625.

Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. Slot-gated modeling for joint slot filling and intent prediction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 753–757.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680.

E Haihong, Peiqing Niu, Zhongfu Chen, and Meina Song. 2019. A novel bi-directional interrelated model for joint intent detection and slot filling. In

- Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5467–5471.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. 2019. Unicoder: A universal language encoder by pre-training with multiple cross-lingual tasks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2485–2494.
- Joo-Kyung Kim, Young-Bum Kim, Ruhi Sarikaya, and Eric Fosler-Lussier. 2017. Cross-lingual transfer learning for pos tagging without cross-lingual resources. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2832–2838.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270.
- Zhaojiang Lin, Zihan Liu, Genta Indra Winata, Samuel Cahyawijaya, Andrea Madotto, Yejin Bang, Etsuko Ishii, and Pascale Fung. 2020. Xpersona: Evaluating multilingual personalized chatbot. *arXiv preprint arXiv:2003.07568*.
- Zihan Liu, Jamin Shin, Yan Xu, Genta Indra Winata, Peng Xu, Andrea Madotto, and Pascale Fung. 2019a. Zero-shot cross-lingual dialogue systems with transferable latent variables. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1297–1303.
- Zihan Liu, Genta Indra Winata, Zhaojiang Lin, Peng Xu, and Pascale Fung. 2019b. Attention-informed mixed-language training for zero-shot cross-lingual task-oriented dialogue systems. *arXiv preprint arXiv:1911.09273*.
- Zihan Liu, Genta Indra Winata, Andrea Madotto, and Pascale Fung. 2020a. Exploring fine-tuning techniques for pre-trained cross-lingual models via continual learning. *arXiv preprint arXiv:2004.14218*.
- Zihan Liu, Genta Indra Winata, Peng Xu, and Pascale Fung. 2020b. **Coach: A coarse-to-fine approach for cross-domain slot filling**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 19–25, Online. Association for Computational Linguistics.
- Stephen Mayhew, Chen-Tse Tsai, and Dan Roth. 2017. Cheap translation for cross-lingual named entity recognition. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 2536–2545.
- Nikola Mrkšić, Ivan Vulić, Diarmuid Ó Séaghdha, Ira Leviant, Roi Reichart, Milica Gašić, Anna Korhonen, and Steve Young. 2017. Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints. *Transactions of the Association for Computational Linguistics*, 5(1):309–324.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. **How multilingual is multilingual BERT?** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Libo Qin, Minheng Ni, Yue Zhang, and Wanxiang Che. 2020. Cosda-ml: Multi-lingual code-switching data augmentation for zero-shot cross-lingual nlp. *arXiv preprint arXiv:2006.06402*.
- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2019. Cross-lingual transfer learning for multilingual task oriented dialog. 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 3795–3805.
- Avirup Sil, Gourab Kundu, Radu Florian, and Wael Hamza. 2018. Neural cross-lingual entity linking. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Shyam Upadhyay, Manaal Faruqi, Gokhan Tür, Hakkani-Tür Dilek, and Larry Heck. 2018a. (almost) zero-shot cross-lingual spoken language understanding. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6034–6038. IEEE.
- Shyam Upadhyay, Nitish Gupta, and Dan Roth. 2018b. Joint multilingual supervision for cross-lingual entity linking. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2486–2495.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. Transferable multi-domain state generator for task-oriented dialogue systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 808–819.
- 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.

- Katherine Yu, Haoran Li, and Barlas Oguz. 2018. Multilingual seq2seq training with similarity loss for cross-lingual document classification. In *Proceedings of The Third Workshop on Representation Learning for NLP*, pages 175–179, Melbourne, Australia. Association for Computational Linguistics.
- Tao Zhang, Kang Liu, and Jun Zhao. 2013. Cross lingual entity linking with bilingual topic model. In *Twenty-Third International Joint Conference on Artificial Intelligence*.
- Yuan Zhang, David Gaddy, Regina Barzilay, and Tommi Jaakkola. 2016. Ten pairs to tag—multilingual pos tagging via coarse mapping between embeddings. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1307–1317.