

Cross-Lingual Low-Resource Set-to-Description Retrieval for Global E-Commerce

Juntao Li,^{1,2*} Chang Liu,^{1,2*} Jian Wang,³ Lidong Bing,³
Hongsong Li,³ Xiaozhong Liu,⁴ Dongyan Zhao,^{1,2} Rui Yan^{1,2†}

¹Center for Data Science, Academy for

Advanced Interdisciplinary Studies, Peking University, Beijing, China

²Wangxuan Institute of Computer Technology, Peking University, Beijing, China

³ DAMO Academy, Alibaba Group

⁴ Indiana University, Bloomington, USA

{lijuntao, liuchang97, zhaody, ruiyan}@pku.edu.cn, binglidong@gmail.com

{eric.wj, hongsong.lhs}@alibaba-inc.com, liu237@indiana.edu

Abstract

With the prosperous of cross-border e-commerce, there is an urgent demand for designing intelligent approaches for assisting e-commerce sellers to offer local products for consumers from all over the world. In this paper, we explore a new task of cross-lingual information retrieval, i.e., cross-lingual set-to-description retrieval in cross-border e-commerce, which involves matching product attribute sets in the source language with persuasive product descriptions in the target language. We manually collect a new and high-quality paired dataset, where each pair contains an unordered product attribute set in the source language and an informative product description in the target language. As the dataset construction process is both time-consuming and costly, the new dataset only comprises of 13.5k pairs, which is a low-resource setting and can be viewed as a challenging testbed for model development and evaluation in cross-border e-commerce. To tackle this cross-lingual set-to-description retrieval task, we propose a novel cross-lingual matching network (CLMN) with the enhancement of context-dependent cross-lingual mapping upon the pre-trained monolingual BERT representations. Experimental results indicate that our proposed CLMN yields impressive results on the challenging task and the context-dependent cross-lingual mapping on BERT yields noticeable improvement over the pre-trained multi-lingual BERT model.

Introduction

Cross-border online shopping has become popular in the past few years, and e-commerce platforms start to pay more effort in improving their user experience. To increase the conversion rate, one critical aspect is how to display products to overseas customers with persuasive and informative descriptions. Fortunately, there already exist various local and cross-national e-commerce platforms, e.g., Amazon, eBay, Taobao, Lazada, with multi-lingual persuasive descriptions of products from all over the world, making it possible for retrieving a feasible product description in the foreign language to a local product. Compared with generation-based methods, retrieved product descriptions are more un-

Product attribute set:

男士, 长袖, 针织衫, 男, 2018, 春秋, 新款, 圆领, 套头, 薄款, 羊毛, 打底衫, 百搭, 潮 (man, long sleeves, sweater, male, 2018, spring and autumn, round collar, sleeve, thin section, wool, bottoming shirt, wild joker, fashion)

Persuasive description:

This handsome knit sweater, the selected fabric makes the hand feel comfortable and delicate, the soft drape shoulder line and the sense sleeves combined with the long contour, contracted and not simple design, very recreational joker, let you handsome and natural and unrestrained appear capable.

Figure 1: An example training pair sampled from the collected dataset, where a Chinese product attribute set is given for retrieving a proper persuasive description in English.

derstandable and accessible to overseas users, and training a retrieval model requires less paired data. Accordingly, we formulate this problem as a cross-lingual information retrieval (CLIR) task for global e-commerce, i.e., ranking foreign product descriptions against a local product with an attribute set (Litschko et al. 2018; Zhang et al. 2019).

Conventional CLIR systems mainly comprise of two components: machine translation and monolingual information retrieval. According to the translation direction, these systems are further categorized into translating the local language to the foreign one and translating the foreign language to the local one. By doing so, the CLIR task is converted to a monolingual setting through the machine translation system (Nie 2010; Rücklé, Swarnkar, and Gurevych 2019). Although these systems are conceptual simple and natural, they are restricted by the performance of machine translation system. Thus, directly modeling CLIR with recent deep neural information retrieval models is promising (Hui et al. 2018; McDonald, Brokos, and Androutsopoulos 2018). However, while performing well, existing deep IR models require a large amount of labeled training data, which is expensive and costly in our cross-lingual retrieval scenario.

In this paper, we explore to leverage a deep neural model to directly solve the low-resource cross-lingual information retrieval task without machine translation process. Herein,

*Equal contribution

†Corresponding author: Rui Yan (ruiyan@pku.edu.cn)

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the critical component of deep IR model is to capture the matching score between product attribute sets in the source language and product descriptions in the target language. As there is not existing open dataset for CLIR in e-commerce, we hire domain experts to collect a heuristic dataset with 13.5K pairs to facilitate the development of cross-lingual retrieval in e-commerce and address the shortage of dataset. Note that since present retrieval/searching systems of cross-border e-commerce platform are not yet applicable to obtain product description candidates in the target language for re-ranking, we mainly focus on designing matching model to solve this low-resource cross-lingual problem directly. Figure 1 illustrates an example in the collected dataset, where an attribute set (i.e., keywords) of a product in the source language has a matched persuasive description in the target language. Specifically, the product attribute sets are collected from a local e-commerce platform, where there exist various overlapping and relations inside the sets to improve the clicking rate. The corresponded heuristic product descriptions in the target language are created by domain experts, and only leverage part of the attribute set considering the overlapping and relations inside the set.

To address the obstacle of lacking paired data and the unordered product attribute sets, we propose a novel model with the enhancement of unsupervised monolingual pre-training, semi-supervised cross-lingual alignment, and a supervised fine-tuning matching model. Concretely, we propose to leverage the deep attention module in neural translation (Vaswani et al. 2017) and response selection (Zhou et al. 2018) for modeling the unordered product attributes and copious cross-lingual correlations between the matched pairs in two different languages. Besides, motivated by the appealing performance of unsupervised monolingual pre-trained models (e.g., ELMo (Peters et al. 2018), BERT (Devlin et al. 2019)), we further incorporate a semi-supervised context-dependent cross-lingual mapping mechanism upon the monolingual BERT representations. By doing so, the data scarcity issue is significantly mitigated.

In a nutshell, our contributions are as follows. (1) We gathered a new and challenging heuristic dataset for cross-lingual retrieval in e-commerce. The dataset will be given by asking. (2) We proposed a cross-lingual matching network (CLMN) for addressing this task, which incorporates deep attention mechanism for modeling unordered attribute sets and cross-lingual correlations, and a context-dependent cross-lingual mapping learning upon BERT. The code is available on <https://github.com/LiuChang97/CLMN>. (3) We conducted extensive experiments on the challenging dataset, including evaluating the performance of fine-tuned multi-lingual BERT representations, exploring the performance of strong machine translation based CLIR methods, comparing with the state-of-the-art directly CLIR models. (4) Experimental results indicate that our proposed CLMN model achieves promising performance on the low-resource cross-lingual retrieval task. With the enhancement of context-dependent mapping upon BERT, our directly cross-lingual retrieval model achieves significant performance improvement over the fine-tuned Multi-lingual BERT and state-of-the-art CLIR models.

Related Work

Information Retrieval Models

Information retrieval models for cross-lingual tasks can be roughly categorized into two groups. The first one is to learn matching signals across different language space directly. For instance, bag-of-words addition has been proved effective (Vulić and Moens 2015) for matching degree calculation, where each text is represented by adding its bilingual word embeddings representation together. With the enhancement of the TF-IDF algorithm, bag-of-words addition model further achieves a performance improvement, and the term-by-term query translation model yields the state-of-the-art performance for unsupervised cross-lingual retrieval (Litschko et al. 2018). Another group of methods transfer the cross-lingual retrieval task to a monolingual retrieval task by using machine translation systems (Schuster et al. 2019), e.g., combining cross-language tree kernel align with neural machine translation system for retrieval (Da San Martino et al. 2017), learning language invariant representations for cross-lingual question re-ranking (Joty et al. 2017).

Low-Resource Cross-Lingual Learning

Existing low-resource cross-lingual learning methods mainly contain three elements, i.e., off-the-shelf monolingual word embedding training algorithm (Bojanowski et al. 2017), cross-lingual mapping learning, and the refinement procedure. In cross-lingual mapping learning, a linear mapping between the source embeddings space and the target embeddings space is learned in an adversarial fashion (Conneau et al. 2018a). To enhance the quality of learned bilingual word embeddings, various refinement strategies are proposed, such as synthetic parallel vocabulary building (Artetxe, Labaka, and Agirre 2017), orthogonal constraint (Smith et al. 2017), cross-domain similarity local scaling (Søgaard, Ruder, and Vulić 2018), self-boosting (Artetxe, Labaka, and Agirre 2018), byte-pair encodings (Sennrich, Haddow, and Birch 2016; Lample et al. 2018). Alternatively, a context-dependent cross-lingual representation mapping based on pre-trained ELMo (Peters et al. 2018) is proposed recently to boost the performance of cross-lingual learning. (Schuster et al. 2019). Unlike previous work, we propose to train a cross-lingual mapping upon the context-dependent monolingual BERT with an effective refinement strategy.

Existing Datasets

There already exists several cross-lingual text retrieval datasets. Most of them are collected from open-domain and comprise of comparable/parallel document pairs, e.g., WaCky translation dataset (Baroni et al. 2009; Joulin et al. 2018), benchmark CLEF corpora (Vulić and Moens 2015), the Askubuntu benchmark corpus in QA task (Dos Santos et al. 2015; Barzilay et al. 2016), Arabic–English language pairs (Da San Martino et al. 2017), text stream alignment (Ge et al. 2018), English–German cross-lingual information retrieval dataset based on English–German Wikipedia (Schamoni et al. 2014). Later on, Sasaki et al. (2018) follow previous dataset construction method (Schamoni et al. 2014) and collect 25 cross-lingual datasets with large scales based on

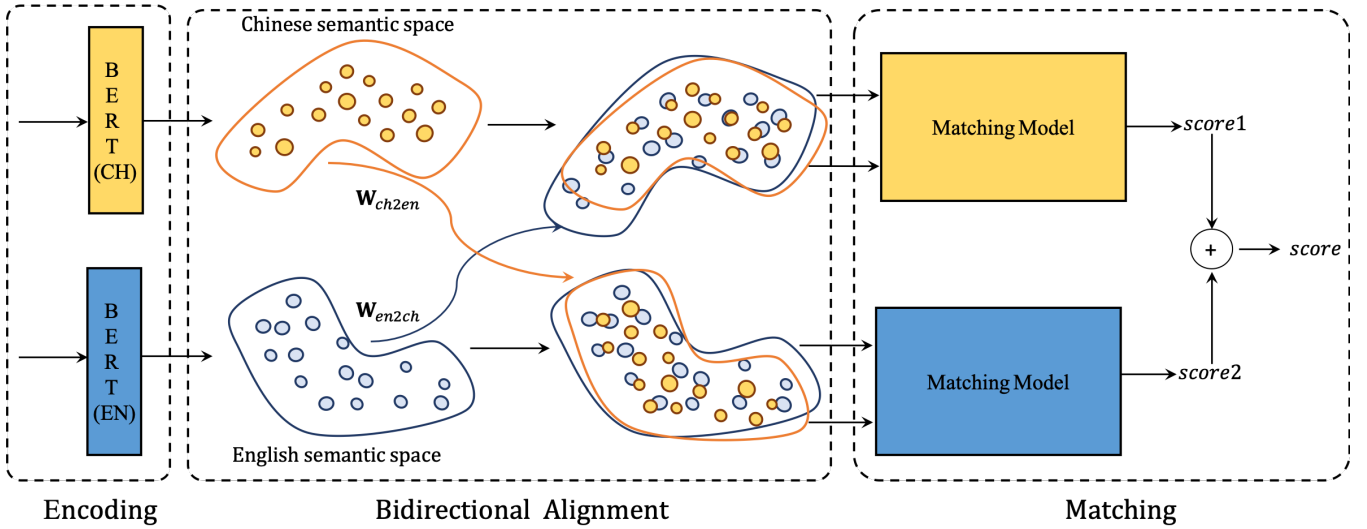


Figure 2: The overall architecture of our proposed CLMN model.

Wikipedia. Rücklé et al. (2019) extend open-domain cross-lingual question retrieval to the task-oriented domain, i.e., constructing a dataset upon StackOverflow. To facilitate the development of cross-lingual information retrieval in cross-border e-Commerce, we, for the first time, create a high-quality heuristic dataset from real commercial applications.

Our Proposed Model

In our cross-lingual set-to-description retrieval task, the main purpose is to train a cross-lingual matching model that computes the relevance score between a product attribute set in the source language and a product description in the target language. To formulate this task, we utilize the following necessary notations. A dataset with paired product attribute sets and descriptions $\mathcal{D} = \{(a_i, d_i, y_i)\}_{i=1}^N$ is first given, where a_i, d_i, y_i represent a product attribute set in the source language, a description in the target language, and the corresponding relevance score $y_i \in \{0, 1\}$ where $y_i = 1$ represents the given description candidate is a proper one for the given attribute set, otherwise $y_i = 0$. Our task is defined as learning a matching model from the given dataset that can output a relevance score for a set-description pair. As shown in Figure 2, our proposed cross-lingual matching network (CLMN) consists of encoding, bidirectional cross-lingual alignment, and matching.

Encoding

As the unsupervised monolingual pre-trained models has achieved promising results on various NLU tasks (Peters et al. 2018; Devlin et al. 2019), we propose to leverage the pre-trained BERT for obtaining the contextualized word representations. In details, we utilize the pre-trained Chinese BERT and the pre-trained English BERT to convert each product attribute set (Chinese) and each product description (English) into two separate monolingual semantic spaces. Their monolingual contextualized representa-

tions are then bidirectionally aligned into the same semantic spaces by context-dependent bilingual mapping.

Bidirectional Cross-lingual Alignment

Inspired by the success of cross-lingual alignment learning (Conneau et al. 2018a), we propose to map the contextualized monolingual embedding representations to a shared bilingual space so as to extract semantic information and matching features of a product attribute set and a description. Different from the case in context-independent word embeddings, the contextualized embedding of a word varies with the surrounding contexts. In order to draw on the effective off-the-shelf cross-lingual alignment algorithms, we bridge this gap by considering the averaged contextualized representations as the anchor (i.e., the static embedding) for each word in a monolingual corpus following Schuster et al. (2019). Then we adopt MUSE (Conneau et al. 2018a) to learn the cross-lingual alignment matrix \mathbf{W} as the initialization of the mapping matrix of our model. To better facilitate the downstream matching task, we propose to update \mathbf{W} during the training process. As demonstrated by Smith et al. (2017), imposing an orthogonal constraint to the mapping matrix leads to better performance. In order to ensure that \mathbf{W} stays close to an orthogonal matrix, we put an additional constraint during parameters updating (Cisse et al. 2017; Conneau et al. 2018a), denoted as:

$$\mathbf{W} \leftarrow (1 + \beta)\mathbf{W} - \beta(\mathbf{W} \cdot \mathbf{W}^T) \cdot \mathbf{W} \quad (1)$$

As illustrated in Figure 2, we conduct bidirectional cross-lingual alignment on the contextualized monolingual embeddings. We first learn the alignment mapping matrix \mathbf{W}_{ch2en} from the Chinese space to the English Space. Then, we learn the alignment mapping matrix \mathbf{W}_{en2ch} from the English space to the Chinese Space. By doing so, we have two groups of representations of a_i and d_i in the bilingual space, i.e., $a_{en,i}$ and $d_{en,i}$ in the English semantic space and $a_{ch,i}$ and $d_{ch,i}$ in the Chinese semantic space.

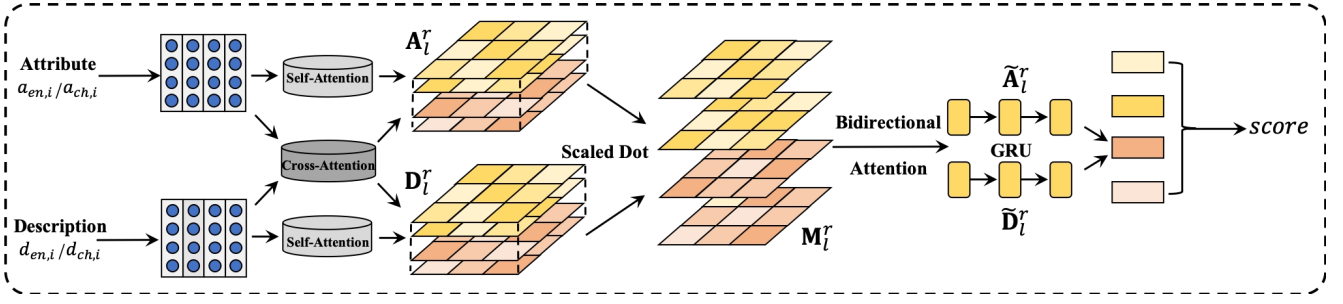


Figure 3: The detailed illustration of the matching model.

Matching Model

Given that the product attribute sets are unordered and there exist abundant correlations between product attributes and descriptions, we propose to use a fully attention-based module, including self-attention and cross-attention, to capture these dependency information. Specifically, the attention-based module is constructed upon the aligned bilingual representations of each product attribute and description. The attentive module takes three inputs, namely the query, the key, and the value, which are denoted as \mathbf{Q} , \mathbf{K} , \mathbf{V} respectively. First, the attentive module uses each word in the query to attend each word in the key through the scaled dot-product attention mechanism. Then, the obtained attention score is functioned upon the value \mathbf{V} to form a new representation of \mathbf{Q} , which is formulated as

$$Att(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = softmax(\frac{\mathbf{Q} \cdot \mathbf{K}^T}{\sqrt{d}}) \cdot \mathbf{V} \quad (2)$$

Thus, each word in the query \mathbf{Q} is represented by the joint meaning of its similar words in \mathbf{V} . In practice, the key \mathbf{K} and the value \mathbf{V} are set to identical. $\mathbf{Q} = \mathbf{V}$ corresponds to the self-attention representation of the query. Otherwise, we can obtain the cross-attention representation of the query. We stack the attention modules L times to better capturing different granularities of representations (Zhou et al. 2018). Accordingly, we obtain the self-attention and the cross-attention representations \mathbf{A}_l^r and \mathbf{D}_l^r , where $1 \leq l \leq L$ and $r \in \{self, cross\}$. To perform interaction between two sets of representations, we first calculate the scaled interaction matrix \mathbf{M}_l^r , denoted as:

$$\mathbf{M}_l^r = \frac{\mathbf{A}_l^r \cdot \mathbf{D}_l^{rT}}{\sqrt{d}} \quad (3)$$

We then use softmax function on both row and column directions of \mathbf{M}_l^r to get bi-directional attentive weights and obtain the bi-directional attentive aligned representations:

$$\begin{aligned} \hat{\mathbf{A}}_l^r &= softmax(\mathbf{M}_l^r) \cdot \mathbf{D}_l^r \\ \hat{\mathbf{D}}_l^r &= softmax(\mathbf{M}_l^{rT}) \cdot \mathbf{A}_l^r \end{aligned} \quad (4)$$

Next, we calculate the interaction between \mathbf{A}_l^r and $\hat{\mathbf{A}}_l^r$ as well as \mathbf{D}_l^r and $\hat{\mathbf{D}}_l^r$ by element-wise multiplication, subtraction, and concatenation, following with a fully connected

layer with ReLU activation denoted as $g(\cdot)$:

$$\begin{aligned} \tilde{\mathbf{A}}_l^r &= g([\mathbf{A}_l^r \odot \hat{\mathbf{A}}_l^r; \mathbf{A}_l^r - \hat{\mathbf{A}}_l^r; \mathbf{A}_l^r; \hat{\mathbf{A}}_l^r]) \\ \tilde{\mathbf{D}}_l^r &= g([\mathbf{D}_l^r \odot \hat{\mathbf{D}}_l^r; \mathbf{D}_l^r - \hat{\mathbf{D}}_l^r; \mathbf{D}_l^r; \hat{\mathbf{D}}_l^r]) \end{aligned} \quad (5)$$

Subsequently, we consider the rows of $\tilde{\mathbf{A}}_l^r$ and $\tilde{\mathbf{D}}_l^r$ as time steps and summarize the interaction information with two GRU layers then concatenate the last hidden states of both as the representation of the overall interaction features of \mathbf{A}_l^r and \mathbf{D}_l^r . Finally, we gather all the interaction information from all the interaction between all the $\mathbf{A}_l^r - \mathbf{D}_l^r$ pairs with a fully connected layer followed by a sigmoid function to obtain the matching scores. We denote the predictions of the matching models in the Chinese semantic space and the English semantic space $score1$ and $score2$, respectively.

Learning and Prediction

Recall that we have two separate matching models with the same structure but different inputs from different aligned semantic spaces, we jointly update the two matching models as well as the two alignment matrices with the same batches of training cases so as to make full use of the forward pass of BERT, which is time-consuming. In learning each of the matching model $f(\cdot)$, the objective is to minimize the cross entropy with dataset \mathcal{D} , formulated as:

$$\begin{aligned} \mathcal{L} &= - \sum_{i=1}^N y_i \log(f(a_i, d_i, y_i)) \\ &\quad + \sum_{i=1}^N (1 - y_i) \log(1 - f(a_i, d_i, y_i)) \end{aligned} \quad (6)$$

where N refers to the number of training pairs, y_i is the label of a product descriptions, $f(\cdot)$ is the matching model.

In prediction, we utilize the addition of $score1$ and $score2$ as the final relevance score between product attribute sets and descriptions.

Experimental Setup

CLIR-EC Dataset

We propose to construct a heuristic dataset CLIR-EC (Cross-Lingual Information Retrieval for e-Commerce) for exploring how to retrieve informative product descriptions in the

	Taobao	eBay	Attri.	Descri.
# Samples	313K	118K	13.5K	13.5K
# Total Words	17.6M	8.92M	233K	639K
Average Length	56.2	75.8	17.2	47.3
Vocabulary Size	163K	50.8K	16.6K	15.5K

Table 1: The statistical results of datasets used in experiments. **Descri.** refers to product descriptions. **Attri.** is product attribute sets.

target language given product attribute sets in the source language. Specifically, we hired 20 annotators with both proficient language skills (with master’s degree majored in translation) and domain knowledge (at least one year experience for e-commerce data annotation) to write a proper informative product description in the target language given a product attribute set in the source language. As there exist multi-level overlaps and correlations inside a product attribute, we asked annotators to choose part of the pivotal attributes of a set for creating a fluent and persuasive product description. We also employed 3 native speakers to conduct proofreading and editing on the collected product descriptions. By doing so, each product attribute set in the source language is paired with a product description in the target language, which is used for training the CLMN model. Totally, we collected 13.5 K training pairs, and the collection process takes 35 days. More details can be seen in Table 1. We randomly split the alignment dataset into 10,500/1,000/2,000 for training/validating/testing the CLMN model. We also launch another human evaluation to verify the quality of the manually created dataset. Concretely, we employed 5 evaluators with proficient language skills and rich cross-national shopping experience to judge the created product descriptions from the following five perspectives: 1) **information integrity**, i.e., whether the descriptions include the most pivotal information in product attributes; 2) **grammatically formed**, i.e., whether the descriptions are fluent and grammatical; 3) **matching degree**, i.e., does the description correlate well with any given attributes; 4) **sentiment polarities**, i.e., whether the descriptions present a positive sentiment to users. Each annotator is asked to rate a product description from these aspects on a scale of 1 (very bad) to 5 (excellent). We randomly sample 500 pairs for quality verification. Table 2 presents the results of human verification, where the evaluation results are very promising.

Other Datasets

Besides, we also automatically collect two large monolingual corpora from e-commerce platforms, i.e., eBay¹ with product descriptions in the target language (English), and Taobao² with product attribute sets in the source language (Chinese) for learning cross-lingual alignment, including bilingual word embeddings and context-dependent bilingual word representations upon BERT. Table 1 illustrates the detailed data statistical results. Overall, there are 313K samples and 17.6 million words in the Taobao dataset while

¹<https://www.ebay.com/>

²<https://www.taobao.com/>

Metrics	Scores	Kappa Values
Grammatically Formed	4.77	0.71
Information Integrity	4.58	0.63
Sentimental Polarities	4.79	0.80
Matching Degree	4.53	0.75

Table 2: Evaluation results of data quality verification, where kappa values refer to the inter-annotator agreement.

there are 118K samples and 8.92 million words in the eBay dataset. The average document length of Taobao and eBay corpora are 75.8 and 40.8, respectively. To preprocess these corpora, we use NLTK for tokenization. All words in both datasets are reserved, which results in a vocabulary size of 162,506 for Taobao and 59,233 for eBay.

Comparison Methods

To thoroughly evaluate the performance of our proposed CLMN model, we leverage three groups of baselines, i.e., advanced unsupervised cross-lingual IR models, the combination of machine translation system and monolingual IR models, directly cross-lingual retrieval models.

Unsupervised CLIR Models We utilize the state-of-the-art unsupervised cross-lingual retrieval model and its extensions (Litschko et al. 2018), including term-by-term query translation model (**TbTQT**), bilingual word embedding aggregation model (**BWE-AGG**), inverse document frequency weighted BWE-AGG (**BWE-IDF**), **Bigram matching**.

Translation Based CLIR Models We follow the recent proposed CLIR model (Rücklé, Swarnkar, and Gurevych 2019) to transfer the cross-lingual information retrieval task to a monolingual information retrieval one by a machine translation model. In our settings, we use the SOTA commercial translation system (Google API) to translate the unordered product attribute sets in the source language to the target language, and then we utilize the obtained paired data to train the following monolingual IR models for retrieval.

SMN, a strong model for response selection (Wu et al. 2017). It first learns a matching vector between the translated attribute sets and the product descriptions with the CNN network. Then, the learned matching vectors are aggregated by an RNN to calculate the final matching score.

DAM, a strong text matching network with only attention modules, which is adapted from the response selection task (Zhou et al. 2018). This model is used for studying whether attention modules can model unordered attribute sets and their matching degrees with descriptions.

CSRAN (Tay, Tuan, and Hui 2018), a state-of-the-art monolingual IR model. This model first adopts stacked recurrent encoders with refinement to learn representations. Then, it calculates interactions based on the attentively aligned representations and aggregates the interaction features via a multi-layered LSTM to obtain matching scores.

Directly CLIR Models We also leverage the state-of-the-art directly cross-lingual IR models for comparison.

ML-BERT (Devlin et al. 2019), the multilingual version of BERT³, which is trained on 104 languages and has achieved the SOTA performance in cross-lingual text pair classification (Conneau et al. 2018b). We further fine-tune the pre-trained multilingual BERT on our collected datasets.

POSIT-DRMM (Zhang et al. 2019), the recent proposed cross-lingual document retrieval model, which is designed for addressing the low-resource issue in CLIR. This model incorporates bilingual representations to capture and aggregate matching signals between an input query in the source language and a document in the target language.

Implementation Details

Following the conventional settings in related research (Zhou et al. 2018) for training the proposed CLMN model, we use the CLIR-EC dataset to prepare the supervised positive sample pairs and negative sample pairs. Each positive sample pair consists of a product attribute set and the matched product description while a negative sample pair comprises of a product attribute set and a mismatched product description sampled from all of the descriptions in CLIR-EC dataset. The proportion of the positive sample and the negative sample in training and validation sets is 1:1, while the corresponding proportion of testing set is 1:9. We limit the length of product attributes and descriptions to 50 words and 100 words, respectively. The pre-trained static monolingual and bilingual word embeddings on Taobao and eBay corpora used by all the non-BERT models have the dimension size 300. For our BERT-based CLMN model, we use the 768-dimension outputs of the second last BERT layer and keep the dimension of the encoder outputs as well as interaction outputs the same as the dimension of BERT. We use $L = 2$ stacked attention modules for attributes and descriptions. The parameters of encoders for attributes and descriptions are shared. We set the mini-batch size to 50 and adopt Adam optimizer (Kingma and Ba 2014) with the initial learning rate of $3e-4$ to for training. The best performance of each model on the validation set is chosen.

Evaluation Metrics

To evaluate the model performance, we follow the conventional settings in related work (Wu et al. 2017; Zhou et al. 2018; Zhang et al. 2019). Specifically, we first calculate the matching scores between a product attribute set and product description candidates, and then rank the matching scores of all candidates to calculate the following automatic metrics, including mean reciprocal rank (MRR) (Voorhees and others 1999), and recall at position k in n candidates ($R_n@k$).

Experimental Results

Table 3 summarizes the evaluation results of different cross-lingual retrieval models. Overall, our proposed CLMN model achieves the best performance on five evaluation metrics. For simplicity, we mainly discuss the performance of $R_{10}@1$ in the following part. In details, the advanced unsupervised CLIR models only achieve a fair performance without the collected CLIR-EC dataset. Through utilizing paired

Model	$R_2@1$	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$	MRR
Unsupervised models					
TbTQT	86.0	40.1	64.7	94.5	61.7
BWE-AGG	64.1	20.4	33.6	65.5	40.3
BWE-IDF	62.0	19.0	32.6	66.4	39.3
Bigram	64.0	19.4	33.4	65.3	39.6
Translation+Monolingual information retrieval models					
SMN	94.7	73.6	88.8	98.4	84.2
DAM	95.8	78.5	91.5	98.8	87.4
CSRN	96.0	79.8	92.5	99.0	88.2
Cross-lingual information retrieval model					
ML-BERT	97.2	83.8	95.0	99.3	90.8
POSIT-DRMM	95.4	74.3	90.7	99.3	85.1
CLMN ⁻	97.1	83.5	94.7	99.6	90.5
CLMN	97.8	86.8	95.5	99.8	92.3

Table 3: Evaluation results of baselines and our models. CLMN⁻ refers to CLMN without using pre-trained BERT.

data, translation based CLIR models significantly outperform the SOTA unsupervised methods and achieves a decent performance on each metric. We can observe from Table 3 that the attention-based translation CLIR model DAM achieves a noticeable improvement upon the SMN model, which demonstrates the effectiveness of attention module for modeling the unordered product attribute sets. Moreover, directly modeling the cross-lingual information retrieval task can further introduce a significant performance improvement over the translation based models. Both the recent proposed POSIT-DRMM and ML-BERT models yield a large increase over the strong translation based models. We also observe that our proposed CLMN⁻ matching model are better than the other matching models and even achieves comparable results with the fine-tuned multilingual BERT. With the enhancement of monolingual BERT, our proposed CLMN model outperforms the fine-tuned multi-lingual BERT model.

Ablation Study

Table 4 presents the experimental results for ablation study. Through analyzing, we observe that monolingual BERT is superior to train domain-specific word embeddings in our task. All models trained on the monolingual BERT achieve impressing results in Table 4, which are significantly better than domain-specific word embeddings pre-training. We attribute such an observation to the relatively small size of the domain-specific monolingual data, which is far from ideal to train a decent model. As for the influence of cross-lingual alignment, it can effectively capture the correlations between two different languages and thus can introduce performance improvement. Besides, we also notice that the context-dependent mapping strategy is more effective in English bilingual space. As reflected in Table 4, CLMN-ch2en model yields a observable improvement over CLMN-en2ch model on two different settings. This observation is perhaps owing to that the relatively small size of the unordered Chinese product attribute can be easily mapped to another language space compared with a relatively long product description with lexical and syntactic structures.

³<https://github.com/google-research/bert>

Model	$R_2@1$	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$	MRR
Initialize with word embedding					
CLMN ⁻ -mono	95.8	78.3	91.6	99.2	87.3
CLMN ⁻ -en2ch	96.3	81.5	93.1	99.5	89.2
CLMN ⁻ -ch2en	96.7	81.8	93.4	99.5	89.5
CLMN ⁻	97.1	83.5	94.7	99.6	90.5
Initialize with BERT					
CLMN-mono	96.5	80.7	93.8	99.8	89.1
CLMN-en2ch	97.2	84.6	95.4	99.7	91.4
CLMN-ch2en	97.5	85.5	94.8	99.8	91.7
CLMN	97.8	86.8	95.5	99.8	92.3

Table 4: Ablation study of cross-lingual alignment, where *mono* refers to without cross-lingual alignment learning, *en2ch* refers to calculating the matching score in the Chinese bilingual space, and *ch2en* is to compute the relevance score in the English Bilingual Space.

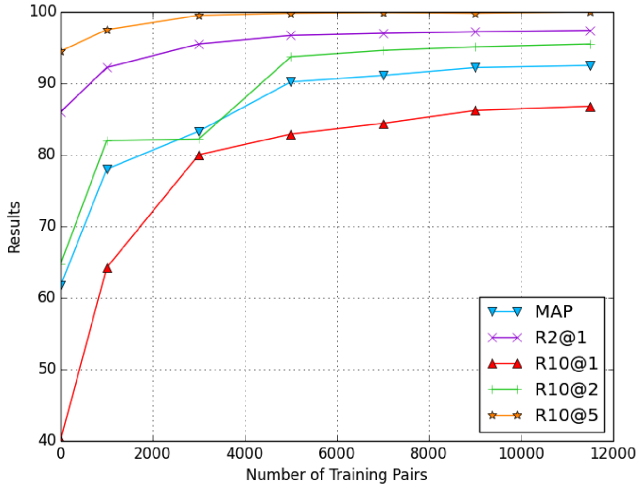


Figure 4: Trade-off between the size of CLIR-EC corpus and model performance.

The Effect of Paired Data Size

We also launch experiments to study the influence of the paired data size for directly cross-lingual retrieving. As illustrated in Figure 4, when the number of training pairs is less than 4000, the performance of cross-lingual information retrieval is positively correlated with the increasement of cross-lingual training pairs. When the number of cross-lingual training pairs is large than 4000, model performance only gains a slight improvement with more training pairs.

Bad Case Analysis

To learn the limitations of our proposed CLMN model, we thoroughly analyze typical bad cases from the test set that fails in $R_{10}@5$. Figure 5 presents the typical example that our proposed model fails in the cross-lingual retrieval task. We can observe that the main error pattern is caused by the commonly occurred phrases in e-commerce platforms, such as “gentleman”, “leisure”, “Korean-style”, “workplace”, “male”. The system pays more attention to these words and leaves the key information unobserved, e.g., the trained model retrieves a product description of “suit” for

Product attributes: 黑色, 酒红色, 领带, 男士, 商务, 韩版, 6cm, 窄, 时尚, 学生, 英伦, 上班, 职业, 休闲, 结婚 (black, claret, tie , male, commercial, Korean-style, 6cm, narrow, fashion, on work, professional, leisure, wedding)
Positive descriptions: VORLAD / Walleton 6cm tie , this tie is divided into various colors, each color has its own unique style, it can easily match any suits to help you succeed, high-quality colored fabric has delicate texture, the bright color shows the high quality of fashion.
Negative descriptions: This Korean-style suit has a strong gentleman's style. Its version is very good, and the design keeps up with the trend. It breaks through the style of traditional suits and is very popular with young people. Whether it is the workplace, every day, or leisure, it can be seen in the image of a gentleman. The quality of the suit fabric is high and the craft quality is outstanding. The color is very nice. If you like, come to buy it!

Figure 5: A typical bad case example in our task that the proposed CLMN model fails in $R_{10}@5$.

the product of “tie”. Such an error pattern is related to the search behaviors of customers on an e-commerce platform. Usually, users are not exactly sure about their target products. They prefer to input a large concept such as “man”, “Korean-style”, rather than “6cm black tie”. To improve the search conversion rate, sellers on e-commerce platforms add these general keywords to their product attributes. As a result, the potential solutions of this issue rely on how to filtrate this general information and augment the matching signals of important product attributes.

Conclusion

In this paper, we propose to address the task of cross-lingual set-to-description retrieval in e-commerce, which involves automatically retrieving a persuasive product description in the target language given a product attribute set in the source language. We collected a high-quality heuristic dataset with 13.5K pairs created by experts and two large monolingual auxiliary corpora for sharing. We further propose a novel cross-lingual matching network (CLMN) aligned with a refined context-dependent cross-lingual mapping upon pre-trained BERT to address this low-resource cross-lingual retrieval task directly. Experimental results indicate that our proposed CLMN is effective for solving this challenging task. In our future work, we will evaluate our model on other open datasets for cross-lingual information retrieval.

Acknowledgments

We thank the reviewers for their constructive comments. This work was supported by the National Key Research and Development Program of China (No. 2017YFC0804001), the National Science Foundation of China (NSFC No. 61876196 and NSFC No. 61672058). Rui Yan was sponsored by Alibaba Innovative Research (AIR) Grant and Beijing Academy of Artificial Intelligence (BAAI).

References

- Artetxe, M.; Labaka, G.; and Agirre, E. 2017. Learning Bilingual Word Embeddings with (almost) no Bilingual Data. In *ACL*, volume 1, 451–462.
- Artetxe, M.; Labaka, G.; and Agirre, E. 2018. A Robust Self-Learning Method for Fully Unsupervised Cross-Lingual Mappings of Word Embeddings. In *ACL*, volume 1, 789–798.
- Baroni, M.; Bernardini, S.; Ferraresi, A.; and Zanchetta, E. 2009. The wacky wide web: a collection of very large linguistically processed web-crawled corpora. *Language resources and evaluation* 43(3):209–226.
- Barzilay, T. L. H. J. R.; Jaakkola, T.; Tymoshenko, K.; and Marquez, A. M. L. 2016. Semi-supervised question retrieval with gated convolutions. In *NAACL-HLT*, 1279–1289.
- Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2017. Enriching Word Vectors with Subword Information. *TACL* 5:135–146.
- Cisse, M.; Bojanowski, P.; Grave, E.; Dauphin, Y.; and Usunier, N. 2017. Parseval networks: Improving robustness to adversarial examples. In *ICML*, 854–863. JMLR. org.
- Conneau, A.; Lample, G.; Ranzato, M.; Denoyer, L.; and Jégou, H. 2018a. Word Translation Without Parallel Data. *ICLR*.
- Conneau, A.; Lample, G.; Rinott, R.; Williams, A.; Bowman, S. R.; Schwenk, H.; and Stoyanov, V. 2018b. Xnli: Evaluating cross-lingual sentence representations. *arXiv preprint arXiv:1809.05053*.
- Da San Martino, G.; Romeo, S.; Barroón-Cedeño, A.; Joty, S.; Maàrquez, L.; Moschitti, A.; and Nakov, P. 2017. Cross-Language Question Re-anking. In *SIGIR*, 1145–1148. ACM.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of Deep Bidirectional Transformers for Language Understanding. *NAACL-HLT*.
- Dos Santos, C.; Barbosa, L.; Bogdanova, D.; and Zadrozny, B. 2015. Learning hybrid representations to retrieve semantically equivalent questions. In *ACL-IJCNLP(Short Papers)*, volume 2, 694–699.
- Ge, T.; Dou, Q.; Ji, H.; Cui, L.; Chang, B.; Sui, Z.; Wei, F.; and Zhou, M. 2018. Fine-grained coordinated cross-lingual text stream alignment for endless language knowledge acquisition. In *EMNLP*, 2496–2506.
- Hui, K.; Yates, A.; Berberich, K.; and De Melo, G. 2018. Copacrr: A context-aware neural ir model for ad-hoc retrieval. In *Proceedings of the eleventh ACM international conference on web search and data mining*, 279–287. ACM.
- Joty, S.; Nakov, P.; Maàrquez, L.; and Jaradat, I. 2017. Cross-language learning with adversarial neural networks. In *CoNLL*, 226–237.
- Joulin, A.; Bojanowski, P.; Mikolov, T.; Jégou, H.; and Grave, E. 2018. Loss in translation: Learning bilingual word mapping with a retrieval criterion. In *EMNLP*, 2979–2984.
- Kingma, D. P., and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*.
- Lample, G.; Ott, M.; Conneau, A.; Denoyer, L.; et al. 2018. Phrase-Based & Neural Unsupervised Machine Translation. In *EMNLP*, 5039–5049.
- Litschko, R.; Glavaš, G.; Ponzetto, S. P.; and Vulić, I. 2018. Unsupervised Cross-Lingual Information Retrieval Using Monolingual Data Only. *SIGIR*.
- McDonald, R.; Brokos, G.; and Androutsopoulos, I. 2018. Deep relevance ranking using enhanced document-query interactions. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 1849–1860.
- Nie, J.-Y. 2010. Cross-language information retrieval. *Synthesis Lectures on Human Language Technologies* 3(1):1–125.
- Peters, M. E.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep contextualized word representations. In *NAACL-HLT*, 2227–2237.
- Rücklé, A.; Swarnkar, K.; and Gurevych, I. 2019. Improved cross-lingual question retrieval for community question answering. In *The World Wide Web Conference*, 3179–3186. ACM.
- Sasaki, S.; Sun, S.; Schamoni, S.; Duh, K.; and Inui, K. 2018. Cross-lingual learning-to-rank with shared representations. In *NAACL-HLT(Short Papers)*, 458–463.
- Schamoni, S.; Hieber, F.; Sokolov, A.; and Riezler, S. 2014. Learning translational and knowledge-based similarities from relevance rankings for cross-language retrieval. In *ACL(Short Papers)*, volume 2, 488–494.
- Schuster, T.; Ram, O.; Barzilay, R.; and Globerson, A. 2019. Cross-lingual alignment of contextual word embeddings, with applications to zero-shot dependency parsing. *arXiv preprint arXiv:1902.09492*.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Neural Machine Translation of Rare Words with Subword Units. In *ACL*, volume 1, 1715–1725.
- Smith, S. L.; Turban, D. H.; Hamblin, S.; and Hammerla, N. Y. 2017. Offline Bilingual Word Vectors, Orthogonal Transformations and the Inverted Softmax. *arXiv preprint arXiv:1702.03859*.
- Søgaard, A.; Ruder, S.; and Vulić, I. 2018. On the Limitations of Unsupervised Bilingual Dictionary Induction. In *ACL*, volume 1, 778–788.
- Tay, Y.; Tuan, L. A.; and Hui, S. C. 2018. Co-stack Residual Affinity Networks with Multi-level Attention Refinement for Matching Text Sequences. *EMNLP*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention Is All You Need. In *NIPS*, 5998–6008.
- Voorhees, E. M., et al. 1999. The TREC-8 Question Answering Track Report. In *Trec*, volume 99, 77–82. Citeseer.
- Vulić, I., and Moens, M.-F. 2015. Monolingual and Cross-Lingual Information Retrieval Models Based on (bilingual) Word Embeddings. In *SIGIR*, 363–372. ACM.
- Wu, Y.; Wu, W.; Xing, C.; Zhou, M.; and Li, Z. 2017. Sequential Matching Network: A New Architecture for Multi-turn Response Selection in Retrieval-based Chatbots. In *ACL*, volume 1, 496–505.
- Zhang, R.; Westerfield, C.; Shim, S.; Bingham, G.; Fabbri, A.; Verma, N.; Hu, W.; and Radev, D. 2019. Improving low-resource cross-lingual document retrieval by reranking with deep bilingual representations. *ACL*.
- Zhou, X.; Li, L.; Dong, D.; Liu, Y.; Chen, Y.; Zhao, W. X.; Yu, D.; and Wu, H. 2018. Multi-Turn Response Selection for Chatbots with Deep Attention Matching Network. In *ACL*, volume 1, 1118–1127.