

Learning to Represent Bilingual Dictionaries

Muhao Chen^{1*}, Yingtao Tian^{2*}, Haochen Chen²,
Kai-Wei Chang¹, Steven Skiena² & Carlo Zaniolo¹

¹University of California, Los Angeles, CA, USA

²The State University of New York, Stony Brook, NY, USA

muhaochen@ucla.edu; {kwchang, zaniolo}@cs.ucla.edu;
{yittian, haocchen, skiena}@cs.stonybrook.edu

Abstract

Bilingual word embeddings have been widely used to capture the correspondence of lexical semantics in different human languages. However, the cross-lingual correspondence between sentences and words is less studied, despite that this correspondence can significantly benefit many applications such as cross-lingual semantic search and textual inference. To bridge this gap, we propose a neural embedding model that leverages bilingual dictionaries¹. The proposed model is trained to map the lexical definitions to the cross-lingual target words, for which we explore with different sentence encoding techniques. To enhance the learning process on limited resources, our model adopts several critical learning strategies, including multi-task learning on different bridges of languages, and joint learning of the dictionary model with a bilingual word embedding model. We conduct experiments on two new tasks. In the cross-lingual reverse dictionary retrieval task, we demonstrate that our model is capable of comprehending bilingual concepts based on descriptions, and the proposed learning strategies are effective. In the bilingual paraphrase identification task, we show that our model effectively associates sentences in different languages via a shared embedding space, and outperforms existing approaches in identifying bilingual paraphrases.

1 Introduction

Cross-lingual semantic representation learning has attracted significant attention recently. Various approaches have been proposed to align words of different languages in a shared embedding space (Ruder et al., 2017). By offering task-invariant se-

matic transfers, these approaches critically support many cross-lingual NLP tasks including neural machine translations (NMT) (Devlin et al., 2014), bilingual document classification (Zhou et al., 2016), knowledge alignment (Chen et al., 2018b) and entity linking (Upadhyay et al., 2018).

While many existing approaches have been proposed to associate lexical semantics between languages (Chandar et al., 2014; Gouws et al., 2015; Luong et al., 2015a), modeling the correspondence between lexical and sentential semantics across different languages is still an unresolved challenge. We argue that learning to represent such cross-lingual and multi-granular correspondence is well desired and natural for multiple reasons. One reason is that, learning word-to-word correspondence has a natural limitation, considering that many words do not have direct translations in another language. For example, *schadenfreude* in German, which means *a feeling of joy that comes from knowing the troubles of other people*, has no proper English counterpart word. To appropriately learn the representations of such words in bilingual embeddings, we need to capture their meanings based on the definitions.

Besides, modeling such correspondence is also highly beneficial to many application scenarios. One example is cross-lingual semantic search of concepts (Hill et al., 2016), where the lexemes or concepts are retrieved based on sentential descriptions (see Fig. 1). Others include discourse relation detection in bilingual dialogue utterances (Jiang et al., 2018), multilingual text summarization (Nenkova et al., 2012), and educational applications for foreign language learners. Finally, it is natural in foreign language learning that a human learns foreign words by looking up their meanings in the native language (Hulstijn et al., 1996). Therefore, learning such correspondence essentially mimics human learning behaviors.

* Both authors contributed equally to this work.

¹We refer the term *dictionary* to its common meaning, i.e. lexical definitions of words. Note that this is different from some papers on bilingual settings that refer dictionaries to seed lexicons for one-to-one word mappings.

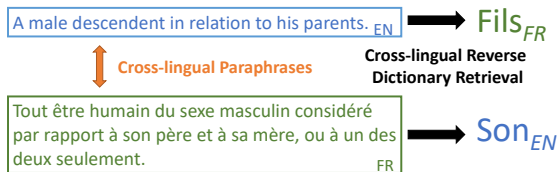


Figure 1: An example illustrating the two cross-lingual tasks. The *cross-lingual reverse dictionary retrieval* finds cross-lingual target words based on descriptions. In terms of *cross-lingual paraphrases*, the French sentence (which means *any male being considered in relation to his father and mother, or only one of them*) describes the same meaning as the English sentence, but has much more content details.

However, realizing such a representation learning model is a non-trivial task, inasmuch as it requires a comprehensive learning process to effectively compose the semantics of arbitrary-length sentences in one language, and associate that with single words in another language. Consequently, this objective also demands high-quality cross-lingual alignment that bridges between single and sequences of words. Such alignment information is generally not available in the parallel and seed-lexicon that are utilized by bilingual word embeddings (Ruder et al., 2017).

To incorporate the representations of bilingual lexical and sentential semantics, we propose an approach to capture *the mapping from the definitions to the corresponding foreign words* by leveraging *bilingual dictionaries*. The proposed model **BiLDRL (Bilingual Dictionary Representation Learning)** first constructs a word embedding space with pre-trained bilingual word embeddings. Based on cross-lingual word definitions, a sentence encoder is trained to realize the mapping from literal descriptions to target words in the bilingual word embedding space, for which we investigate with multiple encoding techniques. To enhance cross-lingual learning on limited resources, **BiLDRL** conducts multi-task learning on different directions of a language pair. Moreover, **BiLDRL** enforces a joint learning strategy of bilingual word embeddings and the sentence encoder, which seeks to gradually adjust the embedding space to better suit the representation of cross-lingual word definitions.

To show the applicability of **BiLDRL**, we conduct experiments on two useful cross-lingual tasks (see Fig. 1). (i) *Cross-lingual reverse dictionary retrieval* seeks to retrieve words or concepts given descriptions in another language. This task is use-

ful to help users find foreign words based on the notions or descriptions, and is especially beneficial to users such as translators, foreigner language learners and technical writers using non-native languages. We show that **BiLDRL** achieves promising results on this task, while bilingual multi-task learning and joint learning dramatically enhance the performance. (ii) *Bilingual paraphrase identification* asks whether two sentences in different languages essentially express the same meaning, which is critical to question answering or dialogue systems that apprehend multilingual utterances (Bannard and Callison-Burch, 2005). This task is challenging, as it requires a model to comprehend cross-lingual paraphrases that are inconsistent in grammar, content details and word orders. **BiLDRL** maps sentences to the lexicon embedding space. This process reduces the problem to evaluate the similarity of lexicon embeddings, which can be easily solved by a simple classifier. **BiLDRL** performs well with even a small amount of data, and significantly outperforms previous approaches.

2 Related Work

We discuss two lines of relevant work.

Bilingual word embeddings. Various approaches have been proposed for training bilingual word embeddings. These approaches span in two families: off-line mappings and joint training.

The off-line mapping based approach fixes the structures of pre-trained monolingual embeddings, and induces bilingual projections based on seed lexicons (Mikolov et al., 2013a). Some variants of this approach improve the quality of projections by adding constraints such as orthogonality of transforms, normalization and mean centering of embeddings (Xing et al., 2015; Artetxe et al., 2016; Vulić et al., 2016). Others adopt canonical correlation analysis to map separate monolingual embeddings to a shared embedding space (Faruqui and Dyer, 2014; Doval et al., 2018).

Unlike off-line mappings, joint training models simultaneously update word embeddings and cross-lingual alignment. In doing so, such approaches generally capture more precise cross-lingual semantic transfer (Ruder et al., 2017; Upadhyay et al., 2018). While a few such models still maintain separated embedding spaces for each language (Artetxe et al., 2017), more of them maintain a unified space for both languages. The

cross-lingual semantic transfer by these models is captured from parallel corpora with sentential or document-level alignment, using techniques such as bilingual bag-of-words distances (BilBOWA) (Gouws et al., 2015), Skip-Gram (Coulmance et al., 2015) and sparse tensor factorization (Vyas and Carpuat, 2016).

Neural sentence modeling. Neural sentence models seek to capture phrasal or sentential semantics from word sequences. They often adopt encoding techniques such as recurrent neural encoders (RNN) (Kiros et al., 2015), convolutional encoders (CNN) (Chen et al., 2018a), and attentive encoders (Rocktäschel et al., 2016) to represent the composed semantics of a sentence as an embedding vector. Recent works have focused on apprehending pairwise correspondence of sentential semantics by adopting multiple neural sentence models in one learning architecture, including Siamese models for detecting discourse relations of sentences (Sha et al., 2016), and sequence-to-sequence models for tasks like style transfer (Shen et al., 2017), text summarization (Chopra et al., 2016) and translation (Wu et al., 2016).

On the other hand, fewer efforts have been put to characterizing the associations between sentential and lexical semantics. Hill et al. (2016) and Ji et al. (2017) learn off-line mappings between monolingual descriptions and lexemes to capture such associations. Eisner et al. (2016) adopt a similar approach to capture emojis based on descriptions. At the best of our knowledge, there has been no previous approach to learn to discover the correspondence of sentential and lexical semantics in a multilingual scenario. This is exactly the focus of our work, in which the proposed strategies of multi-task learning and joint learning are critical to the corresponding learning process under limited resources. Utilizing such correspondence, our approach also sheds light on addressing discourse relation detection in a multilingual scenario.

3 Modeling Bilingual Dictionaries

We hereby begin our modeling with the formalization of bilingual dictionaries. We use \mathcal{L} to denote the set of languages. For a language $l \in \mathcal{L}$, V_l denotes its vocabulary, where for each word $w \in V_l$, bold-faced $\mathbf{w} \in \mathbb{R}^k$ denotes its embedding vector. A l_i - l_j bilingual dictionary $D(l_i, l_j)$ (or simply D_{ij}) contains dictionary entries $(w^i, S_w^j) \in D_{ij}$, in which $w^i \in V_{l_i}$, and $S_w^j = w_1^j \dots w_n^j$ ($w^j \in$

V_{l_j}) is a cross-lingual definition that describes the word w^i with a sequence of words in language l_j . For example, a French-English dictionary $D(\text{Fr}, \text{En})$ could include a French word *appétite* accompanied by its English definition *desire for, or relish of food or drink*. Note that, for a word w^i , multiple definitions in l_j may coexist.

BILDRL is constructed and improved through three stages, as depicted in Fig. 2. A sentence encoder is first used to learn from a bilingual dictionary the association between words and definitions. Then in a pre-trained bilingual word embedding space, multi-task learning is conducted on both directions of a language pair. Lastly, joint learning with word embeddings is enforced to simultaneously adjust the embedding space during the training of the dictionary model, which further enhances the cross-lingual learning process.

It is noteworthy that, NMT (Wu et al., 2016) is considered as an ostensibly relevant method to ours. NMT does not apply to our problem setting because it has major differences from our work in those perspectives: (i) In terms of data modalities, NMT has to bridge between corpora of the same granularity, i.e. either between sentences or between lexicons. This is unlike BILDRL that captures multi-granular correspondence of semantics across different modalities, i.e. sentences and words; (ii) As for learning strategies, NMT relies on an encoder-decoder architecture using end-to-end training (Luong et al., 2015b), while BILDRL employs joint learning of a dictionary-based sentence encoder and a bilingual embedding space.

3.1 Encoders for Lexical Definitions

BILDRL models a dictionary using a neural sentence encoder $E(S)$, which composes the meaning of the sentence into a latent vector representation. We hereby introduce this model component, which is designed to be a GRU encoder with self-attention. Besides that, we also investigate other widely-used neural sequence encoders.

3.1.1 Attentive GRU Encoder

The GRU encoder is a computationally efficient alternative of the LSTM (Cho et al., 2014). Each unit consists of a reset gate \mathbf{r}_t and an update gate \mathbf{z}_t to track the state of the sequence. Given the vector representation \mathbf{w}_t of an incoming item w_t , GRU updates the hidden state $\mathbf{h}_t^{(1)}$ as a linear combination of the previous state $\mathbf{h}_{t-1}^{(1)}$ and the candidate

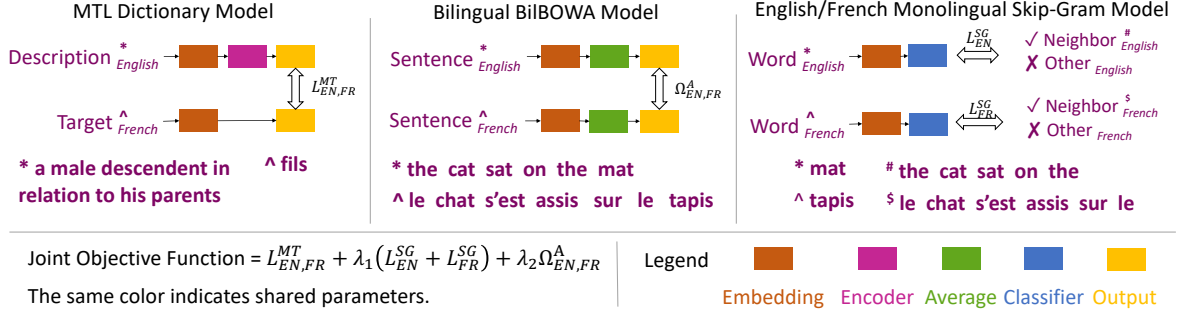


Figure 2: Joint learning architecture of BiLDRL.

state $\tilde{\mathbf{h}}_t^{(1)}$ of the new item w_t as below.

$$\mathbf{h}_t^{(1)} = \mathbf{z}_t \odot \tilde{\mathbf{h}}_t^{(1)} + (1 - \mathbf{z}_t) \odot \mathbf{h}_{t-1}^{(1)}.$$

The update gate \mathbf{z}_t balances between the information of the previous sequence and the new item, where \mathbf{M}_z and \mathbf{N}_z are two weight matrices, \mathbf{b}_z is a bias vector, and σ is the sigmoid function.

$$\mathbf{z}_t = \sigma \left(\mathbf{M}_z \mathbf{x}_t + \mathbf{N}_z \mathbf{h}_{t-1}^{(1)} + \mathbf{b}_z \right).$$

The candidate state $\tilde{\mathbf{h}}_t^{(1)}$ is calculated similarly to those in a traditional recurrent unit as below. The reset gate \mathbf{r}_t thereof controls how much information of the past sequence should contribute to the candidate state.

$$\tilde{\mathbf{h}}_t^{(1)} = \tanh \left(\mathbf{M}_s \mathbf{w}_t + \mathbf{r}_t \odot (\mathbf{N}_s \mathbf{h}_{t-1}^{(1)}) + \mathbf{b}_s \right)$$

$$\mathbf{r}_t = \sigma \left(\mathbf{M}_r \mathbf{w}_t + \mathbf{N}_r \mathbf{h}_{t-1}^{(1)} + \mathbf{b}_r \right).$$

While a GRU encoder can stack multiple of the above GRU layers, without an attention mechanism, the last state $\mathbf{h}_S^{(1)}$ of the last layer represents the overall meaning of the encoded sentence S .

The self-attention mechanism (Conneau et al., 2017) seeks to highlight the important units in an input sentence when capturing its overall meaning, which is calculated as below:

$$\mathbf{u}_t = \tanh \left(\mathbf{M}_a \mathbf{h}_t^{(1)} + \mathbf{b}_a \right)$$

$$a_t = \frac{\exp(\mathbf{u}_t^\top \mathbf{u}_S)}{\sum_{w_m \in S} \exp(\mathbf{u}_m^\top \mathbf{u}_S)}$$

$$\mathbf{h}_t^{(2)} = |S| a_t \mathbf{u}_t.$$

\mathbf{u}_t is the intermediary representation of GRU output $\mathbf{h}_t^{(1)}$, and $\mathbf{u}_S = \tanh(\mathbf{M}_a \mathbf{h}_S^{(1)} + \mathbf{b}_a)$ is that of the last GRU output $\mathbf{h}_S^{(1)}$. \mathbf{u}_S can be seen as a high-level representation of the input sequence.

By measuring the similarity of each \mathbf{u}_t with \mathbf{u}_S , the normalized attention weight a_t , which highlights an input that contributes significantly to the overall meaning, is produced through a softmax. Note that a scalar $|S|$ is multiplied along with a_t to \mathbf{u}_t , so as to keep the weighted representation $\mathbf{h}_t^{(2)}$ from losing the scale of $\mathbf{h}_t^{(1)}$. The sentence encoding is calculated as the average of the last attention layer $E^{(1)}(S) = \frac{1}{|S|} \sum_{t=1}^{|S|} a_t \mathbf{h}_t^{(2)}$.

3.1.2 Other Encoders

We also experiment with other widely used neural sentence modeling techniques², which are however outperformed by the attentive GRU in our tasks. These techniques include the vanilla GRU, CNN (Kalchbrenner et al., 2014), and linear bag-of-words (BOW) (Hill et al., 2016). We briefly introduce the later two techniques in the following.

A convolutional encoder applies a kernel $\mathbf{M}_c \in \mathbb{R}^{h \times k}$ to produce the latent representation $\mathbf{h}_t^{(3)} = \tanh(\mathbf{M}_c \mathbf{w}_{t:t+h-1} + \mathbf{b}_c)$ from each h -gram of the input vector sequence $\mathbf{w}_{t:t+h-1}$, for which h is the kernel size and \mathbf{b}_c is a bias vector. A sequence of latent vectors $\mathbf{H}^{(3)} = [\mathbf{h}_1^{(3)}, \mathbf{h}_2^{(3)}, \dots, \mathbf{h}_{|S|-h+1}^{(3)}]$ is produced from the input, where each latent vector leverages the significant local semantic features from each h -gram. Following convention (Liu et al., 2017), we apply dynamic max-pooling to extract robust features from the convolution outputs, and use the mean-pooling results of the last layer to represent the sentential semantics.

The Linear bag-of-words (BOW) encoder (Ji et al., 2017; Hill et al., 2016) is realized by the

²Note that recent advances in monolingual contextualized embeddings like multilingual ELMo (Peters et al., 2018; Che et al., 2018) and M-BERT (Pires et al., 2019; Devlin et al., 2018) can also be supported to represent sentences for our setting. We leave them as future work, as they require non-trivial adaption to both multilingual settings and joint training, and extensive pre-training on external corpora.

sum of projected word embeddings of the input sentence, i.e. $E^{(2)}(S) = \sum_{t=1}^{|S|} \mathbf{M}_b \mathbf{w}_t$.

3.2 Basic Learning Objective

The objective of learning the dictionary model is to map the encodings of cross-lingual word definitions to the target word embeddings. This is realized by minimizing the following L_2 loss,

$$L_{ij}^{\text{ST}} = \frac{1}{|D_{ij}|} \sum_{(w^i, S_w^j) \in D_{ij}} \|E_{ij}(S_w^j) - \mathbf{w}^i\|_2^2$$

in which E_{ij} is the dictionary model that maps from descriptions in l_j to words in l_i .

The above defines the basic model variants of BiLDRL that learns on a single dictionary. For word representations in the learning process, BiLDRL initializes the embedding space using pre-trained word embeddings. Note that, without adopting the joint learning strategy in Section 3.4, the learning process does not update word embeddings that are used to represent the definitions and target words. While other forms of loss such as cosine proximity (Hill et al., 2016) and hinge loss (Ji et al., 2017) may also be used in the learning process, we find that L_2 loss consistently leads to better performance in our experiments.

3.3 Bilingual Multi-task Learning

In cases where entries in a bilingual dictionary are not amply provided, learning the above bilingual dictionary on one ordered language pair may fall short in insufficiency of alignment information. One practical solution is to conduct a bilingual multi-task learning process. In detail, given a language pair (l_i, l_j) , we learn the dictionary model E_{ij} on both dictionaries D_{ij} and D_{ji} with shared parameters. Correspondingly, we rewrite the previous learning objective function as below, in which $D = D_{ij} \cup D_{ji}$.

$$L_{ij}^{\text{MT}} = \frac{1}{|D|} \sum_{(w, S_w) \in D} \|E_{ij}(S_w) - \mathbf{w}\|_2^2.$$

This strategy non-trivially requests the same dictionary model to represent semantic transfer in two directions of the language pair. To fulfill such a request, we initialize the embedding space using the BilBOWA embeddings (Gouws et al., 2015), which provide a unified embedding space that resolves both monolingual and cross-lingual semantic relatedness of words. In practice, we find this

simple multi-task strategy to bring significant improvement to our cross-lingual tasks. Note that, besides BilBOWA, other joint-training bilingual embeddings in a unified space (Doval et al., 2018) can also support this strategy, for which we leave the comparison to future work.

3.4 Joint Learning Objective

While above learning strategies are based on a fixed embedding space, we lastly propose a joint learning strategy. During the training process, this strategy simultaneously updates the embedding space based on both the dictionary model and the bilingual word embedding model. The learning is through asynchronous minimization of the following joint objective function,

$$J = L_{ij}^{\text{MT}} + \lambda_1(L_i^{\text{SG}} + L_j^{\text{SG}}) + \lambda_2\Omega_{ij}^{\text{A}},$$

where λ_1 and λ_2 are two positive hyperparameters. L_i^{SG} and L_j^{SG} are the original Skip-Gram losses (Mikolov et al., 2013b) to separately obtain word embeddings on monolingual corpora of l_i and l_j . Ω_{ij}^{A} , termed as below, is the alignment loss to minimize bag-of-words distances for aligned sentence pairs (S^i, S^j) in parallel corpora C_{ij} .

$$\Omega_{ij}^{\text{A}} = \frac{1}{|C_{ij}|} \sum_{(S^i, S^j) \in C_{ij}} d_S(S^i, S^j)$$

$$d_S(S^i, S^j) = \left\| \frac{1}{|S^i|} \sum_{w_m^i \in S^i} \mathbf{w}_m^i - \frac{1}{|S^j|} \sum_{w_n^j \in S^j} \mathbf{w}_n^j \right\|_2^2$$

The joint learning process adapts the embedding space to better suit the dictionary model, which is shown to further enhance the cross-lingual learning of BiLDRL.

3.5 Training

To initialize the embedding space, we pre-trained BilBOWA on the parallel corpora Europarl v7 (Koehn, 2005) and monolingual corpora of tokenized Wikipedia dump (Al-Rfou et al., 2013). For models without joint learning, we use AMSGrad (Reddi et al., 2018) to optimize the parameters. Each model without bilingual multi-task learning thereof, is trained on batched samples from each individual dictionary. Multi-task learning models are trained on batched samples from two dictionaries. Within each batch, entries of different directions of languages can be mixed together. For joint learning, we conduct an efficient

multi-threaded asynchronous training (Mnih et al., 2016) of AMSGrad. In detail, after initializing the embedding space based on pre-trained BiBOWA, parameter updating based on the four components of J occurs across four worker threads. Two monolingual threads select batches of monolingual contexts from the Wikipedia dump of two languages for Skip-Gram, one alignment thread randomly samples parallel sentences from Europarl, and one dictionary thread extracts samples of entries for a bilingual multi-task dictionary model. Each thread makes a batched update to model parameters asynchronously for each term of J . The asynchronous training of all threads keeps going until the dictionary thread finishes its epochs.

4 Experiments

We present experiments on two multilingual tasks: the cross-lingual reverse dictionary retrieval task and the bilingual paraphrase identification task.

4.1 Datasets

The experiment of cross-lingual reverse dictionary retrieval is conducted on a trilingual dataset *Wikt3l*. This dataset is extracted from Wiktionary³, which is one of the largest freely available multilingual dictionary resources on the Web. *Wikt3l* contains dictionary entries of language pairs (English, French) and (English, Spanish), which form En-Fr, Fr-En, En-Es and Es-En dictionaries on four bridges of languages in total. Two types of cross-lingual definitions are extracted from Wiktionary: (i) cross-lingual definitions provided under the *Translations* sections of Wiktionary pages; (ii) monolingual definitions for words that are linked to a cross-lingual counterpart with an inter-language link⁴ of Wiktionary. We exclude all the definitions of stop words in constructing the dataset, and list the statistics in Table 1.

Since existing datasets for paraphrase identification are merely monolingual, we contribute with another dataset *WBP3l* for cross-lingual sentential paraphrase identification. This dataset contains 6,000 pairs of bilingual sentence pairs respectively for En-Fr and En-Es settings. Within each bilingual setting, positive cases are formed as pairs of descriptions aligned by inter-language links, which exclude the word descriptions in

³<https://www.wiktionary.org/>

⁴An inter-language link matches the entries of counterpart words between language versions of Wiktionary.

Dictionary	En-Fr	Fr-En	En-Es	Es-En
#Target words	15,666	16,857	8,004	16,986
#Definitions	50,412	58,808	20,930	56,610

Table 1: Statistics of the bilingual dictionary dataset *Wikt3l*.

Positive Examples
En: <i>Being remote in space.</i>
Fr: <i>Se trouvant à une grande distance.</i>
En: <i>The interdisciplinary science that applies theories and methods of the physical sciences to questions of biology.</i>
Es: <i>Ciencia que emplea y desarrolla las teorías y métodos de la física en la investigación de los sistemas biológicos.</i>
Negative Examples
En: <i>A person who secedes or supports secession from a political union.</i>
Fr: <i>Contrôle politique exercé par une grande puissance sur une contre inféodée.</i>
En: <i>The fear of closed, tight places.</i>
Es: <i>Pérdida o disminución considerables de la memoria.</i>

Table 2: Examples of bilingual paraphrases from *WBP3l*.

Wikt3l for training *BiLDRL*. To generate negative examples, given a source word, we first find its 15 nearest neighbors in the embedding space. Within the nearest neighbors, we use ConceptNet (Speer et al., 2017) to filter out the synonyms of the source word, so as to prevent from generating false negative cases. Then we randomly pick one word from the filtered neighbors and pair its cross-lingual definition with the English definition of the source word to create a negative case. This process ensures that each negative case is endowed with limited dissimilarity of sentence meanings, which makes the decision more challenging. For each language setting, we randomly select 70% for training, 5% for validation, and the rest 25% for testing. Note that each language setting of this dataset thereof, matches with the quantity and partitioning of sentence pairs in the widely-used Microsoft Research Paraphrase Corpus benchmark for monolingual paraphrase identification (Yin et al., 2016; Das and Smith, 2009). Several examples from the dataset are shown in Table 2. The datasets and the processing scripts are available at https://github.com/muhaochen/bilingual_dictionaries.

4.2 Cross-lingual Reverse Dictionary Retrieval

The objective of this task is to enable cross-lingual semantic retrieval of words based on descriptions. Besides comparing variants of *BiLDRL* that adopt different sentence encoders and learning strate-

Languages Metric	En-Fr			Fr-En			En-Es			Es-En		
	$P@1$	$P@10$	MRR	$P@1$	$P@10$	MRR	$P@1$	$P@10$	MRR	$P@1$	$P@10$	MRR
BOW	0.8	3.4	0.011	0.4	2.2	0.006	0.4	2.4	0.007	0.4	2.6	0.007
CNN	6.0	12.4	0.070	6.4	14.8	0.072	3.8	7.2	0.045	7.0	16.8	0.088
GRU	35.6	46.0	0.380	38.8	49.8	0.410	47.8	59.0	0.496	57.6	67.2	0.604
ATT	38.8	47.4	0.411	39.8	50.2	0.425	51.6	59.2	0.534	60.4	68.4	0.629
GRU-mono	21.8	33.2	0.242	27.8	37.0	0.297	34.4	41.2	0.358	36.8	47.2	0.392
ATT-mono	22.8	33.6	0.249	27.4	39.0	0.298	34.6	42.2	0.358	39.4	48.6	0.414
GRU-MTL	43.4	49.2	0.452	44.4	52.8	0.467	50.4	60.0	0.530	63.6	71.8	0.659
ATT-MTL	46.8	56.6	0.487	47.6	56.6	0.497	55.8	62.2	0.575	66.4	75.0	0.687
ATT-joint	63.6	69.4	0.654	68.2	75.4	0.706	69.0	72.8	0.704	78.6	83.4	0.803

Table 3: Cross-lingual reverse dictionary retrieval results by BiLDR variants. We report $P@1$, $P@10$, and MRR on four groups of models: (i) basic dictionary models that adopt four different encoding techniques (BOW, CNN, GRU and ATT); (ii) models with the two best encoding techniques that enforce the monolingual retrieval approach by Hill et al. (2016) (GRU-mono and ATT-mono); (iii) models adopting bilingual multi-task learning (GRU-MTL and ATT-MTL); (iv) joint learning that employs the best dictionary model of ATT-MTL (ATT-joint).

gies, we also compare with the monolingual retrieval approach proposed by Hill et al. (2016). Instead of directly associating cross-lingual word definitions, this approach learns definition-to-word mappings in a monolingual scenario. When it applies to the multilingual setting, given a lexical definition, it first retrieves the corresponding word in the source language. Then, it looks for semantically related words in the target language using bilingual word embeddings. As discussed in Section 3, NMT does not apply to this task due that it cannot capture the multi-granular correspondence between a sentence and a word.

Evaluation Protocol. Before training the models, we randomly select 500 word definitions from each dictionary respectively as test cases, and exclude these definitions from the training data. Each of the basic BiLDR variants are trained on one bilingual dictionary. The monolingual retrieval models are trained to fit the target words in the original languages of the word definitions, which are also provided in Wiktionary. BiLDR variants with multi-task or joint learning use both dictionaries of the same language pair. In the test phase, for each test case $(w^i, S_w^j) \in D_{ij}$, the prediction performs a kNN search from the definition encoding $E_{ij}(S_w^j)$, and record the rank of w^i within the vocabulary of l_i . We limit the vocabularies to all words that appear in the Wikt31 dataset, which involve around 45k English words, 44k French words and 36k Spanish words. To prevent the surface information of the target word from appearing in the definition, we have also masked out any translation of the target word occurring in the definition using a wildcard token $\langle \text{concept} \rangle$. We aggregate three metrics on test cases: the accuracy $P@1$ (%), the proportion of

ranks no larger than 10 $P@10$ (%), and mean reciprocal rank MRR .

We pre-train BilBOWA based on the original configuration by Gouws et al. (2015) and obtain 50-dimensional initialization of bilingual word embedding spaces respectively for the English-French and English-Spanish settings. For CNN, GRU, and attentive GRU (ATT) encoders, we stack five of each corresponding encoding layers with hidden-sizes of 200, and two affine layers are applied to the final output for dimension reduction. This encoder architecture consistently represents the best performance through our tuning. Through comprehensive hyperparameter tuning, we fix the learning rate α to 0.0005, the exponential decay rates of AMSGrad β_1 and β_2 to 0.9 and 0.999, coefficients λ_1 and λ_2 to both 0.1, and batch size to 64. Kernel-size and pooling-size are both set to 2 for CNN. Word definitions are zero-padded (short ones) or truncated (long ones) to the sequence length of 15, since most definitions (over 92%) are within 15 words in the dataset. Training is limited to 1,000 epochs for all models as well as the dictionary thread of asynchronous joint learning, in which all models are able to converge.

Results. Results are reported in Table 3 in four groups. The first group compares four different encoding techniques for the basic dictionary models. GRU thereof consistently outperforms CNN and BOW, since the latter two fail to capture the important sequential information for descriptions. ATT that weighs among the hidden states has notable improvements over GRU. While we equip the two better encoding techniques with the monolingual retrieval approach (GRU-mono and ATT-mono), we find that the way of learning the dictionary models towards monolingual targets and

retrieving cross-lingual related words incurs more impreciseness to the task. For models of the third group that conduct multi-task learning in two directions of a language pair, the results show significant enhancement of performance in both directions. For the final group of results, we incorporate the best variant of multi-task models into the joint learning architecture, which leads to compelling improvement of the task on all settings. This demonstrates that properly adapting the word embeddings in joint with the bilingual dictionary model efficaciously constructs the embedding space that suits better the representation of both bilingual lexical and sentential semantics.

In general, this experiment has identified the proper encoding techniques of the dictionary model. The proposed strategies of multi-task and joint learning effectively contribute to the precise characterization of the cross-lingual correspondence of lexical and sentential semantics, which have led to very promising capability of cross-lingual reverse dictionary retrieval.

4.3 Bilingual Paraphrase Identification

The bilingual paraphrase identification problem⁵ is a binary classification task with the goal to decide whether two sentences in different languages express the same meanings. BiDRL provides an effective solution by transferring sentential meanings to word-level representations and learning a simple classifier. We evaluate three variants of BiDRL on this task using WBP31: the multi-task BiDRL with GRU encoders (BiDRL-GRU-MTL), the multi-task BiDRL with attentive GRU encoders (BiDRL-ATT-MTL), and the joint learning BiDRL with with attentive GRU encoders (BiDRL-ATT-joint). We compare against several baselines of neural sentence pair models that are proposed for monolingual paraphrase identification. These models include Siamese structures of CNN (BiCNN) (Yin and Schütze, 2015), RNN (BiLSTM) (Mueller and Thyagarajan, 2016), attentive CNN (ABCNN) (Yin et al., 2016), attentive GRU (BiATT) (Rocktäschel et al., 2016), and BOW (BiBOW). To support the reasoning of cross-lingual semantics, we provide the baselines with the same BiBOWA embeddings.

⁵Paraphrases have similar meanings, but can largely differ in content details and word orders. Hence, they are essentially different from translations. We have found that even the well-recognized Google NMT frequently caused distortions to short sentence meanings, and led to results that were close to random guess by the baseline classifiers after translation.

Languages	En&Fr		En&Es	
Metrics	<i>Acc.</i>	<i>F1</i>	<i>Acc.</i>	<i>F1</i>
BiBOW	54.93	0.622	56.27	0.623
BiCNN	54.33	0.625	53.80	0.611
ABCNN	56.73	0.644	58.83	0.655
BiLSTM	59.60	0.662	57.60	0.637
BiATT	61.47	0.699	61.27	0.689
BiDRL-GRU-MTL	64.80	0.732	63.33	0.722
BiDRL-ATT-MTL	65.27	0.735	66.07	0.735
BiDRL-ATT-joint	68.53	0.785	67.13	0.759

Table 4: Accuracy and F1-scores of bilingual paraphrase identification. For BiDRL, the results by three model variants are reported: BiDRL-GRU-MTL and BiDRL-ATT-MTL are models with bilingual multi-task learning, and BiDRL-ATT-joint is the best ATT-based dictionary model variant deployed with both multi-task and joint learning.

Evaluation protocol. BiDRL transfers each sentence into a vector in the word embedding space. Then, for each sentence pair in the train set, a Multi-layer Perceptron (MLP) with a binary softmax loss is trained on the subtraction of two vectors as a downstream classifier. Baseline models are trained end-to-end, each of which directly uses a parallel pair of encoders with shared parameters and an MLP that is stacked to the subtraction of two sentence vectors. Note that some works use concatenation (Yin and Schütze, 2015) or Manhattan distances (Mueller and Thyagarajan, 2016) of sentence vectors instead of their subtraction (Jiang et al., 2018), which we find to be less effective on small amount of data.

We apply the configurations of the sentence encoders from the last experiment to corresponding baselines, so as to show the performance under controlled variables. Training of a classifier is terminated by early-stopping based on the validation set. Following convention (Hu et al., 2014; Yin et al., 2016), we report the accuracy and F1 scores.

Results. This task is challenging due to the heterogeneity of cross-lingual paraphrases and limitedness of learning resources. The results in Table 4 show that all the baselines, where BiATT consistently outperforms the others, merely reaches slightly over 60% of accuracy on both En-Fr and En-Es settings. We believe that it comes down to the fact that sentences of different languages are often drastically heterogenous in both lexical semantics and the sentence grammar that governs the composition of words. Hence, it is not surprising that previous neural sentence pair models, which capture the semantic relation of bilingual sentences directly from all participating words, fall short at the multilingual task. BiDRL, how-

ever, effectively leverages the correspondence of lexical and sentential semantics to simplify the task to an easier entailment task in the lexicon space, for which the multi-task learning BiLDRL-ATT-MTL outperforms the best baseline respectively by 3.80% and 4.80% of accuracy in both language settings, while BiLDRL-ATT-joint, employing the joint learning, further improves the task by another satisfying 3.26% and 1.06% of accuracy. Both also show notable increment in F1.

5 Conclusion and Future Work

In this paper, we propose a neural embedding model BiLDRL that captures the correspondence of cross-lingual lexical and sentential semantics. We experiment with multiple forms of neural models and identify the best technique. The two learning strategies, bilingual multi-task learning and joint learning, are effective at enhancing the cross-lingual learning with limited resources, and also achieve promising performance on cross-lingual reverse dictionary retrieval and bilingual paraphrase identification tasks by associating lexical and sentential semantics. An important direction of future work is to explore whether the word-sentence alignment can improve bilingual word embeddings. Applying BiLDRL to bilingual question answering and semantic search systems is another important direction.

6 Acknowledgement

We thank the anonymous reviewers for their insightful comments. This work was supported in part by National Science Foundation Grant IIS-1760523.

References

Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2013. Plyglot: Distributed word representations for multilingual nlp. In *The SIGNLL Conference on Computational Natural Language Learning*.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2016. Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2289–2294.

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), volume 1, pages 451–462.

- Colin Bannard and Chris Callison-Burch. 2005. Paraphrasing with bilingual parallel corpora. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 597–604. Association for Computational Linguistics.
- Sarath Chandar, Stanislas Lauly, Hugo Larochelle, Mitesh Khapra, Balaraman Ravindran, Vikas C Raykar, and Amrita Saha. 2014. An autoencoder approach to learning bilingual word representations. In *Advances in Neural Information Processing Systems*, pages 1853–1861.
- Wanxiang Che, Yijia Liu, Yuxuan Wang, Bo Zheng, and Ting Liu. 2018. Towards better UD parsing: Deep contextualized word embeddings, ensemble, and treebank concatenation. In *CoNLL Shared Task*, pages 55–64. ACL.
- Muhao Chen, Chang Ping Meng, Gang Huang, and Carlo Zaniolo. 2018a. Neural article pair modeling for wikipedia sub-article machine. In *Proceedings of European Conference of Machine Learning*.
- Muhao Chen, Yingtao Tian, Kai-Wei Chang, Steven Skiena, and Carlo Zaniolo. 2018b. Co-training embeddings of knowledge graphs and entity descriptions for cross-lingual entity alignment. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734.
- Sumit Chopra, Michael Auli, and Alexander M Rush. 2016. Abstractive sentence summarization with attentive recurrent neural networks. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 93–98.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680.
- Jocelyn Coulmance, Jean-Marc Marty, Guillaume Wenzek, and Amine Benhalloum. 2015. Transgram, fast cross-lingual word-embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1109–1113.

- Dipanjan Das and Noah A Smith. 2009. Paraphrase identification as probabilistic quasi-synchronous recognition. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing*, pages 468–476. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Jacob Devlin, Rabih Zbib, Zhongqiang Huang, Thomas Lamar, Richard Schwartz, and John Makhoul. 2014. Fast and robust neural network joint models for statistical machine translation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1370–1380.
- Yerai Doval, Jose Camacho-Collados, Luis Espinosa Anke, and Steven Schockaert. 2018. Improving cross-lingual word embeddings by meeting in the middle. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 294–304.
- Ben Eisner, Tim Rocktäschel, Isabelle Augenstein, Matko Bošnjak, and Sebastian Riedel. 2016. emoji2vec: Learning emoji representations from their description. *arXiv preprint arXiv:1609.08359*.
- Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 462–471.
- Stephan Gouws, Yoshua Bengio, and Greg Corrado. 2015. Bilbowa: Fast bilingual distributed representations without word alignments. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 748–756.
- Felix Hill, KyungHyun Cho, Anna Korhonen, and Yoshua Bengio. 2016. Learning to understand phrases by embedding the dictionary. *Transactions of the Association for Computational Linguistics*, 4:17–30.
- Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. In *Advances in neural information processing systems*, pages 2042–2050.
- Jan H Hulstijn, Merel Hollander, and Tine Greidanus. 1996. Incidental vocabulary learning by advanced foreign language students: The influence of marginal glosses, dictionary use, and reoccurrence of unknown words. *The modern language journal*, 80(3):327–339.
- Guoliang Ji, Kang Liu, Shizhu He, Jun Zhao, et al. 2017. Distant supervision for relation extraction with sentence-level attention and entity descriptions. In *Proceedings of the AAAI International Conference on Artificial Intelligence*, pages 3060–3066.
- Jyun-Yu Jiang, Francine Chen, Yan-Ying Chen, and Wei Wang. 2018. Learning to disentangle interleaved conversational threads with a siamese hierarchical network and similarity ranking. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 1812–1822.
- Nal Kalchbrenner, Edward Grefenstette, Phil Blunsom, Dimitri Kartsaklis, Nal Kalchbrenner, Mehrnoosh Sadrzadeh, Nal Kalchbrenner, Phil Blunsom, Nal Kalchbrenner, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 212–217. Association for Computational Linguistics.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86.
- Jingzhou Liu, Wei-Cheng Chang, Yuexin Wu, and Yiming Yang. 2017. Deep learning for extreme multi-label text classification. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 115–124. ACM.
- Thang Luong, Hieu Pham, and Christopher D Manning. 2015a. Bilingual word representations with monolingual quality in mind. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 151–159.
- Thang Luong, Hieu Pham, and Christopher D Manning. 2015b. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013a. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- Tomas Mikolov, Ilya Sutskever, et al. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous methods for deep reinforcement learning.

- In *International Conference on Machine Learning*, pages 1928–1937.
- Jonas Mueller and Aditya Thyagarajan. 2016. Siamese recurrent architectures for learning sentence similarity. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2786–2792.
- Ani Nenkova, Kathleen McKeown, et al. 2012. A survey of text summarization techniques. In *Mining text data*, pages 43–76. Springer.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 2227–2237.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual bert? In *ACL*.
- Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. 2018. On the convergence of adam and beyond. In *International Conference on Learning Representations*.
- Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, and Phil Blunsom. 2016. Reasoning about entailment with neural attention. In *International Conference on Learning Representations*.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2017. A survey of cross-lingual word embedding models. *Journal of Artificial Intelligence Research*.
- Lei Sha, Baobao Chang, et al. 2016. Reading and thinking: Re-read lstm unit for textual entailment recognition. In *Proceedings of the International Conference on Computational Linguistics*.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *Advances in Neural Information Processing Systems*, pages 6833–6844.
- Robert Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Shyam Upadhyay, Nitish Gupta, and Dan Roth. 2018. Joint multilingual supervision for cross-lingual entity linking. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2486–2495.
- Ivan Vulić, Anna Korhonen, et al. 2016. On the role of seed lexicons in learning bilingual word embeddings. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 247–257.
- Yogarshi Vyas and Marine Carpuat. 2016. Sparse bilingual word representations for cross-lingual lexical entailment. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1187–1197.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1006–1011.
- Wenpeng Yin and Hinrich Schütze. 2015. Convolutional neural network for paraphrase identification. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Wenpeng Yin, Hinrich Schütze, et al. 2016. Abcnn: Attention-based convolutional neural network for modeling sentence pairs. *Transactions of the Association for Computational Linguistics*, 4(1).
- Xinjie Zhou, Xiaojun Wan, and Jianguo Xiao. 2016. Cross-lingual sentiment classification with bilingual document representation learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1403–1412.