

# Emu: Enhancing Multilingual Sentence Embeddings with Semantic Specialization

Wataru Hirota,<sup>1\*</sup> Yoshihiko Suhara,<sup>2</sup> Behzad Golshan,<sup>2</sup> Wang-Chiew Tan<sup>2</sup>

<sup>1</sup>Osaka University, <sup>2</sup>Megagon Labs  
w-hirota@ist.osaka-u.ac.jp, {yoshi, behzad, wangchiew}@megagon.ai

## Abstract

We present EMU, a system that semantically enhances multilingual sentence embeddings. Our framework fine-tunes pre-trained multilingual sentence embeddings using two main components: a semantic classifier and a language discriminator. The semantic classifier improves the semantic similarity of related sentences, whereas the language discriminator enhances the multilinguality of the embeddings via multilingual adversarial training. Our experimental results based on several language pairs show that our specialized embeddings outperform the state-of-the-art multilingual sentence embedding model on the task of cross-lingual intent classification using only monolingual labeled data.

## Introduction

Learning multilingual sentence representations (Ruder et al. 2019) is a key technique for building NLP applications with multilingual support. A primary advantage of multilingual sentence embeddings is that they enable us to train a single classifier based on a single language (e.g., English) and then apply it to other languages without using training models for those languages (e.g., German.) Furthermore, recent advances in multilingual sentence embedding techniques (Artetxe and Schwenk 2019b; Chidambaram et al. 2019) have shown to exhibit competitive performance on several downstream NLP tasks, compared to the two-stage approach that relies on machine translation followed by monolingual sentence embedding techniques.

The main challenge of multilingual sentence embeddings is that they are sensitive to textual similarity (*textual similarity bias*) which negatively affects the semantic similarity of sentence embeddings (Zhu, Li, and de Melo 2018). The following example illustrates this point:

- S1: *What time is the pool open tonight?*  
S2: *What time are the stores on 5th open tonight?*  
S3: *When does the pool open this evening?*

S1 and S3 have similar intents. They ask for the opening hours of the pool in the evening. S2 has a different intent:

\*This work was done during an internship at Megagon Labs. Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

it asks about the opening hour of stores. We expect embeddings of sentences of the same intent to be closer (e.g., to have higher cosine similarity) to one another than embeddings of sentences with different intents.

We tested several pre-trained (multilingual) sentence embedding models (Pagliardini, Gupta, and Jaggi 2018; Conneau et al. 2017; Artetxe and Schwenk 2019b; Chidambaram et al. 2019) in both monolingual and cross-lingual settings. Somewhat surprisingly, every model provided lower similarity scores between S1 and S3 (compared to S1 and S2, or S2 and S3). This is mainly because S1 and S2 are more textually similar (because both sentences contain “what time” and “tonight”) compared to S1 and S3. This example highlights that general-purpose multilingual sentence embeddings exhibit textual similarity bias, which is a fundamental limitation as they may not correctly capture the semantic similarity of sentences.

Motivated by the need for sentence embeddings that better reflect the semantics of sentence, we examine *multilingual semantic specialization*, which tailors pre-trained multilingual sentence embeddings to handle semantic similarity. Although prior work has developed *semantic specialization* methods for word embeddings (Mrkšić et al. 2017) and semantic and linguistic properties of sentence embeddings (Zhu, Li, and de Melo 2018; Conneau et al. 2018a), no prior work has considered semantic specialization of multilingual sentence embeddings.

In this paper, we develop a “lightweight” approach for semantic specialization of multilingual embeddings that can be applied to any base model. Our approach fine-tunes a pre-trained multilingual sentence embedding model based on a classification task that considers semantic similarity. This aligns with common techniques of pre-training methods for NLP (Howard and Ruder 2018; Peters et al. 2018; Devlin et al. 2019). We explore several loss functions to determine which is appropriate for the semantic specialization of cross-lingual sentence embeddings. We found that naive choices of loss functions such as the softmax loss, which is a common choice for classification, may suffer from significant degradation of the original multilingual sentence embedding model.

We also design EMU to specialize multilingual sentence

embeddings using only monolingual training data as it is expensive to collect parallel training data in multiple languages. Our solution incorporates *language adversarial training* to enhance the multilinguality of sentence embeddings. Specifically, we implemented a language discriminator that tries to identify the language of an input sentence given its embedding and optimizes multilingual sentence embeddings to confuse the language discriminator.

We conducted experiments on three cross-lingual intent classification tasks that involves 6 languages. The results show that EMU successfully specializes the state-of-the-art multilingual sentence embedding techniques, namely LASER, using only monolingual training data with unlabeled data in other languages. It outperforms the original LASER model and monolingual sentence embeddings with machine translation by up to 47.7% and 86.2% respectively.

The contributions of the paper are as follows:

- We developed EMU, a system that semantically enhances pre-trained multilingual sentence embeddings<sup>1</sup>. EMU incorporates multilingual adversarial training on top of fine-tuning to enhance multilinguality without using parallel sentences.
- We experimented with several loss functions and show that the two loss functions, namely  $L_2$  constrained softmax and center loss, outperform common loss functions used for fine-tuning.
- We show that EMU successfully specializes multilingual sentence embedding using only monolingual labeled data.

## Multilingual Semantic Specialization

The architecture of EMU is depicted in Figure 1. There are three main components, which we detail next: multilingual encoder  $E$ , semantic classifier  $C$ , and language discriminator  $D$ . The solid lines show the flow of the forward propagation for fine-tuning  $C$  and  $E$ , and the dotted lines are that for  $D$ . These arrows become reversed during the backpropagation. The semantic classifier and language discriminator are only used for fine-tuning.

After fine-tuning, EMU uses the fine-tuned multilingual encoder to obtain sentence embeddings for input sentences. More specifically, we expect the similarity (e.g., cosine similarity) between two related sentences in any languages to be closer to each other. We consider cosine similarity as it is the most common choice and can be calculated efficiently (Wang et al. 2017).

## Multilingual Encoder

A multilingual encoder is a language-agnostic sentence encoder that converts sentences in any language into embedding vectors in a common space. EMU is flexible with the choice of multilingual encoders and their architectures. The only requirement of this component is that it encodes a sentence in any language into a sentence embedding.

In this paper, we use LASER (Artetxe and Schwenk 2019b) as a base multilingual sentence embedding model. LASER is a multilingual sentence embedding model that

<sup>1</sup>Our code is available at <https://github.com/megagonlabs/emu>.

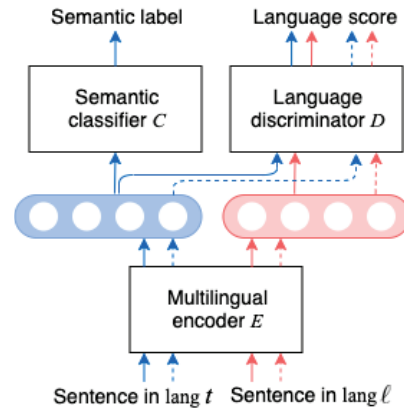


Figure 1: Architecture of EMU.

covers more than 93 languages with more than 23 different alphabets. It is an encoder-decoder model that shares the same BiLSTM encoder with max-pooling and uses byte pair encoding (BPE) (Sennrich, Haddow, and Birch 2016) to accept sentences in any languages as input. The model is trained on a set of bilingual translation tasks and is shown to have the state-of-the-art performance on cross-lingual NLP tasks including bitext mining. We use LASER instead of multilingual models for BERT (Devlin et al. 2019) because (1) LASER outperformed the BERT model on the XNLI task (Artetxe and Schwenk 2019b) and (2) a LASER model can be used as a sentence encoder without any changes<sup>2</sup>.

## Semantic Classifier

The semantic classifier categorizes input sentences into groups that share the same intent, such as “seeking pool information” or “seeking restaurant information”. We expect the semantic classifier to enhance multilingual sentence embeddings to better reflect the semantic similarity of related sentences, where the semantic similarity is calculated as the cosine similarity between the embeddings of the two sentences.

Additionally, we expect that learned embeddings retain semantic similarity with respect to cosine similarity. Thus, we propose the use of  $L_2$ -constrained softmax loss (Ranjan, Castillo, and Chellappa 2017) and center loss (Wen et al. 2016), which are known to be effective for image recognition tasks. To the best of our knowledge, we are the first to apply these loss functions for fine-tuning embedding models. We describe these loss functions next.

**$L_2$ -constrained softmax loss**  $L_2$ -constrained softmax loss (Ranjan, Castillo, and Chellappa 2017) considers hard constraints on the norm of embedding vectors on top of the

<sup>2</sup>A BERT model needs to be fine-tuned to use the first vector corresponding to the class symbol [CLS] as a sentence embedding. A BERT variant for sentence embeddings also needs supervision to train a model (Reimers and Gurevych 2019).

softmax loss:

$$\begin{aligned} \text{minimize} \quad & -\frac{1}{M} \sum_{i=1}^M \log \frac{e^{W_{y_i}^T u_i + b_{y_i}}}{\sum_{j=1}^C e^{W_j^T u_i + b_j}} \\ \text{subject to} \quad & \|u_i\|_2 = \alpha, \forall i = 1, \dots, M, \end{aligned}$$

where  $M$  denotes the number of classes, and  $u_i$  and  $y_i$  are  $i$ -th sentence embedding vector and its true label respectively.

The  $L_2$  constraint ensures that embedding vectors are distributed on the hypersphere with the size of  $\alpha$ . Therefore, the Euclidean distance between two vectors on the hypersphere is approximately close to its cosine distance. This property is helpful for specializing sentence embeddings to learn semantic similarity in the form of cosine similarity. Note that this  $L_2$ -constraint is different from the  $L_2$  regularization term applied to the weight parameters of the output layer. In that case, the regularization term will be considered in the loss function.

To implement  $L_2$ -constrained softmax loss, the model additionally inserts an  $L_2$ -normalized layer that normalizes the encoder output  $u$  (i.e.,  $\frac{u}{\|u\|}$ ) followed by a layer that scales with a hyper-parameter  $\alpha$ . The scaled vectors are then fed into the output layer, where the model evaluates the softmax loss.

**Center loss** The center loss (Wen et al. 2016) was originally developed for face recognition tasks to stabilize deep features learned from data. The center loss is described as follows:

$$L_{center} = \frac{1}{2} \sum_{i=1}^m \|u_i - c_{y_i}\|_2^2, \quad (1)$$

where  $c_{y_i}$  denotes the centroid of sentence embedding vectors of class  $y_i$ . The loss function forces the embedding vector of  $i$ -th sample toward the centroid of the true category. Our motivation to use this loss function is to enhance the *intra-class compactness* of sentence embeddings. That is, we want to ensure that the sentence embeddings that have the same intent form compact clusters because other loss functions, such as the softmax loss, does not have this functionality. The center loss works as *cross-lingual* center loss; it enforces sentences, in any language, that belong to the same intent as a same cluster if multilingual training data are available.

We consider combining the center loss with another function with a hyper-parameter  $\lambda$ :

$$L_C = L_{L_2-sm} + \lambda L_{center}, \quad (2)$$

where  $L_{L_2-sm}$  denotes the  $L_2$ -constrained softmax loss function.

## Language Discriminator

The semantic classifier does not directly consider *multilinguality*, so the model, which is fine-tuned on a single language, may now perform worse on other languages. To avoid this problem, we incorporate *multilingual adversarial learning* into the framework. Specifically, the language discriminator  $D$  aims to identify the language of an input sentence

Table 1: Statistics of the datasets.

	<i>HotelQA</i>	<i>ATIS</i>	<i>Quora</i>
# of classes	28	13	50
# of training data	676	1,195	1,059
# of test data	144	252	353
Vocab. size (en)	977	626	1,308

given its embedding, whereas the multilingual sentence encoder  $E$  incorporates an additional loss function to “confuse”  $D$ . The idea was inspired by related work that used adversarial learning for multilingual NLP models (Chen et al. 2018; Chen and Cardie 2018). We hypothesize and our experiments show that incorporating adversarial learning also enhances the multilinguality of sentence embeddings.

The language discriminator is trained to determine whether the languages of two input embeddings are different. Simultaneously, the other part of the model is trained to confuse the discriminator. In our implementation, we use Wasserstein GAN (Arjovsky, Chintala, and Bottou 2017) because it is known to be more robust than the original GAN (Goodfellow et al. 2014).

Algorithm 1 shows a single training step of EMU. Each step consists of two training routines for language discriminator  $D_t$  and the other components (multilingual sentence encoder  $E$  and semantic classifier  $C$ ). Target language  $t$  denotes the language used for training (e.g., English).  $t$  is randomly chosen from a training language set if multiple languages are used for training. *Adversarial languages*  $\mathcal{L}$  is a set of languages that are used to retrieve adversarial sentences. To train language discriminator  $D_t$ , training sentences in language  $t$  and adversarial sentences from randomly chosen language  $\ell \in \mathcal{L}$  are used to evaluate  $L_{D_t}$ . Formally, the loss function for any training language  $t$  is described as

$$L_{D_t} = L_d(1, D_t(u^t)) + L_d(0, D_t(v^\ell)), \quad (3)$$

where  $L_d(\cdot, \cdot)$  is the cross entropy loss,  $u_t$  and  $v_\ell$  are embedding vectors (encoded by  $E$ ) of sentences in language  $t$  and language  $\ell$  ( $t \neq \ell$ ). Our design implements a language discriminator for each training language  $t$ . For instance, language discriminator  $D_{t=en}$  aims to predict whether an input multilingual sentence embedding belongs to English.

Next, labeled sentences in language  $t$  and adversarial sentences  $\ell$  are sampled to update the parameters of  $E$  and  $C$  with the fixed parameters of  $D_t$ . The overall loss function  $L_{C+D_t}$  now takes into account the loss value of  $D_t$  so that the multilingual encoder  $E$  can generate multilingual sentences embeddings for sentences in languages  $t$  and  $\ell$ , which cannot be classified by the language discriminator  $D_t$ . We use hyper-parameter  $\gamma$  to balance the loss functions:

$$L_{C+D_t} = L_C - \gamma L_{D_t}. \quad (4)$$

## Evaluation

We evaluated EMU based on the cross-lingual intent classification task. The task is to detect the intent of an input sentence in a source language (e.g., German) based on *labeled*

---

**Algorithm 1** Single Training Step of EMU

---

**Require:** Training lang  $t$ , adversarial langs  $\mathcal{L}$ , iteration number  $k$ , clipping interval  $c$ .

```
1: for 1 to  $k$  do
2:   Sample training sentences as  $x^t$ 
3:   Sample adversarial language  $\ell$  from  $\mathcal{L}$ 
4:   Sample adversarial sentences as  $x^\ell$ 
5:    $u^t \leftarrow E(x^t)$ ;  $v^\ell \leftarrow E(x^\ell)$ 
6:   Evaluate loss  $L_{D_t}(u^t, v^\ell)$  ▷ Eq. 3
7:   Update  $D_t$  parameters
8:   Clip  $D_t$  parameters to  $[-c, c]$ 
9:   Sample training sentences and labels as  $x^t$  and  $y^t$ 
10:  Sample adversarial language  $\ell$  from  $\mathcal{L}$ 
11:  Sample adversarial sentences as  $x^\ell$ 
12:   $u^s \leftarrow E(x^t)$ ;  $v^\ell \leftarrow E(x^\ell)$ 
13:  Evaluate loss  $L_{C+D_t}(u^t, v^\ell, y^t)$  ▷ Eq. 4
14:  Update  $E$  and  $C$  parameters
```

---

sentences associated with intent labels in a target language (e.g., English.) We consider similarity-based intent detection, which categorizes an input sentence based on the label of the *nearest neighbor* sentence that has the highest cosine similarity against the input sentence. We adopted this evaluation method since it is widely used in search-based QA systems (Paşca 2003) and works robustly especially if training data are sparse. An intuitive alternative for intent detection is to directly use the trained semantic classifier (see Figure 1). We evaluated the classification results using the semantic classifier but the performance was poor. Therefore, we excluded the results from the tables.

## Dataset

We used three datasets for evaluation. Some statistics of these datasets are shown in Table 1.

**HotelQA** is a real-world private corpus of 820 questions collected via a multi-channel communication platform for hotel guests and hotel staff. Questions are always made by guests and have ground truth labels for 28 intent classes (e.g., check-in, pool.) The utterances are professionally translated into 5 non-English languages (German (de), Spanish (es), French (fr), Japanese (ja), and Chinese (zh).) We split the dataset into training and test sets so that the sentences used for fine-tuning do not appear in the test set.

**ATIS** (Hemphill, Godfrey, and Doddington 1990) is a publicly available corpus for spoken dialog systems and is widely used for intent classification research. The dataset consists of more than 5k sentences and 22 intent labels are assigned to each sentence. We excluded the “flights” class from the dataset since the class accounts for about 75% of the dataset. We also ensured that each class has at least 5 sentences in each of train and test datasets. As a result, 13 classes remained in the dataset. Similar to previous studies (Conneau et al. 2018b; Glavas et al. 2019), we used Google Translate to generate corresponding translations in the same 5 non-English languages as *HotelQA*.

**Quora**<sup>3</sup> is a publicly available paraphrase detection dataset that contains over 400k questions with duplicate labels. Each row is a pair of questions with a duplicate label. Duplicate questions can be considered sentences that belong to the same intent. Therefore, we created a graph where each node is a question and an edge between two nodes denotes that these questions are considered duplicate. By doing this, we can consider each disjoint clique in the graph as a single intent class. Specifically, we filtered only complete subgraphs whose size (i.e., # of nodes) is less than 30 to avoid having extremely large clusters that are too general. We chose the 50 largest clusters after the filtering. The original dataset contains only English sentences. We used Google Translate to translate into the same 5 languages in the same manner as *ATIS*.

## Baselines

**MT + sent2vec** We consider the two-stage approach that uses machine translation and monolingual sentence embeddings in a pipeline<sup>4</sup> We used Google Translate for translation and sent2vec (Pagliardini, Gupta, and Jaggi 2018) as a baseline method<sup>5</sup>.

**Softmax loss** Softmax loss is the most common loss function for classification, and thus a natural choice for fine-tuning the embeddings. We used the softmax loss function to train the semantic classifier and adjust the embeddings.

**Contrastive loss** Contrastive loss (Chopra et al. 2005) is a widely used pairwise loss function for metric learning. The loss function minimizes the squared distance between two embeddings if the labels are the same, and it maximizes the margin (we used  $m = 2.0$ ) between two samples otherwise. For contrastive loss, we use the Siamese (i.e., dual-encoder) architecture (Chopra et al. 2005) that takes two input sentences that will be fed into a shared encoder (i.e., multilingual encoder  $E$ ) to obtain sentence embeddings.

**N-pair loss** As another metric learning method, we used the N-pair sampling cosine loss (Yang et al. 2019), which first samples one positive sample and  $N - 1$  negative samples and then minimizes a cosine similarity-based loss function.

## Experimental Settings

For each dataset, we used only English training data to fine-tune the models with EMU and the baseline methods. To train EMU’s language discriminator, we used unlabeled training data in other non-English languages (i.e., de, es, fr, ja, zh.)

**Emu variants** To verify the effect of the language discriminator and the center loss, we also evaluated EMU without the language discriminator (EMU w/o LD) and EMU without the language discriminator or the center loss (EMU w/o

---

<sup>3</sup><https://data.quora.com/First-Quora-Dataset-Release-Question-Pairs>

<sup>4</sup>The non-English sentences obtained through MT from English had to be translated back to English.

<sup>5</sup>We tested the official implementation of InferSent (Conneau et al. 2017), finding that performance was unstable and often significantly lower than that of sent2vec. Thus, we decided to use sent2vec in the experiments.

Table 2: Experimental results (Acc@1) on three dataset. The highest performance (excluding EMU-PARALLEL) is in bold and the highest performance by EMU-PARALLEL is underlined. \*, \*\*, and \*\*\* denote  $p$ -value < 0.01, 0.05, and 0.10 respectively based on the binomial proportion confidence intervals of Acc@1 values against the baseline methods.

(a) <i>HotelQA</i>												
			EN → *					* → EN				
	Method	en-en	de	es	fr	ja	zh	de	es	fr	ja	zh
Baseline	MT + sent2vec	48.6	41.0	35.4	34.7	47.2	43.1	46.5	47.2	44.4	48.6	41.7
	LASER (original)	55.6	45.1	48.6	48.6	47.9	45.1	43.8	45.8	50.7	44.4	49.3
	Contrastive loss	34.0	19.4	12.5	22.9	25.7	21.5	24.3	18.8	25.7	24.3	20.1
	N-pair loss	27.8	20.8	22.9	21.5	20.8	21.5	24.3	24.3	25.7	25.0	20.1
	Softmax loss	30.6	13.9	13.9	7.6	8.3	7.6	13.2	24.3	16.0	20.8	13.9
Proposed	EMU	<b>78.5</b> ***	<b>66.7</b> ***	<b>66.0</b> ***	63.2***	63.2***	<b>62.5</b> ***	<b>56.9</b> ***	<b>62.5</b> ***	<b>58.3</b> ***	53.5**	<b>59.7</b> **
	EMU w/o LD	76.4***	63.2***	59.7***	<b>65.3</b> ***	<b>66.7</b> ***	56.9***	55.6***	61.8***	56.9	54.2**	58.3**
	EMU w/o LD+CL	77.1***	62.5***	60.4***	61.8***	68.1***	58.3***	58.3***	<b>62.5</b> ***	56.9***	<b>55.6</b> ***	58.3**
	EMU-PARALLEL	<u>79.2</u> ***	<u>68.1</u> ***	65.3***	<u>65.3</u> ***	59.0***	61.8***	<u>59.7</u> ***	61.1***	<u>59.0</u> **	48.6	58.3**
(b) <i>ATIS</i>												
			EN → *					* → EN				
	Method	en-en	de	es	fr	ja	zh	de	es	fr	ja	zh
Baseline	MT + sent2vec	90.5	87.3	89.7	87.7	2.4	7.1	84.9	84.9	86.1	80.6	81.7
	LASER (original)	88.5	86.5	84.1	81.3	85.3	87.7	87.7	87.7	85.7	86.5	<b>87.7</b>
	Contrastive loss	83.3	62.3	67.9	63.9	44.0	52.8	66.7	69.4	64.3	57.9	59.1
	N-pair loss	81.0	57.1	49.2	52.0	30.2	42.1	58.3	57.9	55.2	41.3	41.3
	Softmax loss	90.5	48.0	63.5	52.0	56.0	52.0	50.0	46.0	45.2	35.7	39.7
Proposed	EMU	<b>98.8</b> ***	<b>98.4</b> ***	<b>98.4</b> ***	93.7***	<b>98.4</b> ***	<b>97.6</b> ***	<b>95.6</b> ***	<b>94.8</b> ***	<b>94.0</b> ***	<b>89.3</b>	84.5
	EMU w/o LD	97.6***	98.0***	96.8***	94.8***	95.6***	93.7***	92.5***	83.3	88.5	84.5	87.3
	EMU w/o LD+CL	98.4***	97.6***	98.0***	<b>96.4</b> ***	94.0***	95.6***	91.7**	84.5	87.7	82.1	86.1
	EMU-PARALLEL	<u>99.2</u> ***	<u>98.4</u> ***	98.0***	95.2***	97.2***	96.0***	95.2***	<u>95.6</u> ***	93.7***	77.8	83.7
(c) <i>Quora</i>												
			EN → *					* → EN				
	Method	en-en	de	es	fr	ja	zh	de	es	fr	ja	zh
Baseline	MT + sent2vec	77.6	74.0	75.8	73.5	1.8	72.6	70.4	70.4	69.5	70.4	71.3
	LASER (original)	88.8	83.9	<b>86.5</b>	<b>85.2</b>	<b>86.5</b>	85.2	79.8	79.8	<b>84.3</b>	<b>88.3</b>	<b>85.7</b>
	Contrastive loss	65.0	35.4	43.0	42.6	27.8	26.0	50.2	59.2	54.3	50.7	49.8
	N-pair loss	61.4	23.8	40.4	35.9	12.6	26.5	50.2	53.4	45.7	50.7	52.0
	Softmax loss	75.8	20.2	35.4	30.5	12.1	16.1	31.4	39.0	35.9	28.3	26.0
Proposed	EMU	<b>89.7</b>	<b>87.0</b> *	<b>86.5</b>	84.3	<b>86.5</b>	86.1	<b>81.2</b>	<b>87.0</b>	82.5	83.9	84.8
	EMU w/o LD	89.7	83.9	85.7	83.4	82.1	<b>87.4</b>	77.6	84.3	83.4	86.5	83.9
	EMU w/o LD+CL	88.3	75.3	80.3	75.8	70.9	78.5	72.6	82.1	75.3	81.6	80.7
	EMU-PARALLEL	86.5	78.0	79.4	76.2	74.4	78.5	78.9	82.1	78.0	80.3	82.1

LD+CL) as a part of an ablation study. Finally, we evaluated EMU-PARALLEL, which uses *parallel sentences* instead of randomly sampled sentences for cross-lingual adversarial training.

**Hyper-parameters** We used the official implementation of LASER<sup>6</sup> and the pre-trained models including BPE. We implemented our proposed method and the baseline methods using PyTorch. We used an initial learning rate of  $10^{-3}$  and optimized the model with Adam. We used a batch size of 16. For our proposed methods, we set  $\alpha = 50$  and  $\lambda = 10^{-4}$ . All the models were trained for 3 epochs. The architecture of language discriminator  $D$  has two 900-dimensional fully-connected layers with a dropout rate of 0.2. The hyper-parameters were  $\gamma = 10^{-4}$ ,  $k = 5$ ,  $c = 0.01$  respectively. The language discriminator was also optimized with Adam with an initial learning rate of  $5.0 \times 10^{-4}$ .

<sup>6</sup><https://github.com/facebookresearch/LASER>

**Evaluation Metric** We used the leave-one-out evaluation method on the test data. For each sentence, we consider the other sentences in the test data as labeled sentences to find the nearest neighbor to predict the label. The idea is to exclude the direct translation of an input sentence in the target language to make the nearest neighbor search more challenging and to simulate the real-world setting where parallel sentences are missing. We used Acc@1 (the ratio of test sentences that are correctly categorized into the intent classes) as our evaluation metric.

## Results and Discussion

Table 2 shows the experimental results on these three datasets. In Table 2 (a), EMU achieved the best performance for all the 11 tasks (en-fr, en-ja, and ja-en by EMU w/o LD and en-ja by EMU w/o LD+CL.) EMU outperformed the baseline methods including the original LASER model. In Table 2 (b), EMU achieved the best performance for 10 tasks

Table 3: Relative performance (Acc@1 on *HotelQA*) of EMU w/o LD models trained on different training languages against the original LASER model for each language pair.

Training data	en-en	en-de	en-fr	de-en	de-de	de-fr	fr-en	fr-de	fr-fr
En only	+37.5%	+40.0%	+34.3%	+27.0%	+10.0%	+1.7%	+12.3%	+12.7%	+11.1%
De only	+26.2%	+47.7%	+10.0%	+49.2%	+10.0%	+25.0%	+9.6%	+7.9%	+9.9%
Fr only	+30.0%	+33.8%	+28.6%	+17.5%	+8.7%	+16.7%	+31.5%	+15.9%	+17.3%
En + De	+37.5%	+58.5%	+27.1%	+50.8%	+17.5%	+23.3%	+9.6%	+12.7%	+14.8%
En + Fr	+40.0%	+60.0%	<b>+50.0%</b>	+46.0%	+12.5%	+33.3%	<b>+35.6%</b>	+25.4%	+23.5%
De + Fr	+28.7%	+50.8%	+37.1%	+55.6%	+12.5%	+46.7%	+31.5%	+25.4%	+17.3%
En + De + Fr	<b>+41.2%</b>	<b>+63.1%</b>	+47.1%	<b>+60.3%</b>	<b>+20.0%</b>	<b>+56.7%</b>	+31.5%	<b>+34.9%</b>	<b>+25.9%</b>

Table 4: Relative performance of Acc@1 on *HotelQA* of EMU w/o LD against the original LASER model for each language pair.

*→	en	de	es	fr	zh	ja
en	+37.5%	+40.0%	+22.9%	+34.3%	+39.1%	+26.1%
de	+27.0%	+10.0%	+10.0%	+1.7%	+14.5%	+20.9%
es	+34.8%	+0.0%	+11.5%	+5.0%	+21.3%	+8.0%
fr	+12.3%	+12.7%	+23.2%	+11.1%	+13.2%	+7.1%
zh	+21.9%	+34.5%	+11.4%	+9.6%	+9.1%	+10.1%
ja	+18.3%	+31.0%	+23.4%	+20.3%	+32.3%	+22.1%

(en-fr by EMU w/o LD+CL.) The original LASER model showed the best performance for zh-en and all of the EMU methods degraded the performance for the task. In Table 2 (c), EMU achieved the best performance for 7 tasks (en-zh by EMU w/o LD), whereas the original LASER model achieved the best performance for the rest of the tasks. From the results, EMU consistently outperformed the baseline methods, including the original LASER model. At the same time, EMU failed to improve the performance of the five tasks, namely zh-en on *ATIS* (Table 2 (b)) and en-fr, fr-en, ja-en, ja-zh on *Quora* (Table 2 (c)). We would like to emphasize that the EMU models were trained using labeled data only in English. The EMU also used unlabeled data in non-English languages. Therefore, it is noteworthy that our framework successfully specializes multilingual sentence embeddings for multiple language pairs, which involve English, using only English labeled data. The results support that EMU is effective in semantically specializing multilingual sentence embeddings.

For all the tasks, we observe that the baseline fine-tuning methods (i.e., contrastive loss, N-pair loss, softmax loss) do not improve the performance but instead decrease the accuracy values compared to the original LASER performance. The results indicate that fine-tuning multilingual sentence embeddings is sensitive to the choice of loss functions, and  $L_2$ -constrained softmax loss is the best choice among the loss functions.

The original LASER model consistently performs better on all datasets for the en-en task compared to the other tasks. This is partially due to the higher quality sentence embeddings in English. More specifically, the LASER model was trained on MT tasks, translating text from 93 languages to either English or Spanish as target languages with English having the most training data in the dataset (Artetxe and Schwenk 2019b).

Table 5: Ablation study of EMU. Each value denotes the average percentage point (pp) drop after removing the component. Negative values denote *improvements* after removing the component. \*\*and \*\*\*denote  $p$ -values  $< 0.05$  and  $< 0.01$  (Wilcoxon signed ranked test) respectively.

Component	<i>HotelQA</i>	<i>ATIS</i>	<i>Quora</i>
Language Discriminator	1.45	2.81**	1.05
Center loss	-0.44	0.04	6.00***

MT+sent2vec shows significantly low Acc@1 values for the en-ja and en-zh tasks on *ATIS*, and for the en-ja task on *Quora*. Investigating this trend, we observed that back-translation of sentences that were translated from en into ja/zh results in the following types of degradation: (1) missing words, especially interrogative pronouns (e.g., what, when, which etc.) and verbs, (2) significant changes in the word order. As discussed above, sent2vec embeddings are also susceptible to this type of perturbation. From the evaluation perspective, we can consider the en-en performance of MT+sent2vec as an upper bound on its performance for all en-\* and \*-en tasks, assuming that MT exactly translates back to the original sentence in English. Nevertheless, EMU consistently outperforms MT+sent2vec.

**Ablation study** We conducted an ablation study to quantitatively evaluate the contribution of each component of EMU, namely, the language discriminator and the center loss. First, we compared EMU w/o LD with EMU to verify the effect of the language discriminator, and then compared EMU w/o LD and EMU w/o LD+CL to determine the effect of the center loss.

Table 5 shows the average percentage point drop (i.e., the degree of contributions) of each component. The language discriminator had a significant contribution of 2.81 points on *ATIS*. The contributions were 1.45 points and 1.05 points on *HotelQA* and *Quora* respectively. Similarly, the center loss had a significant impact on *Quora*, whereas it had almost no effect on *ATIS* and had a negative impact on *HotelQA*.

**Sentence embedding visualization** We conducted a qualitative analysis to observe how our framework with the language discriminator specialized multilingual sentence embeddings and enhanced the multilinguality. We filtered English and German sentences from the test data of the *ATIS* dataset and visualized sentence embeddings of (a) the original LASER model, (b) the softmax loss, (c) EMU w/o LD,




Figure 2: Visualizations of the sentence embeddings of English (○) and German (×) test data of the *ATIS* dataset. We used *t*-SNE to convert the sentence embeddings into the 2d space. Each point is a sentence and the color denotes the intent class. The plots are: (a) the original LASER embeddings, (b) softmax loss, (c) EMU w/o LD, (d) EMU.

and (d) EMU into the same 2D space using *t*-SNE.

Figure 2 shows visualizations of these methods. Figure 2(a) shows that the original LASER sentence embeddings have multilinguality, as the sentences in the same intent in English and German were embedded close to each other. Figure 2(b) shows that fine-tuning the model with the softmax loss function broke not only the intent clusters but also spoiled the multilinguality. In Figure 2(c), EMU w/o LD successfully specialized the sentence embeddings, whereas multilinguality was degraded as the sentence embeddings of the same intent classes were separated compared to the original LASER model. Finally, EMU (with the language discriminator) moved sentence embeddings of the same intent in English and German close to each other, as shown in Figure 2(d).

From the results, we observe that incorporating the language discriminator enriches the multilinguality in the embedding space.

**Do we need parallel sentences for Emu?** We compared EMU to EMU-PARALLEL, which uses parallel sentences instead of randomly sampled sentences, to verify whether using parallel sentences makes multilingual adversarial learning more effective. The results are shown in Tables 2 (a)-(c). Compared to EMU, EMU-PARALLEL showed lower Acc@1 values on the three datasets. The decreases were -0.5 points, -1.2 points, and -5.9 points on *HotelQA*, *ATIS*, and *Quora* respectively. The differences are not statistically significant except for *Quora*. The results show that the language discriminator of EMU does not need any cost-expensive parallel corpus but can improve performance using *unlabeled and non-parallel* sentences in other languages.

**What language(s) should we use for training?** We also investigated how the performance changes by fine-tuning with training data in multiple languages other than English. To understand the insights more closely, we turned off the language discriminator in this analysis to ensure that EMU uses data only in specified languages. We summarize the relative performance of EMU w/o LD against the original LASER model on the *HotelQA* dataset. As discussed above, the accuracy values of tasks that involve English in at least one side (i.e., source language, target language, or both) show larger improvements than the other pairs that only involve non-English languages. This is likely because sentence em-

beddings of those languages were not appropriately fine-tuned compared to those of English because training data in those languages were not used.

Therefore, we hypothesized that using training data in the same language for a target and/or source language would be the best choice. To test the hypothesis, we chose English, German, and French as source/target languages and conducted additional experiments on the *HotelQA* dataset. The experimental settings, including the hyper-parameters, followed the main experiments, with only the training data used for fine-tuning being different.

Table 3 shows the results. When only using training data in a single language (i.e., En only, De only, Fr only), the target language was the best training data for monolingual intent classification tasks because this method achieved the best performance in the en-en, de-de, and fr-fr tasks respectively. Similarly, using the source and target languages as training data was the best configuration for methods that trained in two languages. That is, En+De achieved the best performance for the en-de and de-en tasks. En+Fr (De+Fr) also achieved the best performance for the en-fr (de-fr) and fr-en (fr-de.) Finally, the method that used training data in the three languages (En+De+Fr) showed the best accuracy values for 7 out of 9 tasks. The degradation in those two tasks occurred when En+De+Fr incorporated a language that was neither the source nor target languages (i.e., en-fr and fr-en.)

From the results, we conclude that we should focus on creating training data in a target or source language to obtain the best performance with EMU and use our budget effectively.

## Related Work

Multilingual embedding techniques (Ruder et al. 2019) have been well studied, and most of the prior work has focused on word embeddings. However, relatively fewer techniques have been developed for multilingual sentence embeddings. This is because such techniques (Hermann and Blunsom 2014; Artetxe and Schwenk 2019b) require parallel sentences for training multilingual sentence embeddings and some use both sentence-level and word-level alignment information (Luong, Pham, and Manning 2015). Ruckle et al. (Rücklé et al. 2018) developed an unsupervised sentence

embedding method based on concatenating and aggregating cross-lingual word embeddings. They also confirm that the method performs well on cross-lingual as well as monolingual settings. Schwenk and Douze (Schwenk and Douze 2017) used machine translation tasks to learn multilingual sentence representations. This idea has been further expanded in LASER (Artetxe and Schwenk 2019b; 2019a), a recently developed system which trains a language-agnostic sentence embeddings model with a large number of translation tasks on a large-scale parallel corpora.

Similar to the center loss used in this paper, two techniques have incorporated cluster-level information (Huang et al. 2018; Doval et al. 2018) to enhance the compactness of word clusters to improve the quality of multilingual word embedding models. None of them have directly used the centroid of each class to calculate loss values for training.

Adversarial learning (Goodfellow et al. 2014) is a common technique that has been used for many NLP tasks, including (monolingual) sentence embeddings (Patro et al. 2018) and multilingual word embeddings (Conneau et al. 2018b; Chen and Cardie 2018). (Chen et al. 2018) developed a technique that uses a language discriminator to train a cross-lingual sentiment classifier. Our framework is similar in the use of a language discriminator, but our novelty is that it uses a language discriminator for learning multilingual sentence embeddings instead of cross-lingual transfer. Joty et al. (Joty et al. 2017) used a language discriminator to train a model for cross-lingual question similarity calculation. Their setting differs from ours as their method requires parallel sentences in different languages and pair-wise similarity labels instead of class labels.

There is a line of work in post-processing word embedding models called *word embedding specialization* (Faruqui et al. 2015; Kiela, Hill, and Clark 2015; Mrkšić et al. 2017). Prior work specialized word embeddings with different external resources such as semantic information (Faruqui et al. 2015). The common approaches are (1) a post-hoc learning (Faruqui et al. 2015) that uses additional loss function to tune pre-trained embeddings, (2) learning an additional model (Glavaš and Vulić 2018; Vulić et al. 2018), and (3) the fine-tuning approach (Abdalla, Sahlgren, and Hirst 2019), which is similar to our fine-tuning approach. However, to the best of our knowledge, we are the first to approach semantic specialization of multilingual sentence embeddings.

## Conclusion

We have presented EMU, a semantic specialization framework for multilingual sentence embeddings. EMU incorporates multilingual adversarial training on top of fine-tuning to enhance multilinguality without using parallel sentences.

Our experimental results show that EMU outperformed the baseline methods including state-of-the-art multilingual sentence embeddings, LASER, and monolingual sentence embeddings after machine translation with respect to multiple language pairs. The results also show that EMU can successfully train a model using only monolingual labeled data and unlabeled data in other languages.

## Acknowledgments

We thank Sorami Hisamoto for sharing his literature survey on cross-lingual embedding techniques and Tom Mitchell for helpful comments as well as the anonymous reviewers for their constructive feedback.

## References

- Abdalla, M.; Sahlgren, M.; and Hirst, G. 2019. Enriching word embeddings with a regressor instead of labeled corpora. In *Proc. AAAI '19*, 6188–6195.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *Proc. ICML '17*, volume 70, 214–223.
- Artetxe, M., and Schwenk, H. 2019a. Margin-based parallel corpus mining with multilingual sentence embeddings. In *Proc. ACL '19*, 3197–3203.
- Artetxe, M., and Schwenk, H. 2019b. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics* 7:597–610.
- Chen, X., and Cardie, C. 2018. Unsupervised multilingual word embeddings. In *Proc. EMNLP '18*, 261–270.
- Chen, X.; Sun, Y.; Athiwaratkun, B.; Cardie, C.; and Weinberger, K. 2018. Adversarial deep averaging networks for cross-lingual sentiment classification. *Transactions of the Association for Computational Linguistics* 6:557–570.
- Chidambaram, M.; Yang, Y.; Cer, D.; Yuan, S.; Sung, Y.; Strope, B.; and Kurzweil, R. 2019. Learning cross-lingual sentence representations via a multi-task dual-encoder model. In *Proc. RepL4NLP '19*, 250–259.
- Chopra, S.; Hadsell, R.; LeCun, Y.; et al. 2005. Learning a similarity metric discriminatively, with application to face verification. In *Proc. CVPR '05*, 539–546.
- Conneau, A.; Kiela, D.; Schwenk, H.; Barrault, L.; and Bordes, A. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proc. EMNLP '17*, 670–680.
- Conneau, A.; Kruszewski, G.; Lample, G.; Barrault, L.; and Baroni, M. 2018a. What you can cram into a single  $\&\#\&$  vector: Probing sentence embeddings for linguistic properties. In *Proc. ACL '18*, 2126–2136.
- Conneau, A.; Lample, G.; Ranzato, M.; Denoyer, L.; and Jégou, H. 2018b. Word translation without parallel data. In *Proc. ICLR '18*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. NAACL-HLT '19*, 4171–4186.
- Doval, Y.; Camacho-Collados, J.; Espinosa Anke, L.; and Schockaert, S. 2018. Improving cross-lingual word embeddings by meeting in the middle. In *Proc. EMNLP '18*, 294–304.
- Faruqui, M.; Dodge, J.; Jauhar, S. K.; Dyer, C.; Hovy, E.; and Smith, N. A. 2015. Retrofitting word vectors to semantic lexicons. In *Proc. NAACL-HLT '15*, 1606–1615.



- Glavaš, G., and Vulić, I. 2018. Explicit retrofitting of distributional word vectors. In *Proc. ACL '18*, 34–45.
- Glavas, G.; Litschko, R.; Ruder, S.; and Vulic, I. 2019. How to (properly) evaluate cross-lingual word embeddings: On strong baselines, comparative analyses, and some misconceptions. In *Proc. ACL '19*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Proc. NIPS '14*, 2672–2680.
- Hemphill, C. T.; Godfrey, J. J.; and Doddington, G. R. 1990. The ATIS spoken language systems pilot corpus. In *Proc. the Workshop on Speech and Natural Language, HLT '90*, 96–101.
- Hermann, K. M., and Blunsom, P. 2014. Multilingual models for compositional distributed semantics. In *Proc. ACL '14*, 58–68.
- Howard, J., and Ruder, S. 2018. Universal language model fine-tuning for text classification. In *Proc. ACL '18*, 328–339.
- Huang, L.; Cho, K.; Zhang, B.; Ji, H.; and Knight, K. 2018. Multi-lingual common semantic space construction via cluster-consistent word embedding. In *Proc. EMNLP '18*.
- Joty, S.; Nakov, P.; Màrquez, L.; and Jaradat, I. 2017. Cross-language learning with adversarial neural networks. In *Proc. CoNLL '17*, 226–237.
- Kiela, D.; Hill, F.; and Clark, S. 2015. Specializing word embeddings for similarity or relatedness. In *Proc. EMNLP '15*, 2044–2048.
- Luong, T.; Pham, H.; and Manning, C. D. 2015. Bilingual word representations with monolingual quality in mind. In *Proc. RepL4NLP '15*, 151–159.
- Mrkšić, N.; Vulić, I.; Ó Séaghdha, D.; Leviant, I.; Reichart, R.; Gašić, M.; Korhonen, A.; and Young, S. 2017. Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints. *Transactions of the Association for Computational Linguistics* 5:309–324.
- Pagliardini, M.; Gupta, P.; and Jaggi, M. 2018. Unsupervised learning of sentence embeddings using compositional n-gram features. In *NAACL-HLT '18*.
- Paşca, M. 2003. *Open-domain question answering from large text collections*. MIT Press.
- Patro, B. N.; Kurmi, V. K.; Kumar, S.; and Namboodiri, V. P. 2018. Learning semantic sentence embeddings using pairwise discriminator. In *Proc. COLING '18*.
- Peters, M.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep contextualized word representations. In *Proc. NAACL-HLT '18*, 2227–2237.
- Ranjan, R.; Castillo, C. D.; and Chellappa, R. 2017. L2-constrained softmax loss for discriminative face verification. *arXiv preprint arXiv:1703.09507* abs/1703.09507.
- Reimers, N., and Gurevych, I. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proc. EMNLP-IJCNLP '19*, 3973–3983.
- Rücklé, A.; Eger, S.; Peyrard, M.; and Gurevych, I. 2018. Concatenated power mean word embeddings as universal cross-lingual sentence representations. *arXiv preprint arXiv:1803.01400*.
- Ruder, S.; Vulić, I.; Søgaard, A.; and Faruqui, M. 2019. *Cross-Lingual Word Embeddings*. Morgan & Claypool Publishers.
- Schwenk, H., and Douze, M. 2017. Learning joint multi-lingual sentence representations with neural machine translation. In *Proc. RepL4NLP '17*, 157–167.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Neural machine translation of rare words with subword units. In *Proc. ACL '16*, 1715–1725.
- Vulić, I.; Glavaš, G.; Mrkšić, N.; and Korhonen, A. 2018. Post-specialisation: Retrofitting vectors of words unseen in lexical resources. In *Proc. NAACL-HLT '19*.
- Wang, J.; Zhang, T.; Sebe, N.; Shen, H. T.; et al. 2017. A survey on learning to hash. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40(4):769–790.
- Wen, Y.; Zhang, K.; Li, Z.; and Qiao, Y. 2016. A discriminative feature learning approach for deep face recognition. In *Proc. ECCV '16*, 499–515.
- Yang, Y.; Abrego, G. H.; Yuan, S.; Guo, M.; Shen, Q.; Cer, D.; Sung, Y.-h.; Stroppe, B.; and Kurzweil, R. 2019. Improving multilingual sentence embedding using bi-directional dual encoder with additive margin softmax. In *Proc. IJCAI '19*, 5370–5378.
- Zhu, X.; Li, T.; and de Melo, G. 2018. Exploring semantic properties of sentence embeddings. In *Proc. ACL '18*, 632–637.