

# Cross-Lingual Phrase Retrieval

Heqi Zheng<sup>12\*</sup>, Xiao Zhang<sup>1\*</sup>, Zewen Chi<sup>1\*</sup>, Heyan Huang<sup>12</sup>

Tan Yan<sup>1</sup>, Tian Lan<sup>1</sup>, Wei Wei<sup>3</sup>, Xian-Ling Mao<sup>1†</sup>

<sup>1</sup>School of Computer Science and Technology, Beijing Institute of Technology

<sup>2</sup>Beijing Engineering Research Center of High Volume Language Information Processing and Cloud Computing Applications

<sup>3</sup>Huazhong University of Science and Technology

{hikiz, yotta, czw, hhy63, yt, lt, maoxl}@bit.edu.cn, weiw@hust.edu.cn

## Abstract

Cross-lingual retrieval aims to retrieve relevant text across languages. Current methods typically achieve cross-lingual retrieval by learning language-agnostic text representations in word or sentence level. However, how to learn phrase representations for cross-lingual phrase retrieval is still an open problem. In this paper, we propose XPR, a cross-lingual phrase retriever that extracts phrase representations from unlabeled example sentences. Moreover, we create a large-scale cross-lingual phrase retrieval dataset, which contains 65K bilingual phrase pairs and 4.2M example sentences in 8 English-centric language pairs. Experimental results show that XPR outperforms state-of-the-art baselines which utilize word-level or sentence-level representations. XPR also shows impressive zero-shot transferability that enables the model to perform retrieval in an unseen language pair during training. Our dataset, code, and trained models are publicly available at [github.com/cwszz/XPR/](https://github.com/cwszz/XPR/).

## 1 Introduction

Phrase retrieval aims to retrieve relevant phrases from a large phrase set, which is a critical part of information retrieval. Recent studies on phrase retrieval learn dense representations of phrases, and achieve promising results in entity linking, slot filling, and open-domain question answering tasks (Gillick et al., 2019; Lee et al., 2021a,b). Nonetheless, most of the studies focus on monolingual scenarios, leaving the cross-lingual phrase retrieval unexplored.

Various methods have been proposed to perform cross-lingual text retrieval, which learns cross-lingual word or sentence representations shared

across languages. Cross-lingual word representation methods typically train word embeddings on each language separately, and then learn an embedding mapping between the embedding spaces of different languages (Mikolov et al., 2013; Dinu et al., 2014). Then, the bilingual word pairs can be retrieved between vocabularies using nearest neighbor search, which is also known as bilingual lexicon induction (Artetxe et al., 2018; Lample et al., 2018). Cross-lingual sentence retrieval is typically achieved by learning a sentence encoder on multilingual text corpora with self-supervised pretraining tasks (Conneau and Lample, 2019; Conneau et al., 2020), or large-scale parallel corpora (Artetxe and Schwenk, 2019), or both (Chi et al., 2021b). The trained sentence encoders produce language-agnostic sentence representations, which enables sentences to be retrieved across languages.

Despite the effectiveness of word-level and sentence-level methods, how to learn phrase representations for cross-lingual phrase retrieval is still an open problem. Learning cross-lingual phrase representations is challenging in two aspects. First, a phrase is a conceptual unit containing multiple words, so it is necessary to model the interaction between words, which is not considered in word-level methods. Second, a phrase contains fewer words with less information compared to sentences, which prevents sentence encoders from taking the advantage of the ability of understanding full-length sentences.

Thus, in this paper, we propose a novel cross-lingual phrase retriever named as XPR. Unlike previous cross-lingual retrieval methods that directly encode the input text, XPR produces phrase representations using example sentences, which can be collected from unlabeled text corpora. Initialized with a pretrained cross-lingual language

\* Co-first authors with equal contributions.

† Corresponding author.

model, XPR can either directly serve as an unsupervised retriever, or be further trained to produce better-aligned phrase representations. Besides, we propose the cross-lingual phrase contrast (XPCO) loss for training XPR, where the model is trained to distinguish bilingual phrase pairs from negative examples. Furthermore, we create a cross-lingual phrase retrieval dataset, namely WikiXPR. WikiXPR contains 65K bilingual phrase pairs of eight language pairs, and provides example sentences for each phrase.

We conduct a comprehensive evaluation of XPR on WikiXPR under four evaluation settings, i.e., unsupervised, supervised, zero-shot transfer, and multilingual supervised. Our XPR model substantially outperforms the retrieval baselines based on cross-lingual word embeddings and cross-lingual sentence encoders. XPR also shows impressive zero-shot transferability that enables the model to be trained in a language pair and directly perform phrase retrieval for other language pairs. Moreover, we present an in-depth analysis on XPR, showing that using example sentences improves both the learned XPR model and the phrase representations.

Our contributions are summarized as follows:

- We propose XPR, a novel cross-lingual phrase retriever that utilizes example sentences to produce phrase representations.
- We propose the cross-lingual phrase contrast loss for training XPR.
- We demonstrate the effectiveness of XPR on eight language pairs under four evaluation settings.
- We create a cross-lingual phrase retrieval dataset, which provides 65K bilingual phrase pairs with 4.2M example sentences in 8 language pairs.

## 2 Related Work

**Cross-Lingual Retrieval** Current cross-lingual text retrieval methods focus on word-level and sentence-level scenarios. Word-level cross-lingual retrieval methods typically train word embeddings on each language separately, and then align the word embeddings between languages by learning a mapping function (Mikolov et al., 2013; Dinu et al., 2014; Artetxe et al., 2016, 2018; Lample et al., 2018; Doval et al., 2018;

Joulin et al., 2018). Similarly, cross-lingual sentence retrieval can be achieved by aligning sentence representations across different languages. LASER (Artetxe and Schwenk, 2019) learns a multilingual auto-encoder on multilingual parallel corpora to produce language-agnostic sentence embeddings. Training on parallel corpora, cross-lingual sentence representations can also be learned with neural machine translation (Schwenk, 2018), contrastive learning (Chidambaram et al., 2019; Feng et al., 2020; Chi et al., 2021b), translation span corruption (Chi et al., 2021a), or knowledge distillation (Ham and Kim, 2021). Thanks to the recent language model pretraining technique (Devlin et al., 2019), sentence encoders can be learned on a multilingual unlabeled text corpus without using parallel corpora (Conneau and Lample, 2019; Conneau et al., 2020; Chi et al., 2021c; Goswami et al., 2021).

**Phrase Retrieval** Recent research on phrase retrieval typically learns phrase representations. Seo et al. (2019) propose to treat phrases as the smallest retrieval unit for open-domain question answering, where the phrases are encoded as indexable query-agnostic representations. The retrieval methods can be further improved with self-supervised pretraining, leading to better performance on open-domain question answering (Lee et al., 2021a,b). Additionally, DEER (Gillick et al., 2019) formulates the entity linking task as an entity phrase retrieval problem. However, these works study phrase retrieval in a monolingual scenario while we focus on cross-lingual phrase retrieval.

**Contrastive Learning** Contrastive learning learns representations by a contrastive loss that encourages the positive data pairs to be more similar than other data pairs. It has shown to be effective for learning representations of a wide range of modalities including visual representations (He et al., 2020; Chen et al., 2020a; Grill et al., 2020), sentence representations (Kong et al., 2020; Chi et al., 2021b; Gao et al., 2021), audio representations (Saeed et al., 2021), etc. Different from previous work that performs contrastive learning at sentence level, we introduce contrastive learning to learn phrase representations.

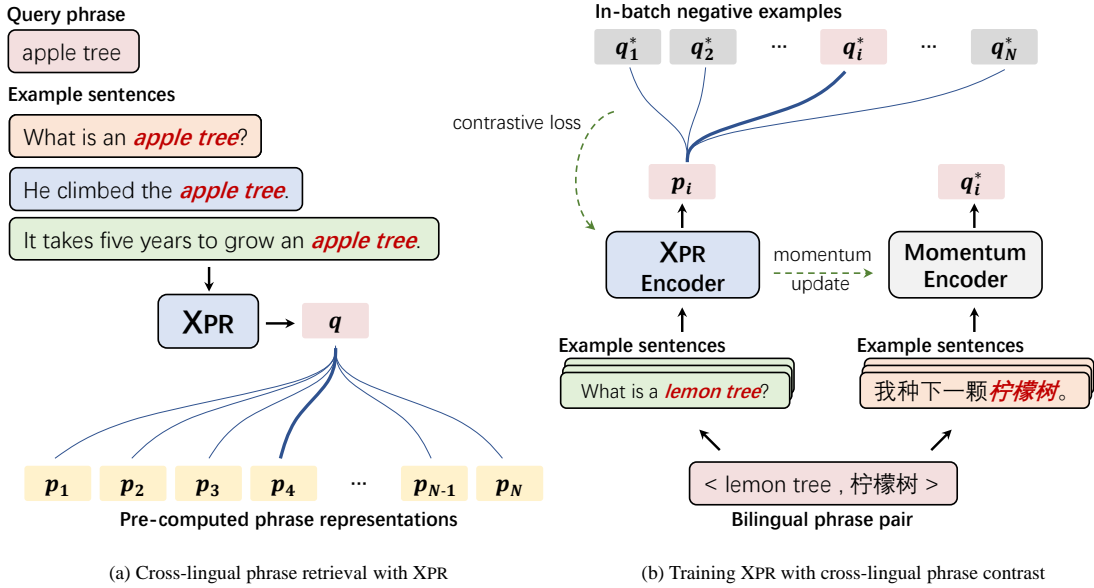


Figure 1: The overview of XPR. (a) The cross-lingual phrase retrieval procedure of XPR. XPR first extracts phrase representations from the example sentences sampled from an unlabeled text corpus, and then performs cross-lingual phrase retrieval with nearest search. (b) The cross-lingual phrase contrast (XPCO) loss. Notice that the example sentences are sampled separately for each phrase, i.e., training without parallel sentences.

### 3 Methods

Figure 1 shows the overview of XPR. In this section, we first introduce the model architecture of XPR, and then present the cross-lingual phrase contrast loss. Finally, we show the training procedure of XPR.

#### 3.1 Model Architecture

The model architecture of XPR is a Transformer (Vaswani et al., 2017) encoder shared across different languages. XPR can be initialized with pretrained cross-lingual language models, which have shown to produce well-aligned sentence representations (Hu et al., 2020; Chi et al., 2021b).

Given a phrase  $p$  and an example sentence  $x = w_1, \dots, w_n$  with  $n$  tokens that contain the phrase. We denote the start and end indices of  $p$  as  $s$  and  $e$ , i.e.,  $p = w_s, \dots, w_e$ . XPR first encodes  $x$  into a sequence of contextualized token representations<sup>1</sup>

$$h_1, \dots, h_n = \text{Transformer}(w_1, \dots, w_n). \quad (1)$$

Then, the phrase is represented as the average of the phrase tokens

$$x = \frac{1}{e - s + 1} \sum_{i=s}^e h_i. \quad (2)$$

<sup>1</sup>Following Chi et al. (2021b), we take the hidden vectors from a specific hidden layer as the token representations rather than only the last layer.

In general, a phrase can have more than one example sentence. Considering  $m$  example sentences  $\mathcal{X} = x_1, \dots, x_m$  for the phrase  $p$ , XPR encodes the sentences separately, and uses the average of the phrase representations as the final phrase representation, i.e.,  $\sum_{x \in \mathcal{X}} x/m$ . Notice that XPR does not introduce additional parameters beyond the original Transformer encoder. Thus, after the initialization with a pretrained cross-lingual language model, XPR can directly serve as an unsupervised cross-lingual phrase retriever.

#### 3.2 Cross-Lingual Phrase Contrast Loss

Recent work (Chen et al., 2020a; Kong et al., 2020) has demonstrated the effectiveness of contrastive learning framework for learning visual and text representations. To learn language-agnostic phrase representations, we propose the cross-lingual phrase contrast (XPCO) loss, where the goal is to distinguish the bilingual phrase pairs from negative examples.

Formally, consider a mini-batch  $\mathcal{B} = \{\mathcal{P}, \mathcal{Q}\}$  of bilingual phrase pairs, where  $\mathcal{P} = \{p\}^N$  and  $\mathcal{Q} = \{q\}^N$  stand for  $N$  phrases in a language and their translations in another language, respectively. For each phrase  $p \in \mathcal{P}$ , we sample example sentences  $\mathcal{X}$  for  $p$ , and compute the phrase representation  $u$  as described in Section 3.1. Following Chen et al. (2020a), we apply a projection head over  $u$  that consists of two linear lay-

---

**Algorithm 1** Training procedure of XPR

---

**Input:** Bilingual phrase pair corpus  $\mathcal{D}$ , unlabeled text corpus  $\mathcal{U}$ , learning rate  $\tau$ , momentum coefficient  $\mu$

**Output:** XPR parameters  $\theta$

```

1: Initialize  $\theta, \theta_m$ 
2: while not converged do
3:    $(\mathcal{P}, \mathcal{Q}) \sim \mathcal{D}$ 
4:   for  $i = 1, 2, \dots, N$  do
5:      $\mathcal{X} \sim \mathcal{U}$  s.t.  $p_i \subset x, \forall x \in \mathcal{X}$ 
6:      $\mathcal{Y} \sim \mathcal{U}$  s.t.  $q_i \subset y, \forall y \in \mathcal{Y}$ 
7:      $\mathbf{p}_i = f(p_i, \mathcal{X}; \theta)$ 
8:      $\mathbf{p}_i^* = f(p_i, \mathcal{X}; \theta_m)$ 
9:      $\mathbf{q}_i = f(q_i, \mathcal{Y}; \theta)$ 
10:     $\mathbf{q}_i^* = f(q_i, \mathcal{Y}; \theta_m)$ 
11:   end for
12:    $\mathbf{g} \leftarrow \nabla_{\theta} \mathcal{L}_{\text{XPCO}}$ 
13:    $\theta \leftarrow \theta - \tau \mathbf{g}$ 
14:    $\theta_m \leftarrow \mu \theta_m + (1 - \mu) \theta$ 
15: end while

```

---

ers with a ReLU in between and a  $l_2$  normalization followed. For simplicity, we denote the above operation that converts an input phrase  $p$  to a normalized vector as  $\mathbf{p} = f(p, \mathcal{X}; \theta)$ , where  $\theta$  stands for the parameters of the encoder and the projection head. For each phrase  $q \in \mathcal{Q}$ , we employ a momentum encoder (He et al., 2020) to encode  $q$ :  $\mathbf{q}^* = f(q, \mathcal{Y}; \theta_m)$ , where  $\mathcal{Y}$  represents the example sentences of  $q$ , and  $\theta_m$  represents the parameters of the momentum encoder.

For the  $i$ -th phrase  $p_i \in \mathcal{P}$ ,  $q_i \in \mathcal{Q}$  is its corresponding positive example and the other  $N - 1$  phrases are treated as negative examples. The contrastive loss in the direction of  $\mathcal{P} \rightarrow \mathcal{Q}$  is defined as

$$\mathcal{L}(\mathcal{P} \rightarrow \mathcal{Q}) = - \sum_{i=1}^N \log \frac{\exp(\mathbf{p}_i^\top \mathbf{q}_i^*/T)}{\sum_{j=1}^N \exp(\mathbf{p}_i^\top \mathbf{q}_j^*/T)} \quad (3)$$

Similarly, we employ an additional contrastive loss in the direction of  $\mathcal{Q} \rightarrow \mathcal{P}$ . The XPCO loss combines both directions, which is defined as

$$\mathcal{L}_{\text{XPCO}} = \mathcal{L}(\mathcal{P} \rightarrow \mathcal{Q}) + \mathcal{L}(\mathcal{Q} \rightarrow \mathcal{P}) \quad (4)$$

where  $T$  is the softmax temperature.

### 3.3 Training Procedure of XPR

Algorithm 1 illustrates the training procedure of XPR. We initialize the XPR encoder  $\theta$  and the

momentum encoder  $\theta_m$  with a pretrained cross-lingual language model. For each training step, we first sample a mini-batch of bilingual phrase pairs  $(\mathcal{P}, \mathcal{Q})$  from the bilingual phrase pair corpus  $\mathcal{D}$ , and then sample example sentences  $\mathcal{X}$  and  $\mathcal{Y}$  for  $\mathcal{P}$  and  $\mathcal{Q}$ , respectively. Each example sentence  $x \in \mathcal{X}$  should contain the phrase  $p_i$ , which is denoted as  $p_i \subset x$ . With the phrase representations produced by the two encoders, we compute the XPCO loss, and update  $\theta$  with gradient descent. Notice that we do not perform back-propagation in the momentum encoder, which is learned by a momentum update (He et al., 2020) with a momentum coefficient of  $\mu$ .

### 3.4 Phrase Retrieval with XPR

Given a phrase set  $\mathcal{P} = \{p\}^N$  with  $N$  candidate phrases, the goal is to find  $p \in \mathcal{P}$  with the same meaning of a query phrase  $q$ . With the trained XPR encoder  $\theta$ , we first sample example sentences candidate phrases and then compute their representations  $\{\mathbf{p}\}^N$  with  $f(\cdot; \theta)$ . Then, for a query phrase  $q$ , we can find the corresponding phrase by:

$$\hat{p} = \arg \max_{p_i} \{\mathbf{p}_i^\top \mathbf{q}\} \quad (5)$$

In practice, the representations of candidate phrases can be pre-computed for reuse. Moreover, although the example sentence number is limited during training, we can use more example sentences to obtain better phrase representation for retrieval.

## 4 WikiXPR: Cross-Lingual Phrase Retrieval Dataset

To evaluate our model, we create WikiXPR, a cross-lingual phrase retrieval dataset extracted from Wikipedia. WikiXPR consists of bilingual phrase pairs in eight English-centric language pairs, and contains large-scale example sentences for the phrases, which enable models to leverage contextual information to better understand phrases. In what follows, we describe how we construct the WikiXPR dataset.

### 4.1 Phrase Pair Mining

Manually translating phrases is expensive when building a large-scale bilingual phrase pair corpus. Therefore, we leverage the link information within Wikipedia for mining bilingual phrase pairs. Specifically, we first extract inter-language

	ar-en	de-en	es-en	fr-en	ja-en	ko-en	ru-en	zh-en	Total
Train	4222	1931	1333	1315	14745	2138	5229	8326	39239
Dev	1408	644	445	438	4915	713	1743	2775	13081
Test	1407	644	445	438	4915	713	1743	2775	13080

Table 1: The number of bilingual phrase pairs for each language pair in WikiXPR.

linked wiki entries from dbpedia<sup>2</sup>. We treat English as the pivot language, and choose a range of diverse languages to build our datasets, so that the models can be evaluated with different language families and scripts. We filter out time expressions, and the phrase pairs with low edit distance using ROUGE-L (Lin, 2004) as the distance measure. The phrase pairs with bidirectional ROUGE-L values higher than 0.5 are removed.

## 4.2 Example Sentence Retrieval

In addition to phrase pairs in diverse languages, XPR also provides example sentences for each phrase, which aims to facilitate the research on phrase representation learning with example sentences. For each phrase, we retrieve example sentences from an unlabeled text corpus. In specific, we first extract raw sentences from Wikipedia dumps as our unlabeled text corpus. Then, we build sentence indices with the Elasticsearch<sup>3</sup> searching engine. For each phrase, we retain the searched sentences with at least 10 more characters than the phrase as the results. Besides, we only retain 32 example sentences for each phrase to keep a reasonable size for the resulting example sentence corpus.

## 4.3 The Resulting WikiXPR Dataset

As shown in Table 1, we present the number of bilingual phrase pairs for each language pair in WikiXPR. The resulting WikiXPR dataset consists of 65,400 phrase pairs in eight language pairs, and 4.2M example sentences in total. For each phrase, WikiXPR provides 32 example sentences extracted from Wikipedia text. WikiXPR is split into training, dev, and test sets by 3:1:1, so WikiXPR can be used for diverse evaluation settings including the supervised setting, cross-lingual zero-shot transfer, etc. See detailed statistics in Appendix A.

<sup>2</sup>[downloads.dbpedia.org/2014/en/](https://downloads.dbpedia.org/2014/en/)

<sup>3</sup>[www.elastic.co](http://www.elastic.co)

# 5 Experiments

In this section, we first present four evaluation settings for cross-lingual phrase retrieval, and describe the models to be compared. Then, we present the experimental results.

## 5.1 Evaluation Settings

We conduct experiments on the cross-lingual phrase retrieval task on our WikiXPR dataset. Detailed description of WikiXPR can be found in Section 4. Since collecting or annotating parallel sentences can be expensive especially for low-resource languages, we only consider unlabeled text and the bilingual pairs provided by WikiXPR in our experiments. According to the difference in the training resource, we present the following four evaluation settings.

**Unsupervised** Under the unsupervised setting, the retrieval model should not use any bilingual phrase pairs or other cross-lingual supervision such as bilingual dictionaries and parallel corpus. The language representations are typically learned from unlabeled text corpora.

**Supervised** In the supervised setting, the retrieval model is trained on and tested on bilingual phrase pairs for each language pair separately, e.g., training and testing with English-French phrase pairs.

**Zero-Shot Transfer** Zero-shot transfer is a widely-used setting in cross-lingual understanding tasks (Conneau and Lample (2019); Wu and Dredze (2019)), where models are trained in a source language but evaluated on other languages. We introduce this setting to the cross-lingual phrase retrieval task, e.g., training a model with English-French phrase pairs but performing retrieval between English and Chinese phrases.

**Multilingual Supervised** In this setting, the retrieval model is able to use training data in multiple languages, e.g., training a model using a combined training set over all languages in WikiXPR and testing it for each language.

## 5.2 Baselines

Considering the lack of methods for cross-lingual phrase retrieval, we develop the following two baselines in our experiments:

Model	ar-en	de-en	en-es	en-fr	en-ja	en-ko	en-ru	en-zh	Avg
<i>Unsupervised</i>									
CLWE	2.74	0.78	0.00	1.02	0.34	0.28	1.32	0.12	0.83
CLSE	9.70	19.10	29.21	20.89	4.83	11.50	16.98	8.76	15.12
XPR	<b>14.71</b>	<b>28.96</b>	<b>42.25</b>	<b>39.38</b>	<b>7.34</b>	<b>15.22</b>	<b>24.24</b>	<b>11.26</b>	<b>22.92</b>
<i>Supervised</i>									
CLWE	56.14	33.62	63.71	51.26	31.62	50.14	38.67	30.02	44.40
CLSE	20.58	18.79	36.06	26.60	16.73	24.58	21.32	17.69	22.79
XPR	<b>88.63</b>	<b>81.44</b>	<b>84.53</b>	<b>80.18</b>	<b>87.32</b>	<b>80.83</b>	<b>91.00</b>	<b>77.62</b>	<b>83.94</b>
<i>Zero-shot transfer</i>									
CLWE	0.04	0.32	0.22	0.23	0.00	2.24	0.09	30.02	4.15
CLSE	6.18	10.25	16.07	10.39	6.73	9.75	8.35	17.69	10.68
XPR	<b>74.12</b>	<b>73.60</b>	<b>82.54</b>	<b>77.36</b>	<b>73.04</b>	<b>78.52</b>	<b>79.10</b>	<b>77.62</b>	<b>76.99</b>
<i>Multilingual supervised</i>									
CLWE	12.33	1.87	6.63	3.77	18.46	4.00	9.84	11.19	8.51
CLSE	11.98	19.64	29.44	21.58	11.91	14.73	18.01	14.50	17.72
XPR	<b>91.90</b>	<b>82.76</b>	<b>90.79</b>	<b>85.16</b>	<b>90.16</b>	<b>88.22</b>	<b>93.09</b>	<b>86.47</b>	<b>88.57</b>

Table 2: Accuracy@1 on WikiXPR cross-lingual phrase retrieval under four evaluation settings. Results are averaged over three random seeds in both the  $xx \rightarrow en$  and  $en \rightarrow xx$  directions, where ‘xx’ denotes one of the eight non-English languages.

**CLWE** Cross-lingual word embeddings (CLWE) encode words from various languages into a shared embedding space. For each word in a phrase, we first represent it with the pretrained fastText multilingual word vectors (Grave et al., 2018), and then map it to a shared embedding space via the VECMAP<sup>4</sup> (Artetxe et al., 2018) tool. Notice that VECMAP can be applied to both unsupervised and supervised scenarios. Finally, the retrieval is achieved by the nearest search using an average word vector as the phrase representation.

**CLSE** Cross-lingual sentence encoders (CLSE) produce language-agnostic sentence representations for the input text sequence. We use XLM-R<sub>base</sub> (Conneau et al., 2020) as the sentence encoder, which is pretrained on a large-scale multilingual text corpus. For the unsupervised setting, we use the averaged hidden vector from a specific middle layer as the phrase representation. For the other settings, we follow Wang et al. (2019), which learns an orthogonal mapping between the feature spaces of the training phrase pairs. As LASER (Artetxe and Schwenk, 2019) and LaBSE (Feng et al., 2020) utilize parallel corpora, we do not use them in our experiments.

As for our model XPR described in Section 3, we initialize XPR with XLM-R<sub>base</sub> (Conneau et al., 2020) for a fair comparison. For each step, we use a batch of 256 phrase pairs and 4 example sen-

tences for each phrase. The model is optimized with the Adam (Kingma and Ba, 2015) optimizer with a learning rate of  $2 \times 10^{-5}$  for 100 epochs. The learning rate is scheduled with 1% warm-up steps and a linear decay during training.

### 5.3 Experimental Results

Table 2 compares the three cross-lingual phrase retrieval models on our WikiXPR dataset under four different evaluation settings.

**Unsupervised Results** As present in Table 2, XPR obtains the best performance over all languages without any cross-lingual supervision, achieving an average accuracy@1 of 22.92. On the contrary, CLWE and CLSE only obtain 0.83 and 15.12, respectively. It indicates that XPR successfully leverage example sentences to produce better phrase representations. Besides, the performance varies in different languages. We observe that the retrieval between English and European languages can be easier than other language pairs when using CLSE and XPR. It is worth mentioning that CLWE and CLSE are proven to be effective for bilingual lexicon induction and cross-lingual sentence retrieval, respectively (Lample et al., 2018; Hu et al., 2020). Nonetheless, they do not perform as well as on word or sentence level tasks, indicating that they are not directly applicable to cross-lingual phrase retrieval.

<sup>4</sup>[github.com/artetxem/vecmap](https://github.com/artetxem/vecmap)

**Supervised Results** Under the supervised setting, XPR achieves an average accuracy of 83.94, largely outperforming the other two models over all evaluation language pairs. Comparing the results between the unsupervised and the supervised settings, all the three models greatly benefit from the training data. In particular, XPR pushes the average result from 7.34 to 87.32 for the en-ja phrase retrieval. The results suggest that the bilingual phrase pairs can help to learn cross-lingual alignment for both word-level and sentence-level representations. We find that using training data brings more gains for CLWE than CLSE, showing that the contextualized phrase representations in CLSE can be harder to align.

**Zero-shot Transfer** In zero-shot transfer, the models are trained using an en-xx dataset but evaluated on all language pairs. The table only presents the results of the model trained on en-zh data. Detailed results of other transfer directions can be found in Appendix B. Although the XPR model only learns on en-zh training data, it performs surprisingly well on other languages. On en-es and en-ko, XPR even produces comparable results to the results in the supervised setting. Comparing the results to the unsupervised setting, XPR pushes the average accuracy from 22.92 to 76.99. This demonstrates the strong cross-lingual transferability of XPR, which allows our model to be applied to low-resource languages without training data. On the contrary, CLSE fails to leverage the en-zh training data for the retrieval in other languages, resulting in a consistent performance drop.

**Multilingual Supervised** In the multilingual supervised setting, XPR obtains the best results over all models and settings, achieving an average accuracy of 88.57. Compared to the supervised setting, using the combined training data leads to consistent improvement over all languages, which demonstrates that XPR successfully leverage the supervision signals from both the same and different languages.

#### 5.4 Ablation Studies

We conduct ablation studies by removing main components from XPR. In specific, we compare three variants of XPR that are trained without example sentences, momentum update, or projection head, respectively. The evaluation results are shown in Table 3.

Model	en-fr	en-ko	en-ru	Avg
<i>Supervised</i>				
XPR	<b>80.18</b>	<b>80.83</b>	<b>91.00</b>	<b>84.00</b>
–Example Sentence	60.84	63.39	83.39	69.21
–Momentum Update	78.54	80.43	90.50	83.16
–Projection Head	77.28	77.70	89.16	81.38
<i>Zero-shot transfer</i>				
XPR	<b>77.36</b>	<b>78.52</b>	<b>79.10</b>	<b>78.33</b>
–Example Sentence	60.50	60.03	60.90	60.48

Table 3: Ablation results of XPR on WikiXPR cross-lingual phrase retrieval. We report the average accuracy@1 scores that are averaged in both the xx→en and en→xx directions. In zero-shot transfer, the models are trained using the en-zh data but evaluated on three other language pairs.




Figure 2: Effects of the example sentence number. We train XPR models on the en-fr set of WikiXPR under two settings: 1) Training and inference with various numbers of example sentences for each phrase, 2) Training with 32 example sentences for each phrase but inference with various numbers of example sentences.

**Example Sentence** We first investigate whether using example sentences helps cross-lingual phrase retrieval. During training, we remove the example sentences from XPR, i.e., the model extracts the phrase representation only from the input phrase itself. As shown in Table 3, removing example sentences substantially harms the performance of XPR for both the supervised and zero-shot transfer settings. Notice that example sentences are not parallel across languages, but they still make the resulting phrase representations from different languages better aligned. Besides, compared to the supervised setting, the gains are even larger for zero-shot transfer, improving the average accuracy from 60.48 to 78.33. The above results demonstrate that using example sentences not only learns better phrase representations, but also encourages cross-lingual alignment.

**Projection Head** We train a XPR model without the projection head, i.e., directly using the aver-

Layer	1	2	3	4	5	6	7	8	9	10	11	12
XPR (Unsupervised)	29.22	29.22	31.39	35.73	36.64	37.67	35.50	36.07	35.84	37.10	<b>39.38</b>	29.11
XPR (Supervised)	30.94	30.06	36.45	42.13	48.13	51.52	51.29	55.82	59.51	70.85	78.01	<b>80.18</b>

Table 4: Effects of using phrase representations from different hidden layers of XPR. We report the accuracy@1 on the en-fr set of WikiXPR, where the results are averaged in both the fr→en and en→fr directions.

Objective	en-fr	en-ko	en-ru	Avg
XPCO	<b>80.18</b>	<b>80.83</b>	<b>91.00</b>	<b>84.00</b>
MoCo	77.28	74.05	87.95	79.76

Table 5: Comparison of XPCO and MoCo for training XPR. We report the accuracy@1 on WikiXPR, where the results are averaged in both the xx→en and en→xx directions.

age of the hidden vectors as the phrase representation. As shown in Table 3, the projection head provides consistent gains on the three language pairs, showing the effectiveness of the projection head in contrastive learning. The results also agree with the finding in visual representation learning (Chen et al., 2020a,b).

**Momentum Update** We study the effects of momentum update used in XPR. It shows that the momentum update strategy slightly improves the results on all of the three evaluation language pairs, providing 0.84 accuracy improvement.

## 5.5 Effects of Example Sentence Number

We study the effects of the example sentence number used in XPR. We conduct an evaluation on the en-fr set of WikiXPR, under two settings where the example sentence number varies during training or inference: 1) Training and inference with various numbers of example sentences for each phrase, 2) Training with 32 example sentences for each phrase but inference with various numbers of example sentences.

Figure 2 illustrates the evaluation results. It shows a trend that using more example sentences during inference notably improves the performance in both settings. The gain is larger when using fewer example sentences, demonstrating the effectiveness of using multiple example sentences for producing phrase representations. Comparing the results between the two settings, we find that the model moderately benefits from a large number of example sentences if we use a lower sentence number for inference. Although using more example sentences during training provides gains,

the heavier computation load should be taken into consideration.

## 5.6 Effects of Layer

Recent work (Chi et al., 2021b,c) has shown that a middle layer can produce better-aligned sentence representations than the last layer, resulting in higher cross-lingual sentence retrieval performance. We investigate which hidden layer of XPR produces phrase representations that achieve higher retrieval accuracy. To this end, we evaluate XPR using representations from various hidden layers on the en-fr set of WikiXPR.

As shown in Table 4, we present the evaluation results of XPR under both the unsupervised and the supervised settings. For the unsupervised XPR, we observe that Layer-11 produces the best results while the last layer even performs worse than the first layer. Differently, the supervised XPR obtains the best results on Layer-12, indicating that our XPCO loss encourages the model to fully utilize the last few layers. Moreover, it shows that using representations from higher layers of the supervised XPR leads to consistent improvement.

## 5.7 Comparison of Contrast Losses

We explore whether using momentum contrast (MoCo; He et al. 2020) trains our XPR model better, which is proven to be effective for cross-lingual language model pretraining (Chi et al., 2021b). In specific, we train a variant of XPR with MoCo, which maintains more negative examples encoded by the momentum encoder with a queue with a length of 1024. The evaluation results are presented in Table 5. XPCO consistently outperforms MoCo on the three language pairs, suggesting that the negative examples stored in the queue can be out-of-date for contrastive learning.

## 6 Conclusion

In this work, we propose a cross-lingual phrase retriever XPR, which outperforms the baseline retrievers on a range of diverse languages. Moreover, we create a cross-lingual phrase retrieval



dataset that contains diverse languages with large-scale example sentences. For future work, we would like to improve XPR by: 1) extending XPR to asymmetric retrieval scenarios such as open-domain question answering, 2) exploring how to utilize parallel corpora for training XPR.

## 7 Ethical considerations

XPR is designed as a cross-lingual phrase retriever that retrieve relevant phrases across different languages. We believe XPR would help the communication between the people who speak different languages. Besides, our work can facilitate the research on multilingual natural language processing (NLP), which helps to build NLP applications for low-resource languages. In addition, we construct the WikiXPR dataset using open-source data from Wikipedia and dbpedia.

## Acknowledgements

The work is supported by National Key R&D Plan (No. 2018YFB1005100), National Natural Science Foundation of China (No. U19B2020, 62172039, 61732005, 61602197 and L1924068), the funds of Beijing Advanced Innovation Center for Language Resources (No. TYZ19005), and in part by CCF-AFSG Research Fund under Grant No.RF20210005, and in part by the fund of Joint Laboratory of HUST and Pingan Property & Casualty Research (HPL). We would like to acknowledge Qian Liu for the helpful discussions.

## References

- 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. 2018. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 789–798.
- Mikel Artetxe and Holger Schwenk. 2019. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7:597–610.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020a. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. 2020b. [Big self-supervised models are strong semi-supervised learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 22243–22255. Curran Associates, Inc.
- Zewen Chi, Li Dong, Shuming Ma, Shaohan Huang, Saksham Singhal, Xian-Ling Mao, Heyan Huang, Xia Song, and Furu Wei. 2021a. [mT6: Multilingual pretrained text-to-text transformer with translation pairs](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1671–1683, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021b. [InfoXLM: An information-theoretic framework for cross-lingual language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online. Association for Computational Linguistics.
- Zewen Chi, Shaohan Huang, Li Dong, Shuming Ma, Saksham Singhal, Payal Bajaj, Xia Song, and Furu Wei. 2021c. [XLM-E: Cross-lingual language model pre-training via ELECTRA](#). *ArXiv*, abs/2106.16138.
- Muthu Chidambaram, Yinfei Yang, Daniel Cer, Steve Yuan, Yunhsuan Sung, Brian Strope, and Ray Kurzweil. 2019. [Learning cross-lingual sentence representations via a multi-task dual-encoder model](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 250–259, Florence, Italy. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, pages 7057–7067. Curran Associates, Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference*

- of the North American Chapter of the Association for Computational Linguistics: *Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Georgiana Dinu, Angeliki Lazaridou, and Marco Baroni. 2014. Improving zero-shot learning by mitigating the hubness problem. *arXiv preprint arXiv:1412.6568*.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Matthew Cer, N. Arivazhagan, and Wei Wang. 2020. Language-agnostic bert sentence embedding. *ArXiv*, abs/2007.01852.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldridge, Eugene Ie, and Diego Garcia-Olano. 2019. [Learning dense representations for entity retrieval](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 528–537, Hong Kong, China. Association for Computational Linguistics.
- Koustava Goswami, Sourav Dutta, Haytham Assem, Theodorus Franssen, and John P. McCrae. 2021. [Cross-lingual sentence embedding using multi-task learning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9099–9113, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. Learning word vectors for 157 languages. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. 2020. [Bootstrap your own latent - a new approach to self-supervised learning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 21271–21284. Curran Associates, Inc.
- Jiyeon Ham and Eun-Sol Kim. 2021. [Semantic alignment with calibrated similarity for multilingual sentence embedding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1781–1791, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. Momentum contrast for unsupervised visual representation learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization. *arXiv preprint arXiv:2003.11080*.
- Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. 2018. [Loss in translation: Learning bilingual word mapping with a retrieval criterion](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2984, Brussels, Belgium. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations*, San Diego, CA.
- Lingpeng Kong, Cyprien de Masson d’Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. 2020. [A mutual information maximization perspective of language representation learning](#). In *International Conference on Learning Representations*.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. [Word translation without parallel data](#). In *International Conference on Learning Representations*.
- Jinhyuk Lee, Mujeen Sung, Jaewoo Kang, and Danqi Chen. 2021a. [Learning dense representations of phrases at scale](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6634–6647, Online. Association for Computational Linguistics.
- Jinhyuk Lee, Alexander Wettig, and Danqi Chen. 2021b. Phrase retrieval learns passage retrieval, too. *ArXiv*, abs/2109.08133.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *ACL 2004*.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- Aaqib Saeed, David Grangier, and Neil Zeghidour. 2021. Contrastive learning of general-purpose audio

representations. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3875–3879. IEEE.

Holger Schwenk. 2018. [Filtering and mining parallel data in a joint multilingual space](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 228–234, Melbourne, Australia. Association for Computational Linguistics.

Minjoon Seo, Jinhyuk Lee, Tom Kwiatkowski, Ankur Parikh, Ali Farhadi, and Hannaneh Hajishirzi. 2019. [Real-time open-domain question answering with dense-sparse phrase index](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4430–4441, Florence, Italy. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, pages 5998–6008. Curran Associates, Inc.

Yuxuan Wang, Wanxiang Che, Jiang Guo, Yijia Liu, and Ting Liu. 2019. [Cross-lingual BERT transformation for zero-shot dependency parsing](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5721–5727, Hong Kong, China. Association for Computational Linguistics.

Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 833–844, Hong Kong, China. Association for Computational Linguistics.

## Appendix

### A Additional WikiXPR Statistics

Table 6 presents the detailed statistics of WikiXPR, including the number of phrase pairs, average phrase length, and the average length of example sentences.

### B Detailed Results of Zero-Shot Transfer

Table 7 presents the evaluation results of XPR on WikiXPR under the zero-shot transfer setting, where the XPR model is trained in a source language pair but evaluated on target language pairs.

	<b>ar-en</b>	<b>de-en</b>	<b>en-es</b>	<b>en-fr</b>	<b>en-ja</b>	<b>en-ko</b>	<b>en-ru</b>	<b>en-zh</b>
<i>Train</i>								
#Phrase pairs	4222	1931	1333	1315	14745	2138	5229	8326
Avg phrase length (xx/en)	2.32 / 2.55	1.88 / 2.72	3.26 / 2.79	3.05 / 2.81	7.73 / 2.48	5.18 / 2.35	1.95 / 2.55	4.47 / 2.48
Avg example sentence length (xx/en)	38.28 / 25.36	33.93 / 35.61	42.13 / 26.24	34.75 / 27.63	116.57 / 32.01	130.63 / 25.18	32.26 / 27.56	97.30 / 33.25
<i>Dev</i>								
#Phrase pairs	1408	644	445	438	4915	713	1743	2775
Avg phrase length (xx/en)	2.32 / 2.55	1.89 / 2.64	3.31 / 2.72	3.13 / 2.88	7.68 / 2.45	5.32 / 2.35	1.96 / 2.56	4.49 / 2.46
Avg example sentence length (xx/en)	38.04 / 24.58	34.10 / 35.53	43.02 / 26.74	36.39 / 29.25	116.28 / 32.00	128.20 / 24.50	31.40 / 26.06	98.54 / 33.84
<i>Test</i>								
#Phrase pairs	1407	644	445	438	4915	713	1743	2775
Avg phrase length (xx/en)	2.33 / 2.57	1.89 / 2.68	3.24 / 2.83	3.04 / 2.75	7.75 / 2.47	5.13 / 2.34	1.93 / 2.56	4.48 / 2.50
Avg example sentence length (xx/en)	27.75 / 24.94	33.87 / 35.94	42.82 / 24.99	34.65 / 28.62	117.46 / 31.94	130.61 / 25.24	31.63 / 27.02	97.64 / 33.00

Table 6: Statistics of WikiXPR. For each language pair, we present the number of phrase pairs, average phrase length, and average length of example sentences. Notice that the length is counted by characters for Japanese (ja), Korean (ko), and Chinese (zh), and counted by words for the other languages.

<b>src \ trg</b>	<b>ar-en</b>	<b>de-en</b>	<b>en-es</b>	<b>en-fr</b>	<b>en-ja</b>	<b>en-ko</b>	<b>en-ru</b>	<b>en-zh</b>
<b>ar-en</b>	88.63	68.01	79.55	73.52	63.82	73.82	83.58	55.74
<b>de-en</b>	53.80	81.44	84.16	78.84	46.93	61.92	70.52	47.61
<b>en-es</b>	51.96	70.99	84.53	79.15	42.47	57.76	67.65	43.60
<b>en-fr</b>	51.15	71.17	83.60	80.18	41.40	57.20	65.88	43.16
<b>en-ja</b>	78.93	76.19	82.70	77.85	87.32	84.66	87.25	74.24
<b>en-ko</b>	67.41	66.05	76.78	72.34	62.34	80.83	77.62	56.93
<b>en-ru</b>	73.48	70.58	81.23	76.03	64.33	73.58	91.00	54.62
<b>en-zh</b>	74.12	73.60	82.54	77.36	73.04	78.52	79.10	77.62

Table 7: Evaluation results of XPR on WikiXPR under the zero-shot transfer setting. ‘src’ denotes the source language pair for training. ‘trg’ denotes the target language pair for evaluation.