

Retrieval-Augmented Multilingual Keyphrase Generation with Retriever-Generator Iterative Training*

Yifan Gao¹, Qingyu Yin^{1†}, Zheng Li^{1‡}, Rui Meng²,
Tong Zhao¹, Bing Yin¹, Irwin King³, Michael R. Lyu³

¹Amazon Inc. ²Salesforce Research ³Chinese University of Hong Kong

¹{yifangao, qingyy, amzzhe}@amazon.com ²ruimeng@salesforce.com

Abstract

Keyphrase generation is the task of automatically predicting keyphrases given a piece of long text. Despite its recent flourishing, keyphrase generation on non-English languages haven't been vastly investigated. In this paper, we call attention to a new setting named multilingual keyphrase generation and we contribute two new datasets, EcommerceMKP and AcademicMKP, covering six languages. Technically, we propose a retrieval-augmented method for multilingual keyphrase generation to mitigate the data shortage problem in non-English languages. The retrieval-augmented model leverages keyphrase annotations in English datasets to facilitate generating keyphrases in low-resource languages. Given a non-English passage, a cross-lingual dense passage retrieval module finds relevant English passages. Then the associated English keyphrases serve as external knowledge for keyphrase generation in the current language. Moreover, we develop a retriever-generator iterative training algorithm to mine pseudo parallel passage pairs to strengthen the cross-lingual passage retriever. Comprehensive experiments and ablations show that the proposed approach outperforms all baselines.¹

1 Introduction

Keyphrases are single or multi-word lexical units that best summarize a piece of text. As such, they are of great importance for indexing, categorizing, and mining in many information retrieval and natural language processing tasks (Jones and Staveley, 1999; Frank et al., 1999; Hulth and Megyesi, 2006; Dave et al., 2003). Keyphrase generation is the task of automatically predicting keyphrases given

a piece of long text. Existing works on keyphrase generation mostly focus on English datasets (Galina et al., 2019; Meng et al., 2017) while keyphrase generation for languages other than English is still under-explored. Since search engines usually provide services to customers using different languages, multilingual keyphrase generation becomes a significant problem while it is still unknown how well existing keyphrase generation approaches perform in non-English languages.

Nevertheless, there are two challenges we will face regarding multilingual keyphrase generation. First, to the best of our knowledge, there is no large-scale dataset publicly available for training and benchmarking multilingual keyphrase generation models. Building keyphrase datasets at a sufficient scale is difficult and costly. Second, compared with the existing datasets in English, which can contain millions of data examples and cover a wide diversity of topics, the data resources in non-English languages are inherently scarce. For example, in the domain of e-commerce, marketplaces using English have abundant customer queries to be used for keyphrase mining, while queries in some languages are relatively less than in English, which is probably because of a smaller size of user population or a shorter operation time.

We start tackling these challenges by contributing two new datasets for multilingual keyphrase, which cover six languages and two domains. The first dataset EcommerceMKP is collected from a real-world major e-commerce website. The product descriptions are used as the source text while the target keyphrases are collected from user search queries. This dataset contains a total of 73k data examples, covering four different languages (Spanish, German, Italian and French). The second multilingual keyphrase dataset AcademicMKP lies in the academic domain, in which titles and abstracts are used as the source text and the author-provided keyphrases are deemed targets. A total of 2,693 aca-

* Qingyu Yin and Zheng Li are corresponding authors. This work was mainly done while Rui Meng was a Ph.D. student at University of Pittsburgh.

¹The datasets are released at https://github.com/Yifan-Gao/multilingual_keyphrase_generation.

demographic papers in Chinese and Korean are included in AcademicMKP.

To overcome the resource scarcity challenge in training multilingual models, we propose a retrieval-based method to leverage the keyphrase knowledge in large-scale English datasets. By investigating multilingual keyphrase data, we observe that data in different languages may talk about similar topics. Therefore, we conjecture that passage-keyphrases pairs in English can be of help as an external knowledge base for multilingual keyphrase generation. To be specific, given a passage in low-resource language *XX*, we propose to use a retrieval model to find multiple top-related passages in English. These retrieved English passages provide high-quality English keyphrases that can be used as hints for generating keyphrases in other languages. After that, the generator takes the code-mixed inputs, including the passage in language *XX* and retrieved English keyphrases, and predicts keyphrases in language *XX*.

In the cross-lingual retrieval training, parallel passage-keyphrases pairs between English and other languages are extremely limited. For example, in the e-commerce domain, only a small fraction of products have both English and non-English descriptions (being sold in multiple countries). Such a data scarcity issue weakens the ability of cross-lingual knowledge acquisition from high-resource English keyphrases as intermediary, and finally hinders the potential of the retrieval-augmented keyphrase generation. To mitigate the problem, we propose a retriever-generator iterative training (RGIT) algorithm to automatically mine pseudo training parallel pairs from unlabeled data. Concretely, the retriever can dynamically adjust in terms of the current variations of generation performance between the proposed retrieval-augmented generator and the base one without the aid of retrieved English keyphrases. Starting from insufficient seed parallel pairs, if the retrieved pseudo passage-keyphrases pairs in the current iteration can bring in higher generation results as the generator’s feedback, those pseudo parallel data will be regarded as high quality and incorporated into the seed ones to further boost the retriever. Such cycle providing positive effects can be repeated until the increasing generation performance stopped.

We conduct extensive experiments on EcommerceMKP and AcademicMKP and demonstrate that large-scale English datasets do provide use-

ful knowledge for multilingual keyphrase generation. The proposed retrieval-augmented method outperforms traditional extraction-based models, sequence-to-sequence neural models, and its variants. Moreover, the RGIT algorithm boosts the retrieval performance significantly by mining over 20k pseudo-parallel passage pairs. We also conduct detailed analyses to investigate the effectiveness of retriever-generator iterative training.

2 Related Work

Keyphrase Generation. The advance of neural language generation enables models to freely generate keyphrases according to the phrase importance and semantics, rather than extracting a list of sub-strings from the text (Witten et al., 1999a; Liu et al., 2011; Wang et al., 2016). Meng et al. (2017) propose the first keyphrase generation model Copy-RNN, which not only generates words based on a vocabulary but also points to words in the source text — overcoming the barrier of predicting absent keyphrases. Following this idea, Chen et al. (2018); Zhao and Zhang (2019); Ahmad et al. (2021) leverage the attention mechanism to reduce duplication and improve coverage. Chen et al. (2019b); Ye and Wang (2018); Wang et al. (2019); Liang et al. (2021) propose to leverage extra structure information (e.g., title, topic) to guide the generation. Chan et al. (2019); Luo et al. (2021) propose a model using reinforcement learning, and Swaminathan et al. (2020) propose using GAN for KPG. Chen et al. (2020) introduce hierarchical decoding and exclusion mechanism to prevent models from generating duplicate phrases. Ye et al. (2021b) propose to dynamically align target phrases to eliminate the influence of order, as highlighted by Meng et al. (2021). Mu et al. (2020); Liu et al. (2020a); Park and Caragea (2020) use pre-trained language models for better representations of documents.

Retrieval Augmented Text Generation (RAG) recently shows great power in knowledge-intensive NLP tasks such as open-domain question answering, fact checking and entity linking (Lewis et al., 2020; Petroni et al., 2021; Guu et al., 2020). In RAG, a retriever (either sparse (Lee et al., 2019) or dense (Karpukhin et al., 2020)) searches for useful non-parametric knowledge from a knowledge base, then a generator combines the non-parametric retrieved knowledge with its parametric knowledge, learned during pre-training, for solving the task. Different from these tasks, keyphrase gen-

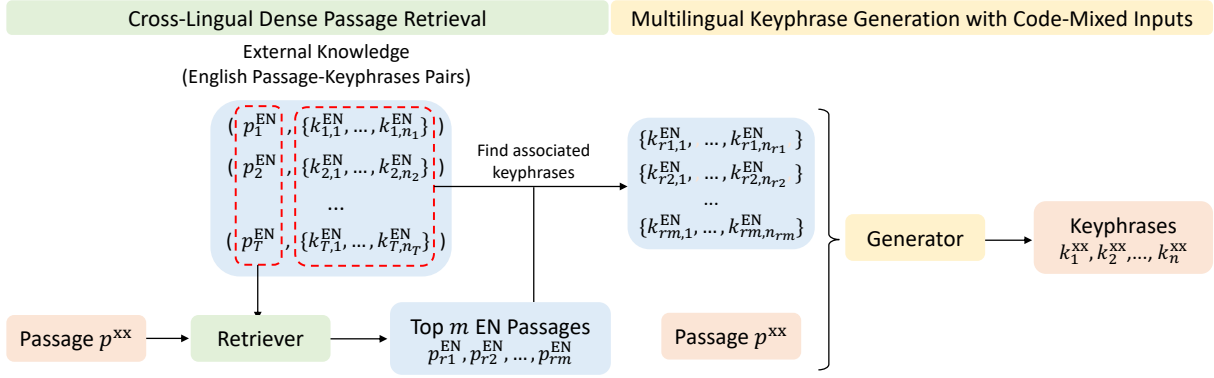


Figure 1: Overview of our Retrieval-Augmented Multilingual Keyphrase Generation (RAMKG) framework. p^{XX} , k_i^{XX} denote a passage and keyphrases in language XX ($XX \in \{DE, ES, FR, IT, KO, ZH\}$). p^{EN} , k_i^{EN} denote relevant passages and keyphrases retrieved from the English dataset.

eration is not a knowledge-intensive task but we treat the English passage-keyphrase training data as our knowledge. Similar approaches have been investigated in neural machine translation (Gu et al., 2018; Cai et al., 2021), dialogue (Weston et al., 2018), and knowledge-base QA (Das et al., 2021). In keyphrase generation, Chen et al. (2019a); Ye et al. (2021a); Kim et al. (2021) retrieve similar documents from training data to produce more accurate keyphrases. However, their retrieval module is a non-parametric model and cannot be generalized in the multilingual setting due to the large vocabulary gap between languages.

3 Task Definition

In this paper, we aim to tackle the keyphrase generation task in a multilingual setting, which means one model of desire is capable of generating keyphrases in any language that it has been trained with. The benefits of having a single keyphrase generation model for multiple languages are threefold: (1) Collecting keyphrase annotation for individual language can be prohibitively expensive; (2) Training and deploying separate models for each language is laborious; (3) Joint training of multiple languages can alleviate the resource scarcity by utilizing rich monolingual data.

Formally, we define the multilingual keyphrase generation task as follows. Given a piece of text p^{XX} in language XX, our goal is to predict its corresponding keyphrases $k_1^{XX}, k_2^{XX}, \dots, k_n^{XX}$ in language XX, where n is the total number of target keyphrases for this text p^{XX} . In this study, XX can be German (DE), Spanish (ES), Italian (IT), French (FR), Korean (KO) or Chinese (ZH).

4 Model

Scarcity of resources is one of the topmost challenges for multilingual tasks, which is also the case for multilingual keyphrase generation. One may find it difficult to collect enough text data in languages other than English, much less the annotation of keyphrases in specific domains. To overcome this problem, we propose a retrieval-augmented approach to make use of the relatively rich resources in English. The motivation for our proposed retrieval-augmented approach comes from an observation from data: texts and keyphrases expressed in different languages usually share common topics or knowledge concepts. For example, in e-commerce websites, it is often the case that the same products are sold in different marketplaces/countries. Thus these products as well as their keyphrases, though exhibited in different languages, can share a high semantic similarity. In other words, given a text in language p^{AA} , if we could find a similar text in language p^{BB} , its associated keyphrases k^{BB} may serve as a good hint for the to-be-generated keyphrases k^{AA} in language AA. Since English has the most abundant text-keyphrases pairs in both e-commerce and academic papers domains, its resource can be treated as a non-parametric keyphrase knowledge base, which provides texts in English covering a wide range of topics and concepts, as well as the associated high-quality keyphrases.

As shown in Fig. 1, our framework consists of a retrieval step and a generation step:

1. *Retrieval Step*: given a source passage p^{XX} in language XX, the cross-lingual retriever first finds m semantically relevant English passages $p_1^{EN}, p_2^{EN}, \dots, p_m^{EN}$. Each retrieved En-

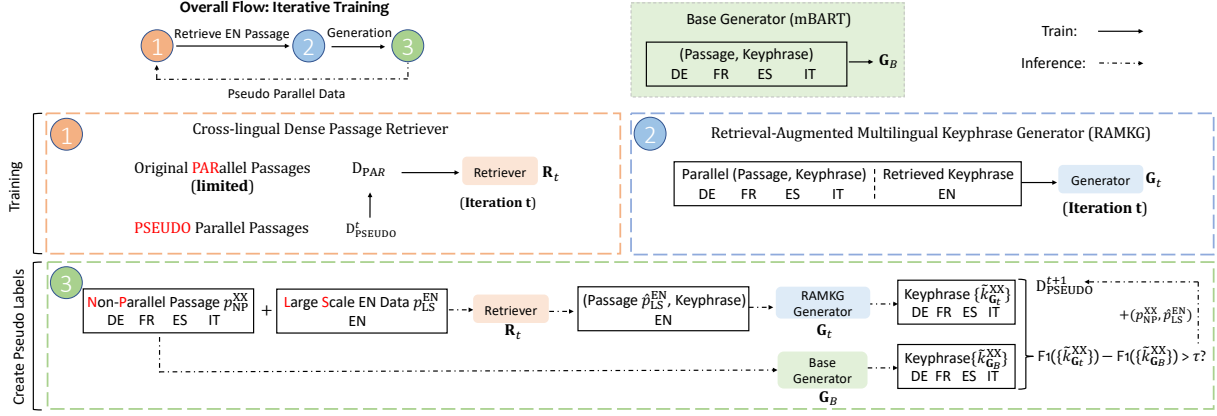


Figure 2: Retriever-Generator Iterative Training for Parallel Passage Mining

English passage p_j^{EN} has its associated n_j English keyphrases $k_{j,1}^{\text{EN}}, k_{j,2}^{\text{EN}}, \dots, k_{j,n_j}^{\text{EN}}$. These retrieved English keyphrases are taken as external knowledge for keyphrase generation in step 2.

2. *Generation Step*: taking the source text p^{XX} in language XX and all retrieved English keyphrases $\{k_{1,1}^{\text{EN}}, \dots, k_{1,n_1}^{\text{EN}}\}, \dots, \{k_{m,1}^{\text{EN}}, \dots, k_{m,n_m}^{\text{EN}}\}$ as inputs, the generation module concatenates them as a sequence and generates keyphrases in target language XX.

4.1 Cross-Lingual Dense Passage Retrieval

The cross-lingual retriever includes a passage encoder $E_P(\cdot)$ and a query encoder $E_Q(\cdot)$. The passage encoder $E_P(\cdot)$ maps millions of English passages into d -dimensional vectors and builds indices for all English passages using FAISS (Johnson et al., 2021) offline. At inference time, the passage in language XX goes through the query encoder $E_Q(\cdot)$ and is converted into a d -dimensional vector. Then the cross-lingual retriever performs a KNN search to retrieve m English passages whose vectors are closest to the query vector measured by the dot product similarity: $\text{sim}(p^{\text{XX}}, p^{\text{EN}_j}) = E_Q(p^{\text{XX}})^\top E_P(p_j^{\text{EN}})$.

Passage Encoder. Since naive lexical similarity can hardly handle text matching across languages, we resort to a BERT-based dense passage retriever (Karpukhin et al., 2020), expecting the contextualized semantic matching can retrieve similar passages accurately and robustly. In order to meet the demand of multilingual representation, we utilize multilingual pre-trained model mBERT (Devlin et al., 2019) to encode passages into 768-dimensional vectors.

Training. Since the output vectors of mBERT are not aligned across languages, we need extra alignment training to ensure that similar passages in different languages can be mapped into near regions in the high-dimensional space. Given a passage p_i^{XX} in language XX, we take its corresponding English passage $p_i^{\text{EN}+}$ as the positive example and randomly select n negative passages $p_{i,1}^{\text{EN}-}, \dots, p_{i,n}^{\text{EN}-}$ in the English corpus. The dense retriever is trained by optimizing the negative log likelihood loss of the positive English passage.

In the e-commerce domain, we select the positive passage according to product metadata. For a product sold in both EN and XX marketplaces, we regard its bilingual product descriptions $(p_i^{\text{XX}}, p_i^{\text{EN}+})$ as a parallel passage pair, i.e., positive training example. For the domain of academic paper, we notice that papers with parallel text is very rare. Therefore, we develop an automatic approach to mine parallel abstract pairs of English and the target language. Specifically, we adopt an off-the-shelf bi-text mining tool named Sentence Transformers (Reimers and Gurevych, 2019) to mine pseudo parallel pairs. Given two datasets in different languages, we encode passages using LaBSE (Feng et al., 2020), the current best method for learning language-agnostic sentence embeddings for 109 languages, and then parallel passages can be extracted through nearest-neighbor retrieval and filtered by setting a fixed threshold over a margin-based similarity score, as proposed in (Artetxe and Schwenk, 2019).

4.2 Multilingual Keyphrase Generation with Code-Mixed Inputs

Given the top m retrieved English passages $p_1^{\text{EN}}, \dots, p_m^{\text{EN}}$, we find their associated keyphrases in the

dataset: $\{k_{1,1}^{\text{EN}}, \dots, k_{1,n_1}^{\text{EN}}\}, \dots, \{k_{m,1}^{\text{EN}}, \dots, k_{m,n_m}^{\text{EN}}\}$. We utilize mBART (Liu et al., 2020b), a multilingual denoising pre-trained sequence-to-sequence language model, to integrate information from multiple languages. Different from machine translation which maps a sentence in the source language to a target language, our multilingual keyphrase generation model takes code-mixed inputs – a combination of retrieved English keyphrases $\{k_{1,1}^{\text{EN}}, \dots, k_{1,n_1}^{\text{EN}}\}, \dots, \{k_{m,1}^{\text{EN}}, \dots, k_{m,n_m}^{\text{EN}}\}$ from m retrieved English passages and the source passage p^{XX} in the target language XX.

We concatenate retrieved English keyphrases with a delimiter token [SEP], and add special tokens to separate different inputs: [ENKPS] for retrieved keyphrases and [CTX] for the source passage. Besides, we follow the fine-tuning setup of mBART by adding the language identifier [XX] (e.g. [DE] for German) at the end of the input sequence to denote the current input language:

[ENKPS] $k_{1,1}^{\text{EN}}$ [SEP] ... [SEP] k_{m,n_m}^{EN} [CTX] p^{XX} [XX]. The training target is a sequence of concatenated keyphrases $k_1^{\text{XX}}, \dots, k_n^{\text{XX}}$, separated by a special token [SEP]. [XX], the language identifier of the current language, is also added at the beginning of the target sequence to indicate the target language:

[XX] k_1^{XX} [SEP] k_2^{XX} [SEP] ... [SEP] k_n^{XX} .

4.3 Retriever-Generator Iterative Training

In spite of having utilized parallel passage pairs to align the multilingual representations of the retrieval module, it remains a concern because the parallel passage pairs between English and non-English languages account for only a small portion of the whole multilingual dataset. For example, in a popular e-commerce platform, only a small percentage of products (less than 10%) have both English and non-English descriptions. Without enough quality parallel pairs, the cross-lingual dense passage retriever may not work well to find relevant English passages. Consequently, associated English keyphrases may provide little help for multilingual keyphrase generation.

To make the multilingual keyphrase generation generalize better to any target languages or domains without reliance on numerous parallel passage pairs, we propose an iterative training method to mine parallel passages which requires only a small number of initial parallel pairs of bootstrap the process. Since our ultimate goal of the retriever model is to provide useful external knowledge for

Algorithm 1 Parallel Passage Mining via Iterative Training

```

1: Input: (1) Parallel data  $D_{\text{PAR}} = \{(p_{\text{PAR}}^{\text{EN}}, p_{\text{PAR}}^{\text{XX}}, k_{\text{PAR}}^{\text{EN}}, k_{\text{PAR}}^{\text{XX}})\}$ , (2)
   Non-parallel data  $D_{\text{NP}} = \{(p_{\text{NP}}^{\text{XX}}, k_{\text{NP}}^{\text{XX}})\}$ , (3) Large-scale English
   corpus  $D_{\text{LS}} = \{(p_{\text{LS}}^{\text{EN}}, k_{\text{LS}}^{\text{EN}})\}$ .
2: Output: Pseudo parallel passage pairs  $D_{\text{PSEUDO}}$ 
3:  $\triangleright$  0. Train Seq2Seq baseline w/o retrieved keyphrases
4:  $\mathbf{G}_B \leftarrow \text{train}(\{(p_{\text{PAR}}^{\text{XX}}, k_{\text{PAR}}^{\text{XX}})\} \in D_{\text{PAR}})$ 
5:  $D_{\text{PSEUDO}}^0 \leftarrow \{\}$  // Pseudo parallel passage pairs
6: for  $t \in \{0 \dots T - 1\}$  do
7:    $\triangleright$  1. Train retriever on pseudo and parallel data
8:    $\mathbf{R}_t \leftarrow \text{train}(\{(p_{\text{PAR}}^{\text{EN}}, p_{\text{PAR}}^{\text{XX}})\} \in D_{\text{PSEUDO}}^t \cup D_{\text{PAR}})$ 
9:    $\triangleright$  2. Train retrieval-augmented generator on  $D_{\text{PAR}}$ 
10:  for each  $(p_{\text{PAR}}^{\text{XX}}, k_{\text{PAR}}^{\text{XX}}) \in D_{\text{PAR}}$  do
11:     $\hat{p}_{\text{LS}}^{\text{EN}} \leftarrow \mathbf{R}_t(p_{\text{PAR}}^{\text{XX}}, D_{\text{LS}})$  // Retrieve EN passages
12:     $\{\hat{k}_{\text{LS}}^{\text{EN}}\} \leftarrow \hat{p}_{\text{LS}}^{\text{EN}}$  // Find associated EN keyphrases
13:  // Train retrieval-augmented generator with EN keyphrases
14:   $\mathbf{G}_t \leftarrow \text{train}(\{(p_{\text{PAR}}^{\text{XX}}, k_{\text{PAR}}^{\text{XX}})\} \in D_{\text{PAR}}, \{\hat{k}_{\text{LS}}^{\text{EN}}\})$ 
15:   $\triangleright$  3. Create pseudo parallel passage pairs
16:   $D_{\text{PSEUDO}}^{t+1} \leftarrow \{\}$ 
17:  for each  $p_{\text{NP}}^{\text{XX}} \in D_{\text{NP}}$  do
18:     $\hat{p}_{\text{LS}}^{\text{EN}} \leftarrow \mathbf{R}_t(p_{\text{NP}}^{\text{XX}}, D_{\text{LS}})$  // Retrieve EN passages
19:     $\{\hat{k}_{\text{LS}}^{\text{EN}}\} \leftarrow \hat{p}_{\text{LS}}^{\text{EN}}$  // Find associated EN keyphrases
20:    // Predict keyphrases w/o and w/ EN keyphrases
21:     $\{\hat{k}_{\text{NP}}^{\text{XX}}, \mathbf{G}_B\} \leftarrow \mathbf{G}_B(p_{\text{NP}}^{\text{XX}})$ 
22:     $\{\hat{k}_{\text{NP}}^{\text{XX}}, \mathbf{G}_t\} \leftarrow \mathbf{G}_t(p_{\text{NP}}^{\text{XX}}, \{\hat{k}_{\text{LS}}^{\text{EN}}\})$ 
23:    // If adding EN keyphrases leads to better keyphrase pre-
       dictions, the retrieved EN passages are taken as positive examples
24:    if  $\text{F1}(\{\hat{k}_{\text{NP}}^{\text{XX}}, \mathbf{G}_t\}) - \text{F1}(\{\hat{k}_{\text{NP}}^{\text{XX}}, \mathbf{G}_B\}) > \tau$  then
25:       $D_{\text{PSEUDO}}^{t+1} \leftarrow D_{\text{PSEUDO}}^{t+1} \cup \{(p_{\text{NP}}^{\text{XX}}, p_{\text{NP}}^{\text{XX}})\}$ 
26: return  $D_{\text{PSEUDO}}^T$ 

```

multilingual keyphrase generation, we mine parallel passage pairs (English and a non-English language) according to whether the retrieved English passage-keyphrases pairs could help the keyphrase generation for the target non-English language XX. For example, let p_a^{EN} and p_b^{EN} be two retrieved English passages for a passage p^{XX} in target language, if the associated keyphrases of p_a^{EN} provide more useful information for generating the keyphrases of p^{XX} than p_b^{EN} , then $(p_a^{\text{EN}}, p^{\text{XX}})$ would be considered as a better parallel passage pair. That said, we expect the mined pseudo parallel passage pairs to be of high quality according to the retrieval score, at the same time they can be directly helpful for training the generation module.

The proposed iterative training approach is sketched in Algo. 1 and Fig. 2. Given a Large-Scale keyphrase dataset in English $D_{\text{LS}} = \{(p_{\text{LS}}^{\text{EN}}, k_{\text{LS}}^{\text{EN}})\}$ and a smaller one $D_{\text{XX}} = \{(p^{\text{XX}}, k^{\text{XX}})\}$ in target language XX, we denote the set of annotated parallel examples (bilingual passages in English and other languages) as **PARALLEL split** $D_{\text{PAR}} = \{(p_{\text{PAR}}^{\text{EN}}, p_{\text{PAR}}^{\text{XX}}, k_{\text{PAR}}^{\text{EN}}, k_{\text{PAR}}^{\text{XX}})\}$, in which $\{(p_{\text{PAR}}^{\text{XX}}, k_{\text{PAR}}^{\text{XX}})\}$ comes from the XX dataset while $\{(p_{\text{PAR}}^{\text{EN}}, k_{\text{PAR}}^{\text{EN}})\}$ comes from the English dataset. The remaining data examples in the target dataset

Language	Train Size	Dev Size	Test Size	Passage Length (Avg/Std/Mid)	#Keyphrases (Avg/Std/Mid)	Absent Kps%
AcademicMKP Dataset						
Chinese (ZH)	1,110	158	319	217/48/207	5/1/5	27.2%
Korean (KO)	774	110	222	115/31/111	4/1/4	37.7%
Total	1,884	268	541	171/57/155	4/1/4	31.3%
EcommerceMKP Dataset						
German (DE)	23,997	1,411	2,825	157/79/141	10/5/8	57.1%
Spanish (ES)	12,222	718	1,440	159/84/139	9/5/7	54.6%
French (FR)	16,986	998	2,000	163/84/144	9/5/8	63.0%
Italian (IT)	9,163	538	1,081	167/84/152	8/3/7	42.6%
Total	62,368	3,665	7,346	161/82/143	9/5/7	56.4%

Table 1: AcademicMKP & EcommerceMKP Dataset

D_{XX} have no annotated corresponding English examples in D_{LS} (the pairs may exist but are not known yet), and we name this set as the **Non-Parallel split** $D_{NP} = \{(p_{NP}^{XX}, k_{NP}^{XX})\}$. We firstly fine-tune a mBART using only keyphrases data of target language $\{(p_{PAR}^{XX}, k_{PAR}^{XX})\}$ in D_{PAR} (Line 4). Then we start a loop to mine pseudo parallel passage pairs for refining the passage retriever. Each iteration is expected to bring in a higher quality of pseudo passage pairs, resulting in a better performance of retriever, with three steps:

1. We train a retriever R_t using existing available EN-XX passage pairs $\{(p_{PAR}^{EN}, p_{PAR}^{XX})\}$ from both parallel passage data D_{PAR} and most up-to-date pseudo passage data D_{PSEUDO}^t (Line 8).
2. We train a retrieval-augmented model G_t using multilingual passages $\{(p_{PAR}^{XX}, k_{PAR}^{XX})\}$ from the parallel data D_{PAR} and retrieved English keyphrases $\{\hat{k}_{LS}^{EN}\}$ from the English dataset D_{LS} (Line 14). To get the retrieved English keyphrases for each passage p_{PAR}^{XX} , we take the retriever R_t trained in step (1) to do a KNN search for passages \hat{p}_{LS}^{EN} in EN D_{LS} and find their associated keyphrases $\{\hat{k}_{LS}^{EN}\}$ (Line 10-12).
3. For each passage p_{NP}^{XX} in the non-aligned dataset D_{NP} , we also retrieve English passages \hat{p}_{LS}^{EN} and keyphrases $\{\hat{k}_{LS}^{EN}\}$ from D_{LS} (Line 18-19). Then the retrieved English passage \hat{p}_{LS}^{EN} will be taken as the parallel text to p_{NP}^{XX} if its associated keyphrases $\{\hat{k}_{LS}^{EN}\}$ provide *useful* information. The *usefulness* is measured by the keyphrase generation performance (F-score) between the retrieval-augmented generation model G_t and the base model G_0 that does not use EN keyphrases (Line 21-25).

After T iterations, we train the retriever on the pseudo data D_{PSEUDO}^T and fine-tune it on the parallel data D_{PAR} . Then we treat it as our final retriever and train the generation model in Sec. 4.2.

5 Datasets

EcommerceMKP Dataset is collected from a popular E-commerce shopping platform. There are four languages we consider for building EcommerceMKP: German (DE), French (FR), Spanish (ES) and Italian (IT). The title, product description, and bullet description provided by manufacturers are concatenated and treated as source input. The keyphrases of each product are selected from search queries under the following protocol. First, given a product, we only keep search queries that lead to purchases and treat them as effective queries. Then phrases are chunked from these effective queries using AutoPhrase (Shang et al., 2018) and further ranked by their frequency. Our assumption is: the more times a phrase appears in effective search queries of a product, the more important a phrase is. Finally a threshold is set to filter out unimportant phrases. Under this protocol, we receive 73k examples over four languages. The statistics are shown in Table 1.

We collect the passages and keyphrases under the same protocol for the English (EN) dataset and name it as EcommerceMKP-EN. In total the English dataset contains 3 million passage-keyphrases pairs. To obtain the parallel passage pairs for training the cross-lingual dense passage retriever, we pair the product descriptions in different languages according to the product identification information. We select a total of 1,247 parallel passages from EcommerceMKP training set which include 480 passages for DE-EN, 244 for ES-EN, 340 for FR-EN, and 183 for IT-EN. Besides, we keep 1,000 parallel passage pairs in the DEV set of EcommerceMKP to evaluate the performance of retrieval and bi-text mining.

AcademicMKP Dataset is collected from the academic domain. We take the title and abstract of each paper as the source text and author-provided keywords as the target output. All papers are sampled from Microsoft Academic Graph (Sinha et al., 2015), a web-scale academic entity graph that contains multiple types of scholarly entities and relationships: field of study, author, institution, abstract, venue, and keywords. We use Spacy (<https://spacy.io/>) to detect the language of abstracts and keyphrases, and choose two languages Chinese (ZH) and Korean (KO) to construct the AcademicMKP dataset. Since Microsoft Academic Graph (MAG) is automatically crawled and

Model	German (DE)			Spanish (ES)			French (FR)			Italian (IT)			Average		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1
Unsupervised Statistical Keyphrase Extraction															
KP-Miner	9.13	2.35	3.34	14.56	4.59	6.22	7.78	2.76	3.62	22.51	7.71	10.31	11.80	3.69	5.01
YAKE	2.31	26.54	4.17	3.26	36.86	5.89	2.47	27.29	4.43	3.87	50.12	7.10	2.77	32.24	5.01
Unsupervised Graph-based Keyphrase Extraction															
TextRank	5.77	9.00	6.39	7.15	11.59	7.97	5.45	8.33	5.94	8.18	15.27	9.85	6.31	10.25	7.09
TopicalPageRank	3.59	14.10	5.34	5.07	22.08	7.74	3.75	15.49	5.65	5.22	25.43	8.26	4.16	17.71	6.33
PositionRank	5.24	11.67	6.90	7.86	19.36	10.75	5.56	13.50	7.56	7.98	21.46	11.34	6.24	15.12	8.49
MultipartiteRank	5.09	8.76	6.15	7.56	14.25	9.46	5.31	9.86	6.59	7.72	15.93	10.12	6.02	11.19	7.50
Supervised Feature-based Keyphrase Extraction															
Kea	9.64	17.73	11.92	12.69	24.64	16.11	8.81	16.76	11.05	14.16	30.18	18.81	10.68	20.65	13.52
Neural-based Supervised Keyphrase Generation															
CopyRNN	13.48	7.59	9.08	21.11	12.40	14.78	14.79	8.48	10.04	34.66	20.07	24.28	18.45	10.61	12.69
Transformer	29.92	25.40	26.22	34.01	30.57	30.83	29.17	24.03	25.18	44.05	43.34	42.56	32.60	28.68	29.25
mBART (monolingual)	44.91	39.59	40.04	47.70	44.79	44.17	42.64	36.81	37.72	57.98	58.16	56.25	46.76	42.58	42.60
mBART (multilingual)	45.78	40.93	41.09	48.43	44.99	44.57	43.21	38.40	38.72	60.37	58.91	57.87	47.75	43.68	43.60
mBART + EN Joint Train	45.91	39.90	40.64	49.27	43.92	44.52	43.21	37.48	38.26	59.21	57.03	56.42	47.79	42.55	43.08
mBART + EN Pretrain	45.77	40.76	41.05	48.34	44.94	44.58	42.96	38.10	38.45	60.24	58.51	57.59	47.64	43.46	43.47
RAMKG (Ours)	46.88	41.90	42.14	49.35	45.89	45.49	43.79	39.48	39.61	60.89	59.51	58.43	48.59	44.61	44.50
RAMKG + RGIT (Ours)	48.11	43.05	43.30	50.54	47.04	46.64	45.07	40.77	40.86	62.35	60.75	59.87	49.86	45.81	45.73

Table 2: Main results on the EcommerceMKP dataset. The best results are in bold. (RGIT: Iterative-training)

Model	Chinese (ZH)			Korean (KO)			Average		
	P	R	F1	P	R	F1	P	R	F1
mBART (mono.)	32.52	31.50	31.50	24.57	26.46	24.96	29.26	29.43	28.81
mBART (multi.)	32.48	32.27	31.85	27.03	26.93	26.44	30.25	30.08	29.63
mBART + Joint	31.10	30.38	30.23	27.36	26.85	26.58	29.57	28.93	28.73
mBART + Pretrain	32.72	29.77	30.66	27.56	25.30	27.56	30.60	27.94	28.68
RAMKG (Ours)	33.40	32.66	32.45	28.30	28.17	27.68	31.31	30.82	30.49
RAMKG + RGIT (Ours)	34.38	33.05	33.15	29.33	27.87	28.00	32.31	30.92	31.04

Table 3: Results on AcademicMKP (mono: monolingual, multi: multilingual, Joint: EN Joint Train, Pretrain: EN Pretrain, RGIT: Iterative-training).

constructed, we find some of its data is extremely noisy. For example, keyphrases might be missing or contain incorrect information such as titles, author names, and publication venues. Some of the abstracts are incomplete. Therefore, we hire three annotators to manually examine the samples from MAG dataset. Data examples with incomplete abstracts are removed. We further manually verify the metadata of all examples and correct their keyphrase information if needed. Finally, 2,693 high-quality data examples of scientific papers in the computer science domain are collected to constitute AcademicMKP.

Besides the multilingual AcademicMKP dataset, we use KP20K (Meng et al., 2017) as the English data for retrieval-augmented generation. KP20K has 560k abstract-keyphrases pairs collected from various online digital libraries in computer science domain. The threshold is set as 1.03 for passage mining and we receive 841 parallel passage pairs from AcademicMKP training set, in which 433 ZH-EN passage pairs and 384 for KO-EN.

6 Experimental Setup

6.1 Evaluation Metrics

Keyphrase Generation. Let the ground truth keyphrases be $Y : k_1, k_2, \dots, k_n$ and the predicted keyphrases be $\tilde{Y} : \tilde{k}_1, \tilde{k}_2, \dots, \tilde{k}_M$, we compute the precision ($P@M$), recall ($R@M$) and F-score ($F_1@M$) between Y and \tilde{Y} as $P@M = \frac{|Y \cap \tilde{Y}|}{|\tilde{Y}|}$, $R@M = \frac{|Y \cap \tilde{Y}|}{|Y|}$, $F_1@M = \frac{2 \times P \times R}{P + R}$, where $|Y|$ denotes the number of keyphrases in the gold set Y . We only consider exact match of two keyphrases (with some post-processing such as lowercase) for $|Y \cap \tilde{Y}|$. Then the average are computed for all languages in the test set.

Passage Retrieval. The quantity of retrieved English passages directly influences how much external knowledge could be utilized for keyphrase generation. Therefore, we evaluate the top-k recall (k=1,2,5,10,20) on the DEV set for evaluating retrieval performance.

6.2 Baselines and Ablations

We consider following baselines and ablations: **1) Unsupervised Statistical Keyphrase Extraction:** KP-Miner (El-Beltagy and Rafea, 2010), YAKE (Campos et al., 2020); **2) Unsupervised Graph-based Keyphrase Extraction:** TextRank (Mihalcea and Tarau, 2004), TopicalPageRank (Sterckx et al., 2015), PositionRank (Florescu and Caragea, 2017), MultipartiteRank (Boudin, 2018); **3) Supervised Feature-based Keyphrase Extraction:** KEA (Witten et al., 1999b); **4) Neural**

Supervised Keyphrase Generation: CopyRNN, Transformer; **5) mBART (monolingual):** separately trained 6 mBART models on each language; **6) mBART (multilingual):** a single mBART model on all languages; **7) mBART + EN Joint Train:** a mBART model jointly trained on the multilingual data and English data (KP20K (Meng et al., 2017) for AcademicMKP; EcommerceMKP-EN for EcommerceMKP). **8) mBART + EN Pre-train:** a mBART firstly pre-trained on the English data and then fine-tuned on the multilingual data. **9) RAMKG (Ours):** The Retrieval-Augmented Multilingual Keyphrase Generation model (Sec. 4.1 & 4.2). **10) RAMKG + RGIT (Ours):** RAMKG improved with retriever-generator iterative training (RGIT) (Sec. 4.3).

7 Results and Analyses

7.1 Main Results

Main results are shown in Table 2 & 3 for EcommerceMKP and AcademicMKP respectively, and we make the following observations:

- The unsupervised approaches, both statistical-based and graph-based, have robust results across all languages. PositionRank performs the best among all unsupervised approaches.
- The supervised approaches consistently outperform unsupervised approaches. The feature-based approach KEA receives a high recall by predicting more keyphrases while the CopyRNN receives a high precision. Different from the results on English keyphrase generation where Transformer and CopyRNN are comparable, the Transformer beats the CopyRNN by a large margin in the multilingual scenario.
- We observe that jointly training on all languages (mBART multilingual) receives better results than separately training on each language (mBART monolingual). This implies the ability of locating and summarizing key information is transferable across languages.
- Comparing different approaches using external large-scale English data, we find that our proposed RAMKG outperforms both “EN Joint Train” and “EN Pretrain”. This is because the retrieval-augmented approach provides auxiliary knowledge information as part of the input to the generation module, while the other two variants have to “infuse” the knowledge learned from English data to model parameters. Moreover, “EN Joint Train” and “EN Pretrain” have

Recall @ Top K	1	2	5	10	20
RAMKG	26.4%	36.8%	50.1%	59.2%	67.3%
RAMKG + RGIT	45.8%	58.3%	72.4%	79.5%	85.2%

Table 4: Retrieval recall on EcommerceMKP DEV set.

Iteration	1	2	3	4	5	6
# Pseudo Passages	20,288	21,402	19,942	21,241	22,343	21,557
Label Accuracy %	28.0%	37.9%	40.5%	44.2%	45.1%	47.0%

Table 5: Number of pseudo parallel passages and their accuracy on EcommerceMKP DEV set in different iterations of parallel passages mining.

no positive effect on AcademicMKP dataset (Table 3). Compared with multilingual and English data are from the same website, there is still a domain gap between papers (multilingual) in AcademicMKP and papers (English) in KP20K.

- The retriever-generator self-training (RAMKG + Iter) alleviates the data scarcity issue with the help of stronger retriever: since the retriever can find more relevant English keyphrases, it leads to a general improvement on keyphrase performance across languages.

7.2 Effect of Iterative Training

Retrieval Results We investigate the effect of retriever-generator iterative training by comparing the retrieval recall for models trained w/o and w/ the mined pseudo parallel passage pairs. Results on the DEV set of EcommerceMKP are shown in the Table 4. With additional mined pseudo parallel passage pairs, the retriever improves the Recall@5 from 50.1% to 72.4%. And therefore, the better retrieved English keyphrases lead to a better generation performance (44.50 vs. 45.73 in Table 2).

Quantity and Quality of Pseudo Parallel Passages We show the quantity and quality of mined pseudo parallel pairs in Table 5. After each iteration of passage mining, our algorithm can consistently find around 20k passage pairs from EcommerceMKP training set, which are nearly 20 times of the initial data. To assess the quality of mined passage pairs, we examine the label accuracy using the 1,000 parallel passage pairs from the DEV set of EcommerceMKP. Results in Table 5 show that while the passage mining finds a similar number of pseudo passage pairs, the labelling accuracy does increase from 28.0% to 47.0%. This is because the better pseudo parallel data improves the retriever, and the stronger retriever results in a better generator, which in turn leads to more relevant passages.

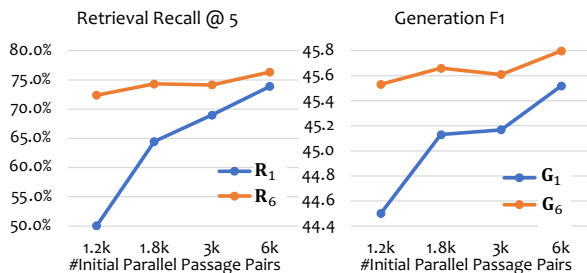


Figure 3: Performance of passage retrieval and keyphrase generation on EcommerceMKP, with different number of initial parallel data for iterative training.

<p>Product Description (German): Steiff 113437 Soft Cuddly Friends Honey Teddybär, grau, 38 cm. Bereits der Name des Soft Cuddly Friends Honey Teddybär sagt es schon aus: der 38 cm große Freund mit seinem honigsüßen Lächeln begeistert alle Kinderherzen ...</p> <p>(Translation in English): Steiff 113437 Soft Cuddly Friends Honey teddy bear, gray, 38 cm. The name of the Soft Cuddly Friends Honey Teddy bear already says it all: the 38 cm tall friend with his honey-sweet smile delights all children’s hearts ...</p> <p>Gold Keyphrases (German): <i>steiff kuscheltier; steiff teddy; soft cuddly friend; steiff; baer; grau.</i></p> <p>(Translation in English): steiff cuddly toy; steiff teddy; soft cuddly friend; steiff; bear; grey.</p> <p>Retrieved English Keyphrases: steiff teddy bear; teddy bear; my first; grey; honey; sweetheart; steiff bear; pink; vintage; steiff stuffed animal; steiff; terry; soft; jimmy.</p> <p>Predicted Keyphrases (German): <i>steiff kuscheltier; steiff teddy; soft cuddly friend; steiff; baer; grau; jimmy.</i></p> <p>(Translation in English): steiff cuddly toy; steiff teddy; soft cuddly friend; steiff; bear; grey; jimmy.</p>
--

Figure 4: Case study on the EcommerceMKP dataset. The present keyphrases (keyphrases shown in the description) are in **bold** while absent keyphrases are in *italics*. Correct predictions are in **green** while wrong predictions are in **red**.

Initializing with Different Amount of Parallel Data

To investigate the impact of initial parallel passage on the training, we conduct experiments by varying the number of parallel passage pairs on EcommerceMKP, from 1.2k (default setting) to 6k instances. We compare the single-round training (i.e., training with initial data) and iterative training after six rounds (in which round we generally obtain the best retrieval recall), on both passage retrieval (R_1 & R_6) and keyphrase generation (G_1 & G_6). Results are shown in Fig. 3. We observe that (1) the score of iterative training consistently increase when more annotated parallel data is available; (2) our iterative training demonstrates great robustness with limited parallel data (e.g. 1.2k pairs), while the benefit gradually diminishes while more parallel data becomes available.

7.3 Case Study

Fig. 4 exhibits an example of our model’s prediction. Given a product description in German, the model retrieves several English keyphrases and generates keyphrases in German accurately (translations in English are also provided). Through this example, we find that the retrieved English keyphrases do provide certain useful information such as “steiff teddy bear”, “grey” and “soft”, while it also brings some noise such as “my first”, “sweetheart” and “vintage”. Although there is a wrong prediction “jimmy” caused by the retrieved English keyphrases, the improvement in results shows that the benefits of retrieved knowledge outweigh the noise it introduces. Moreover, the retrieved keyphrases are only regarded as a supplement to the original passage, and the generator can automatically focus on the informative parts from both inputs through self-attention. Our retrieval-augmented multilingual keyphrase generation can tolerate some noise from the retrieved English keyphrases and predict better keyphrases based on these external knowledge.

8 Conclusion

In this study, we investigate a novel task setting – multilingual keyphrase generation – and contribute two new multilingual keyphrase generation datasets covering multiple domains and languages. Furthermore, we propose a retrieval-augmented multilingual keyphrase generation framework with retriever-generator iterative training. Results show that the proposed approach outperforms a wide range of baselines.

References

Wasi Ahmad, Xiao Bai, Soomin Lee, and Kai-Wei Chang. 2021. [Select, extract and generate: Neural keyphrase generation with layer-wise coverage attention](#). 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 1389–1404, Online. Association for Computational Linguistics.

Mikel Artetxe and Holger Schwenk. 2019. [Margin-based parallel corpus mining with multilingual sentence embeddings](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3197–3203, Florence, Italy. Association for Computational Linguistics.

- Florian Boudin. 2018. [Unsupervised keyphrase extraction with multipartite graphs](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 667–672, New Orleans, Louisiana. Association for Computational Linguistics.
- Deng Cai, Yan Wang, Huayang Li, Wai Lam, and Lemao Liu. 2021. [Neural machine translation with monolingual translation memory](#). 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 7307–7318, Online. Association for Computational Linguistics.
- Ricardo Campos, Vítor Mangaravite, Arian Pasquali, Alípio Jorge, Célio Nunes, and Adam Jatowt. 2020. [Yake! keyword extraction from single documents using multiple local features](#). *Inf. Sci.*, 509:257–289.
- Hou Pong Chan, Wang Chen, Lu Wang, and Irwin King. 2019. [Neural keyphrase generation via reinforcement learning with adaptive rewards](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2163–2174, Florence, Italy. Association for Computational Linguistics.
- Jun Chen, Xiaoming Zhang, Yu Wu, Zhao Yan, and Zhoujun Li. 2018. [Keyphrase generation with correlation constraints](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4057–4066, Brussels, Belgium. Association for Computational Linguistics.
- Wang Chen, Hou Pong Chan, Piji Li, Lidong Bing, and Irwin King. 2019a. [An integrated approach for keyphrase generation via exploring the power of retrieval and extraction](#). 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 2846–2856, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wang Chen, Hou Pong Chan, Piji Li, and Irwin King. 2020. [Exclusive hierarchical decoding for deep keyphrase generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1095–1105, Online. Association for Computational Linguistics.
- Wang Chen, Yifan Gao, Jiani Zhang, Irwin King, and Michael R. Lyu. 2019b. [Title-guided encoding for keyphrase generation](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6268–6275. AAAI Press.
- Rajarshi Das, Manzil Zaheer, Dung Thai, Ameya Godbole, Ethan Perez, Jay Yoon Lee, Lizhen Tan, Lazaros Polymenakos, and Andrew McCallum. 2021. [Case-based reasoning for natural language queries over knowledge bases](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9594–9611, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kushal Dave, Steve Lawrence, and David M. Pennock. 2003. [Mining the peanut gallery: opinion extraction and semantic classification of product reviews](#). In *Proceedings of the Twelfth International World Wide Web Conference, WWW 2003, Budapest, Hungary, May 20-24, 2003*, pages 519–528. ACM.
- 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.
- Samhaa R. El-Beltagy and Ahmed Rafea. 2010. [KP-miner: Participation in SemEval-2](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 190–193, Uppsala, Sweden. Association for Computational Linguistics.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. [Language-agnostic BERT sentence embedding](#). *CoRR*, abs/2007.01852.
- Corina Florescu and Cornelia Caragea. 2017. [PositionRank: An unsupervised approach to keyphrase extraction from scholarly documents](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1105–1115, Vancouver, Canada. Association for Computational Linguistics.
- Eibe Frank, Gordon W. Paynter, Ian H. Witten, Carl Gutwin, and Craig G. Nevill-Manning. 1999. [Domain-specific keyphrase extraction](#). In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, IJCAI 99, Stockholm, Sweden, July 31 - August 6, 1999. 2 Volumes, 1450 pages*, pages 668–673. Morgan Kaufmann.
- Ygor Gallina, Florian Boudin, and Beatrice Daille. 2019. [KPTimes: A large-scale dataset for keyphrase generation on news documents](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 130–135, Tokyo, Japan. Association for Computational Linguistics.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor O. K. Li. 2018. [Search engine guided neural machine translation](#). In *AAAI*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [REALM: Retrieval-augmented language model pre-training](#). *International Conference on Machine Learning (ICML)*.

- Anette Hulth and Beáta B. Megyesi. 2006. [A study on automatically extracted keywords in text categorization](#). In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 537–544, Sydney, Australia. Association for Computational Linguistics.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. [Billion-scale similarity search with gpus](#). *IEEE Trans. Big Data*, 7(3):535–547.
- Steve Jones and Mark S. Staveley. 1999. [Phrasier: A system for interactive document retrieval using keyphrases](#). In *SIGIR '99: Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, August 15-19, 1999, Berkeley, CA, USA*, pages 160–167. ACM.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Jihyuk Kim, Myeongho Jeong, Seungtaek Choi, and Seung-won Hwang. 2021. [Structure-augmented keyphrase generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2657–2667, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Dong-Hyun Lee. 2013. Pseudo-label : The simple and efficient semi-supervised learning method for deep neural networks.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Xinnian Liang, Shuangzhi Wu, Mu Li, and Zhoujun Li. 2021. [Unsupervised keyphrase extraction by jointly modeling local and global context](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 155–164, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Rui Liu, Zheng Lin, Peng Fu, and Weiping Wang. 2020a. Reinforced keyphrase generation with bert-based sentence scorer. In *2020 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking (ISPA/BDCLOUD/SocialCom/SustainCom)*, pages 1–8. IEEE.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020b. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Zhiyuan Liu, Xinxiong Chen, Yabin Zheng, and Maosong Sun. 2011. Automatic keyphrase extraction by bridging vocabulary gap. *the Fifteenth Conference on Computational Natural Language Learning*.
- Yichao Luo, Yige Xu, Jiacheng Ye, Xipeng Qiu, and Qi Zhang. 2021. [Keyphrase generation with fine-grained evaluation-guided reinforcement learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 497–507, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Rui Meng, Xingdi Yuan, Tong Wang, Sanqiang Zhao, Adam Trischler, and Daqing He. 2021. [An empirical study on neural keyphrase generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4985–5007, Online. Association for Computational Linguistics.
- Rui Meng, Sanqiang Zhao, Shuguang Han, Daqing He, Peter Brusilovsky, and Yu Chi. 2017. [Deep keyphrase generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 582–592, Vancouver, Canada. Association for Computational Linguistics.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing order into text](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.
- Funan Mu, Zhenting Yu, LiFeng Wang, Yequan Wang, Qingyu Yin, Yibo Sun, Liqun Liu, Teng Ma, Jing Tang, and Xing Zhou. 2020. Keyphrase extraction with span-based feature representations. *arXiv preprint arXiv:2002.05407*.
- Seoyeon Park and Cornelia Caragea. 2020. [Scientific keyphrase identification and classification by pre-trained language models intermediate task transfer](#)

- learning. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5409–5419, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. **KILT: a benchmark for knowledge intensive language tasks**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2523–2544, Online. Association for Computational Linguistics.
- Hieu Pham, Qizhe Xie, Zihang Dai, and Quoc V. Le. 2021. Meta pseudo labels. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11552–11563.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Jingbo Shang, Jialu Liu, Meng Jiang, Xiang Ren, Clare R Voss, and Jiawei Han. 2018. Automated phrase mining from massive text corpora. *IEEE Transactions on Knowledge and Data Engineering*, 30(10):1825–1837.
- Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darin Eide, Bo-June Paul Hsu, and Kuansan Wang. 2015. **An overview of microsoft academic service (MAS) and applications**. In *Proceedings of the 24th International Conference on World Wide Web Companion, WWW 2015, Florence, Italy, May 18-22, 2015 - Companion Volume*, pages 243–246. ACM.
- Lucas Sterckx, Thomas Demeester, Johannes Deleu, and Chris Develder. 2015. **Topical word importance for fast keyphrase extraction**. In *Proceedings of the 24th International Conference on World Wide Web Companion, WWW 2015, Florence, Italy, May 18-22, 2015 - Companion Volume*, pages 121–122. ACM.
- Avinash Swaminathan, Haimin Zhang, Debanjan Mahata, Rakesh Gosangi, Rajiv Shah, and Amanda Stent. 2020. A preliminary exploration of gans for keyphrase generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8021–8030.
- Minmei Wang, Bo Zhao, and Yihua Huang. 2016. Ptr: Phrase-based topical ranking for automatic keyphrase extraction in scientific publications. *23rd International Conference, ICONIP 2016*.
- Yue Wang, Jing Li, Hou Pong Chan, Irwin King, Michael R. Lyu, and Shuming Shi. 2019. **Topic-aware neural keyphrase generation for social media language**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2516–2526, Florence, Italy. Association for Computational Linguistics.
- Jason Weston, Emily Dinan, and Alexander Miller. 2018. **Retrieve and refine: Improved sequence generation models for dialogue**. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 87–92, Brussels, Belgium. Association for Computational Linguistics.
- Ian H. Witten, Gordon W. Paynter, Eibe Frank, Carl Gutwin, and Craig G. Nevill-Manning. 1999a. Kea: Practical automatic keyphrase extraction. In *Proceedings of the Fourth ACM Conference on Digital Libraries, DL '99*, pages 254–255, New York, NY, USA. ACM.
- Ian H. Witten, Gordon W. Paynter, Eibe Frank, Carl Gutwin, and Craig G. Nevill-Manning. 1999b. **KEA: practical automatic keyphrase extraction**. In *Proceedings of the Fourth ACM conference on Digital Libraries, August 11-14, 1999, Berkeley, CA, USA*, pages 254–255. ACM.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Hai Ye and Lu Wang. 2018. **Semi-supervised learning for neural keyphrase generation**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4142–4153, Brussels, Belgium. Association for Computational Linguistics.
- Jiacheng Ye, Ruijian Cai, Tao Gui, and Qi Zhang. 2021a. Heterogeneous graph neural networks for keyphrase generation. *arXiv preprint arXiv:2109.04703*.
- Jiacheng Ye, Tao Gui, Yichao Luo, Yige Xu, and Qi Zhang. 2021b. **One2Set: Generating diverse keyphrases as a set**. 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 4598–4608, Online. Association for Computational Linguistics.
- Jing Zhao and Yuxiang Zhang. 2019. Incorporating linguistic constraints into keyphrase generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5224–5233.

A Appendix

A.1 Implementation Details

In the cross-lingual dense passage retriever, we use “bert-base-multilingual-cased” model (Wolf et al., 2020) to initialize the query and passage encoders and fine-tune it for 15 epochs with a batch size of 32. We share the parameters between the query encoder $E_Q(\cdot)$ and the passage encoder $E_P(\cdot)$ and map English and non-English passages into the same embedding space. Empirical results show the encoder with parameter sharing can perform slightly better. The positive examples are the corresponding English passages while we randomly sample 100 passages as negative examples in training. For the retrieval-augmented keyphrase generator, we fine-tune “mbart-large-cc25” (Wolf et al., 2020) for 10 epochs with Adam (Kingma and Ba, 2015) optimizer with a learning rate of $1e-4$, a batch size of 8, a warm-up rate of 50 training steps. Similar to most Seq2Seq models, we train the mBART-based generation module by optimizing the negative log-likelihood loss of the ground-truth keyphrase sequence, and use beam search decoding with a beam size of 5 during inference. The number of retrieved keyphrases m for retrieval-augmented generation is a hyperparameter and is tuned on the development set. We use keyphrases from $m = 1$ English passages for AcademicMKP dataset and $m = 5$ for EcommerceMKP dataset. During inference, we set the maximum target sequence length as 128 and set the beam decoding size as 5. For parallel passage mining via iterative training, we continue the iterative process until the retrieval recall does not improve. The total number of iterations (T in Algo. 1) are 6 and 3 on EcommerceMKP and AcademicMKP respectively. The threshold τ in Line 23 for Algo. 1 is set as 5.

A.2 Variants of Retrieval Targets

There exists a misalignment between the retriever and the generator model. The retriever retrieves similar passages while the generator utilizes the associated keyphrases of these passages (not the retrieved passages) as external knowledge for generation. Therefore, a good retriever does not necessarily guarantee the good quality and usefulness of these keyphrases.

We tried two different retrieval targets which might close the misalignment. Given a non-English passage, we tried to either directly retrieve English keyphrases (RAMKG-P2K) or retrieve the

Model	Retrieval Results			Generation Results		
	Recall@1	Recall@2	Recall@5	P	R	F1
P2P	26.00%	36.78%	50.05%	48.51	44.71	44.50
P2K	2.86%	4.70%	9.09%	46.91	43.13	42.95
P2PK	25.25%	35.45%	49.51%	48.50	44.39	44.34

Table 6: Results of RAMKG variants with different retrieval targets.

τ	0	5	10	15
Recall	62.28%	64.01%	63.67%	62.97%

Table 7: Influence of the threshold τ on the retriever-generator self-training algorithm.

concatenated sequences of passage-keyphrase pair (RAMKG-P2PK). We find that (1) RAMKG-P2K that directly retrieves keyphrases has poor retrieval performance. This is because it is hard to capture the similarity between non-English passages and English keyphrases; (2) RAMKG-P2PK has slightly worse results than only retrieval EN passages, which implies that additionally adding keyphrases in the retrieval targets does not bring any benefit.

Results are shown in Table 6. RAMKG (P2P) is our original model which retrieves English passages given a non-english passage. Results tell us that 1) directly retrieval of keyphrases have poor retrieval performance. This is because it is hard to capture the similarity between non-english passages and english keyphrases; 2) RAMKG (P2PK) has slightly worse results than the model , which implies that additionally adding keyphrases in the retrieval targets does not bring any benefit.

A.3 Discussion on Retriever-Generator Iterative training (RGIT) Algorithm

Difference between RGIT and Self-Training.

Our approach shares some similarities with self-training (Lee, 2013; Pham et al., 2021) but there are some differences. In self-training, the teacher and student models are in the same architecture and focus on the same training objectives. In our proposed retriever-generator iterative training, the retriever and generator are two different models and optimized by different objectives.

Threshold Tuning. In this section, we investigate the impact of the chosen threshold τ (line 24) in our proposed retriever-generator iterative training. We tune the threshold (tau) on AcademicMKP and results are shown in Table 7. Results show

that $\tau=5$ receives the best retrieval performance. $\tau=0$ brings more pseudo parallel passage pairs but introduces more noise. Larger τ (10/15) reduces the number of pseudo pairs, making the iterative training less effective.