

# DictBERT: Dictionary Description Knowledge Enhanced Language Model Pre-training via Contrastive Learning

Qianglong Chen<sup>1,2</sup>, Feng-Lin Li<sup>2</sup>, Guohai Xu<sup>2</sup>, Ming Yan<sup>2</sup>, Ji Zhang<sup>2</sup>, Yin Zhang<sup>1\*</sup>

<sup>1</sup>College of Computer Science and Technology, Zhejiang University, China

<sup>2</sup>Alibaba Group, China

{chenqianglong, zhangyin98}@zju.edu.cn, {guohai.xgh, ym119608, zj122146}@alibaba-inc.com, maillifenglin@gmail.com

## Abstract

Although pre-trained language models (PLMs) have achieved state-of-the-art performance on various natural language processing (NLP) tasks, they are shown to be lacking in knowledge when dealing with knowledge driven tasks. Despite the many efforts made for injecting knowledge into PLMs, this problem remains open. To address the challenge, we propose **DictBERT**, a novel approach that enhances PLMs with dictionary knowledge which is easier to acquire than knowledge graph (KG). During pre-training, we present two novel pre-training tasks to inject dictionary knowledge into PLMs via contrastive learning: *dictionary entry prediction* and *entry description discrimination*. In fine-tuning, we use the pre-trained DictBERT as a plugin knowledge base (KB) to retrieve implicit knowledge for identified entries in an input sequence, and infuse the retrieved knowledge into the input to enhance its representation via a novel extra-hop attention mechanism. We evaluate our approach on a variety of knowledge driven and language understanding tasks, including NER, relation extraction, CommonsenseQA, OpenBookQA and GLUE. Experimental results demonstrate that our model can significantly improve typical PLMs: it gains a substantial improvement of 0.5%, 2.9%, 9.0%, 7.1% and 3.3% on BERT-large respectively, and is also effective on RoBERTa-large.

## 1 Introduction

Pre-trained language models (PLMs) such as BERT [Devlin *et al.*, 2019], RoBERTa [Liu *et al.*, 2019] and ALBERT [Lan *et al.*, 2019] have been prevailing in both academic and industrial community due to their state-of-the-art performance on various natural language processing (NLP) tasks. However, as they capture only a general language representation learned from large-scale corpora, they are shown to be lacking in knowledge when dealing with knowledge driven tasks [Talmor *et al.*, 2019; Mihaylov *et al.*, 2018]. To address this challenge, many efforts, such as ERNIE-THU [Zhang *et al.*,

2019], KEPLER [Wang *et al.*, 2021c], KnowBERT [Peters *et al.*, 2019], K-Adapter [Wang *et al.*, 2021b] and ERICA [Qin *et al.*, 2021] have been made for injecting knowledge into PLMs for further improvement.

However, existing knowledge enhanced PLMs (i.e., K-PLMs) still suffer from several deficiencies. First, few methods pay attention to knowledge itself, including what type of knowledge is needed and the feasibility of acquiring such knowledge. On the one hand, some models take for granted the use of knowledge graph (KG), which is difficult to acquire in practice and shown to be less effective than dictionary knowledge [Xu *et al.*, 2021; Chen *et al.*, 2020]. On the other hand, many methods use Wikipedia, which is easier to access but often noisy and of low knowledge density. Second, current K-PLMs mainly focus on one or two types of knowledge-driven tasks. Although they are shown to be useful on a few specific tasks, their language understanding ability was either not further validated on GLUE [Liu *et al.*, 2020; Wang *et al.*, 2021b] or not improved [Zhang *et al.*, 2019]. That is, the application scope of such K-PLMs is limited.

Inspired by the hint that dictionary knowledge can be even more effective than structured knowledge [Chen *et al.*, 2020], we leverage dictionary sources as external knowledge to enhance PLMs. In our experience, this enjoys several benefits. First, it is consistent with human reading habit and cognitive process. In the process of reading, when encountering unfamiliar words, people usually consult dictionaries or encyclopedias. Second, compared with long Wikipedia texts, dictionary knowledge is more concise and of high knowledge density. Third, dictionary knowledge is much easier to access, which is of key importance for applying K-PLMs in practice. Even in the case of lacking a dictionary, it can be acquired through simply constructing a generator to summarize the description explaining a word.

Correspondingly, we propose **DictBERT**, an effective approach that enhances PLMs with dictionary knowledge via contrastive learning. In the pre-training stage, we inject dictionary knowledge into PLMs through two novel pre-training tasks: *dictionary entry prediction*, in which we use a description to predict its masked entry and learn entry representations from descriptive texts; and *entry description discrimination*, where we use contrastive learning to improve the robustness of entry representations by constructing positive and negative samples with dictionary synonyms and antonyms.

\*Corresponding Author: Yin Zhang.

During fine-tuning, we first identify dictionary entries from a given input, then use DictBERT as a plugin KB to retrieve corresponding entry information. For the fusion of retrieved entry information and original input, we propose a novel extra-hop attention mechanism to enhance its representation for downstream tasks.

The main contributions of our paper are as follows:

- We propose DictBERT, a novel approach enhancing PLMs with dictionary knowledge, which is able to effectively not only integrate external knowledge into but also improve the language understanding ability of PLMs.
- For pre-training, we present two novel pre-training tasks with contrastive learning, namely dictionary entry prediction and entry description discrimination, for injecting dictionary knowledge into PLMs. For fine-tuning, we present three knowledge infusion mechanisms to utilize the retrieved knowledge from pre-trained DictBERT for improving downstream tasks.
- We conducted a series of experiments on NER, relation extraction (RE), CommonsenseQA, OpenBookQA and GLUE. Experimental results show that our model can significantly improve typical PLMs (BERT-large and RoBERTa-large).

## 2 Related Work

**Knowledge Enhanced PLMs.** To alleviate the problem of lacking knowledge for PLMs, a popular approach is to inject factual knowledge through infusing pre-trained entity embeddings [Zhang *et al.*, 2019; Peters *et al.*, 2019] or incorporating symbolic knowledge triples [Liu *et al.*, 2020]. One problem of using pre-trained entity embeddings, as pointed out by CoLAKE [Sun *et al.*, 2020], is the separation between entity embedding and language embedding. To tackle this problem, we use textual entry descriptions to predict their masked entries in the entry prediction pre-training task. Being different from KnowBERT [Peters *et al.*, 2019] and KEPLER [Wang *et al.*, 2021c] that use structured KGs, we use semi-structured dictionary knowledge. Inspired by K-Adapter [Wang *et al.*, 2021b], we also use the PLM enhanced with dictionary knowledge as a plugin for downstream tasks. It should be noted that Dict-BERT [Yu *et al.*, 2021] and our work are at the same period. There are many differences between them in pre-training and fine-tuning, and our results are superior to those of Dict-BERT.

**Contrastive Learning.** The main idea of contrastive learning is to improve the robustness of representations through bringing closer positive samples and pushing away negative ones. It has been widely used to obtain better sentence representations [Logeswaran and Lee, 2018; Wu *et al.*, 2020; Wang *et al.*, 2021a; Qin *et al.*, 2021]. While [Logeswaran and Lee, 2018] take a sentence B that follows A as a positive example and randomly chooses a sentence C from other documents as a negative, [Wang *et al.*, 2021a] construct positive and negative examples via replacing representative tokens in a sentence with WordNet, and [Wu *et al.*, 2020] present multiple sentence-level data augmentation strategies. In addition, [Qin *et al.*, 2021] leverage contrastive pre-training to

<b>Entry word</b>	forest
<b>Description</b>	a large area of land covered with trees and plants, usually larger than a wood, or the trees and plants themselves
<b>Synonyms</b>	jungle, woodland
<b>Antonyms</b>	desert, wasteland

Table 1: A sample of dictionary entries.

improve the ability of PLMs on capturing relational facts in texts. Being differently, we use synonyms and antonyms in dictionary to construct contrastive pairs, and use contrastive pre-training to learn a better dictionary entry representation for downstream tasks.

## 3 The Proposed Approach

### 3.1 Dictionary Description Knowledge

A dictionary is a resource that lists the words of a language, clarifies their meanings through explanatory descriptions, and often specifies their pronunciation, origin, usage, synonyms and antonyms, etc. Table 1 shows an example about the entry word “forest”. In this paper, we use four kinds of information for pre-training: each entry, its description(s), synonym(s) and antonym(s). We leverage dictionary entry words and their meanings (i.e., explanatory descriptions) for knowledge injection pre-training. Also, in order to improve the robustness of entry representation, we use the synonyms and antonyms of an entry word for contrastive learning.

### 3.2 Pre-training DictBERT

As shown in Figure 1, we use two novel pre-training tasks: (1) dictionary entry prediction and (2) entry description discrimination, to capture the different aspects of dictionary knowledge through further training a PLM.

**Dictionary Entry Prediction.** For entry word prediction, we follow the design of masked language modeling (MLM) in BERT [Devlin *et al.*, 2019], but impose constraints on the tokens to be masked. Originally, given an input sequence, the MLM task randomly masks a certain percentage of the input tokens with a special [MASK] symbol, and then tries to recover them. Inspired by [Tsukagoshi *et al.*, 2021], to effectively learn entry representations, we take as input the concatenation of each entry word  $e = \{t_1, \dots, t_i, \dots, t_m\}$  and its description  $desc = \{w_1, w_2, \dots, w_n\}$  in a dictionary  $D$ , perform masking only on the tokens of entry  $e$  in a chosen input sample  $s = \{[CLS]e[SEP]desc[SEP]\}$ , and at last predict the masked entry tokens based on the corresponding description  $desc$ . Note that if an entry  $e$  consists of multiple tokens, all of the component tokens will be masked. In the case of polysemy, where an entry  $e$  has multiple meanings (i.e., descriptions), we construct an input sample for each meaning in a similar way. We formulate the entry token prediction as:

$$P(t_1, \dots, t_i, \dots, t_m | s \setminus \{t_1, \dots, t_i, \dots, t_m\}) \quad (1)$$

where the  $t_i$  is the  $i$ -th token of  $e$ , and  $s \setminus \{t_1, \dots, t_i, \dots, t_m\}$  denotes the sample  $s$  with entry tokens  $t_{i..m}$  being masked. We initialize our model with the pre-trained checkpoint of BERT-large and keep MLM as one of our objectives, which uses the cross-entropy loss as loss function  $L_{dep}$ .

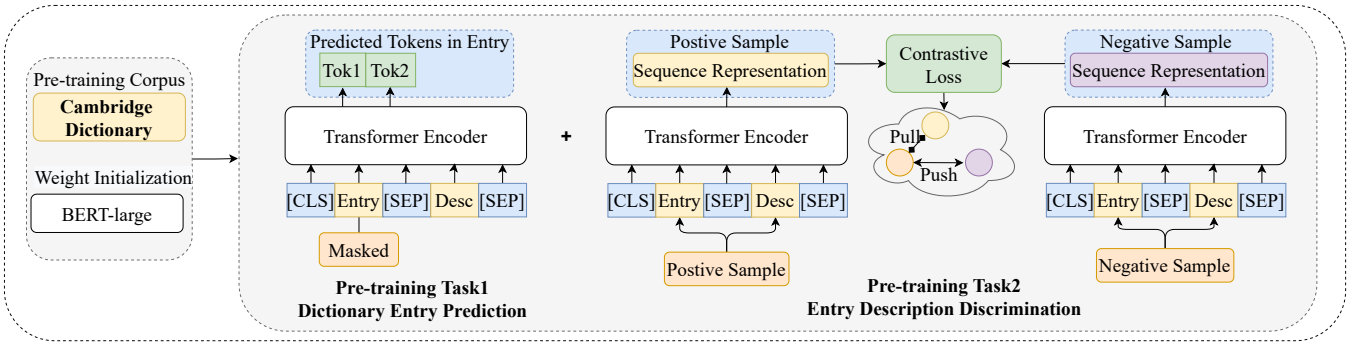


Figure 1: DictBERT pre-training. We take the Cambridge dictionary as our knowledge source. The pre-training tasks include **dictionary entry prediction** and **entry description discrimination**. In the former task, we mask only the entry tokens in the chosen input samples. In the latter task, we try to obtain better entry representations through contrastive learning.

<b>Positive</b>	[CLS] woodland [SEP] Land covered with wood or trees [SEP]
<b>Negative</b>	[CLS] desert [SEP] arid land with little or no vegetation [SEP]

Table 2: The positive sample “woodland” and negative example “desert” of the entry “forest”.

**Entry Description Discrimination.** To better capture the semantics of dictionary entries, we introduce entry description discrimination, which tries to improve the robustness of entry representations through contrastive learning. Specifically, we construct positive (resp. negative) samples as follows: given an entry word  $e$  and its description  $desc$ , we obtain its synonyms  $D_s = \{e_{syn}\}$  (resp. antonyms  $D_a = \{e_{ant}\}$ ) from the dictionary source, and treat the concatenation of each  $e_{syn}$  (resp.  $e_{ant}$ ) and its description  $desc_{syn}$  (resp.  $desc_{ant}$ ) as a positive (resp. negative) sample. Take the entry “forest” in Table 1 for example, “woodland” and “desert” are one of its synonyms and antonyms, respectively. The corresponding positive and negative samples are shown in Table 2. In our experiments, we use the same number of (e.g., 5) positive and negative samples. Note that we currently only utilize the antonyms of an entry word to construct strict negative samples, and will explore the construction of negative samples through random selection in the future.

We use  $h_{ori}$ ,  $h_{syn}$ ,  $h_{ant}$  to indicate the representations of the original, the positive, and the negative input sample. To bring closer  $h_{ori}$  and  $h_{syn}$ , and push away  $h_{ori}$  and  $h_{ant}$ , we develop a contrastive objective, where  $(e_{ori}, e_{syn})$  is considered a positive pair and  $(e_{ori}, e_{ant})$  is considered negative. We use  $h^c$ , which denotes the hidden state of the special symbol [CLS], to indicate the representation of an input sample. We define a contrastive objective  $L_{edd}$ :

$$h_{ori}^c = \text{Encoder}(e_{ori}) \quad (2)$$

$$h_{syn}^c = \text{Encoder}(e_{syn}) \quad (3)$$

$$h_{ant}^c = \text{Encoder}(e_{ant}) \quad (4)$$

$$f(h_i^c, h_j^c) = \exp(h_i^c h_j^c) \quad (5)$$

$$L_{edd} = - \sum_{e \in D} \log \frac{f(e_{ori}, e_{syn})}{f(e_{ori}, e_{syn}) + f(e_{ori}, e_{ant})} \quad (6)$$

where the  $f(x, y)$  denotes the exponentiation of the dot product between hidden states  $x$  and  $y$ .

We sum the dictionary entry prediction task loss and the entry description discrimination task loss, and finally obtain the overall loss function  $L$ :

$$L = \lambda_1 L_{dep} + \lambda_2 L_{edd} \quad (7)$$

where  $L_{dep}$  and  $L_{edd}$  denote the loss functions of the two tasks. In our experiments, we set  $\lambda_1 = 0.4$  and  $\lambda_2 = 0.6$ .

### 3.3 Fine-tuning with DictBERT

Inspired by [Petroni *et al.*, 2019], we use DictBERT as a plugin with a backbone PLM during fine-tuning (i.e., it is frozen). In this way, we can enjoy the flexibility of training different DictBERTs for different dictionaries and avoid the catastrophic forgetting problem of continuous training. Specifically, we first identify dictionary entries from a given input, then use DictBERT as a KB to retrieve corresponding entry information (i.e., entry embeddings), and finally inject the retrieved entry information into the original input to get an enhanced representation for downstream tasks. In the case an input consists of more than one sequence (e.g., NLI), we process each input sequence individually and then feed them into the downstream task specific layer for subsequent processing. To better leverage the retrieved implicit knowledge on downstream tasks, we introduce three different kinds of knowledge infusion mechanisms (See Figure 2): (1) pooled output concatenation, (2) extra-hop attention and (3) layer-wise extra-hop attention.

**Pooled Output Concatenation.** As shown in Figure 2 (a), we directly concatenate the pooled output of the backbone BERT (i.e.,  $h^c$ ) and the sum of entry embeddings retrieved from DictBERT (i.e.,  $\hat{h}$ ). Then, we feed the concatenation (i.e.,  $[h^c; \hat{h}]$ ) into a task specific layer for downstream tasks.

**Extra-hop Attention.** The simplest way to incorporate identified entries into original text is to sum up their embeddings and concatenate the summation with the text representation. However, this method can not tell which entry is more

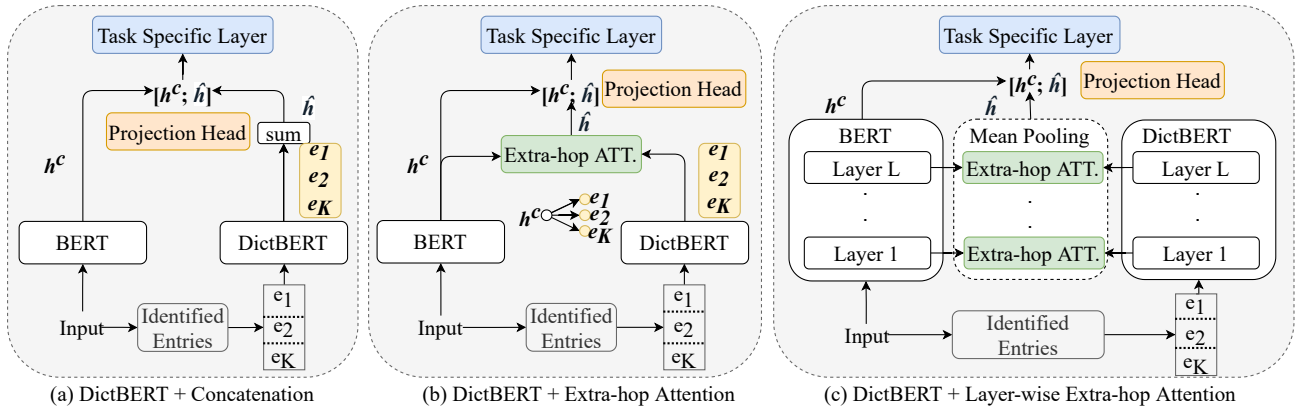


Figure 2: Downstream task fine-tuning with DictBERT, where we present three different knowledge infusion mechanisms.

important, and which sense is more suitable in the case of polyseme entries. Therefore, we further propose an extra-hop attention mechanism to address this deficiency. As shown in Figure 2 (b), we follow Transformer-XH [Zhao *et al.*, 2020] to use  $h^c$ , the hidden state of the [CLS] token in an input query as the ‘‘attention hub’’, which attends to each entry word identified in the same input. With the attentive weights, our method focuses on more important entries or meanings when integrating them as external knowledge into the original input query. The extra hop attention mechanism is formulated as:

$$\hat{h} = \sum_{i=1}^K ATT(h^c, e_i) \quad (8)$$

where  $e_i$  denotes the DictBERT output of  $i_{th}$  identified entry,  $K$  is the number of identified entries in the input query, and  $\hat{h}$  denotes the weighted sum of retrieved entry embeddings. After we obtain the  $\hat{h}$ , we use  $[h^c; \hat{h}]$  for the final inference.

**Layer-wise Extra-hop Attention.** To further improve performance, we extend the extra-hop attention at the last layer to each inner layer, making it become layer-wise. As shown in Figure 2 (c), we compute the attention score at each layer, and finally use their mean for implicit entry knowledge injection. Specifically, the layer-wise extra-hop attention can be formulated as:

$$\hat{h}_l = \sum_{i=1}^K ATT(h_l, e_i^l) \quad (9)$$

$$\hat{h} = \frac{1}{L} \sum_{l=1}^L \hat{h}_l \quad (10)$$

where  $\hat{h}_l$  denotes the weighted sum of  $l$ -th layer outputs of DictBERT. With the final implicit  $\hat{h}$  obtained via Equation 10, we use  $[h^c; \hat{h}]$  in a similar way for downstream tasks.

## 4 Experiments

### 4.1 Datasets and Tasks

**Pre-training Dictionary Source.** To pre-train DictBERT, we use the Cambridge Dictionary<sup>1</sup>, which includes 315K en-

try words, as our pre-training corpus. We construct input samples for the two pre-training tasks, namely dictionary entry prediction and entry description discrimination, as introduced in the section of the proposed approach.

**CoNLL2003 & TACRED.** We use these two traditional knowledge-driven tasks, CoNLL2003 [Tjong Kim Sang and De Meulder, 2003] and TACRED [Zhang *et al.*, 2017], to have a quick check on the effectiveness of our approach.

**CommonsenseQA & OpenBookQA.** We use CommonsenseQA [Talmor *et al.*, 2019] and OpenBookQA [Mihaylov *et al.*, 2018] to evaluate the ability of DictBERT acting as KBs and providing implicit knowledge to downstream tasks.

**GLUE.** We follow existing knowledge enhanced PLMs such as KEPLER and KnowBERT to use GLUE [Wang *et al.*, 2018] to evaluate the general natural language understanding capability of our approach.

### 4.2 Experimental Settings

For pre-training, we use the BERT-large-uncased and RoBERTa-large model as backbone and set the learning rate to  $1e^{-5}$ , dropout rate to 0.1, max-length of tokens to 128, batch size to 32, and number of epochs to 10. We use AdamW as the optimizer. For fine-tuning, we adopt cross-entropy loss as the loss function, set batch size to 32 and number of epochs to 30. We run 5 times for each task and report their average.

### 4.3 Baselines

**BERT & RoBERTa.** We adopt BERT-large [Devlin *et al.*, 2019] instead of BERT-base as baseline because the former is more difficult to improve. To be more convincing, we also use the more adequately trained RoBERTa-large [Liu *et al.*, 2019] for comparison in our experiments.

**Enhanced BERT & RoBERTa.** For CommonsenseQA, we use BERT+AMS [Ye *et al.*, 2019], BERT+OMCS, RoBERTa+CSPT, RoBERTa+KE, G-DAUG [Yang *et al.*, 2020] as baselines for comparison. For OpenbookQA, we use AristoBERTv7, AristoRoBERTav7 and BERT Multi-Task as baselines for comparison.

<sup>1</sup><https://dictionary.cambridge.org>

Model	CoNLL2003	TACRED
BERT-large	92.8	70.1
DictBERT + Concat(K)	93.1	72.3
DictBERT + EHA(K)	93.2	72.7
DictBERT + EHA(K+V)	93.3	72.8
DictBERT + LWA(K+V)	<b>93.3</b>	<b>73.0</b>

Table 3: Experimental results on CoNLL2003 (NER) and TACRED (relation extraction).

Model	CSQA	OBQA
BERT-large	56.7	60.4
BERT-large + AMS	62.2	-
BERT-large + OMCS	62.5	-
BERT-large Multi-Task	-	63.8
AristoBERTv7-large	64.6	72.0
RoBERTa-large	72.1	71.8
RoBERTa-large + CSPT	69.6	-
RoBERTa-large + G-DAUG-Combo	72.6	-
RoBERTa-large + KE	73.3	-
AristoRoBERTav7-large	-	77.8
DictBERT + Concat(K)	62.7	64.4
DictBERT + EHA(K)	65.1	66.3
DictBERT + EHA(K+V)	65.4	66.7
DictBERT + LWA(K+V)	65.7	67.5
DictRoBERTa + Concat(K)	75.7	75.2
DictRoBERTa + EHA(K)	77.5	77.6
DictRoBERTa + EHA(K+V)	77.8	78.1
DictRoBERTa + LWA(K+V)	<b>78.5</b>	<b>78.3</b>

Table 4: Experimental results on CommonsenseQA (CSQA) and OpenBookQA (OBQA).

**KnowBERT & KEPLER.** For GLUE, we use KnowBERT and KEPLER as baselines for comparison. KnowBERT [Peters *et al.*, 2019] enhances contextual word representations through embedding structured, human-curated knowledge into BERT-base through entity linking and word-to-entity attention. KEPLER [Wang *et al.*, 2021c] encodes textual entity descriptions with RoBERTa-base as their embeddings, and then jointly optimizes the knowledge embedding and language modeling objectives.

#### 4.4 DictBERT Variants

We evaluate different variants of DictBERT in our experiments. **DictBERT+Concat(K)** uses the concatenation mechanism, **DictBERT+EHA(K)** and **DictBERT+EHA(K+V)** adopt the extra-hop attention mechanism, and **DictBERT+LWA(K+V)** uses layer-wise attention. The symbol **K** indicates the use of entry word to retrieve entry embeddings from DictBERT, **K+V** denotes that we use both entry word and its corresponding description for knowledge retrieval.

#### 4.5 Experimental Results and Analysis

**Traditional Knowledge Driven Task Results.** Firstly, we evaluate DictBERT on NER and relation extraction, the most commonly used knowledge driven tasks. As shown in Table 3, our approach is finally able to improve the performance on CoNLL2003 and TACRED by 0.5% and 2.9% compared

with the strong baseline BERT-large. Further, we can observe that all the three knowledge infusion mechanisms are helpful, and the layer-wise attention achieve the best results. This indicates the identification and explanation of important entities in an input sample are of key importance. Meanwhile, we found that the use of additional entry description (i.e., the K+V setting) can help retrieve better entry embeddings.

**Knowledge Driven QA Task Results.** We further assess DictBERT on knowledge driven QA tasks, namely CommonsenseQA and OpenBookQA, and report the results in Table 4. Compared with BERT-large, our basic setting DictBERT+Concat gains a significant improvement of 6.0% and 4.0% on the two tasks, respectively. Further, we observe that the extra-hop attention brings an evident increase (2.4% and 1.9%), verifying again the importance of identifying attentive weights of entries in an input sample. Lastly, DictBERT+LWA(K+V) achieves the best result on both tasks, bringing a final gain of 9.0% and 7.1% compared to the BERT-large baseline. To be more convincing, we also compare DictRoBERTa with the original RoBERTa-large on CommonsenseQA and OpenBookQA. As shown in Table 4, the conclusion also holds for RoBERTa. Similarly, DictRoBERTa+LWA(K+V) achieves the best results, which can ultimately improve over 6.4% and 6.5%, respectively.

**GLUE Results.** We also evaluate DictBERT on GLUE to examine whether it can improve the general natural language understanding ability of PLMs. Table 5 shows that compared with BERT-large our basic setting DictBERT+Concat achieves an average improvement of 2.4%, indicating the effectiveness of injecting dictionary knowledge for language understanding. Similarly, the extra-hop attention and the use of additional entry description (i.e., the K+V setting) contribute to further improvement, and DictBERT+LWA(K+V) achieves the best results, bringing a final increase of 3.3% on average. With the baseline being RoBERTa-large, our best model can achieve an average increase of 0.9% on GLUE, validating the effectiveness and broad applicability of our approach. As for other K-PLMs, KnowBERT enhanced with WordNet and 470K Wikipedia entities can improve BERT-base by 2.0%, which is smaller than the performance gain (3.3%) brought by our method on BERT-large. KEPLER-wiki can only improve RoBERTa-base by 0.2% when using 5M Wikidata entities for knowledge (entity) embedding and extra 13GB text data for MLM. With the only 5M entity descriptions for MLM, there is an obvious performance drop for KEPLER-OnlyDesc. Therefore, our approach is more effective in improving language understanding ability with external knowledge (we use only 315K dictionary entries).

**Ablation Study.** We perform ablation studies on the different components of DictBERT. Firstly, we evaluate BERT-large+Concat(K) and BERT-large+LWA(K+V), which directly use BERT-large, instead of our pre-trained DictBERT, as the plugin. As we can see, the improvement is rather marginal, confirming the necessity of injecting external knowledge. Secondly, we assess the effectiveness of the two pre-training tasks: DictBERT(DEP)+Concat and DictBERT(DEP+EDD)+Concat. As shown in Table 6, contrastive learning is helpful to some degree (0.4% on average), and

Model	Single-sentence		Similarity Paraphrase			Inference			Avg
	CoLA	SST-2	QQP	STS-B	MRPC	QNLI	MNLI	RTE	
BERT-base	52.1	93.5	88.9	85.8	88.9	90.5	84.6	66.4	81.3
BERT-large	60.5	94.9	89.3	87.6	89.3	92.7	85.9	70.1	83.8
RoBERTa-base	63.6	94.8	91.9	91.2	90.2	92.7	87.5	80.9	86.6
RoBERTa-large	67.8	96.7	90.2	92.0	93.0	95.4	90.2	87.2	89.0
KnowBERT-WordNet+Wiki	54.6	93.6	90.3	89.1	88.2	91.5	85.7	73.8	83.3
KEPLER-only Desc	55.8	94.4	90.8	90.2	88.5	92.4	85.9	78.3	84.5
KEPLER-wiki	63.6	94.5	<b>91.7</b>	91.2	89.3	92.4	87.2	85.2	86.8
DictBERT + Concat(K)	64.1	95.4	90.3	90.5	91.8	95.1	87.1	75.2	86.2
DictBERT + EHA(K)	64.2	95.6	90.3	91.5	92.0	95.7	88.4	76.6	86.8
DictBERT + EHA(K+V)	64.5	95.6	90.4	91.6	92.1	95.8	88.5	76.8	86.9
DictBERT + LWA(K+V)	64.7	95.7	90.4	91.7	92.3	96.0	88.7	77.1	87.1
DictRoBERTa + Concat(K)	68.1	97.1	90.3	91.2	92.5	96.2	90.5	88.1	89.2
DictRoBERTa + EHA(K)	68.3	97.3	90.4	91.7	92.7	96.4	90.8	88.6	89.5
DictRoBERTa + EHA(K+V)	68.5	97.5	90.5	91.8	92.8	96.5	90.9	89.1	89.7
DictRoBERTa + LWA(K+V)	<b>68.6</b>	<b>97.8</b>	90.8	<b>92.1</b>	<b>93.2</b>	<b>96.8</b>	<b>91.1</b>	<b>89.4</b>	<b>89.9</b>

Table 5: Experimental results on the GLUE development set. The parameter of DictBERT is based on BERT-large. For parameter initialization, KnowBERT uses the BERT-base, while KEPLER uses RoBERTa-base.

Model	CSQA	OBQA	GLUE
BERT-large	56.7	60.4	83.8
BERT-large + Concat(K)	57.1	60.6	83.9
BERT-large + LWA(K+V)	57.3	60.7	84.1
<b>Pre-training</b>			
BERT-large (dict corpus)+Concat(K)	61.9	63.8	85.5
DictBERT(DEP)+Concat(K)	62.2	64.1	85.8
DictBERT(DEP+EDD)+Concat(K)	62.7	64.4	86.2
<b>Fine-tuning</b>			
DictBERT-only	62.6	64.1	85.7
DictBERT + Concat(K)	62.7	64.4	86.2
DictBERT + EHA(K)	65.1	66.3	86.8
DictBERT + EHA(K+V)	65.4	66.7	86.9
DictBERT + LWA(K+V)	65.7	67.5	87.1
DictBERT plus + LWA(K+V)	<b>66.1</b>	<b>67.7</b>	<b>87.2</b>

Table 6: Ablation study results on CSQA, OBQA and GLUE.

masking only entry tokens is better than masking tokens of both entries and descriptions (+0.3% for all the three). Finally, we examine the necessity of using DictBERT as a plugin KB instead of directly using it for downstream task fine-tuning (DictBERT-only), and whether the dictionary size matters (DictBERT plus). As shown in Table 6, all of our three knowledge infusion mechanisms can further improve the performance of DictBERT-only, indicating the use of DictBERT as a plugin is rewarding. To assess the effect of dictionary size, we use the union of the Cambridge Dictionary, the Oxford Dictionary and the Wiktionary, which totals more than 1M unique entry words. The results show that DictBERT plus+LWA(K+V) can further improve the performance of the three task sets (+0.23% on average).

#### 4.6 Discussions

Experimental results show that DictBERT can not only integrate external knowledge into but also improve the language understanding ability of PLMs. It is worth mentioning

that DictBERT is assessed on very strong baselines (BERT-large and RoBERTa-large, rather than the base counterparts adopted by many other K-PLMs), which indicates the effectiveness of our method from another side. Last but not least, our approach can be easily applied in practice: dictionary source is relatively easy to acquire, through either crawling or simple generative models. As for computation cost, DictBERT as a KB plugin can be further simplified to be DictBERT as a lookup table, with each entry in a dictionary being mapped to an embedding in advance, largely accelerating the inference speed. Through generating dictionary entry embeddings in advance by using the plugin, the complexity of our approach is similar to that of the backbone PLM.

## 5 Conclusion

In this paper, we propose DictBERT, an effective approach that enhances PLMs with dictionary knowledge through two novel pre-training tasks and an attention-based knowledge infusion mechanism during downstream task fine-tuning. We also demonstrate its effectiveness through an adequate set of experiments. Importantly, our approach can be easily applied in practice. In the future, we are going to further explore more effective pre-training tasks and knowledge infusion mechanisms for injecting knowledge into multilingual pre-trained language models.

## Acknowledgments

We thank the anonymous reviewers for their helpful comments on this paper. This work was supported by National Key R&D Program of China (No. 2018AAA0101900), the NSFC projects (No. 62072399, No. U19B2042, No. 61402403), Chinese Knowledge Center for Engineering Sciences and Technology, MoE Engineering Research Center of Digital Library, Alibaba Group, Alibaba-Zhejiang University Joint Research Institute of Frontier Technologies, and the Fundamental Research Funds for the Central Universities (No. 226-2022-00070).

## References

- [Chen *et al.*, 2020] Qianglong Chen, Feng Ji, Haiqing Chen, and Yin Zhang. Improving commonsense question answering by graph-based iterative retrieval over multiple knowledge sources. In *COLING*, pages 2583–2594, 2020.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, 2019.
- [Lan *et al.*, 2019] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv:1909.11942*, 2019.
- [Liu *et al.*, 2019] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv:1907.11692*, 2019.
- [Liu *et al.*, 2020] Weijie Liu, Peng Zhou, Zhe Zhao, et al. K-bert: Enabling language representation with knowledge graph. In *AAAI*, pages 2901–2908, 2020.
- [Logeswaran and Lee, 2018] Lajanugen Logeswaran and Honglak Lee. An efficient framework for learning sentence representations. In *ICLR*, 2018.
- [Mihaylov *et al.*, 2018] Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*, pages 2381–2391, 2018.
- [Peters *et al.*, 2019] Matthew E. Peters, Mark Neumann, Robert L Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A. Smith. Knowledge enhanced contextual word representations. In *EMNLP*, pages 43–54, 2019.
- [Petroni *et al.*, 2019] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. Language models as knowledge bases? In *EMNLP*, pages 2463–2473, 2019.
- [Qin *et al.*, 2021] Yujia Qin, Yankai Lin, Ryuichi Takanobu, Zhiyuan Liu, Peng Li, Heng Ji, Minlie Huang, Maosong Sun, and Jie Zhou. ERICA: Improving entity and relation understanding for pre-trained language models via contrastive learning. In *ACL*, pages 3350–3363, 2021.
- [Sun *et al.*, 2020] Tianxiang Sun, Yunfan Shao, Xipeng Qiu, Qipeng Guo, Yaru Hu, Xuanjing Huang, and Zheng Zhang. Colake: Contextualized language and knowledge embedding. In *COLING*, pages 3660–3670, 2020.
- [Talmor *et al.*, 2019] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. CommonsenseQA: A question answering challenge targeting commonsense knowledge. In *NAACL-HLT*, pages 4149–4158, 2019.
- [Tjong Kim Sang and De Meulder, 2003] Erik F. Tjong Kim Sang and Fien De Meulder. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *NAACL-HLT*, pages 142–147, 2003.
- [Tsukagoshi *et al.*, 2021] Hayato Tsukagoshi, Ryohei Sasano, and Koichi Takeda. DefSent: Sentence embeddings using definition sentences. In *ACL*, 2021.
- [Wang *et al.*, 2018] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *EMNLP*, 2018.
- [Wang *et al.*, 2021a] Dong Wang, Ning Ding, Piji Li, and Hai-Tao Zheng. Cline: Contrastive learning with semantic negative examples for natural language understanding. In *ACL*, pages 2332–2342, 2021.
- [Wang *et al.*, 2021b] Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Jianshu Ji, Guihong Cao, Daxin Jiang, and Ming Zhou. K-Adapter: Infusing Knowledge into Pre-Trained Models with Adapters. In *Findings of ACL*, pages 1405–1418, 2021.
- [Wang *et al.*, 2021c] Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. Kepler: A unified model for knowledge embedding and pre-trained language representation. *TACL*, 2021.
- [Wu *et al.*, 2020] Zhuofeng Wu, Sinong Wang, Jiatao Gu, Madian Khabsa, Fei Sun, and Hao Ma. Clear: Contrastive learning for sentence representation. *arXiv:2012.15466*, 2020.
- [Xu *et al.*, 2021] Yichong Xu, Chenguang Zhu, Ruochen Xu, Yang Liu, Michael Zeng, and Xuedong Huang. Fusing context into knowledge graph for commonsense question answering. In *Findings of ACL*, pages 1201–1207, 2021.
- [Yang *et al.*, 2020] Yiben Yang, Chaitanya Malaviya, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Jiping Wang, Chandra Bhagavatula, Yejin Choi, and Doug Downey. Generative data augmentation for commonsense reasoning. In *Findings of EMNLP*, pages 1008–1025, 2020.
- [Ye *et al.*, 2019] Zhi-Xiu Ye, Qian Chen, Wen Wang, and Zhen-Hua Ling. Align, mask and select: A simple method for incorporating commonsense knowledge into language representation models. *arXiv:1908.06725*, 2019.
- [Yu *et al.*, 2021] Wenhao Yu, Chenguang Zhu, Yuwei Fang, Donghan Yu, Shuohang Wang, Yichong Xu, Michael Zeng, and Meng Jiang. Dict-bert: Enhancing language model pre-training with dictionary. *arXiv:2110.06490*, 2021.
- [Zhang *et al.*, 2017] Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. Position-aware attention and supervised data improve slot filling. In *EMNLP*, pages 35–45, 2017.
- [Zhang *et al.*, 2019] Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. ERNIE: Enhanced language representation with informative entities. In *ACL*, pages 1441–1451, 2019.
- [Zhao *et al.*, 2020] Chen Zhao, Chenyan Xiong, Corby Rosset, Xia Song, Paul Bennett, and Saurabh Tiwary. Transformer-xh: Multi-evidence reasoning with extra hop attention. In *ICLR*, 2020.