

# A Survey of Knowledge-Enhanced Pre-trained Language Models

Linmei Hu, Zeyi Liu, Ziwang Zhao, Lei Hou, Liqiang Nie, *Senior Member, IEEE* and Juanzi Li

**Abstract**—Pre-trained Language Models (PLMs) which are trained on large text corpus via self-supervised learning method, have yielded promising performance on various tasks in Natural Language Processing (NLP). However, though PLMs with huge parameters can effectively possess rich knowledge learned from massive training text and benefit downstream tasks at the fine-tuning stage, they still have some limitations such as poor reasoning ability due to the lack of external knowledge. Research has been dedicated to incorporating knowledge into PLMs to tackle these issues. In this paper, we present a comprehensive review of Knowledge-Enhanced Pre-trained Language Models (KE-PLMs) to provide a clear insight into this thriving field. We introduce appropriate taxonomies respectively for Natural Language Understanding (NLU) and Natural Language Generation (NLG) to highlight these two main tasks of NLP. For NLU, we divide the types of knowledge into four categories: linguistic knowledge, text knowledge, knowledge graph (KG), and rule knowledge. The KE-PLMs for NLG are categorized into KG-based and retrieval-based methods. Finally, we point out some promising future directions of KE-PLMs.

**Index Terms**—Pre-trained language models, natural language processing, knowledge-enhanced pre-trained language models, natural language understanding, natural language generation.

## 1 INTRODUCTION

WITH the continuous development of deep learning technologies in recent years, Pre-trained Language Models (PLMs) which are trained with unsupervised objectives on massive text corpora, have been widely used in the field of Natural Language Processing (NLP), and yielded state-of-the-art performance on various downstream tasks. Different from traditional supervised learning, PLMs based on self-supervised learning are usually pre-trained on general-purpose large-scale unlabeled data first and then fine-tuned on small-scale labeled data for the specific tasks. Representative work, such as BERT [1], GPT [2], T5 [3], has refreshed benchmark records constantly in many Natural Language Understanding (NLU) and Natural Language Generation (NLG) tasks, successfully promoting the development of NLP.

As the size of PLMs grows larger, PLMs with hundreds of millions of parameters have been extensively demon-

strated to possess the ability of capturing rich linguistic [4], [5], [6] and factual knowledge [7], [8] in certain probings. However, due to the lack of explicit representation of knowledge in the raw data, PLMs suffer from limited performance on downstream tasks. In particular, the prior study has found that traditional pre-training objectives often have weak symbolic reasoning capabilities [9] since PLMs tend to concentrate on the word co-occurrence information. Incorporating knowledge into PLMs can empower their memorization and reasoning [10]. For instance, in the language understanding problem of “The monument to the people’s Heroes sits solemnly on [MASK] square”, traditional PLMs predict the output of the masked position as “the”, while knowledge-enhanced PLMs predict “Tiananmen”, which is more accurate.

For language generation, although existing PLMs are able to obtain rich language information from text corpus and generate correct sentences, almost all of them fail to generate output towards capturing the human commonsense, since they overlook the external world knowledge [11]. In other words, sentences generated by PLMs often conform to the grammatical norm, but not to logic. For example, given a concept set {hand, sink, wash, soap} to generate a sentence, conventional PLMs may generate “hands washing soap on the sink”, while the PLM with extra knowledge generates “man is washing his hands with soap in a sink”, which is more natural and logical.

To address the above issues, explicitly incorporating knowledge into PLMs has been an emerging trend in recent NLP studies. Wei et al. [12] reviewed the knowledge-enhanced PLMs along three taxonomies: types of knowledge sources, knowledge granularity, and application. Yin et al. [13] summarized the recent progress of pre-trained language model-based knowledge-enhanced models (PLMKEs) according to three crucial elements of them:

- L. Hu is with the School of Computer Science and Technology, Beijing Institute of Technology, Beijing 100081, China. (e-mail: hulinmei@bit.edu.cn).
- Z. Liu and Z. Zhao are with the School of Computer Science, Beijing University of Posts and Telecommunications, Beijing 100876, China. (e-mail: liu\_zy@bupt.edu.cn, zhaoziwang@bupt.edu.cn).
- L. Hou is with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China. (e-mail: houlei@tsinghua.edu.cn).
- L. Nie is with the School of Computer Science and Technology, Harbin Institute of Technology (Shenzhen) (e-mail: nieliqiang@gmail.com).
- J. Li is with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China. (e-mail: lijuanzi@tsinghua.edu.cn).

knowledge sources, knowledge-intensive NLP tasks, and knowledge fusion methods. In this work, considering the fact that injecting knowledge into language models can promote both NLU and NLG tasks, and these two areas have different focuses, we aim to present a comprehensive review of Knowledge-Enhanced Pre-trained Language Models (KE-PLMs) in the two areas to provide respective insights of KE-PLMs in NLU and NLG. The main contributions of this survey can be summarized as follows:

- (1) In this survey, we divide KE-PLMs into two main categories according to the downstream tasks: NLU and NLG. Appropriate taxonomies are respectively presented to highlight the focuses of these two different kinds of tasks in NLP.
- (2) For NLU, KE-PLMs are further divided into four sub-categories according to the types of knowledge: linguistic knowledge, text knowledge, knowledge graph (KG), and rule knowledge. For NLG, focused on the knowledge sources, KE-PLMs are further categorized into retrieval-based methods and KG-based methods. Fig. 1 shows our proposed taxonomies for NLU and NLG.
- (3) We discuss some possible directions that may tackle the existing problems and challenges of KE-PLMs in the future.

The rest of this paper is arranged as follows. In section 2, we provide the background of PLMs under the development of training paradigms in NLP. In Section 3, we introduce the taxonomy of KE-PLMs in the field of NLU. In Section 4, we introduce the taxonomy of KE-PLMs in the field of NLG. For both NLU and NLG fields, we discuss the representative work of each leaf category in the taxonomy. In Section 5, we propose the possible research directions of KE-PLMs in the future based on the existing limitations and challenges. Finally, we conclude in Section 6.

## 2 BACKGROUND

**Pre-trained Language Models.** Although the idea of pre-training on a language modeling task is not novel, the training paradigm has shifted to *pre-train and fine-tune*, with the emergence of ELMo [14] and ULMFiT [15]. Both are based on Long Short-Term Memory (LSTM) architecture. They propose to fine-tune the language model layer by layer for downstream tasks, and their performance demonstrates the competitiveness of pre-trained language models.

Unlike the early fully supervised method that learns salient features from limited data [16], language models can be trained on a large amount of raw textual data to obtain general-purpose representations. Then, the pre-trained models will be applied to various downstream tasks by fine-tuning them through task-specific objective functions [17]. For example, UNILM [18] unifies three language modeling objectives, which can be adapted to NLU and NLG tasks simultaneously.

As the Transformer architecture with multi-head self-attention mechanism is put forward [19], all popular language models, including GPT [2], BERT [1], BART [20] and T5 [3] are proposed based on the Transformer. The multi-head self-attention mechanism allows every word to

attend to each other, making the models capture long-range dependencies and learn more expressive representations. These models differ in the model structure and training objectives. In particular, GPT is an autoregressive language model (unidirectional) which predicts the next word given all the previous words. BERT is a masked language model (bidirectional) that aims to predict the “masked” word conditioned on all the other words. BART and T5 are encoder-decoder language models that learn to generate a sequence when given an input sequence.

**Prompt Learning.** Instead of adapting PLMs to different downstream tasks by designing specific objective functions, “*pre-train, prompt, and predict*” which reformulates downstream tasks through textual prompts has taken the place of “*pre-train, fine-tune*” to become the fourth paradigm in NLP [17]. In this paradigm, there is even no need for fine-tuning models using task-specific training objectives, PLMs themselves can directly be employed to predict the output that is desired, breaking through the problem of data constraints and bridging the gap of objective forms between pre-training and fine-tuning [7], [21], [22], [23], [24], [25], [26]. Previous work [17] conducts a comprehensive survey on the emerging field of prompt-based learning.

Though this prompt learning method has achieved promising results via constructing prompt information without changing the structure and parameters of PLMs significantly, it also calls for the necessity of choosing the most appropriate prompt template and verbalizer [22] which may have a great impact on model’s performance [22]. To this end, some work [27], [28], [29], [30], [31] has proposed to use knowledge as prompt to enhance the prompt-tuning process and reduce the cost of template construction and label mapping. Through injecting domain/task-relevant knowledge at the fine-tuning stage, PLMs can better serve downstream tasks and obtain better performance. In this survey, we also investigated the knowledge-enhanced pre-trained language models that incorporate external knowledge via prompt learning.

## 3 KE-PLMS FOR NLU

NLU is a subpart of NLP concerning all the methods which enable machines to understand and interpret the content of textual data. It extracts core semantic information from the unstructured text and applies this information to downstream tasks, thus playing a vital role in applications such as text classification, relation extraction, named entity recognition (NER), and dialogue system. In line with the taxonomy shown in Fig. 1, we divide the knowledge incorporated by KE-PLMs which are designed for NLU tasks into the following four categories according to their different types, that is, linguistic knowledge, text knowledge, knowledge graph, and rule knowledge. For each category, we discuss its representative methods.

### 3.1 Incorporating Linguistic Knowledge into PLMs

Linguistic knowledge, mainly divided into lexical knowledge and syntax tree, is the most common auxiliary feature incorporated into PLMs [32]. Among them, lexical knowledge includes but is not limited to Part-of-Speech

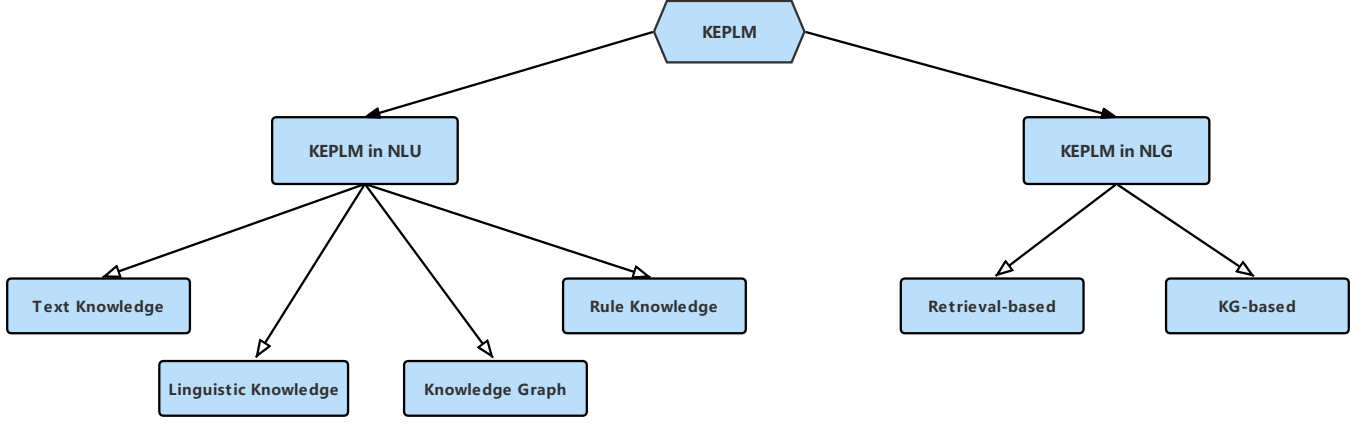


Fig. 1. Taxonomy of Knowledge Enhanced Pre-trained Language Models (KE-PLMs) based on the two core tasks of NLP: Natural Language Understanding (NLU) and Natural Language Generation (NLG).

(POS) tagging, and sentiment tags of words. LIBERT [33] introduces lexical relation classification (LRC) as a new pre-training task on the basis of standard BERT objectives, using synonyms and hypernym-hyponym pairs to predict whether two words are in specific semantic relations, which enhances the modeling ability of PLMs for semantic information. SenseBERT [34] integrates word-supersense (e.g., noun.food, noun.state) and predicts their corresponding supersenses by restoring the masked words, which can explicitly learn the semantic information of words in a given context. SKEP [35] improves the effect of PLM on the sentiment analysis task by integrating sentiment knowledge (sentiment words, polarity, and aspect-sentiment pairs). Sentiprompt [28] incorporates sentiment knowledge about aspects, opinions, and polarities into prompt through the construction of consistency and polarity judgment templates, and explicitly models term relations between aspect and opinion terms, better introducing task-related knowledge for the language models through prompt-tuning methods. LET [36] integrates semantic information of HowNet to improve the Chinese sentence matching task. KEAR [37] combines the knowledge from ConceptNet, dictionary entry definition, and labeled training data to enhance its performance on commonsense knowledge question answering. DictBERT [38] takes dictionary knowledge as the external source, and achieves knowledge enhancement for pre-training tasks by means of contrastive learning.

Considering the different ways of incorporating syntax tree knowledge into PLMs, we divided the KE-PLMs into three categories as illustrated in Fig. 2, including introducing new relevant pre-training tasks [39], adopting new attention mechanism [40], [41], and designing new model structure [42]. For instance, LIMIT-BERT [39] realizes multi-task learning across five linguistic tasks, and simply sums up these task-specific losses together on the basis of a variety of linguistic knowledge such as POS tagging, semantic role labeling (SRL), dependency relations, syntax trees, and etc. in model training. Syntax-BERT [40] provides additional syntax information by constructing a syntax tree parser, and generates a new attention mechanism according to the constructed syntax tree. Syntax-augmented BERT [42] introduces a syntax-based graph neural network to fuse the

syntax information from dependency trees to improve PLM.

Despite the above methods consistently incorporating linguistic knowledge into PLMs, they differ in the stages of knowledge fusion. LIBERT [33], SenseBERT [34], SKEP [35], Dictbert [38], LIMIT-BERT [39], and Syntax-BERT [40] fuse linguistic knowledge in the pre-training stage of PLMs to enhance the representation of the input, while Sentiprompt [28], LET [36], and KEAR [37] fuse knowledge in the fine-tuning stage of PLMs for improving task performance.

### 3.2 Incorporating Text Knowledge into PLMs

Text knowledge is usually retrieved from general-domain text collection (such as WikiText [43], Wiktionary [37]) or large corpus (such as Wikipedia [44]). KNN-LM [43] selects the nearest K neighbors from training samples as knowledge incorporated into PLM, and its earlier idea comes from cache-LM [45] which remains the first K words in the cache. REALM [44] utilizes text corpus to train a text retriever explicitly, exploiting information retrieved from external knowledge bases such as Wikipedia documents to help the prediction of tokens that are masked. ExpBERT [46] and KEAR [37] also incorporate textual descriptions into their models to improve performance. OK-Transformer [47] incorporates large-scale out-of-domain commonsense descriptions to enhance the representation of input text. Kformer [48] obtains some external text knowledge through retrieval, and injects this knowledge into the FFN layer of Transformer. REINA [49] retrieves some training samples similar to the input from external datasets as knowledge to enhance PLM. UniK-QA [50] and UDT-QA [51] use text, knowledge graph and table knowledge together, and transform all the knowledge into text for knowledge enhancement.

In addition to the type of general text knowledge that has been mentioned above, some work also utilizes domain-specific corpora or scientific and technical literature for pre-training tasks. In particular, BioBERT [52] and SciBERT [53] conduct the pre-training process on large-scale scientific domain corpora and achieve promising results on downstream academic NLP tasks. S2ORC-BERT [54] leverages the same method as SciBERT on a larger corpus which contains numerous academic papers covering dozens of

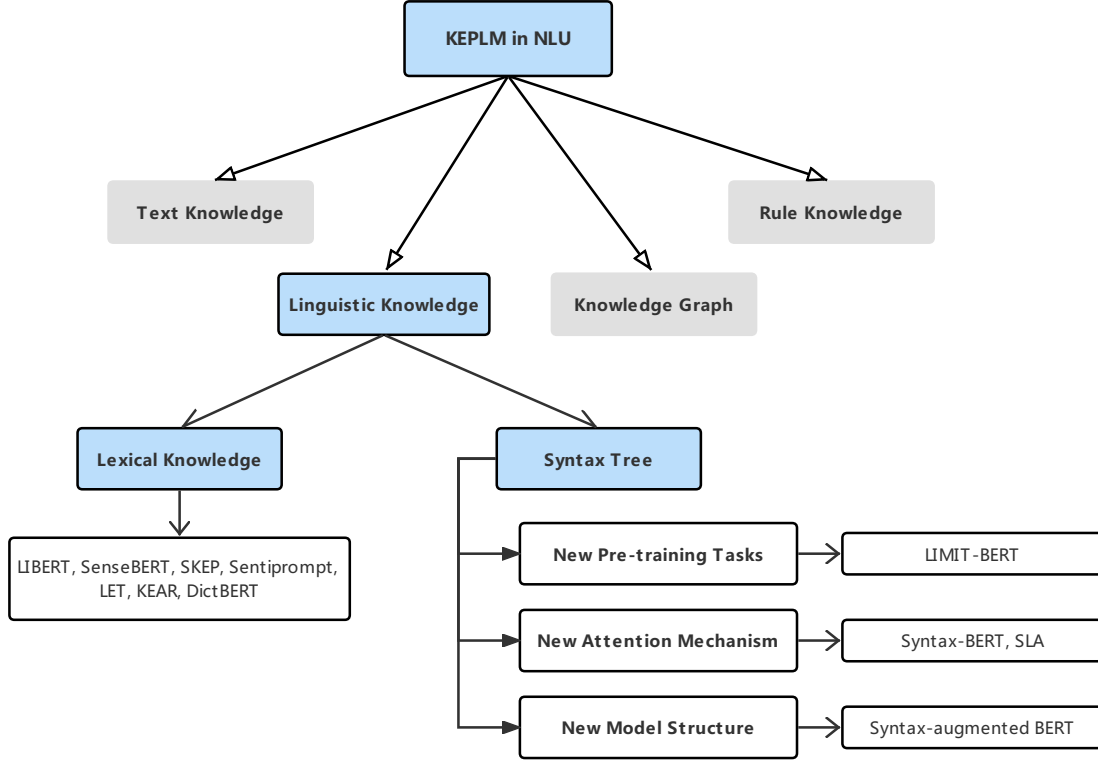


Fig. 2. Categorization of Linguistic knowledge.

academic disciplines, and slightly promotes its performance on several downstream tasks.

As for the stage of fusing knowledge, REALM [44], BioBERT [52], SciBERT [53] and S2ORC-BERT [54] integrate text knowledge in the pre-training stage, while KNN-LM [43], ExpBERT [46], KEAR [37], OK-Transformer [47], Kformer [48], REINA [49], UniK-QA [50], and UDT-QA [51] integrate knowledge in the fine-tuning stage.

### 3.3 Incorporating Knowledge Graph into PLMs

Knowledge graph can be considered as a kind of powerful expression that represents real world knowledge in the structural form of graph, where its nodes represent entities, and edges represent relations between entities [55], [56]. Compared with other types of knowledge, such as text knowledge which is unstructured, knowledge graph often contains more abundant structured information which makes it more applicable to enhance the potential learning capability of models [57], [58], [59], [60]. Here, we further divide KG into entity knowledge and triplet knowledge as shown in Fig. 3.

#### 3.3.1 Entity Knowledge

As the most basic semantic unit, entity plays a vital role in many natural language scenarios, such as machine reading comprehension, NER, sentiment analysis, and etc [61]. Incorporating entity knowledge into PLMs helps to improve the semantic understanding ability of models and their performance in downstream tasks as well. According to the different ways of incorporating entity knowledge, we divide them into three sub-categories as follows.

One is to design entity related pre-training tasks [62], [63], [64], [65], [66]. ERNIE [62] proposes a multi-stage knowledge masking strategy of both entity level and phrase level. WKLM [63] integrates knowledge of external entities through the entity replacement strategy, and determines whether each entity is replaced by performing binary classification prediction. KECP [64] adopts contrastive learning and prompt learning to integrate entity knowledge.

The second is to change the attention mechanism of the model. For example, LUKE [67] introduces an entity-aware self-attention mechanism to capture the entity information in calculating attention scores.

The third is to change the model structure [68], [69], [70], [82]. ERNIE-THU [68] leverages an N-layer T-encoder structure which is similar to the BERT-base model to extract text information, and then uses a proposed K-encoder to integrate entity knowledge. KnowBERT [69] retrieves entity embeddings related to the input text via an entity linker and then updates contextual word representations through word-to-entity attention, allowing the long-range interactions between words and all entities. EaE [70] expands each entity mention embedding in the input text by fusing it with the top K relevant entity embeddings retrieved by a proposed component called an entity memory layer.

Note that several methods may belong to more than one sub-category, and we assign them to the sub-category according to their most significant contributions.

#### 3.3.2 Triplet Knowledge

In addition to the entity knowledge mentioned above, there are also a great number of relational triples in the knowledge graph, which can provide sufficient structured information



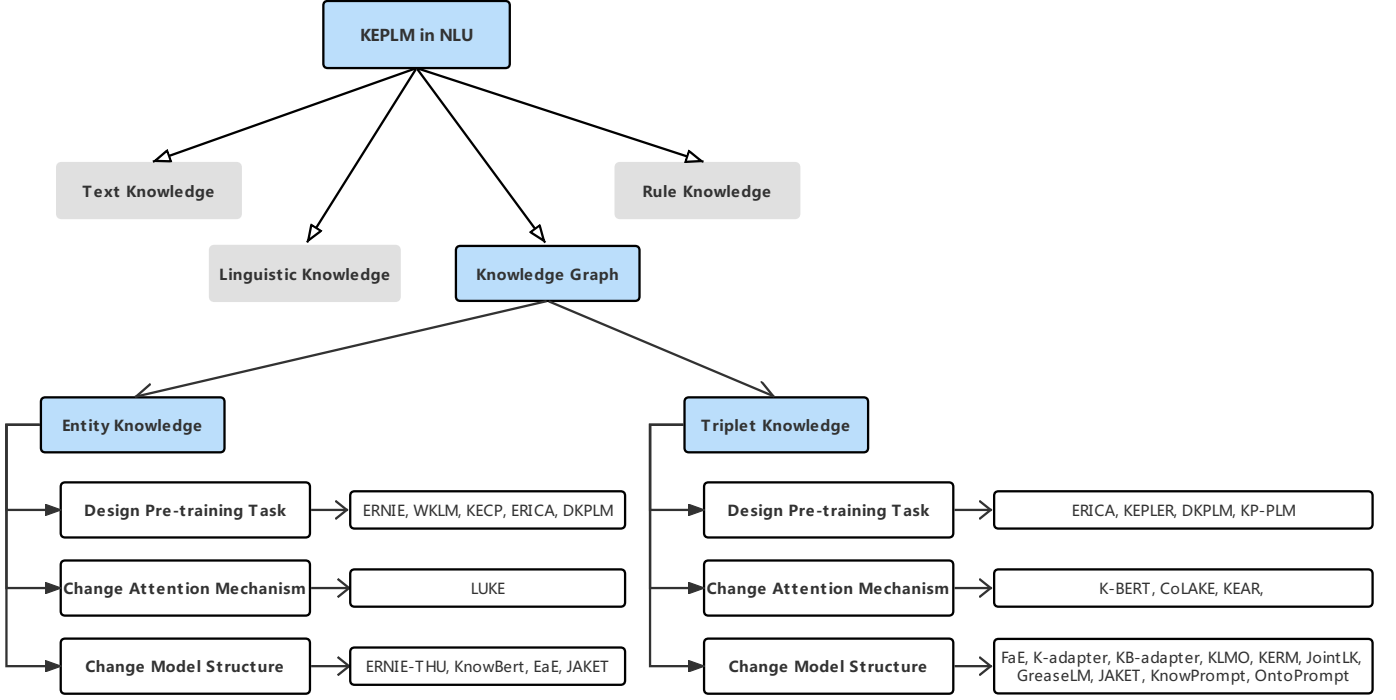


Fig. 3. Categorization of knowledge graph and injection methods of each sub-category.

TABLE 1  
Entity/Triplet related pre-training tasks. Here MLM represents masked language modeling.

Method	Pre-training Tasks/Objectives
ERNIE [62]	Token-level, Phrase-level and Entity-level MLM
WKLM [63]	Entity replacement, MLM
KECP [64]	Token-level MLM, Span-level contrastive learning
ERICA [65]	Entity and relation discrimination tasks, MLM
DKPLM [66]	Relational knowledge decoding, Token-level MLM
KP-PLM [72]	Prompt relevance inspection, Masked prompt modeling
KEPLER [71]	Knowledge embedding, MLM

for PLMs and also improve the semantic understanding ability of them. Similar to the KE-PLMs incorporating entity knowledge, we also divide KE-PLMs in this category into three sub-categories according to their specific ways of knowledge incorporation.

One is to design pre-training tasks related to triplets [65], [71], [66], [72]). ERICA [65] introduces both entity and relation discrimination tasks to deepen PLM’s understanding of entities and relations through contrastive learning. KEPLER [71] trains knowledge embedding and masked language modeling objectives jointly to improve the knowledge representation. DKPLM [66] focuses on long-tail entities, enriching semantic information of low-frequency entities with knowledge graphs. KP-PLM [72] designs two knowledge-aware pre-training tasks to incorporate knowledge triplets into multiple continuous prompts for NLU tasks. In Table 1, we summarize entity and triplet relevant pre-training tasks designed by the existing work.

The second is to change the attention mechanism of the model [73], [74], [37]). K-BERT [73] leverages a knowledge

layer to inject relevant triplets from KG into the input sentence and transform it into a knowledge-rich sentence tree. Then this sentence tree is converted into a visible matrix to control the visible area of each word in the sentence, preventing the sentence from deviating from the correct semantics due to injecting too much knowledge. If K-BERT intends to expand input text into a sentence tree, the core concept of CoLAKE [74] is to expand the input context into word-knowledge graphs (WK graphs), and then feed these constructed WK graphs into masked self-attention to gather information of nodes. KEAR [37] proposes an external attention mechanism to enhance the Transformer architecture, and integrate external knowledge into its prediction process. We illustrate their attention mechanism in Table 2, where  $Q$ ,  $K$ ,  $V$  denote the query, key, and value matrices, respectively.  $d_k$  is the dimension of the key, which is used as the scaling factor.

The third is to change the model structure, which usually introduces a knowledge fusion module [30], [31], [75], [76], [77], [78], [79], [80], [81], [82], as shown in Fig. 4(a). FaE [75] introduces an additional memory module of facts on the basis of EaE [70], so that it can effectively combine the information in the symbolic knowledge graph. Besides the layer structure of the original pre-trained model, K-adaptor [76] and KB-adaptors [77] incorporate knowledge into PLM through external adapter modules. KLMO [78] uses a component named knowledge aggregator to fuse the embeddings of the input text and KG, which applies an entity-level cross-KG attention to interactively model entity segments in text along with entities and relations in KG, as shown in 4(b). KERM [79] designs a knowledge injector module that combines the knowledge between text corpus and KG for passage re-ranking task as shown in Fig. 4(b). JointLK [80] and GreaseLM [81] exploit GNNs for model-

TABLE 2  
Examples of changing attention mechanism for incorporating triplet knowledge.

Method	Attention Mechanism	Formalized expression	Variation in Calculation
K-BERT [73]	Masked Self-Attention	$Attn(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}(\frac{\mathbf{Q}\mathbf{K}^\top + \mathbf{M}}{\sqrt{d_k}})\mathbf{V}$	visible matrix $\mathbf{M}$ for the input sentence tree, $M_{ij} = 0$ if token $i$ and token $j$ are in the same branch, while $M_{ij} = -\infty$ if not
CoLAKE [74]	Masked Self-Attention	$Attn(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}} + \mathbf{M})\mathbf{V}$	mask matrix $\mathbf{M}$ for word-knowledge graph, $M_{ij} = 0$ if node $i$ and node $j$ are connected, while $M_{ij} = -\infty$ if not
KEAR [37]	External Attention	$Attn(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}})\mathbf{V}$	concatenate extra knowledge $\mathbf{K}$ into the input text $\mathbf{H}_0 = [\mathbf{X}; \mathbf{K}]$

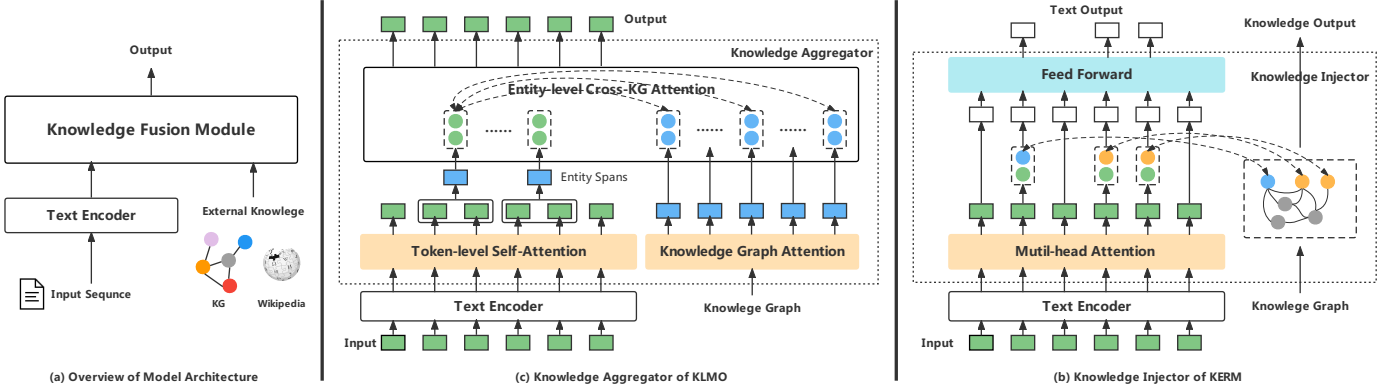


Fig. 4. (a) Incorporating triplet knowledge through a knowledge fusion module. (b) KLMO designs a knowledge aggregator to fuse knowledge into the input token sequence. (c) KERM develops a knowledge injector to integrate knowledge explicitly.

ing extracted knowledge graphs and couple the LM with GNN modules to perform joint reasoning for commonsense reasoning. Besides, KnowPrompt [30] incorporates entity and relation knowledge contained in the relation labels into prompt templates, and inserts these templates into the input text for relation extraction. OntoPrompt [31] linearizes ontology knowledge extracted from external knowledge graph into texts as auxiliary prompts, and introduces a visible matrix to guide the knowledge injection process, avoiding injecting irrelevant or noisy knowledge.

In addition to using knowledge graph as auxiliary information to improve PLMs’ ability of language understanding, some work also learns a knowledge embedding representation while incorporating triplet knowledge, so that PLMs can complete some tasks related to knowledge reasoning, such as entity classification, relation prediction, knowledge graph completion and etc. Representative work includes ERICA [65], KEPLER [71], KLMO [78], FaE [75] mentioned above.

### 3.3.3 Fusion stage

The above mentioned methods ERNIE [62], ERNIE-THU [68], WKLM [63], LUKE [67], EaE [70], FaE [75], ERICA [65], CoLAKE [74], KEPLER [71], KLMO [78], DKPLM [66], KERM [79], JAKET [82] and KP-PLM [72] are pre-fusion methods that fuse knowledge in the pre-training stage, while KnowBert [69], K-BERT [73], K-adapter [76], [37], KECP [64], KB-adapters [77], JointLK [80], GreaseLM [81] are post-fusion methods which fuse knowledge in the fine-tuning stage. The two fusion stages are also called the training stage and reasoning stage, respectively. For exam-

ple, CoLAKE [74] jointly learns the embeddings of entities and relations during the training phase, while K-BERT [73] injects triples from KG during the reasoning phase.

## 3.4 Incorporating Rule Knowledge into PLMs

Logic rules always contain clear logical reasoning processes, and can formalize knowledge from external sources [83]. Incorporating this type of knowledge into PLMs can facilitate the demonstration of reasoning path via its good interpretability. For example, RuleBERT [84] utilizes the Horn rules of existing corpus to establish a training dataset and then fine-tunes the model on it. It adopts a probabilistic answer set programming to predict the probability of events, and tries to learn soft rules from PLM. The results show PLMs that reason with soft rules over natural language could improve their performance for deductive reasoning tasks. Besides, PTR [27] incorporates logic rules to construct task-specific prompts composed with sub-prompts designed manually, so that the model can encode task-related prior knowledge in the prompt-tuning and generate prompts that are more interpretable. Both RuleBERT and PTR incorporate rule knowledge in the fine-tuning stage.

## 4 KE-PLMs FOR NLG

The goal of NLG is to enable machines to generate language texts that can be understood by humans and follow the way where humans express themselves. Incorporating various forms of knowledge into generation models other than input sequences helps to improve the performance of text

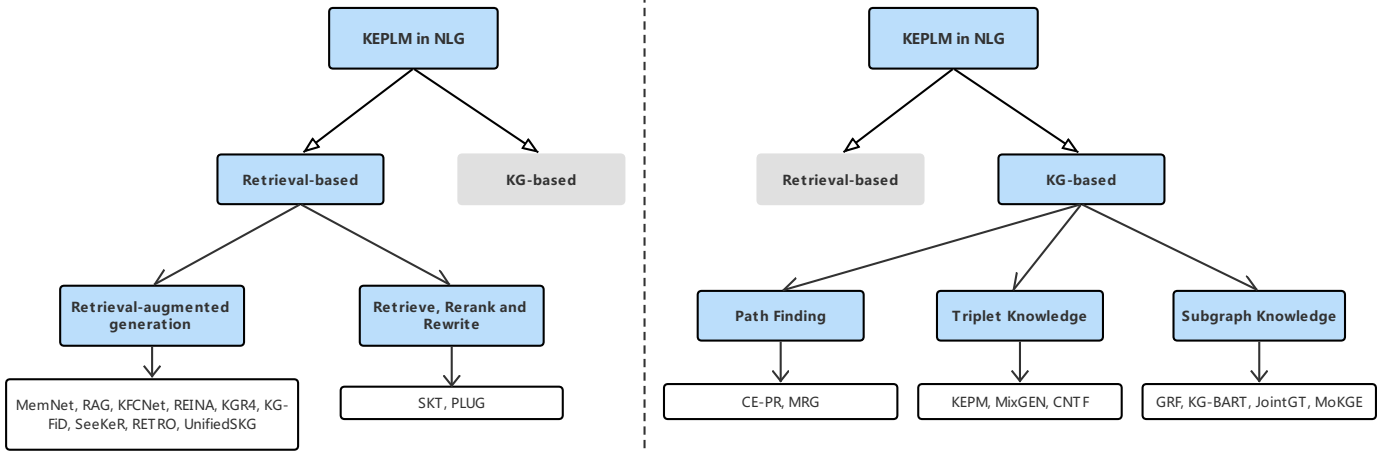


Fig. 5. Further categorization of retrieval-based method and KG-based method. The left figure demonstrates the categorization of retrieval-based method, and the right demonstrates the categorization of KG-based method.

generation tasks. Referring to the survey on knowledge-enhanced text generation [32], we further divide KE-PLMs in the field of NLG into two categories based on their different knowledge sources: one is retrieval-based method and the other is KG-based method.

#### 4.1 Incorporating Retrieval-based Knowledge into PLMs

The retrieval-based methods mainly integrate and utilize additional knowledge related to the input sequences through retrieval. Other than the input sequence itself, the additional knowledge is retrieved from external sources such as online search engines, large data sets, and training sets to guide the generation process. Considering whether the methods focus on re-ranking the retrieved items for generation or not, we divide these retrieval-based methods into two sub-methods as shown in the left of Fig. 5: one is retrieval augmented generation method that aims to improve generation by retrieving related knowledge [49], [85], [86], [87], [88], [89], [90], [91], [92], and the other is retrieve, rerank and rewrite method which focuses on re-ranking retrieved items for generation [93], [94]. The flow charts of these two kinds of methods are shown in Fig. 6. Notably, these KE-PLMs incorporate external knowledge in the fine-tuning stage to improve their performance on downstream tasks.

In the line of retrieval augmented generation method, MemNet [85] proposes a Transformer memory network that can retrieve topic-relevant knowledge based on the dialogue history, and then generates the next dialogue utterance with the help of this retrieved knowledge. RAG [86] leverages a retriever to find the top-K relevant documents given the input sequence, and uses them as additional context when generating predictions. KFCNet [87] first retrieves prototypes that comprise concepts in the given concept set, while keeping these retrieved results semantically similar to the target sentences. Then it applies two contrastive learning modules in both encoder and decoder to capture global target information and learn general features from the multiple retrieved prototypes. REINA [49] retrieves some training samples which are similar to the input text and

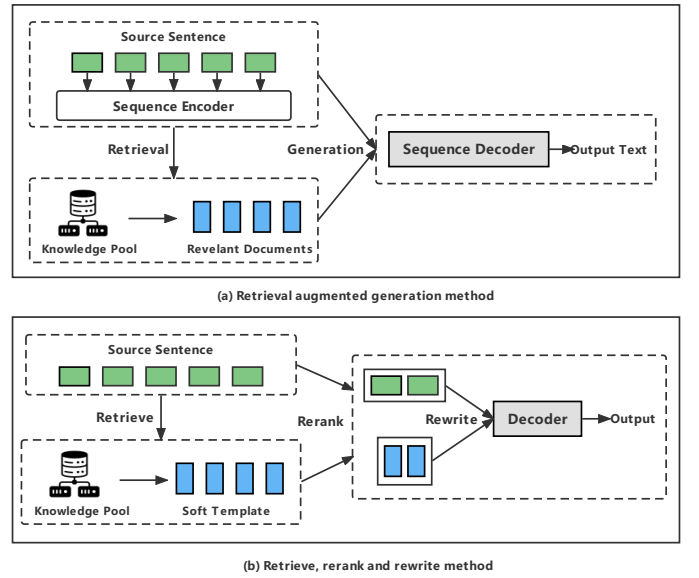


Fig. 6. Flow charts of retrieval-based methods. (a) demonstrates the retrieval augmented generation method. (b) demonstrates the retrieve, rerank and rewrite method. Notably, the retrieved candidate templates will be reranked before rewriting.

takes them as knowledge to improve the effect of machine translation. KGR4 [88] divides the generation of commonsense into four stages, that is, retrieval of commonsense, generation of commonsense by means of generation models, refinement and correction of the generated commonsense statements, and scoring of the generated statements. Best results can be obtained through this process. KG-FiD [89] is based on the Fusion-in-Decoder (FiD) [95] model, but it proposes to solve the validity and efficiency of FiD with the help of KG, which improves the ability of open domain question answering significantly. SeeKeR [90] leverages a set of documents retrieved from the search engine to generate knowledge response. This method can incorporate up-to-date information through its three-module framework (search, knowledge generation, and final response). RETRO [91] retrieves from large text databases and builds trillions

of tokens as retrieve sources to expand the language model. UnifiedSKG [92] unifies six categories of tasks (i.e., semantic parsing, question answering, data-to-text generation, fact verification, conversation, and formal language-to-text translation tasks) into a text-to-text format and introduces linearized knowledge to strengthen their performance.

Representative work of the retrieve, rerank and rewrite method, such as SKT [93], regards knowledge selection as a sequential decision-making process, and uses the sequential latent variable model to improve the accuracy of knowledge selection in multiple rounds of dialogue. PLUG [94] retrieves related knowledge from Wikipedia, dictionary, and knowledge graph, and then ranks them based on statistical and semantic information for knowledge-grounded dialogue generation.

#### 4.2 Incorporating KG-based Knowledge into PLMs

In order to distinguish the granularity of knowledge utilized by different KE-PLMs more specifically, we divide existing work into three categories: knowledge extracted by path finding, triplet knowledge, and subgraph knowledge extracted from KG as shown in the right of Fig. 5.

The first is to extract knowledge by path finding [96], [97]. In this way, the relation path is clearly reasoned to make reliable decisions. CE-PR [96] and MRG [97] mainly perform explicit reasoning on relation paths, significantly improving the effectiveness of text generation. Specifically, CE-PR first retrieves a subgraph given the source concepts, and then scores each triple on this graph, propagates scores along the paths to each node from the source concepts and preserves the nodes with higher scores. MRG leverages the reasoning module to infer paths step by step from the source concepts, and then uses the sentence realization module to generate a complete sentence based on the inferred paths. Fig. 7 demonstrates the inference process of these two methods.

The second sub-category is KE-PLMs based on triplet knowledge [98], [99], [100]. KEPM [98] converts commonsense triplets in the knowledge bases into natural language statements based on templates to provide additional information for story generation. MixGEN [99] adopts an encoder-decoder framework to incorporate expert knowledge (from dataset annotation), explicit knowledge (from knowledge graph), and implicit knowledge (from generative PLM) to produce implications of toxic text. CNTF [100] models the commonsense, named entities and topic-specific knowledge via a multi-hop attention module to facilitate dialogue generation task.

The third one is subgraph knowledge [11], [101], [102], [103]. Different from the first two categories, subgraph contains context information on related concepts, which plays an important role in understanding related concepts and language generation. This kind of methods generally use GNNs to model extracted subgraphs, and then fuse these extracted subgraphs to enhance the natural language generation ability.

For the integration of subgraph knowledge, we further divide them based on the specific position they integrate as shown in Fig. 8. Some work incorporates subgraph knowledge into only encoder to improve the language understanding ability, such as JointGT [102] and MoKGE [103].

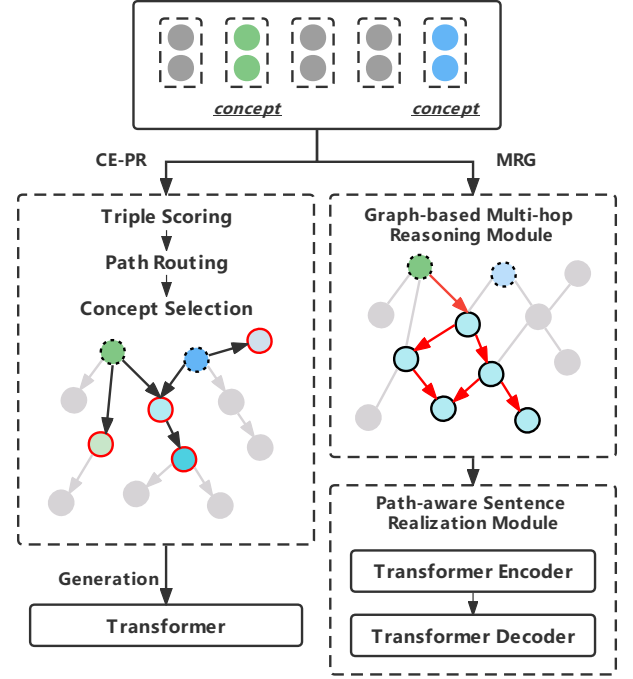


Fig. 7. Comparison of CE-PR and MRG incorporating knowledge extracted by path finding.

Some inject subgraph knowledge into the decoder, such as GRF [101]. In this way, every step of decoding can be traced, so that the model can be better interpreted. Others introduce subgraph knowledge into both encoder and decoder, such as KG-BART [11].

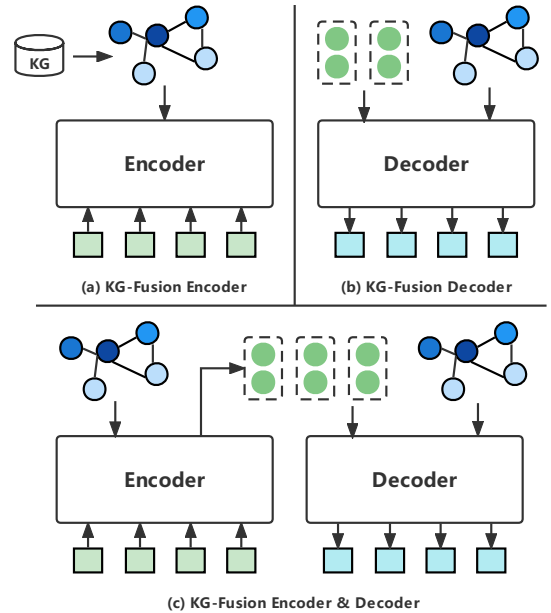


Fig. 8. Integration position of subgraph knowledge. (a) incorporating KG into the encoder. (b) incorporating KG into the decoder. (c) incorporating KG into both the encoder and decoder.

Among the above mentioned KE-PLMs for NLG, JointGT [102], PLUG [94], UnifiedSKG [92] are pre-fusion methods (fuse knowledge during pre-training), and CE-PR [96], MRG [97], KEPM [98], GRF [101], KG-BART [11],

MoKGE [103], MixGEN [99], CNTF [100] are post-fusion methods (fuse knowledge during fine-tuning).

In Table 3 and Table 4, we respectively summarize the existing KE-PLMs for NLU and NLG.

## 5 FUTURE DIRECTIONS

In this section, we propose some possible research directions of KE-PLMs in the future, which may meet the existing problems and challenges.

### 5.1 Integrating Knowledge from Homogeneous and Heterogeneous Sources

Since most of the existing work only utilizes knowledge from a single source, such as knowledge graph or web resource, exploring how to integrate knowledge from heterogeneous sources is still a valuable direction for future research.

As we present in the section above, some prior work has tried to incorporate different types of knowledge to improve the performance of question-answering. For example, UniK-QA [50] integrates external knowledge including text, tables, and relational triplets in the knowledge base. Through the heuristic method of linearizing heterogeneous knowledge sources including knowledge base (KB) triples and semi-structured tables into text, it unifies structured knowledge involved in KBQA and unstructured knowledge involved in TextQA, expanding the sources of external knowledge. UDT-QA [51] introduces structured knowledge such as knowledge graphs and tables into open-domain question answering, and converts them into linear sequences as the input of text generation tasks.

In the field of open-domain question answering, improving the ability of PLMs to integrate multiple knowledge sources can effectively increase the knowledge coverage, so that models can generate more reliable answers.

### 5.2 Exploring Multi-modal Knowledge

Most of the current research focuses merely on text knowledge with fewer multi-modal sources. In fact, images, videos, and audio in addition to textual and tabulated information can also become the knowledge sources of PLMs, which can further improve the performance of KE-PLMs.

Several studies have explored integrating multi-modal knowledge. Representative work includes KB-VLP [104], and ERNIE-VIL [105]. KB-VLP [104] extracts knowledge information from the external knowledge base based on both the input text and image, and uses the knowledge as additional inputs to enhance the model’s ability of semantic alignment and knowledge perception. ERNIE-VIL [105] parses the input description texts of images into structured scene graphs, and designs cross-modal pre-training tasks to pay attention to detailed semantic alignments across vision and language modalities.

Since images and associated text contain rich semantics, the injection of these different modalities of knowledge and concentration on detail semantics can make them complement and enhance each other, which will boost the performance of PLMs on both NLU and NLG tasks.

### 5.3 Providing Interpretability Evidence

Although many existing KE-PLMs have achieved great success on a series of text generation tasks, it should not be ignored that, if the generation process requires common-sense knowledge reasoning, the performance of models will be affected.

Some work has attempted to tackle this problem [44], [76], [86]. For example, GRF [101] utilizes external knowledge graphs for explicit commonsense reasoning, and incorporates rich structural information in order to perform dynamic multi-hop reasoning on multiple relational paths. Reasoning paths obtained in this process provide a theoretical basis for the generation of results. This work suggests that, giving an explicit reasoning path will help improve the interpretability of models and make predictions more rational.

### 5.4 Learning Knowledge in a Continuous Way

Existing work is usually trained on a large number of static or non-updated data in the pre-training stage. But models may forget the original knowledge learned before when facing new tasks, which leaves them vulnerable to a phenomenon called catastrophic forgetting problem [106]. With the continuous growth of knowledge from heterogeneous sources, exploring methods to make models master new knowledge while not forgetting the previous one learned in the past requires continual learning (also called life-long learning) to integrate various knowledge constantly.

ELLE [107] proposes an extension module that maintains the network function to expand the width and depth of the model, so that the model can effectively acquire new knowledge and retain the old to a greater extent at the same time. K-adaptor [76] and KB-adapters [77] adds the adapters with PLMs to store factual and linguistic knowledge, so as to continuously incorporate more knowledge into PLM.

Incorporating knowledge continuously is a promising direction in future research [32]. The application of continuous and increasing pre-training will effectively improve the universality of PLMs, and solve the catastrophic forgetting problem while incorporating more knowledge.

### 5.5 Optimizing the Efficiency of Incorporating Knowledge into Large Models

The scale of pre-trained models and injection of knowledge has become increasingly large in recent years [108], thus bringing severe challenges to the computational efficiency and computational resources that cannot be ignored. Though most of the existing work has achieved good results in various pre-training tasks, few studies mention the cost of knowledge fusion in the process.

In view of this challenge, we propose the following two possible directions that may be worth further exploring: one is to improve the efficiency of knowledge acquisition and filtering, and the other is to optimize the computational burden.

Existing work, such as ZeRO [109], has been explored in the second area. Based on traditional data parallel training mode, ZeRO deeply optimizes the redundant space and eliminates the memory occupied by redundancy through dividing the parameters, gradients, and optimizer states of the model into different processes.

TABLE 3

Summarization of different KEPLMs. Here RC: relation classification; ET: entity typing; QA: question answering; RE: relation extraction; WiC: words in context; NER: named entity recognition; NLI: natural language inference; GLUE: general language understanding evaluation.

Method	Knowledge Type	Fusion in Pre-training	Fusion in Fine-tuning	NLU or NLG tasks
LIBERT [33]	lexical	Yes		lexical simplification, sentence/sentence-pair classification, NLI
LIMIT-BERT [39]	syntax tree	Yes		syntactic parsing, semantic parsing
SenseBERT [34]	lexical	Yes		word supersense disambiguation, WiC
SKEP [35]	lexical	Yes		sentence/aspect-level sentiment classification, opinion role labeling
Sentiprompt [28]	lexical		Yes	triplet extraction, pair extraction aspect term extraction
LET [36]	lexical		Yes	Chinese short text matching
KEAR [37]	lexical, general text, triplet		Yes	commonsense reasoning
Syntax-BERT [40]	syntax tree	Yes		syntactic and semantic compositionality of sentiment classification, NLI, GLUE
DictBERT [38]	lexical	Yes		NER, RE, commonsense reasoning
KNN-LM [43]	general text		Yes	language modeling
REALM [44]	general text	Yes		open-QA
ExpBERT [46]	general text		Yes	RE
OK-Transformer [47]	general text		Yes	commonsense reasoning, text classification
Kformer [48]	general text		Yes	commonsense reasoning, medical QA
REINA [49]	general text, retrieval augmented generation		Yes	summarization, language modeling machine translation, QA
UniK-QA [50]	general text		Yes	multi-source QA
UDT-QA [51]	general text		Yes	open-domain QA
BioBERT [52]	domain-specific text	Yes		biomedical NER, RE, QA
SciBERT [53]	domain-specific text	Yes		sequence tagging, sentence classification, dependency parsing
S2ORC-BERT [54]	domain-specific text	Yes		inline citation detection, bibliography parsing, bibliography linking
ERNIE [62]	entity	Yes		NLI, semantic similarity, NER, sentiment analysis, QA
ERNIE-THU [68]	entity	Yes	Yes	ET, RC
WKLM [63]	entity	Yes		QA, ET
KnowBert [69]	entity	Yes	Yes	RE, WiC, ET
LUKE [67]	entity	Yes		ET, RC, NER, cloze-style/extractive QA
EaE [70]	entity	Yes		open-domain QA, RE
KECP [64]	entity		Yes	extractive QA
ERICA [65]	entity, triplet	Yes		RE, ET, QA
K-BERT [73]	triplet		Yes	sentence classification, QA, NER
CoLAKE [74]	triplet	Yes		ET, RE, knowledge probing, GLUE knowledge graph completion
FaE [75]	entity, triplet	Yes		open-domain QA
K-adapter [76]	general text, triplet		Yes	ET, QA, RC
KB-adapters [77]	entity, triplet		Yes	knowledge-probing using response selection fact memorization, response generation
KLMO [78]	entity, triplet	Yes		ET, RC
KEPLER [71]	entity, triplet	Yes		RC, ET, GLUE, link prediction
DKPLM [66]	entity, triplet	Yes		knowledge probing, RE, ET
JAKET [82]	entity, triplet	Yes		RC, QA over KG, entity classification
KP-PLM [72]	triplet	Yes		knowledge probing, RE, ET
KERM [79]	triplet	Yes		passage re-ranking
JointLK [80]	triplet		Yes	commonsense reasoning
GreaseLM [81]	triplet		Yes	commonsense reasoning
RuleBERT [84]	rule		Yes	rule reasoning
PTR [27]	rule		Yes	RC



TABLE 4  
Summarization of different KEPLMs. Here QA: question answering.

Method	Knowledge Source	Fusion in Pre-training	Fusion in Fine-tuning	NLU or NLG tasks
MemNet [85]	retrieved text		Yes	open-domain dialogue generation
RAG [86]	retrieved text		Yes	open-domain/abstractive QA, question generation, fact verification
SKT [93]	retrieved text		Yes	knowledge-grounded dialogue
KFCNet [87]	retrieved text		Yes	commonsense/keyword generation
KG-FiD [89]	retrieved text		Yes	open-domain QA
KGR4 [88]	retrieved text		Yes	commonsense generation
SeeKeR [90]	retrieved text		Yes	open-domain dialogue, prompt completion
RETRO [91]	retrieved text		Yes	language modelling, QA
UnifiedSKG [92]	structured knowledge	Yes		structured knowledge grounding
PLUG [94]	retrieved text, knowledge graph	Yes		dialogue generation
CE-PR [96]	knowledge graph		Yes	commonsense explanation generation
MRG [97]	knowledge graph		Yes	story/review/description generation
KEPM [98]	knowledge graph		Yes	story generation
MixGEN [99]	knowledge graph, expert knowledge		Yes	toxic text explanation
CNTF [100]	knowledge graph		Yes	dialogue generation
GRF [101]	knowledge graph		Yes	story ending generation, abductive NLG, explanation generation
KG-BART [11]	knowledge graph		Yes	commonsense generation/QA
JointGT [102]	knowledge graph	Yes		KG-to-text generation
MoKGE [103]	knowledge graph		Yes	commonsense explanation generation, abductive commonsense reasoning

## 5.6 Increasing the Variety of Results Generated

It is a vital research direction in NLG to generate alternative outputs or predict all possible results for the real situation, which is also the purpose of output diversity in the generative commonsense reasoning task. Existing work, such as MoKGE [103], uses the diversified knowledge reasoning of commonsense knowledge graph to complete the diversity generation of NLG. Based on the observation of human annotations, the concepts related to original input are associated into the generation process, and the mixture of expert method is used to generate diversified reasonable outputs, thus increasing the diversity of generated results.

## 6 CONCLUSION

In this survey, we present a comprehensive review of KE-PLMs from the perspective of NLU and NLG, and respectively propose proper taxonomies for both NLU and NLG to highlight their different focuses. We also discuss the representative work in the taxonomies. Finally, in view of the existing problems and challenges, we discuss potential future research directions of KE-PLMs, hoping to facilitate relevant research in this promising area.

## REFERENCES

- [1] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *NAACL-HLT*, 2019, pp. 4171–4186.
- [2] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever *et al.*, "Improving language understanding by generative pre-training," 2018.
- [3] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *Journal of Machine Learning Research*, vol. 21, pp. 140:1–140:67, 2020.
- [4] N. F. Liu, M. Gardner, Y. Belinkov, M. E. Peters, and N. A. Smith, "Linguistic knowledge and transferability of contextual representations," in *NAACL-HLT*, 2019, pp. 1073–1094.
- [5] Y. Belinkov and J. R. Glass, "Analysis methods in neural language processing: A survey," in *NAACL-HLT*, J. Burstein, C. Doran, and T. Solorio, Eds., 2019, pp. 3348–3354.
- [6] I. Vulic, E. M. Ponti, R. Litschko, G. Glavas, and A. Korhonen, "Probing pretrained language models for lexical semantics," in *EMNLP*, 2020, pp. 7222–7240.
- [7] F. Petroni, T. Rocktäschel, S. Riedel, P. S. H. Lewis, A. Bakhtin, Y. Wu, and A. H. Miller, "Language models as knowledge bases?" in *EMNLP-IJCNLP*, 2019, pp. 2463–2473.
- [8] S. Li, X. Li, L. Shang, Z. Dong, C. Sun, B. Liu, Z. Ji, X. Jiang, and Q. Liu, "How pre-trained language models capture factual knowledge? A causal-inspired analysis," in *ACL*, 2022, pp. 1720–1732.
- [9] A. Talmor, Y. Elazar, Y. Goldberg, and J. Berant, "olmpics - on what language model pre-training captures," *Transactions of the Association for Computational Linguistics*, vol. 8, pp. 743–758, 2020.
- [10] Y. Sun, S. Wang, S. Feng, S. Ding, C. Pang, J. Shang, J. Liu, X. Chen, Y. Zhao, Y. Lu *et al.*, "ERNIE 3.0: Large-scale knowledge

- enhanced pre-training for language understanding and generation," *arXiv preprint arXiv:2107.02137*, 2021.
- [11] Y. Liu, Y. Wan, L. He, H. Peng, and S. Y. Philip, "Kg-bart: Knowledge graph-augmented bart for generative commonsense reasoning," in *AAAI*, vol. 35, no. 7, 2021, pp. 6418–6425.
  - [12] X. Wei, S. Wang, D. Zhang, P. Bhatia, and A. Arnold, "Knowledge enhanced pretrained language models: A comprehensive survey," *arXiv preprint arXiv:2110.08455*, 2021.
  - [13] D. Yin, L. Dong, H. Cheng, X. Liu, K.-W. Chang, F. Wei, and J. Gao, "A survey of knowledge-intensive NLP with pre-trained language models," *arXiv preprint arXiv:2202.08772*, 2022.
  - [14] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations," in *NAACL-HLT*, 2018, pp. 2227–2237.
  - [15] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," in *ACL*, 2018, pp. 328–339.
  - [16] X. Liu, F. Zhang, Z. Hou, L. Mian, Z. Wang, J. Zhang, and J. Tang, "Self-supervised learning: Generative or contrastive," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
  - [17] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig, "Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing," *arXiv preprint arXiv:2107.13586*, 2021.
  - [18] L. Dong, N. Yang, W. Wang, F. Wei, X. Liu, Y. Wang, J. Gao, M. Zhou, and H. Hon, "Unified language model pre-training for natural language understanding and generation," in *NeurIPS*, 2019, pp. 13 042–13 054.
  - [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017, pp. 5998–6008.
  - [20] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," in *ACL*, 2020, pp. 7871–7880.
  - [21] T. Shin, Y. Razeghi, R. L. L. IV, E. Wallace, and S. Singh, "Autoprompt: Eliciting knowledge from language models with automatically generated prompts," in *EMNLP*, 2020, pp. 4222–4235.
  - [22] T. Schick and H. Schütze, "Exploiting cloze-questions for few-shot text classification and natural language inference," in *EACL*, 2021, pp. 255–269.
  - [23] X. L. Li and P. Liang, "Prefix-tuning: Optimizing continuous prompts for generation," in *ACL-IJCNLP*, 2021, pp. 4582–4597.
  - [24] Y. Gu, X. Han, Z. Liu, and M. Huang, "PPT: pre-trained prompt tuning for few-shot learning," in *ACL*, 2022, pp. 8410–8423.
  - [25] N. Zhang, L. Li, X. Chen, S. Deng, Z. Bi, C. Tan, F. Huang, and H. Chen, "Differentiable prompt makes pre-trained language models better few-shot learners," in *ICLR*, 2022.
  - [26] N. Ding, S. Hu, W. Zhao, Y. Chen, Z. Liu, H. Zheng, and M. Sun, "Openprompt: An open-source framework for prompt-learning," in *ACL*, 2022, pp. 105–113.
  - [27] X. Han, W. Zhao, N. Ding, Z. Liu, and M. Sun, "PTR: prompt tuning with rules for text classification," *arXiv preprint arXiv:2105.11259*, 2021.
  - [28] C. Li, F. Gao, J. Bu, L. Xu, X. Chen, Y. Gu, Z. Shao, Q. Zheng, N. Zhang, Y. Wang *et al.*, "Sentiprompt: Sentiment knowledge enhanced prompt-tuning for aspect-based analysis," *arXiv preprint arXiv:2109.08306*, 2021.
  - [29] S. Hu, N. Ding, H. Wang, Z. Liu, J. Wang, J. Li, W. Wu, and M. Sun, "Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification," in *ACL*, 2022, pp. 2225–2240.
  - [30] X. Chen, N. Zhang, X. Xie, S. Deng, Y. Yao, C. Tan, F. Huang, L. Si, and H. Chen, "Knowprompt: Knowledge-aware prompt-tuning with synergistic optimization for relation extraction," in *ACM*, 2022, pp. 2778–2788.
  - [31] H. Ye, N. Zhang, S. Deng, X. Chen, H. Chen, F. Xiong, X. Chen, and H. Chen, "Ontology-enhanced prompt-tuning for few-shot learning," in *ACM*, 2022, pp. 778–787.
  - [32] W. Yu, C. Zhu, Z. Li, Z. Hu, Q. Wang, H. Ji, and M. Jiang, "A survey of knowledge-enhanced text generation," *ACM Computing Surveys (CSUR)*, 2022.
  - [33] A. Lauscher, I. Vulic, E. M. Ponti, A. Korhonen, and G. Glavas, "Specializing unsupervised pretraining models for word-level semantic similarity," in *COLING*, 2020, pp. 1371–1383.
  - [34] Y. Levine, B. Lenz, O. Dagan, O. Ram, D. Padnos, O. Sharir, S. Shalev-Shwartz, A. Shashua, and Y. Shoham, "Sensebert: Driving some sense into BERT," in *ACL*, 2020, pp. 4656–4667.
  - [35] H. Tian, C. Gao, X. Xiao, H. Liu, B. He, H. Wu, H. Wang, and F. Wu, "SKEP: sentiment knowledge enhanced pre-training for sentiment analysis," in *ACL*, 2020, pp. 4067–4076.
  - [36] B. Lyu, L. Chen, S. Zhu, and K. Yu, "Let: Linguistic knowledge enhanced graph transformer for chinese short text matching," in *AAAI*, vol. 35, no. 15, 2021, pp. 13 498–13 506.
  - [37] Y. Xu, C. Zhu, S. Wang, S. Sun, H. Cheng, X. Liu, J. Gao, P. He, M. Zeng, and X. Huang, "Human parity on commonsenseqa: Augmenting self-attention with external attention," in *IJCAI*, 2022, pp. 2762–2768.
  - [38] Q. Chen, F. Li, G. Xu, M. Yan, J. Zhang, and Y. Zhang, "Dictbert: Dictionary description knowledge enhanced language model pre-training via contrastive learning," in *IJCAI*, 2022, pp. 4086–4092.
  - [39] J. Zhou, Z. Zhang, H. Zhao, and S. Zhang, "LIMIT-BERT : Linguistics informed multi-task BERT," in *EMNLP*, 2020, pp. 4450–4461.
  - [40] J. Bai, Y. Wang, Y. Chen, Y. Yang, J. Bai, J. Yu, and Y. Tong, "Syntaxbert: Improving pre-trained transformers with syntax trees," in *EACL*, 2021, pp. 3011–3020.
  - [41] Z. Li, Q. Zhou, C. Li, K. Xu, and Y. Cao, "Improving BERT with syntax-aware local attention," in *ACL-IJCNLP*, 2021, pp. 645–653.
  - [42] D. S. Sachan, Y. Zhang, P. Qi, and W. L. Hamilton, "Do syntax trees help pre-trained transformers extract information?" in *EACL*, 2021, pp. 2647–2661.
  - [43] U. Khandelwal, O. Levy, D. Jurafsky, L. Zettlemoyer, and M. Lewis, "Generalization through memorization: Nearest neighbor language models," in *ICLR*, 2020.
  - [44] K. Guu, K. Lee, Z. Tung, P. Pasupat, and M. Chang, "Retrieval augmented language model pre-training," in *ICML*, 2020, pp. 3929–3938.
  - [45] E. Grave, A. Joulin, and N. Usunier, "Improving neural language models with a continuous cache," in *ICLR*, 2017.
  - [46] S. Murty, P. W. Koh, and P. Liang, "Expbert: Representation engineering with natural language explanations," in *ACL*, 2020, pp. 2106–2113.
  - [47] W. Cui and X. Chen, "Enhancing natural language representation with large-scale out-of-domain commonsense," in *ACL*, 2022, pp. 1746–1756.
  - [48] Y. Yao, S. Huang, L. Dong, F. Wei, H. Chen, and N. Zhang, "Kformer: Knowledge injection in transformer feed-forward layers," in *NLPCC*, vol. 13551, 2022, pp. 131–143.
  - [49] S. Wang, Y. Xu, Y. Fang, Y. Liu, S. Sun, R. Xu, C. Zhu, and M. Zeng, "Training data is more valuable than you think: A simple and effective method by retrieving from training data," in *ACL*, 2022, pp. 3170–3179.
  - [50] B. Oguz, X. Chen, V. Karpukhin, S. Peshterliev, D. Okhonko, M. S. Schlichtkrull, S. Gupta, Y. Mehdad, and S. Yih, "Unik-qa: Unified representations of structured and unstructured knowledge for open-domain question answering," in *NAACL*, 2022, pp. 1535–1546.
  - [51] K. Ma, H. Cheng, X. Liu, E. Nyberg, and J. Gao, "Open domain question answering with a unified knowledge interface," in *ACL*, 2022, pp. 1605–1620.
  - [52] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, and J. Kang, "Biobert: a pre-trained biomedical language representation model for biomedical text mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, 2020.
  - [53] I. Beltagy, K. Lo, and A. Cohan, "Scibert: A pretrained language model for scientific text," in *EMNLP-IJCNLP*, 2019, pp. 3613–3618.
  - [54] K. Lo, L. L. Wang, M. Neumann, R. Kinney, and D. S. Weld, "S2ORC: the semantic scholar open research corpus," in *ACL*, 2020, pp. 4969–4983.
  - [55] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. D. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier *et al.*, "Knowledge graphs," *ACM Computing Surveys (CSUR)*, vol. 54, no. 4, pp. 1–37, 2021.
  - [56] B. Xue and L. Zou, "Knowledge graph quality management: a comprehensive survey," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
  - [57] Q. Wang, Z. Mao, B. Wang, and L. Guo, "Knowledge graph embedding: A survey of approaches and applications," *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 12, pp. 2724–2743, 2017.

- [58] X. Wang, F. Luo, Q. Wu, and Z. Bao, "How context or knowledge can benefit healthcare question answering," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [59] L. Fang, Y. Luo, K. Feng, K. Zhao, and A. Hu, "A knowledge-enriched ensemble method for word embedding and multi-sense embedding," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [60] Q. Li, D. Wang, S. F. K. Song, Y. Zhang, and G. Yu, "Oerl: Enhanced representation learning via open knowledge graphs," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [61] Y. Zhao, H. Zhou, A. Zhang, R. Xie, Q. Li, and F. Zhuang, "Connecting embeddings based on multiplex relational graph attention networks for knowledge graph entity typing," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [62] Y. Sun, S. Wang, Y. Li, S. Feng, X. Chen, H. Zhang, X. Tian, D. Zhu, H. Tian, and H. Wu, "ERNIE: enhanced representation through knowledge integration," *arXiv preprint arXiv:1904.09223*, 2019.
- [63] W. Xiong, J. Du, W. Y. Wang, and V. Stoyanov, "Pretrained encyclopedia: Weakly supervised knowledge-pretrained language model," in *ICLR*, 2020.
- [64] J. Wang, C. Wang, M. Qiu, Q. Shi, H. Wang, J. Huang, and M. Gao, "Keyp: Knowledge enhanced contrastive prompting for few-shot extractive question answering," *arXiv preprint arXiv:2205.03071*, 2022.
- [65] Y. Qin, Y. Lin, R. Takanobu, Z. Liu, P. Li, H. Ji, M. Huang, M. Sun, and J. Zhou, "ERICA: improving entity and relation understanding for pre-trained language models via contrastive learning," in *ACL-IJCNL*, 2021, pp. 3350–3363.
- [66] T. Zhang, C. Wang, N. Hu, M. Qiu, C. Tang, X. He, and J. Huang, "Dkplm: Decomposable knowledge-enhanced pre-trained language model for natural language understanding," in *AAAI*, vol. 36, no. 10, 2022, pp. 11 703–11 711.
- [67] I. Yamada, A. Asai, H. Shindo, H. Takeda, and Y. Matsumoto, "LUKE: deep contextualized entity representations with entity-aware self-attention," in *EMNLP*, 2020, pp. 6442–6454.
- [68] Z. Zhang, X. Han, Z. Liu, X. Jiang, M. Sun, and Q. Liu, "ERNIE: enhanced language representation with informative entities," in *ACL*, 2019, pp. 1441–1451.
- [69] M. E. Peters, M. Neumann, R. L. L. IV, R. Schwartz, V. Joshi, S. Singh, and N. A. Smith, "Knowledge enhanced contextual word representations," in *EMNLP-IJCNLP*, 2019, pp. 43–54.
- [70] T. Févry, L. B. Soares, N. FitzGerald, E. Choi, and T. Kwiatkowski, "Entities as experts: Sparse memory access with entity supervision," in *EMNLP*, 2020, pp. 4937–4951.
- [71] X. Wang, T. Gao, Z. Zhu, Z. Zhang, Z. Liu, J. Li, and J. Tang, "KEPLER: A unified model for knowledge embedding and pre-trained language representation," *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 176–194, 2021.
- [72] J. Wang, W. Huang, Q. Shi, H. Wang, M. Qiu, X. Li, and M. Gao, "Knowledge prompting in pre-trained language model for natural language understanding," *arXiv preprint arXiv:2210.08536*, 2022.
- [73] W. Liu, P. Zhou, Z. Zhao, Z. Wang, Q. Ju, H. Deng, and P. Wang, "K-bert: Enabling language representation with knowledge graph," in *AAAI*, vol. 34, no. 03, 2020, pp. 2901–2908.
- [74] T. Sun, Y. Shao, X. Qiu, Q. Guo, Y. Hu, X. Huang, and Z. Zhang, "Colake: Contextualized language and knowledge embedding," in *COLING*, 2020, pp. 3660–3670.
- [75] P. Verga, H. Sun, L. B. Soares, and W. W. Cohen, "Facts as experts: Adaptable and interpretable neural memory over symbolic knowledge," *arXiv preprint arXiv:2007.00849*, 2020.
- [76] R. Wang, D. Tang, N. Duan, Z. Wei, X. Huang, J. Ji, G. Cao, D. Jiang, and M. Zhou, "K-adapter: Infusing knowledge into pre-trained models with adapters," in *ACL-IJCNLP*, 2021, pp. 1405–1418.
- [77] D. Emelin, D. Bonadiman, S. Alqahtani, Y. Zhang, and S. Mansour, "Injecting domain knowledge in language models for task-oriented dialogue systems," in *EMNLP*, 2022.
- [78] L. He, S. Zheng, T. Yang, and F. Zhang, "Klmo: Knowledge graph enhanced pretrained language model with fine-grained relationships," in *EMNLP*, 2021, pp. 4536–4542.
- [79] Q. Dong, Y. Liu, S. Cheng, S. Wang, Z. Cheng, S. Niu, and D. Yin, "Incorporating explicit knowledge in pre-trained language models for passage re-ranking," in *SIGIR*, 2022, pp. 1490–1501.
- [80] Y. Sun, Q. Shi, L. Qi, and Y. Zhang, "Jointlk: Joint reasoning with language models and knowledge graphs for commonsense question answering," in *NAACL*, 2022, pp. 5049–5060.
- [81] X. Zhang, A. Bosselut, M. Yasunaga, H. Ren, P. Liang, C. D. Manning, and J. Leskovec, "Greaselm: Graph reasoning enhanced language models," in *ICLR*, 2022.
- [82] D. Yu, C. Zhu, Y. Yang, and M. Zeng, "JAKET: joint pre-training of knowledge graph and language understanding," in *AAAI*, 2022, pp. 11 630–11 638.
- [83] L. Von Rueden, S. Mayer, K. Beckh, B. Georgiev, S. Giesselbach, R. Heese, B. Kirsch, J. Pfrommer, A. Pick, R. Ramamurthy *et al.*, "Informed machine learning—a taxonomy and survey of integrating knowledge into learning systems," *arXiv preprint arXiv:1903.12394*, 2019.
- [84] M. Saeed, N. Ahmadi, P. Nakov, and P. Papotti, "Rulebert: Teaching soft rules to pre-trained language models," in *EMNLP*, 2021, pp. 1460–1476.
- [85] E. Dinan, S. Roller, K. Shuster, A. Fan, M. Auli, and J. Weston, "Wizard of wikipedia: Knowledge-powered conversational agents," in *ICLR*, 2019.
- [86] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel *et al.*, "Retrieval-augmented generation for knowledge-intensive nlp tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020.
- [87] H. Li, Y. Gong, J. Jiao, R. Zhang, T. Baldwin, and N. Duan, "Kfcnet: Knowledge filtering and contrastive learning network for generative commonsense reasoning," *arXiv preprint arXiv:2109.06704*, 2021.
- [88] X. Liu, D. Liu, B. Yang, H. Zhang, J. Ding, W. Yao, W. Luo, H. Zhang, and J. Su, "Kgr4: Retrieval, retrospect, refine and rethink for commonsense generation," in *AAAI*, vol. 36, no. 10, 2022, pp. 11 029–11 037.
- [89] D. Yu, C. Zhu, Y. Fang, W. Yu, S. Wang, Y. Xu, X. Ren, Y. Yang, and M. Zeng, "Kg-fid: Infusing knowledge graph in fusion-in-decoder for open-domain question answering," in *ACL*, 2022, pp. 4961–4974.
- [90] K. Shuster, M. Komeili, L. Adolphs, S. Roller, A. Szlam, and J. Weston, "Language models that seek for knowledge: Modular search & generation for dialogue and prompt completion," *arXiv preprint arXiv:2203.13224*, 2022.
- [91] S. Borgeaud, A. Mensch, J. Hoffmann, T. Cai, E. Rutherford, K. Millican, G. B. Van Den Driessche, J.-B. Lespiau, B. Damoc, A. Clark *et al.*, "Improving language models by retrieving from trillions of tokens," in *ICML*, 2022, pp. 2206–2240.
- [92] T. Xie, C. H. Wu, P. Shi, R. Zhong, T. Scholak, M. Yasunaga, C. Wu, M. Zhong, P. Yin, S. I. Wang, V. Zhong, B. Wang, C. Li, C. Boyle, A. Ni, Z. Yao, D. R. Radev, C. Xiong, L. Kong, R. Zhang, N. A. Smith, L. Zettlemoyer, and T. Yu, "Unifiedskg: Unifying and multi-tasking structured knowledge grounding with text-to-text language models," *arXiv preprint arXiv:2201.05966*, 2022.
- [93] B. Kim, J. Ahn, and G. Kim, "Sequential latent knowledge selection for knowledge-grounded dialogue," in *ICLR*, 2020.
- [94] Y. Li, B. Peng, Y. Shen, Y. Mao, L. Liden, Z. Yu, and J. Gao, "Knowledge-grounded dialogue generation with a unified knowledge representation," in *NAACL*, 2022, pp. 206–218.
- [95] G. Izacard and E. Grave, "Leveraging passage retrieval with generative models for open domain question answering," in *EACL*, 2021, pp. 874–880.
- [96] H. Ji, P. Ke, S. Huang, F. Wei, and M. Huang, "Generating commonsense explanation by extracting bridge concepts from reasoning paths," in *AAACL-IJCNLP*, 2020, pp. 248–257.
- [97] L. Zhao, J. Xu, J. Lin, Y. Zhang, H. Yang, and X. Sun, "Graph-based multi-hop reasoning for long text generation," *arXiv preprint arXiv:2009.13282*, 2020.
- [98] J. Guan, F. Huang, M. Huang, Z. Zhao, and X. Zhu, "A knowledge-enhanced pretraining model for commonsense story generation," *Transactions of the Association for Computational Linguistics*, vol. 8, pp. 93–108, 2020.
- [99] R. Sridhar and D. Yang, "Explaining toxic text via knowledge enhanced text generation," in *NAACL*, 2022, pp. 811–826.
- [100] D. Varshney, A. Prabhakar, and A. Ekbal, "Commonsense and named entity aware knowledge grounded dialogue generation," in *NAACL*, 2022, pp. 1322–1335.
- [101] H. Ji, P. Ke, S. Huang, F. Wei, X. Zhu, and M. Huang, "Language generation with multi-hop reasoning on commonsense knowledge graph," in *EMNLP*, 2020, pp. 725–736.
- [102] P. Ke, H. Ji, Y. Ran, X. Cui, L. Wang, L. Song, X. Zhu, and M. Huang, "Jointgt: Graph-text joint representation learning for

text generation from knowledge graphs,” in *ACL-IJCNLP*, 2021, pp. 2526–2538.

- [103] W. Yu, C. Zhu, L. Qin, Z. Zhang, T. Zhao, and M. Jiang, “Diversifying content generation for commonsense reasoning with mixture of knowledge graph experts,” in *ACL*, 2022, pp. 1896–1906.
- [104] K. Chen, Q. Huang, Y. Bisk, D. McDuff, and J. Gao, “Kb-vlp: Knowledge based vision and language pretraining,” in *ICML workshop*, July 2021.
- [105] F. Yu, J. Tang, W. Yin, Y. Sun, H. Tian, H. Wu, and H. Wang, “Ernie-vil: Knowledge enhanced vision-language representations through scene graphs,” in *AAAI*, vol. 35, no. 4, 2021, pp. 3208–3216.
- [106] G. M. Van de Ven and A. S. Tolias, “Three scenarios for continual learning,” *arXiv preprint arXiv:1904.07734*, 2019.
- [107] Y. Qin, J. Zhang, Y. Lin, Z. Liu, P. Li, M. Sun, and J. Zhou, “ELLE: efficient lifelong pre-training for emerging data,” in *ACL*, 2022, pp. 2789–2810.
- [108] X. Han, Z. Zhang, N. Ding, Y. Gu, X. Liu, Y. Huo, J. Qiu, Y. Yao, A. Zhang, L. Zhang, W. Han, M. Huang, Q. Jin, Y. Lan, Y. Liu, Z. Liu, Z. Lu, X. Qiu, R. Song, J. Tang, J. Wen, J. Yuan, W. X. Zhao, and J. Zhu, “Pre-trained models: Past, present and future,” *AI Open*, vol. 2, pp. 225–250, 2021.
- [109] S. Rajbhandari, J. Rasley, O. Ruwase, and Y. He, “Zero: memory optimizations toward training trillion parameter models,” in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2020, pp. 1–16.



**Linmei Hu** Linmei Hu works as an associate professor in the School of Computer Science and Technology, Beijing Institute of Technology. She received her Ph.D degree in Tsinghua University in 2018. Her research interests include Knowledge Graph, Natural Language Processing and Multimodal Content Analysis. She has published some papers in top conferences and journals, such as ACL, AAAI, EMNLP, ACM SIGKDD, MM, TOIS, and IEEE TKDE.



**Zeyi Liu** Zeyi Liu will receive the B.S. degree from University of Science and Technology Beijing in 2023, and continue to pursue her master's Degree in computer science and technology at Beijing University of Posts and Telecommunications.



**Ziwang Zhao** Ziwang Zhao received the B.S. degree from Xidian University, Xi'an, China, in 2021. He is currently pursuing the Master's Degree in computer science and technology at Beijing University of Posts and Telecommunications. His research interests include Pre-Trained Models, Knowledge Graph and Multimodal.



**Lei Hou** Lei Hou is an assistant professor in Department of Computer Science and Technology, Tsinghua University. He obtained his Ph.D. degree from Tsinghua University in 2016. His research interests include knowledge graph construction and application, news and user-generated content mining. He has published many research papers on top-tier international conferences, such as ACL, EMNLP, AAAI, etc.



**Liqiang Nie** Liqiang Nie, IAPR Fellow, is currently the dean with the School of Computer Science and Technology, Harbin Institute of Technology (Shenzhen). He received his B.Eng. and Ph.D. degree from Xi'an Jiaotong University and National University of Singapore (NUS), respectively. His research interests lie primarily in multimedia content analysis and information retrieval. Dr. Nie has co-/authored more than 100 CCF-A papers and 5 books, with 15k plus Google Scholar citations. He is an AE of IEEE TKDE, IEEE TMM, IEEE TCSVT, ACM ToMM, and Information Science. Meanwhile, he is the regular area chair or SPC of ACM MM, NeurIPS, IJCAI and AAAI. He is a member of ICME steering committee. He has received many awards over the past three years, like ACM MM and SIGIR best paper honorable mention in 2019, the AI 2000 most influential scholars 2020, SIGMM rising star in 2020, MIT TR35 China 2020, DAMO Academy Young Fellow in 2020, SIGIR best student paper in 2021, first prize of the provincial science and technology progress award in 2021 (rank 1), and provincial youth science and technology award in 2022. Some of his research outputs have been integrated into the products of Alibaba, Kwai, and other listed companies.



**Juanzi Li** Juanzi Li is a Professor in Department of Computer Science and Technology at Tsinghua University. She received her Ph.D. degree from Tsinghua University in 2000. Her research interests include knowledge engineering and semantic web, text and social network mining. She has published over 100 research papers in top international journals and conferences such as TKDE, KDD, AAAI, ACL, EMNLP, etc.