

A Survey of Knowledge-Enhanced Text Generation

WENHAO YU, University of Notre Dame, USA

Accepted to ACM Computing Survey (CUSR)

CHENGUANG ZHU, Microsoft Research, USA

ZAITANG LI, The Chinese University of Hong Kong, China

ZHITING HU, University of California at San Diego, USA

QINGYUN WANG, University of Illinois at Urbana-Champaign, USA

HENG JI, University of Illinois at Urbana-Champaign, USA

MENG JIANG, University of Notre Dame, USA

The goal of text-to-text generation is to make machines express like a human in many applications such as conversation, summarization, and translation. It is one of the most important yet challenging tasks in natural language processing (NLP). Various neural encoder-decoder models have been proposed to achieve the goal by learning to map input text to output text. However, the input text alone often provides limited knowledge to generate the desired output, so the performance of text generation is still far from satisfaction in many real-world scenarios. To address this issue, researchers have considered incorporating (i) internal knowledge embedded in the input text and (ii) external knowledge from outside sources such as knowledge base and knowledge graph into the text generation system. This research topic is known as *knowledge-enhanced text generation*. In this survey, we present a comprehensive review of the research on this topic over the past five years. The main content includes two parts: (i) general methods and architectures for integrating knowledge into text generation; (ii) specific techniques and applications according to different forms of knowledge data. This survey can have broad audiences, researchers and practitioners, in academia and industry.

CCS Concepts: • General and reference → Surveys and overviews; • Computing methodologies → Natural language processing; Neural networks.

Additional Key Words and Phrases: Natural language generation, Knowledge-enhanced Methods

§§ Some useful materials related to this survey:

- A tutorial entitled “Knowledge-enriched Natural Language Generation”, at EMNLP 2021. Tutorial abstract, slides, videos can be found at <https://kenlg-tutorial.github.io>.
- A tutorial entitled “Knowledge-augmented Methods for NLP”, to appear at ACL 2022.
- A Github repository with a more complete collection of papers and codes can be found at <https://github.com/wyu97/KENLG-Reading>. It will be frequently updated with new papers.

Authors' addresses: Wenhao Yu, wyu1@nd.edu, University of Notre Dame, Notre Dame, Indiana, USA, 46556; Chenguang Zhu, chezhu@microsoft.com, Microsoft Research, Redmond, Washington, USA, 98052; Zaitang Li, 1155107739@link.cuhk.edu.hk, The Chinese University of Hong Kong, Hong Kong, China, 999077; Zhiting Hu, zhitinghu@gmail.com, University of California at San Diego, San Diego, California, USA, 92092; Qingyun Wang, qingyun4@illinois.edu, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA, 61801; Heng Ji, hengji@illinois.edu, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA, 61801; Meng Jiang, mjiang2@nd.edu, University of Notre Dame, Notre Dame, Indiana, USA, 46556.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

0360-0300/2022/1-ART1 \$15.00

<https://doi.org/10.1145/3512467>

1 INTRODUCTION

Text generation, which is often formally referred as natural language generation (NLG), is one of the most important yet challenging tasks in natural language processing (NLP) [37]. NLG aims at producing understandable text in human language from linguistic or non-linguistic data in a variety of forms such as textual data, numerical data, image data, structured knowledge bases, and knowledge graphs. Among these, text-to-text generation is one of the most important applications and thus often shortly referred as “text generation”. Researchers have developed numerous technologies for this task in a wide range of applications [38, 55, 125]. Text generation takes text (e.g., a sequence, keywords) as input, processes the input text into semantic representations, and generates desired output text. For example, machine translation generates text in a different language based on the source text; summarization generates an abridged version of the source text to include salient information; question answering (QA) generates textual answers to given questions; dialogue system supports chatbots to communicate with humans with generated responses.

With the recent resurgence of deep learning technologies [66], deep neural NLG models have achieved remarkable performance in enabling machines to understand and generate natural language. A basic definition of the text generation task is to generate an expected *output sequence* from a given *input sequence*, called sequence-to-sequence (Seq2Seq). The Seq2Seq task and model were first introduced in 2014 [117]. It maps an input text to an output text under encoder-decoder schemes. The encoder maps the input sequence to a fixed-sized vector, and the decoder maps the vector to the target sequence. Since then, developing NLG systems has rapidly become a hot topic. Various text generation models have been proposed under deep neural encoder-decoder architectures. Popular architectures include recurrent neural network (RNN) encoder-decoder [117], convolutional neural network (CNN) encoder-decoder [39], and Transformer encoder-decoder [122].

Nevertheless, the input text alone contains limited knowledge to support neural generation models to produce the desired output. Meanwhile, the aforementioned methods generally suffer from an inability to well comprehend language, employ memory to retain and recall knowledge, and reason over complex concepts and relational paths; as indicated by their name, they involve encoding an input sequence, providing limited reasoning by transforming their hidden state given the input, and then decoding to an output. Therefore, the performance of generation is still far from satisfaction in many real-world scenarios. For example, in dialogue systems, conditioning on only the input text, a text generation system often produces trivial or non-committal responses of frequent words or phrases in the corpus [139, 159], such as “*Me too.*” or “*Oh my god!*” given the input text “*My skin is so dry.*” These mundane responses lack meaningful content, in contrast to human responses rich in knowledge. In comparison, humans are constantly acquiring, understanding, and storing knowledge from *broader sources* so that they can be employed to understand the current situation in communicating, reading, and writing. For example, in conversations, people often first select *concepts from related topics* (e.g., sports, food), then organize those topics into understandable content to respond; for summarization, people tend to write summaries containing *keywords* used in the input document and perform necessary modifications to ensure grammatical correctness and fluency; in question answering (QA), people use *commonsense* or *professional knowledge* pertained to the question to infer the answer. Therefore, it is often the case that knowledge beyond the input sequence is required to produce informative output text.

1.1 What is Knowledge-enhanced Text Generation?

In general, knowledge is the familiarity, awareness, or understanding that coalesces around a particular subject. In NLG systems, knowledge is an awareness and understanding of the input text and its surrounding context. These knowledge sources can be categorized into internal knowledge and external knowledge (see Figure 1). *Internal knowledge* creation takes place within the input

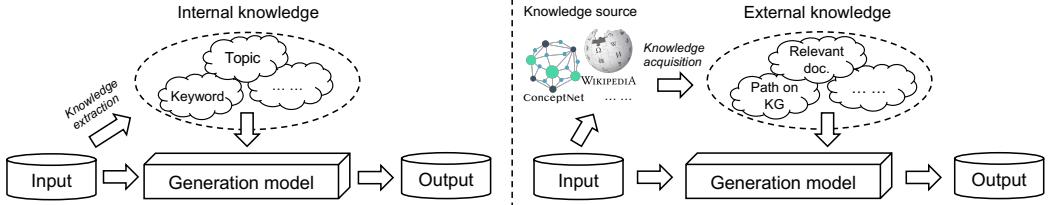


Fig. 1. We divide different knowledge sources into internal knowledge and external knowledge. Internal knowledge creation takes place within the input text(s), while external knowledge acquisition occurs when knowledge is provided from outside sources (e.g., Wikipedia, ConceptNet [115]).

text(s), including but not limited to keyword, topic, linguistic features, and internal graph structure. *External knowledge* acquisition occurs when knowledge is provided from outside sources, including but not limited to knowledge base, external knowledge graph, and grounded text. These sources provide information (e.g., commonsense triples, topic words, reviews, background documents) that can be used as knowledge through various neural representation learning methods, and then applied to enhance the process of text generation. In addition, knowledge introduces interpretability for models with explicit semantics. This research direction of incorporating knowledge into text generation is named as *knowledge-enhanced text generation*.

PROBLEM 1 (KNOWLEDGE-ENHANCED TEXT GENERATION). *Given a text generation problem where the system is given an input sequence X , and aims to generate an output sequence Y . Assume we also have access to additional knowledge denoted as K . Knowledge-enhanced text generation aims to incorporate the knowledge K to enhance the generation of Y given X , through leveraging the dependencies among the input text, knowledge, and output text.*

Many existing knowledge-enhanced text generation systems have demonstrated promising performance on generating informative, logical, and coherent texts. In dialogue systems, a topic-aware Seq2Seq model helped understand the semantic meaning of an input sequence and generate a more informative response such as “*Then hydrate and moisturize your skin.*” to the aforementioned example input “*My skin is so dry.*” In summarization, knowledge graph produced a structured summary and highlight the proximity of relevant concepts, when complex events related with the same entity may span multiple sentences. A knowledge graph enhanced Seq2Seq model generated summaries that were able to correctly answer 10% more topically related questions [54]. In question answering (QA) systems, facts stored in knowledge bases completed missing information in the question and elaborate details to facilitate answer generation [30, 48]. In story generation, using commonsense knowledge acquired from knowledge graph facilitated understanding of the storyline and better narrate following plots step by step, so each step could be reflected as a link on the knowledge graph and the whole story would be a path [46].

1.2 Why a Survey of Knowledge-enhanced Text Generation?

Recent years have witnessed a surge of interests in developing methods for incorporating knowledge in NLG beyond input text. However, there is a lack of comprehensive survey of this research topic. Related surveys have laid the foundation of discussing this topic. For example, Garbacea et al. [37] and Gatt et al. [38] reviewed model architectures for core NLG tasks but did not discuss knowledge-enhanced methods. Ji et al. [58] presented a review on knowledge graph techniques which could be used for enhancing NLG. Wang et al. [125] summarized how to represent structural knowledge such as knowledge base and knowledge graph for reading comprehension and retrieval.

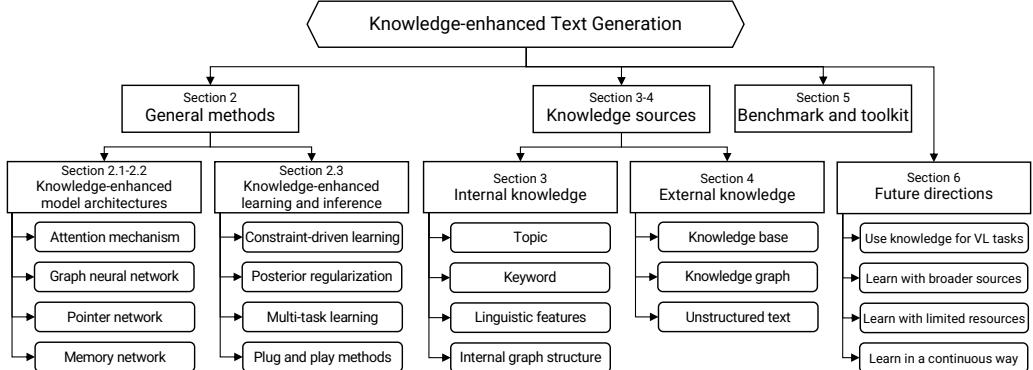


Fig. 2. Categorization of information sources and methods for knowledge-enhanced text generation. Knowledge can be learnt from various information sources, and then integrated into the generation process.

To the best of our knowledge, this is the first survey that presents a comprehensive review of knowledge-enhanced text generation. It aims to provide NLG researchers a synthesis and pointer to related research. Our survey includes a detailed discussion about how NLG can benefit from recent progress in deep learning and artificial intelligence, including technologies such as graph neural network, reinforcement learning, and neural topic modeling.

1.3 What are the Challenges in Knowledge-enhanced Text Generation?

To start with, we note that the *first challenge* in knowledge-enhanced NLG is to *obtain* useful related knowledge from diverse sources. There has been a rising line of work that discovers knowledge from topic, keyword, knowledge base, knowledge graph and knowledge grounded text. The *second challenge* is how to effectively *understand* and *leverage* the acquired knowledge to facilitate text generation. Multiple methods have been explored to improve the encoder-decoder architecture (e.g., attention mechanism, copy and pointing mechanism).

Based on the first challenge, the main content of our survey is divided into two parts: (1) general methods of integrating knowledge into text generation (Section 2); (2) specific methods and applications according to different sources of knowledge enhancement (Sections 3–4). More concretely, since knowledge can be obtained from different sources, we first divide existing knowledge enhanced text generation work into two categories: internal knowledge enhanced and external knowledge enhanced text generation. The division of internal and external knowledge is widely adopted by management science [88], which can be analogous with knowledge enhanced text generation. Based on the second challenge, we categorize recent knowledge-enhanced text generation methods evolved from how knowledge is extracted and incorporated into the process of text generation in each section (named as M1, M2, and etc). Furthermore, we review methods for a variety of natural language generation applications in each section to help practitioners choose, learn, and use the methods. In total, we discuss seven mainstream applications presented in more than 80 papers that were published or released in or after the year of 2016.

As shown in Figure 2, the remainder of this survey is organized as follows. Section 2 presents basic NLG models and general methods of integrating knowledge into text generation. Sections 3 reviews internal knowledge-enhanced NLG methods and applications. The internal knowledge is obtained from topic, keyword, linguistic features and internal graph structures. Sections 4 reviews external knowledge-enhanced NLG methods and applications. The external knowledge sources include knowledge bases, knowledge graphs, and grounded text. Section 5 presents knowledge-enhanced NLG benchmarks. Section 6 discusses future work and concludes the survey.

2 GENERAL METHODS OF INTEGRATING KNOWLEDGE INTO NLG

2.1 The Basic Text Generation Models

Early encoder-decoder frameworks are often based on recurrent neural network (RNN) such as RNN-Seq2Seq [117]. Convolutional neural network (CNN) based encoder-decoder [39] and Transformer encoder-decoder [122] have been increasingly widely used. From a probabilistic perspective, the encoder-decoder frameworks learn the conditional distribution over a variable length sequence conditioned on yet another variable length sequence:

$$P(Y|X) = P(y_1, \dots, y_m | x_1, \dots, x_n) = \prod_{t=1}^m p(y_t | X, y_1, \dots, y_{t-1}). \quad (1)$$

Encoder. The encoder learns to encode a variable length sequence into a fixed length vector representation. RNN encoder reads the input sentence X *sequentially*. CNN encoder performs convolutional operations on a word and its surrounding word(s) in a sequential window. Transformer encoder eschews recurrence and instead relying entirely on the self-attention mechanism to draw global dependencies between different tokens in the input X . We denote them uniformly as:

$$(\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n) = \text{ENCODER}(\mathbf{e}(x_1), \mathbf{e}(x_2), \dots, \mathbf{e}(x_n)), \quad (2)$$

where $\mathbf{e}(x_i)$ is the word embedding of word x_i , \mathbf{h}_i is the contextualized hidden representation of x_i .

Decoder. The decoder is to decode a given fixed length vector representation into a variable length sequence [117]. Specially, the decoder generates an output sequence one token at each time step. At each step the model is auto-regressive, consuming the previously generated tokens as additional input when generating the next token. Formally, the decoding function is represented as:

$$\mathbf{s}_t = \text{DECODER}(\mathbf{s}_{t-1}, \mathbf{e}(y_{t-1})), \quad (3)$$

$$p(y_t | y_{t-1}, y_{t-2}, \dots, y_1) = \text{READOUT}(\mathbf{s}_t), \quad (4)$$

where $\text{READOUT}(\cdot)$ is a nonlinear multi-layered function that outputs the probability of y_t .

Optimization. A generation process is regarded as a sequential multi-label classification problem. It can be directly optimized by the negative log likelihood (NLL) loss. Therefore, the objective of a text generation model via maximum likelihood estimation (MLE) is formulated as:

$$\mathcal{L}_{NLL}(\theta) = -\log p_\theta(Y|X) = -\sum_{t=1}^m \log (p_\theta(y_t | y_{<t}, X)). \quad (5)$$

2.2 Knowledge-enhanced Model Architectures

The most popular idea of incorporating knowledge is designing *specialized architectures* of text generation models that can reflect the particular type of knowledge. In the context of neural networks, several general neural architectures are widely used and customized to bake the knowledge about the problems being tackled into the models.

2.2.1 Attention Mechanism. It is useful to capture the weight of each time step in both encoder and decoder [3]. During the decoding phase, the context vector \mathbf{c}_t is added, so the hidden state \mathbf{s}_t is:

$$\mathbf{s}_t = \text{DECODER}(\mathbf{s}_{t-1}, \mathbf{e}(y_{t-1}), \mathbf{c}_t). \quad (6)$$

Unlike Eq.(3), here the probability is conditioned on the distinct context vector \mathbf{c}_t for target word y_t , and \mathbf{c}_t depends on a sequence of hidden states $\mathbf{H} = \{\mathbf{h}_i\}_{i=1}^n$ that were mapped from input sequence.

In RNN-Seq2Seq decoder, the \mathbf{c}_t is computed as a weighted sum of $\{\mathbf{h}_i\}_{i=1}^n$:

$$\mathbf{c}_t = \sum_{i=1}^n \alpha_{ti} \mathbf{h}_i, \text{ where } \alpha_{ti} = \frac{\exp(\eta(\mathbf{s}_{t-1}, \mathbf{h}_i))}{\sum_{k=1}^n \exp(\eta(\mathbf{s}_{t-1}, \mathbf{h}_k))}, \quad (7)$$

Table 1. NLG methods that incorporates knowledge attention (§2.2.1) and knowledge mode (§2.2.2).

	Topic	Keyword	Knowledge base	Knowledge graph	Grounded text
Knowledge-related attention	[134, 139, 152]	[69, 70, 73]	[34, 48]	[46, 54, 151, 159]	[9, 87]
Knowledge-related mode	[139]	[70]	[48]	[57, 151, 159]	[87, 107]
Knowledge-related memory	[34, 158]	-	[82, 135]	[144]	[61]

where $\eta(\cdot)$ is parametrized as a multi-layer perception to compute a soft alignment. $\eta(\cdot)$ enables the gradient of loss function to be backpropagated. There are six alternatives for the $\eta(\cdot)$ function (see Table 2 in [37]). The probability α_{ti} reflects the importance of the hidden state of input sequence in presence of the previous hidden state \mathbf{s}_{t-1} for deciding the next hidden state.

In Transformer decoder, on top of the two sub-layers in the encoder, the decoder inserts a third sub-layer, which performs multi-head attention over the output of the encoder stack \mathbf{H} . Efficient implementations of the transformer use the cached history matrix \mathbf{S}_t to generate next token. To compare with RNN-Seq2Seq, we summarize the Transformer decoder using recurrent notation:

$$\mathbf{S}_t = \text{TRANSFORMER-DECODER}(\mathbf{S}_{t-1}, \mathbf{e}(y_{t-1}), \mathbf{H}), \quad (8)$$

where $\mathbf{S}_t = [(\mathbf{K}_t^{(1)}, \mathbf{V}_t^{(1)}), \dots, (\mathbf{K}_t^{(l)}, \mathbf{V}_t^{(l)})]$, where $(\mathbf{K}_t^{(i)}, \mathbf{V}_t^{(i)})$ corresponds to the key-value pairs from the i -th layer generated at all time-steps from 0 to t . Instead of noting a specific name, we will use $\text{ENCODER}(\cdot)$ and $\text{DECODER}(\cdot)$ to represent encoder and decoder in the following sections.

Knowledge-related attention. Attention mechanism has been widely used to incorporate knowledge representation in recent knowledge-enhanced NLG work. The general idea is to learn a knowledge-aware context vector (denoted as $\tilde{\mathbf{c}}_t$) by combining both hidden context vector (\mathbf{c}_t) and knowledge context vector (denoted as \mathbf{c}_t^K) into decoder update, such as $\tilde{\mathbf{c}}_t = f_{mlp}(\mathbf{c}_t \oplus \mathbf{c}_t^K)$. The knowledge context vector (\mathbf{c}_t^K) calculates attentions over knowledge representations (e.g., topic vectors, node vectors in knowledge graph). Table 1 summarizes a variety of knowledge attentions, including keyword attention [69, 70, 73], topic attention [79, 134, 139, 152], knowledge base attention [34, 48], knowledge graph attention [54, 63, 151], and grounded text attention [9, 87].

2.2.2 Copy and Pointing Mechanisms. CopyNet and Pointer-generator (PG) are used to choose subsequences in the input sequence and put them at proper places in the output sequence.

CopyNet and PG have a differentiable network architecture [43]. They can be easily trained in an end-to-end manner. In CopyNet and PG, the probability of generating a target token is a combination of the probabilities of two modes, generate-mode and copy-mode. First, they represent unique tokens in the global vocabulary \mathcal{V} and the vocabulary of source sequence \mathcal{V}_X . They build an extended vocabulary $\mathcal{V}_{\text{ext}} = \mathcal{V} \cup \mathcal{V}_X \cup \{\text{unk}\}$. The difference between CopyNet and PG is the way to calculate distribution over the extended vocabulary. CopyNet calculates the distribution by

$$p(y_t) = p_g(y_t) + p_c(y_t), \quad (9)$$

where $p_g(\cdot)$ and $p_c(\cdot)$ stand for the probability of generate-mode and copy-mode. Differently, PG explicitly calculates a switch probability p_m between generate-mode and copy-mode. It recycles the attention distribution to serve as the copy distribution. The distribution over \mathcal{V}_{ext} is calculated by

$$p(y_t) = p_m(g) \cdot p_g(y_t) + (1 - p_m(g)) \cdot p_c(y_t), \quad (10)$$

where $p_m(g)$ indicates the probability of choosing generate-mode, which is obtained by a nonlinear multi-layered (MLP) function. Importantly, CopyNet and pointer-generator network have been used as the *base module* for a lot of knowledge-enhanced NLG work.

Knowledge-related mode. A knowledge-related mode chooses subsequences in the obtained knowledge and puts them at proper places in the output sequence. It helps NLG models to generate words that are not included in the global vocabulary (\mathcal{V}) and input sequence (\mathcal{V}_X). For example, by adding the model of knowledge base, the extended vocabulary (\mathcal{V}_{ext}) adds entities and relations

from the knowledge base, i.e., $\mathcal{V}_{ext} = \mathcal{V} + \mathcal{V}_X + \mathcal{V}_{KB}$. The probability of generating a target token is a combination of the probabilities of three modes: generate-mode, copy-mode and knowledge base-mode. Therefore, knowledge-related mode is not only capable of regular generation of words but also operation of producing appropriate subsequences in knowledge sources. Table 1 summarizes different kinds of knowledge-related modes such as topic mode [139], keyword mode [70], knowledge base mode [48], knowledge graph mode [151, 159], and background mode [87, 107].

2.2.3 Memory Network. Memory networks (MemNNs) are recurrent attention models over a possibly large external memory [116]. They write external memories into several embedding matrices, and use query (generally speaking, the input sequence X) vectors to read memories repeatedly. This approach encodes long dialog history and memorize external information.

Given an input set $\{m_1, \dots, m_i\}$ to be stored in memory. The memories of MemNN are represented by a set of trainable embedding matrices $\mathbf{C} = \{\mathbf{C}^1, \dots, \mathbf{C}^{K+1}\}$, where each \mathbf{C}^k maps tokens to vectors, and a query (i.e., input sequence) vector \mathbf{h}_X^k is used as a reading head. The model loops over K hops and it computes the attention weights at hop k for each memory m_i using:

$$\mathbf{p}_i^k = \text{softmax}((\mathbf{h}_X^k)^\top \mathbf{C}_i^k), \quad (11)$$

where $\mathbf{C}_i^k = \mathbf{C}^k(m_i)$ is the memory content in i -th position, i.e., mapping m_i into a memory vector. Here, \mathbf{p}^k is a soft memory selector that decides the memory relevance with respect to the query vector \mathbf{h}_X^k . Then, the model reads out the memory \mathbf{o}^k by the weighted sum over \mathbf{C}^{k+1} ,

$$\mathbf{o}^k = \sum_i \mathbf{p}_i^k \mathbf{C}_i^{k+1}. \quad (12)$$

Then, the query vector is updated for the next hop by using $\mathbf{h}_X^{k+1} = \mathbf{h}_X^k + \mathbf{o}^k$. The result from the encoding step is the memory vector \mathbf{o}^K and becomes the input for the decoding step.

Knowledge-related memory. Memory augmented encoder-decoder framework has achieved promising progress for many NLG tasks. For example, MemNNs are widely used for encoding dialogue history in task-oriented dialogue systems [106, 135]. Such frameworks enable a decoder to retrieve information from a memory during generation. Recent work explored to model external knowledge with memory network such as knowledge base [82, 144] and topic [34, 158].

2.2.4 Graph Network. Graph network captures the dependence of graphs via message passing between the nodes of graphs. Graph neural networks (GNNs) [138] and graph-to-sequence (Graph2Seq) [6] potentiate to bridge up the gap between graph representation learning and text generation. Knowledge graph, dependency graph, and other graph structures can be integrated into text generation through various GNN algorithms. Here we denote a graph as $\mathcal{G} = (\mathcal{U}, \mathcal{E})$, where \mathcal{U} is the set of entity nodes and \mathcal{E} is the set of (typed) edges. Modern GNNs typically follow a neighborhood aggregation approach, which iteratively updates the representation of a node by aggregating information from its neighboring nodes and edges. After k iterations of aggregation, a node representation captures the structural information within its k -hop neighborhood. Formally, the k -th layer of a node $u \in \mathcal{U}$ is:

$$\mathbf{u}^{(k)} = \text{COMBINE}_k(\mathbf{u}^{(k-1)}, \text{AGGREGATE}_k(\{(\mathbf{u}_i^{(k-1)}, \mathbf{e}_{ij}^{(k-1)}, \mathbf{u}_j^{(k-1)}) : \forall (u_i, e_{ij}, u_j) \in \mathcal{N}(u)\})), \quad (13)$$

where $\mathcal{N}(u) = \{(u_i, e_{ij}, u_j) \in \mathcal{E} | u_i = u \text{ or } u_j = u\}$ denotes the set of edges containing node u , $\mathbf{u}^{(k)}$ and $\mathbf{e}_{ij}^{(k)}$ are feature vectors of a node u and the edge between u_i and u_j at the k -th iteration/layer. The choice of $\text{AGGREGATE}(\cdot)$ and $\text{COMBINE}(\cdot)$ in GNNs is crucial. A number of architectures for $\text{AGGREGATE}(\cdot)$ have been proposed in different GNN works such as GAT [123]. Meanwhile, the $\text{AGGREGATE}(\cdot)$ function used in labeled graphs (e.g., a knowledge graph) is often taken as those GNNs for modeling relational graphs [108]. To obtain the representation of graph \mathcal{G} (denoted as

\mathbf{h}_G), the READOUT(\cdot) function (either a simple permutation invariant function or sophisticated graph-level pooling function) pools node features from the final iteration K ,

$$\mathbf{h}_G = \text{READOUT}(\{\mathbf{u}^{(K)} : u \in \mathcal{U}\}). \quad (14)$$

Applications. Graph network has been commonly used in integrating knowledge in graph structure such as knowledge graph and dependency graph. Graph attention network [123] can be combined with sequence attention and jointly optimized [151, 159]. We will introduce different graph structure knowledge in subsequent sections such as knowledge graph (Section 4.2), dependency graph (Section 3.3.2-3.3.3), and open knowledge graph (OpenKG) (Section 3.4).

2.2.5 Pre-trained Language Models. Pre-trained language models (PLMs) aims to learn universal language representation by conducting self-supervised training on large-scale unlabeled corpora. Recently, substantial PLMs such as BERT [25] and T5 [104] have achieved remarkable performance in various NLP downstream tasks. However, these PLMs suffer from two issues when performing on knowledge-intensive tasks. First, these models struggle to grasp structured world knowledge, such as concepts and relations, which are very important in language understanding. For example, BERT cannot deliver great performance on many commonsense reasoning and QA tasks, in which many of the concepts are directly linked on commonsense knowledge graphs [146]. Second, due to the domain discrepancy between pre-training and fine-tuning, these models do not perform well on domain-specific tasks. For example, BERT can not give full play to its value when dealing with electronic medical record analysis task in the medical field [78].

Recently, a lot of efforts have been made on investigating how to integrate knowledge into PLMs [45, 78, 80, 140, 146, 161]. Specifically, we will introduce some PLMs designed for NLG tasks. Overall, these approaches can be grouped into two categories: The first one is to explicitly inject entity representation into PLMs, where the representations are pre-computed from external sources [80, 155]. For example, KG-BART encoded the graph structure of KGs with knowledge embedding algorithms like TransE [11], and then took the informative entity embeddings as auxiliary input [80]. However, the method of explicitly injecting entity representation into PLMs has been argued that the embedding vectors of words in text and entities in KG are obtained in separate ways, making their vector-space inconsistent [78]. The second one is to implicitly modeling knowledge information into PLMs by performing knowledge-related tasks, such as concept order recovering [161], entity category prediction [146]. For example, CALM proposed a novel contrastive objective for packing more commonsense knowledge into the parameters, and jointly pre-trained both generative and contrastive objectives for enhancing commonsense NLG tasks [161].

2.3 Knowledge-enhanced Learning and Inference

Besides specialized model architectures, one common way of injecting knowledge to generation models is through the supervised knowledge learning. For example, one can encode knowledge into the objective function that guides the model training to acquire desired model behaviors [27, 61]. Such approaches enjoy the flexibility of integrating diverse types of knowledge by expressing them as certain forms of objectives. In general, knowledge-enhanced learning is agnostic to the model architecture, and can be combined with the aforementioned architectures.

2.3.1 Learning with knowledge-related tasks. One could devise learning tasks informed by the knowledge so that the model is trained to acquire the knowledge information.

Knowledge as target. The methods can be mainly divided into two categories as shown in Figure 3. The first category of knowledge-related tasks creates learning targets based on the knowledge, and the model is trained to recover the targets. These tasks can be combined as auxiliary tasks with the text generation task, resulting in a *multi-task learning* setting. For example, knowledge loss is defined as the cross entropy between the predicted and true knowledge sentences, and it is combined

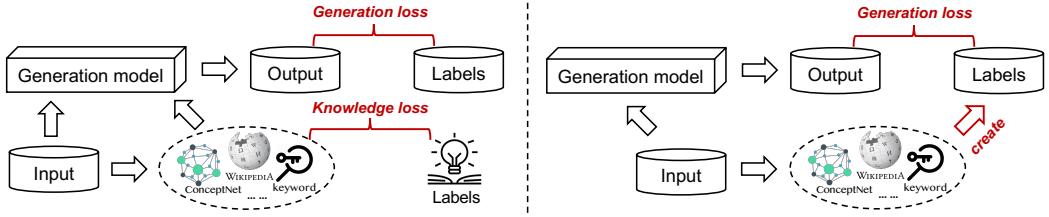


Fig. 3. Incorporating knowledge into text generation by treating knowledge as the target. The first category of methods (left) combine knowledge-related tasks as auxiliary into the text generation task, resulting in a *multi-task learning* setting. The second category of methods (right) create *weakly-supervised* labels from knowledge, enforcing the relevancy between the knowledge and the target sequence.

with the standard conversation generation loss to enhance grounded conversation [27, 61]. Similar tasks include keyword extraction loss [70], template re-ranking loss [13, 129], link prediction loss on knowledge graph [57], path reasoning loss [81], mode loss [137, 159], bag-of-word (BOW) loss [74, 143], etc. The second category of methods directly derive the text generation targets from the knowledge, and use those (typically noisy) targets as supervisions in the standard text generation task. The approach is called *weakly-supervised* learning. Weakly-supervised learning enforces the relevancy between the knowledge and the target sequence. For example, in the problem of aspect based summarization, the work [118] automatically creates target summaries based on external knowledge bases, which are used to train the summarization model in a supervised manner.

Knowledge as condition. The second way of devising knowledge-related tasks is to augment the text generation task by conditioning the generation on the knowledge. That is, the goal is to learn a function $p_\theta(Y|X, K)$, where X is the input sequence, Y is the target text and K is the knowledge. Generally, the knowledge K is first given externally (e.g., style, emotion) or retrieved from external resources (e.g., facts from knowledge base, a document from Wikipedia) or extracted from the given input text (e.g., keywords, topic words). Second, a conditional text generation model is used to incorporate knowledge and generate target output sequence. In practice, knowledge is often remedied by soft enforcing algorithms such as attention mechanism [3] and copy/pointing mechanism [43, 109]. Regarding knowledge as condition is widely used in knowledge-enhanced text generation. For examples, work has been done in making personalized dialogue response by taking account of persona [154] and emotion [158], controlling various aspects of the response such as politeness [96], grounding the responses in external source of knowledge [27, 42, 159] and generating topic-coherent sequence [119, 143]. Besides, using variational autoencoder (VAE) to enforce the generation process conditioned on knowledge is one popular approach to unsupervised NLG. By manipulating latent space for certain attributes, such as topic [132] and style [50], the output sequence can be generated with desired attributes without supervising with parallel data.

2.3.2 Learning with knowledge constraints. Instead of creating training objectives in standalone tasks that encapsulate knowledge, another paradigm of knowledge-enhanced learning is to treat the knowledge as the *constraints* to regularize the text generation training objective.

The posterior regularization (PR) framework was proposed to restrict the space of the model posterior on unlabeled data as a way to guide the model towards desired behavior [35, 164]. PR has been used as a principled framework to impose knowledge constraints on probabilistic models (including deep networks) in general [51, 153]. PR augments any regular training objective $\mathcal{L}(\theta)$ (e.g., negative log-likelihood, as in Eq.(5)) with a constraint term to encode relevant knowledge. Formally, denote the constraint function as $f(X, Y) \in \mathbb{R}$ such that a higher $f(X, Y)$ value indicates

a better generated sequence Y that incorporates the knowledge. PR introduces an auxiliary distribution $q(Y|X)$, and imposes the constraint on q by encouraging a large expected $f(X, Y)$ value: $\mathbb{E}q[f(X, Y)]$. Meanwhile, the model p_θ is encouraged to stay close to q through a KL divergence term. The learning problem is thus a constrained optimization:

$$\max_{\theta, q} \mathcal{L}(\theta) - \text{KL}(q(Y|X) || p_\theta(Y|X)) + \xi \quad (15)$$

$$\text{s.t. } \mathbb{E}q[f(X, Y)] > \xi, \quad (16)$$

where ξ is the slack variable. The PR framework is also related to other constraint-driven learning methods [14, 83]. We refer readers to [35] for more discussions.

2.3.3 Inference with knowledge constraints. Pre-trained language models leverage large amounts of unannotated data with a simple log-likelihood training objective. Controlling language generation by particular knowledge in a pre-trained model is difficult if we do not modify the model architecture to allow for external input knowledge or fine-tuning with specific data [24]. Plug and play language model (PPLM) opened up a new way to control language generation with particular knowledge during inference. At every generation step during inference, the PPLM shifts the history matrix in the direction of the sum of two gradients: one toward higher log-likelihood of the attribute a under the conditional attribute model $p(a|Y)$ and the other toward higher log-likelihood of the unmodified pre-trained generation model $p(Y|X)$ (e.g., GPT). Specifically, the attribute model $p(a|Y)$ makes gradient based updates to $\Delta \mathbf{S}_t$ as follows:

$$\Delta \mathbf{S}_t \leftarrow \Delta \mathbf{S}_t + \frac{\nabla_{\Delta \mathbf{S}_t} \log p(a|\mathbf{S}_t + \Delta \mathbf{S}_t)}{||\nabla_{\Delta \mathbf{S}_t} \log p(a|\mathbf{S}_t + \Delta \mathbf{S}_t)||^\gamma}, \quad (17)$$

where γ is the scaling coefficient for the normalization term; $\Delta \mathbf{S}_t$ is update of history matrix \mathbf{S}_t (see Eq.(8)) and initialized as zero. The update step is repeated multiple times. Subsequently, a forward pass through the generation model is performed to obtain the updated $\tilde{\mathbf{S}}_{t+1}$ as $\tilde{\mathbf{S}}_{t+1} = \text{DECODER}((\mathbf{S}_t + \Delta \mathbf{S}_t), \mathbf{e}(y_t), \mathbf{H})$. The perturbed $\tilde{\mathbf{S}}_{t+1}$ is then used to generate a new logit vector. PPLMs is efficient and flexible to combine differentiable attribute models to steer text generation [102].

3 NLG ENHANCED BY INTERNAL KNOWLEDGE

3.1 NLG Enhanced by Topic

Topic, which can be considered as a representative or compressed form of text, has been often used to maintain the semantic coherence and guide the NLG process. Topic modeling is a powerful tool for finding the high-level content of a document collection in the form of latent topics [10]. A classical topic model, Latent Dirichlet allocation (LDA), has been widely used for inferring a low dimensional representation that captures latent semantics of words and documents [10]. In LDA, each topic is defined as a distribution over words and each document as a mixture distribution over topics. LDA generates words in the documents from topic distribution of document and word distribution of topic. Recent advances of neural techniques open a new way of learning low dimensional representations of words from the tasks of word prediction and context prediction, making neural topic models become a popular choice of finding latent topics from text [12, 47].

Next, we introduce popular NLG applications enhanced by topics:

- **Dialogue system.** A vanilla Seq2Seq often generates trivial or non-committal sentences of frequent words or phrases in the corpus [139]. For example, a chatbot may say “*I do not know*”, “*I see*” too often. Though these off-topic responses are safe to reply to many queries, they are boring with very little information. Such responses may quickly lead the conversation to an end, severely hurting user experience. Thus, on-topic response generation is highly needed.

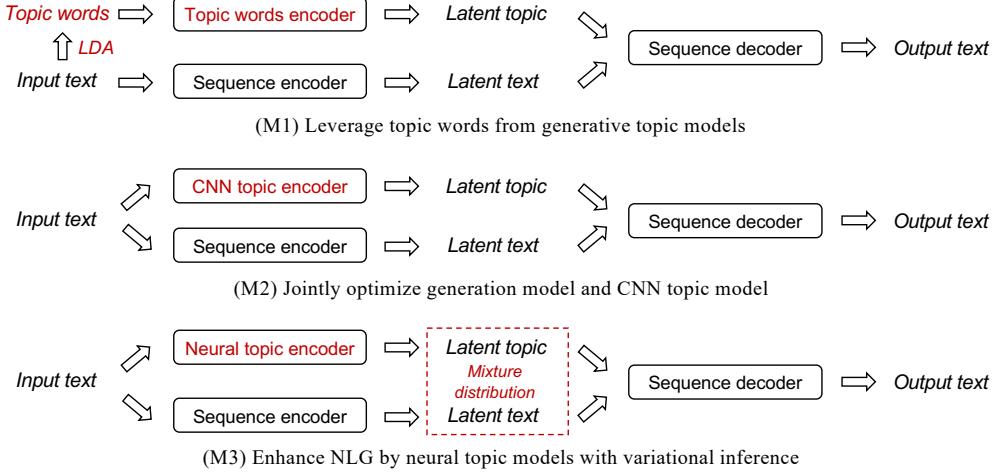


Fig. 4. Three typical methodologies for incorporating topics into NLG. Detailed designs are not included.

- **Machine translation.** Though the input and output languages are different (e.g., translating English to Chinese), the contents are the same, and globally, under the same topic. Therefore, topic can serve as an auxiliary guidance to preserve the semantics information of input text in one language into the output text in the other language.
- **Paraphrase.** Topic information helps understand the potential meaning and determine the semantic range to a certain extent. Naturally, paraphrases concern the same topic, which can serve as an auxiliary guidance to promote the preservation of source semantic.

As shown in Figure 4, we summarize topic-enhanced NLG methods into three methodologies: (M1) leverage topic words from generative topic models; (M2) jointly optimize generation model and CNN topic model; (M3) enhance NLG by neural topic models with variational inference.

3.1.1 M1: Leverage Topic Words from Generative Topic Models. Topics help understand the semantic meaning of sentences and determine the semantic spectrum to a certain extent. To enhance text generation, an effective solution is to first discover topics using generative topic models (e.g., LDA), and then incorporate the topics representations into neural generation models, as illustrated in Figure 4(a). In existing work, there are two mainstream methods to represent topics obtained from generative topic models. The first way is to use the generated topic distributions for each word (i.e., word distributions over topics) in the input sequence [94, 152]. The second way is to assign a specific topic to the input sequence, then picks the top- k words with the highest probabilities under the topic, and use word embeddings (e.g., GloVe) to represent topic words [79, 139]. Explicitly making use of topic words can bring stronger guidance than topic distributions, but the guidance may deviate from the target output sequence when some generated topic words are irrelevant. Zhang et al. proposed the first work of using a topic-informed Seq2Seq model by concatenating the topic distributions with encoder and decoder hidden states [152]. Xing et al. designed a topic-aware Seq2Seq model in order to use topic words as prior knowledge to help dialogue generation [139].

3.1.2 M2: Jointly Optimize Generation Model and CNN Topic Model. The LDA models were separated from the training process of neural generation model and were not able to adapt to the diversity of dependencies between input and output sequences. Therefore, the idea of addressing this issue is to use neural topic models. Convolutional neural networks (CNN) were used to learn latent topic representations through iterative convolution and pooling operations. There are growing interests of using the CNNs to map latent topics implicitly into topic vectors that can be

Table 2. Natural language generation methods that incorporate topic knowledge in text generation. Since most of the methods are tested on different tasks and datasets, we only compare the performance between “w/o topic” setting and “with topic” setting. For evaluation metrics, PPL is short for perplexity (lower is better); B-4 is short for BLEU-4 (higher is better); R-L is short for ROUGE-L (higher is better).

Task	Method	Ref.	Cat.	Framework components		Effect of topic modeling		
				Seq. Enc/Dec	Topic model	Dataset	w/o topic	with topic
Dialogue system	Tp-S2S	[139]	M1	RNN Seq2Seq	LDA topics	Baidu Tieba	(PPL) 147.0	(PPL) 134.6
	PEE	[143]	M3	RNN Seq2Seq	Neural topics	PersonaChat	(B-4) 2.98	(B-4) 3.56
Machine translation	Tp-NMT	[152]	M1	RNN Seq2Seq	LDA topics	NIST	(B-4) 34.76	(B-4) 35.91
	BLT-NMT	[134]	M2	RNN Seq2Seq	CNN topics	NIST	(B-4) 38.97	(B-4) 40.10
Summarization	Tp-CS2S	[94]	M1	CNN Seq2Seq	LDA topics	XSum	(R-L) 25.23	(R-L) 25.75
	TGVAE	[132]	M3	RNN with VAE	Neural topics	Gigawords	(R-L) 32.13	(R-L) 33.02
	VHTM	[33]	M3	RNN with VAE	Neural topics	CNN/DM	(R-L) 36.73	(R-L) 37.18
Paraphrase	TGLM	[36]	M2	RNN Seq2Seq	CNNs topics	Yahoo! Ans	(PPL) 99.13	(PPL) 88.69
	PTA	[79]	M1	RNN Seq2Seq	LDA topics	Quora	(B-4) 28.76	(B-4) 31.75

used to enhance text generation tasks [36, 134]. Empirical analyses showed that convolution-based topic extractors could outperform LDA-based topic models for multiple applications (e.g., dialogue system, text summarization, machine translation). However, theoretical analysis was missing to ensure the quality of the topics captured by the convolutions. And their interpretability is not as satisfactory as the LDA-based topic models.

3.1.3 M3: Enhance NLG by Neural Topic Models with Variational Inference. Neural topic models can be trained efficiently by backpropagation [12]. In neural topic models, Dirichlet distributions can be employed as the prior to generate the parameters of the multinomial distribution θ_d for each document [89]. The generative process of LDA is represented as: (1) $\theta_d \sim \text{Dirichlet}(\alpha)$; (2) $t_i \sim \text{Multinomial}(\theta_d)$; (3) $w_i \sim \text{Multinomial}(\beta_{t_i})$, where d denotes the bag-of-words representation of a document, t_i represents the topic assignment for word w_i , and β_{t_i} represents the topic distribution over words given topic assignment t_i . However, a directed generative model comes up against the problem of establishing low variance gradient estimators. Miao et al. parameterized the multinomial distributions with neural networks and jointly learned the model parameters via variational inference [89]. They created neural structures for constructing topic distributions conditioned on a draw from a multivariate Gaussian distribution, represented as $\theta_d \sim G(\mu_0, \sigma_0^2)$, where $G(\mu_0, \sigma_0^2)$ is composed of a neural network conditioned on an isotropic Gaussian $N(\mu_0, \sigma_0^2)$. Taking a Gaussian prior distribution makes re-parameterization feasible to build an unbiased and low-variance gradient estimator for the variational distribution [26]. Without conjugacy prior, the updates of the parameters are derived directly and easily from the variational lower bound. Formally, a variational lower bound for the document log-likelihood is:

$$\mathcal{J}_{\text{topic}} = \mathbb{E}_{q(\theta|d)} [\log p(d|\beta, \theta)] - \text{KL}(q(\theta|d)||p(\theta|\mu_0, \sigma_0^2)), \quad (18)$$

where $q(\theta|d)$ is the variational distribution approximating the true posterior $p(\theta|d)$. Its lower bound is estimate by sampling θ from $q(\theta|d) = G(\theta|\mu(d), \sigma^2(d))$.

In order to combine neural topic model and neural generation model, the idea is to use the Variational Auto-Encoder (VAE) [26]. It adopts autoregressive networks (e.g., LSTM) both as the encoder and decoder. VAE can learn latent codes z of texts by reconstructing texts with its decoder. It assumes that the generation process is controlled by codes in a continuous latent space. This kind of VAE implementation considers sequential information of texts that can model the linguistic structure of texts. Wang et al. proposed topic guided variational autoencoder (TGVAE), to draw latent code z from a topic-dependent Gaussian Mixture Prior in order to incorporate the topical knowledge into latent variables [132]. The topic-dependent Gaussian Mixture Model (GMM) is

defined as: $p(z|\beta, t) = \sum_{i=1}^T t_i N(\mu(\beta_i), \sigma^2(\beta_i))$, where T is the number of topics, $\mu(d)$ and $\sigma^2(d)$ are functions implemented by MLP. TGVAE uses bag-of-words as input and embeds an input document into a topic vector. The topic vector is then used to reconstruct the bag-of-words input, and the learned topic distribution over words is used to model a topic-dependent prior to generate an output sequence Y from conditioned on an input sequence X . Therefore, to maximize the log-likelihood $\log p(Y, d|X)$, a variational objective function is constructed as:

$$\mathcal{J}_{seq2seq} = \mathbb{E}_{q(z|X)} [\log p(Y|X, z)] - \mathbb{E}_{q(\theta|d)} [\text{KL}(q(z|X)||p(z|\beta, \theta))], \quad (19)$$

where $q(z|X)$ is variational distributions for z . The combined object function is given by:

$$\mathcal{J} = \mathcal{J}_{topic} + \mathcal{J}_{seq2seq}. \quad (20)$$

3.1.4 Discussion and Analysis of Different Methods. For **M1**, topic models (e.g., LDA) has a strict probabilistic explanation since the semantic representations of both words and documents are combined into a unified framework. Besides, topic models can be easily used and integrated into generation frameworks. For example, topic words can be represented as word embeddings; topic embeddings can be integrated into the decoding phase through topic attention. However, LDA models are separated from the training process of generation, so they cannot adapt to the diversity of dependencies between input and output sequences.

For **M2**, it is an end-to-end neural framework that simultaneously learns latent topic representations and generates output sequences. Convolutional neural networks (CNN) are often used to generate the latent topics through iterative convolution and pooling operations. However, theoretical analysis is missing to ensure the quality of the topics captured by the convolutions. And their interpretability is not as good as the LDA-based topic models.

For **M3**, neural topic models combine the advantages of neural networks and probabilistic topic models. They enable back propagation for joint optimization, contributing to more coherent topics, and can be scaled to large data sets. Generally, neural topic models can provide better topic coherence than LDAs [12, 132, 143]. However, neural variational approaches share a same drawback that topic distribution is assumed to be an isotropic Gaussian, which makes them incapable of modeling topic correlations. Existing neural topic models assume that the documents should be i.i.d. to adopt VAE, while they are commonly correlated. The correlations are critical for topic modeling.

3.2 NLG Enhanced by Keywords

Keyword (aka., key phrase, key term) is often referred as a sequence of one or more words, providing a compact representation of the content of a document. The mainstream methods of keyword acquisition for documents can be divided into two categories [112]: keyword assignment and keyword extraction. Keyword assignment means that keywords are chosen from a controlled vocabulary of terms or predefined taxonomy. Keyword extraction selects the most representative words explicitly presented in the document, which is independent from any vocabulary. Keyword extraction techniques (e.g., TF-IDF, TextRank, PMI) have been widely used over decades. Many NLG tasks can benefit from incorporating such a condensed form of essential content in a document to maintain the semantic coherence and guide the generation process.

Next, we introduce popular NLG applications enhanced by keywords:

- **Dialogue system.** Keywords help enlighten and drive the generated responses to be informative and avoid generating universally relevant replies which carry little semantics. Besides, recent work introduced personalized information into the generation of dialogue to help deliver better dialogue response such as emotion [71, 114, 158], and persona [154, 157].
- **Summarization.** Vanilla Seq2Seq models often suffer when the generation process is hard to control and often misses salient information [69]. Making use of keywords as explicit

Table 3. Natural language generation methods that incorporate keyword in text generation.

(a) (M1) Descriptions and quantitative comparisons between three methods for emotional dialogue systems.

Task	Method	Ref.	Assignment method	Experiments on NLPCC dataset		
				BLEU	D-1/D-2	Emotion w/s
Dialogue system	Seq2Seq	[3]	Seq2Seq attention <i>without</i> using keywords	1.50	0.38/1.20	33.5/37.1
	E-SCBA	[71]	MLP classifier to 7 emotions (categories)	1.69	0.54/4.84	72.0/51.2
	EmoChat	[158]	E-SCBA + two memory modules for decoding	1.68	0.90/7.35	76.5/58.0
	EmoDS	[114]	MLP classifier after decoding (discriminator)	1.73	1.13/8.67	81.0/68.7

(b) (M2) As most methods are tested on different tasks and datasets, we only compare the performance between “w/o keyword” setting and “with keyword” setting. Besides, HM is short for human evaluation.

Task	Method	Ref.	Extraction method	Keyword labels	Effect of keyword		
					Dataset	w/o keyword	with keyword
Summarization	KIGN	[69]	TextRank	Unsupervised	CNN/DM Gigaword	(R-2) 15.66 (R-2) 23.61	(R-2) 17.12 (R-2) 23.93
	ComGen	[73]	PMI and TFIDF	Unsupervised	Tencent Gigaword	(HM) 5.77 (R-2) 23.61	(HM) 7.19 (R-2) 25.06
	KGAS	[70]	BiLSTM-Softmax	$w(X) \cap w(Y)$			
Question generation	Selector	[22]	BiLSTM-Softmax	$w(X) \cap w(Y)$	SQuAD	(B-4) 14.72	(B-4) 15.87
	Prior	[133]	BiLSTM-Softmax	$w(X) \cap w(Y)$	SQuAD	(B-4) 14.72	(B-4) 15.34

guidance can provide significant clues of the main points about the document [69, 70]. It is closer to the way that humans write summaries: make sentences to contain the keywords, and then perform necessary modifications to ensure the fluency and grammatically correctness.

- **Question generation.** It aims to generate questions from a given answer and its relevant context. Given an answer and its associated context, it is possible to raise multiple questions with different focuses on the context and various means of expression.

Researchers have developed a great line of keyword-enhanced NLG methods. These methods can be categorized into two methodologies: (M1) Incorporate keyword assignment into text generation; (M2) Incorporate keyword extraction into text generation.

3.2.1 M1: Incorporate Keyword Assignment into Text Generation. When assigning a keyword to an input document, the set of possible keywords is bounded by a pre-defined vocabulary [112]. The keyword assignment is typically implemented by a classifier that maps the input document to a word in the pre-defined vocabulary [23, 71, 114, 158]. Unfortunately, some NLG scenarios do not hold an appropriate pre-defined vocabulary, so keyword assignment cannot be widely used to enhance NLG tasks. One applicable scenario is to use a pre-determined domain specific vocabulary to maintain relevance between the input and the output sequence [23]. Another scenario is to generate dialogue with specific attributes such as persona [113, 143], emotion [71, 114, 158].

M1.1: Adding assigned keyword into the decoder. A straightforward method of keyword assignment is to assign the words from pre-defined vocabulary and use them as the keywords [113, 143]. Sometimes, the input sequence does not have an explicit keyword, but we can find one from the pre-defined vocabulary. For example, a dialogue utterance “*If you had stopped him that day, things would have been different.*” expresses sadness but it does not have the word “sad.” To address this issue, Li et al. propose a method to predict an emotion category by fitting the sum of hidden states from encoder into a classifier [71]. Then, the response will be generated with the guidance of the emotion category. In order to dynamically track how much the emotion is expressed in the generated sequence, Zhou et al. propose a memory module to capture the emotion dynamics during decoding [158]. Each category is initialized with an emotion state vector before the decoding phase starts. At each step, the emotion state decays by a certain amount. Once the decoding process is completed, the emotion state decays to zero, indicating that the emotion is completely expressed.

M1.2: Assigning keyword for generated sequence. As mentioned in [114], explicitly incorporating emotional keywords suffers from expressing a certain emotion overwhelmingly. Instead, Song et al. propose to increase the intensity of the emotional experiences not by using emotional words explicitly, but by implicitly combining neutral words in distinct ways on emotion [114]. Specifically, they use an emotion classifier to build a sentence-level emotion discriminator, which helps to recognize the responses that express a certain emotion but not explicitly contain too many literal emotional words. The discriminator is connected to the end of the decoder.

3.2.2 M2: Incorporate Keyword Extraction into Text Generation. Keyword extraction selects salient words from input documents [112]. Recent work has used statistical keyword extraction techniques (e.g., PMI [73], TextRank [69]), and neural-based keyword extraction techniques (e.g., BiLSTM [70]). The process of incorporating extracted keywords into generation is much like the process discussed in Section 3.2.1. It takes keywords as an additional input into decoder. Recent work improves encoding phase by adding another sequence encoder to represent keywords [69, 70]. Then, the contextualized keywords representation is fed into the decoder together with input sequence representation. To advance the keyword extraction, Li et al. propose to use multi-task learning for training a keyword extractor network and generating summaries [22, 70]. Because both summarization and keyword extraction aim to select important information from input document, these two tasks can benefit from sharing parameters to improve the capacity of capturing the gist of the input text. In practice, they take overlapping words between the input document and the ground-truth summary as keywords, and adopt a BiLSTM-Softmax as keyword extractor. Similar idea has also been used in question generation tasks [22, 133]. They use overlapping words between the input answer context and the ground-truth question as keywords.

3.2.3 Discussion and Analysis of Different Methods.

Pros and cons. For **M1**, the primary advantage of keyword assignment is that the quality of keywords is guaranteed, because irrelevant keywords are not included in the pre-defined vocabulary. Another advantage is that even if two semantically similar documents do not have common words, they can still be assigned with the same keyword. However, there are mainly two drawbacks. On one hand, it is expensive to create and maintain dictionaries in new domains. So, the dictionaries might not be available. On the other hand, potential keywords occurring in the document would be unfortunately ignored if they were not in the vocabulary. Therefore, keyword assignment is suitable for the task that requires specific categories of keywords to guide the generated sentences with these key information. For example, dialogue systems generate responses with specific attitudes.

For **M2**, keyword extraction selects the most representative words explicitly presented in the document, which is independent from any vocabulary. It is easy to use but has two drawbacks. First, it cannot guarantee consistency because similar documents may still be represented by different keywords if they do not share the same set of words. Second, when an input document does not have a proper representative word, and unfortunately, the keyword extractor selects an irrelevant word from the document as a keyword, this wrong guidance will mislead the generation. Therefore, keyword extraction is suitable for the task that the output sequence needs to keep important information in the input sequence such as document summarization and paraphrase.

Quantitative analysis. Table 3 summarizes tasks and datasets used in keyword-enhanced NLG work. Comparing with keyword-enhanced methods (E-SCBA [71]) and the basic Seq2Seq attention model, keyword-enhanced methods can greatly improve both generation quality (evaluated by BLEU) and emotional expression (evaluated by emotion-w and emotion-s) on the NLPCC dataset. Besides, as shown in Table 3(a), EmoDS [114] achieved the best performance among three M1 methods, which indicates taking keyword assignment as a discriminant task can make better

improvement than assigning keyword before the sentence decoding. For M2 methods, since most methods were evaluated on different tasks, we can only compare the performance between “without using keyword” and “using keyword”. As shown in Table 3(b), leveraging extracted keywords from input sequence into Seq2Seq model can improve the generation quality on summarization and question generation tasks. Comparing with KGAS [70] and KIGN [69], we can observe using BiLSTM-Softmax to extract keyword (a supervised manner by using overlapping words between X and Y as labels) can make better performance than using TextRank (an unsupervised manner).

3.3 NLG Enhanced by Linguistic Features

Feature enriched encoder means that the encoder not only reads the input sequence, but also incorporates auxiliary hand-crafted features [110, 149, 160]. Linguistic features are the most common hand-crafted features, such as part-of-speech (POS) tags, dependency parsing, and semantic parsing.

3.3.1 POS tags and NER tags. Part-of-speech tagging (POS) assigns token tags to indicate the token’s grammatical categories and part of speech such as *noun* (N), *verb* (V), *adjective* (A). Named-entity recognition (NER) classifies named entities mentioned in unstructured text into pre-defined categories such as *person* (P), *location* (L), *organization* (O). CoreNLP is the most common used tool [84]. In spite of homonymy and word formation processes, the same surface word form may be shared between several word types. Incorporating NER tags and POS tags can detect named entities and understand input sequence better, hence, further improve NLG [28, 93, 160].

3.3.2 Syntactic dependency graph. Syntactic dependency graph is a directed acyclic graph representing syntactic relations between words [4]. For example, in the sentence “The monkey eats a banana”, “monkey” is the subject of the predicate “eats”, and “banana” is the object. Enhancing sequence representations by utilizing dependency information captures source long-distance dependency constraints and parent-child relation for different words [1, 4, 15]. In NLG tasks, dependency information is often modeled in three different ways as follows: (i) linearized representation: linearize dependency graph and then use sequence model to obtain syntax-aware representation [1]; (ii) path-based representation: calculate attention weights based on the linear distance between a word and the aligned center position, i.e., the greater distance a word to the center position on the dependency graph is, the smaller contribution of the word to the context vector is [15]; and (iii) graph-based representation: use GNNs to aggregate information from dependency relations [4].

3.3.3 Semantic dependency graph. Semantic dependency graph represents *predicate-argument* relations between content words in a sentence and have various semantic representation schemes (e.g., DM) based on different annotation systems. Nodes in a semantic dependency graph are extracted by semantic role labeling (SRL) or dependency parsing, and connected by different intra-semantic and inter-semantic relations [98]. Since semantic dependency graph introduces a higher level of information abstraction that captures commonalities between different realizations of the same underlying predicate-argument structures, it has been widely used to improve text generation [59, 75, 98]. Jin et al. propose a semantic dependency guided summarization model [59]. They incorporate the semantic dependency graph and the input text by stacking encoders to guide summary generation process. The stacked encoders consist of a sequence encoder and a graph encoder, in which the sentence encoder first reads the input text through stacked multi-head self-attention, and then the graph encoder captures semantic relationships and incorporates the semantic graph structure into the contextual-level representation.

3.4 NLG Enhanced by Open Knowledge Graphs

For those KGs (e.g., ConceptNet) constructed based on data beyond the input text, we refer them as *external KGs*. On the contrary, an *internal KG* is defined as a KG constructed solely based on the input text. In this section, we will discuss incorporating internal KG to help NLG [30, 54].

Internal KG plays an important role in understanding the input sequence especially when it is of great length. By constructing an internal KG intermediary, redundant information can be merged or discarded, producing a substantially compressed form to represent the input document [30]. Besides, representations on KGs can produce a structured summary and highlight the proximity of relevant concepts, when complex events related with the same entity may span multiple sentences [54]. One of the mainstream methods of constructing an internal KG is using open information extraction (OpenIE). Unlike traditional information extraction (IE) methods, OpenIE is not limited to a small set of target entities and relations known in advance, but rather extracts all types of entities and relations found in input text [95]. In this way, OpenIE facilitates the domain independent discovery of relations extracted from text and scales to large heterogeneous corpora.

After obtaining an internal KG, the next step is to learn the representation of the internal KG and integrate it into the generation model. For example, Zhu et al. use a graph attention network (GAT) to obtain the representation of each node, and fuse that into a transformer-based encoder-decoder architecture via attention [163]. Their method generates abstractive summaries with higher factual correctness. Huang et al. extend by first encoding each paragraph as a sub-KG using GAT, and then connecting all sub-KGs with a Bi-LSTM [54]. This process models topic transitions and recurrences, which enables the identification of notable content, thus benefiting summarization.

4 NLG ENHANCED BY EXTERNAL KNOWLEDGE

4.1 NLG Enhanced by Knowledge Base

One of the biggest challenges in NLG is to discover the dependencies of elements within a sequence and/or across input and output sequences. The dependencies are actually various types of *knowledge* such as commonsense, factual events, and semantic relationship. Knowledge base (KB) is a popular technology that collects, stores, and manages large-scale information for knowledge-based systems like search engines. It has a great number of triples composed of subjects, predicates, and objects. People also call them “facts” or “factual triplets”. Recently, researchers have been designing methods to use KB as external knowledge for learning the dependencies easier, faster, and better.

Next, we introduce popular NLG applications enhanced by knowledge base:

- **Question answering.** It is often difficult to generate proper answers only based on a given question. This is because, depending on what the question is looking for, a good answer may have different forms. It may complete the question precisely with the missing information. It may elaborate details of some part of the question. It may need reasoning and inference based on some facts and/or commonsense. So, only incorporating input question into neural generation models often fails the task due to the lack of commonsense/factual knowledge [8]. Related structured information of commonsense and facts can be retrieved from KBs.
- **Dialogue system.** The needs of KB in generating conversations or dialogues are relevant with QA but differ from two aspects. First, a conversation or dialogue could be open discussions when started by an open topic like “*Do you have any recommendations?*” Second, responding an utterance in a certain step needs to recall previous contexts to determine involved entities. KB will play an important role to recognize dependencies in the long-range contexts.

To handle different kinds of relationships between KB and input/output sequences, these methods can be categorized into two methodologies which is shown in Figure 5: (M1) design supervised tasks around KB for joint optimization; (M2) enhance incorporation by selecting KB or facts.

4.1.1 M1: Design Supervised Tasks around KB for Joint Optimization. Knowledge bases (KBs) that acquire, store, and represent factual knowledge can be used to enhance text generation. However, designing effective incorporation to achieve a desired enhancement is challenging because a vanilla Seq2Seq often fails to represent discrete isolated concepts though they perform well to

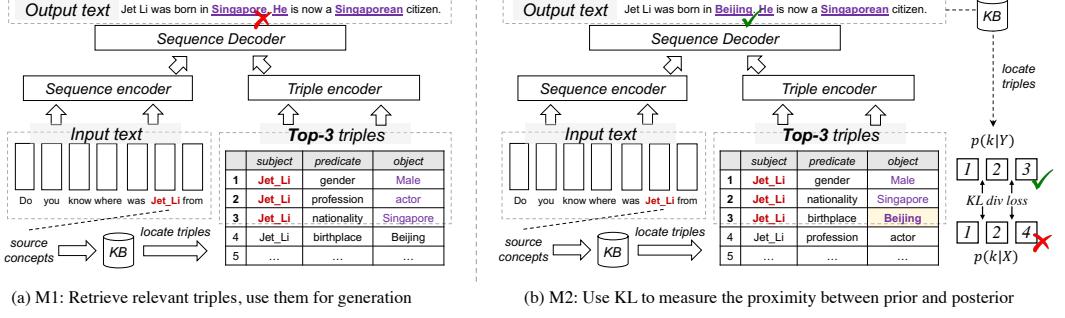


Fig. 5. The left figure demonstrates retrieving relevant triples, then using them for generation; the right figure demonstrate using KL to measure the proximity between prior and posterior distribution.

learn smooth shared patterns (e.g., language diversity). To fully utilize the knowledge bases, the idea is to jointly train neural models on multiple tasks. For example, the target task is answer sequence generation, and additional tasks include question understanding and fact retrieval in the KB. Knowledge can be shared across a unified encoder-decoder framework design. Typically, question understanding and fact retrieval are relevant and useful tasks, because a question could be parsed to match (e.g., string matching, entity linking, named entity recognition) its subject and predicate with the components of a fact triple in KB, and the answer is the object of the triple. KBCopy was the first work to generate responses using factual knowledge bases [29]. During the generation, KBCopy is able to copy words from the KBs. However, the directly copying relevant words from KBs is extremely challenging. CoreQA used both copying and retrieving mechanisms to generate answer sequences with an end-to-end fashion [48]. Specifically, it had a retrieval module to understand the question and find related facts from the KB. Then, the question and all retrieved facts are transformed into latent representations by two separate encoders. During the decoding phase, the integrated representations are fed into the decoder by performing a joint attention on both input sequence and retrieved facts. Figure 5(a) demonstrates a general pipeline that first retrieves relevant triples from KBs, then leverages the top-ranked triples into the generation process.

4.1.2 M2: Enhance Incorporation by Selecting KB or Facts in KB. Ideally, the relevance of the facts is satisfactory with the input and output sequence dependencies, however, it is not always true in real cases. Lian et al. addressed the issue of selecting relevant facts from KBs based on retrieval models (e.g. semantic similarity) might not effectively achieve appropriate knowledge selection [74]. The reason is that different kinds of selected knowledge facts can be used to generate diverse responses for the same input utterance. Given a specific utterance and response pair, the posterior distribution over knowledge base from both the utterance and the response may provide extra guidance on knowledge selection. The challenge lies in the discrepancy between the prior and posterior distributions. Specifically, the model learns to select effective knowledge only based on the prior distribution, so it is hard to obtain the correct posterior distribution during inference.

To tackle this issue, the work of Lian et al. [74] and Wu et al. [137] (shown in Figure 5(b)) approximated the posterior distribution using the prior distribution in order to select appropriate knowledge even without posterior information. They introduced an auxiliary loss, called Kullback-Leibler divergence loss (KLDivLoss), to measure the proximity between the prior distribution and the posterior distribution. The KLDivLoss is defined as follows:

$$\mathcal{L}_{KLDiv}(\theta) = \sum_{i=1}^N p(k = k_i|X, Y) \log \frac{p(k = k_i|X, Y)}{p(k = k_i|X)}, \quad (21)$$

Table 4. M2-based methods can retrieve more precise triples, and further improve the generation performance.

Method	Cat.	Ref.	Chinese Weibo (large) [137]				Chinese Weibo (small) [136]			
			Entity score		Generation score		Entity score		Generation score	
			Match	Recall	BLEU-2	Dist-2	Match	Recall	BLEU-2	Dist-2
GenDS	M1	[165]	0.97	0.37	3.42	4.27	0.75	0.26	2.09	1.66
CCM	M1	[159]	1.09	0.37	4.75	4.87	0.99	0.28	3.26	2.59
ConKADI	M2	[136]	-	-	-	-	1.48	0.38	5.06	23.93
TaFact	M2	[137]	1.81	0.47	5.07	23.56	-	-	-	-

where N is the number of retrieved facts. When minimizing KLDivLoss, the posterior distribution $p(k|X, Y)$ can be regarded as labels to apply the prior distribution $p(k|X)$ for approximating $p(k|X, Y)$. Finally, the total loss is written as the sum of the KLDivLoss and NLL (generation) loss.

4.1.3 Discussion and Analysis of Different Methods. The relevance between triples in KBs and input sequences plays a central role in discovering knowledge for sequence generation. Methods in **M1** typically follows the process that parses input sequence, retrieves relevant facts, and subsequently, a knowledge-aware output can be generated based on the input sequence and previously retrieved facts. Even though the improvement by modeling KB with memory network [82], existing KG-enhanced methods still suffer from effectively selecting precise triples.

Methods of **M2** improve the selection of facts, in which the ground-truth responses used as the posterior context knowledge to supervise the training of the prior fact probability distribution. Wu et al. used exact match and recall to measure whether the retrieved triples is used to generate the target outputs [136]. Table 4 shows the entity recall scores of M1-based methods and M2-based methods reported in [136, 137]. We observe that compared to M1-based methods, M2-based methods can greatly improve the accuracy of triple retrieval, as well as the generation quality.

There are still remaining challenges in KB-enhanced methods. One is that retrieved facts may contain noisy information, making the generation unstable [61]. This problem is extremely harmful in NLG tasks, e.g., KB-based question answering and task-oriented dialogue system, since the information in KB is usually the expected entities in the response.

4.2 NLG Enhanced by Knowledge Graph

Knowledge graph (KG), as a type of structured human knowledge, has attracted great attention from both academia and industry. A KG is a structured representation of facts (a.k.a. knowledge triplets) consisting of entities*, relations, and semantic descriptions [58]. The terms of “knowledge base” and “knowledge graph” can be interchangeably used, but they do not have to be synonymous. The knowledge graph is organized as a graph, so the connections between entities are first-class citizens in it. In the KG, people can easily traverse links to discover how entities are interconnected to express certain knowledge. Recent advances in artificial intelligence research have demonstrated the effectiveness of using KGs in various applications like recommendation systems [127].

Next, we introduce popular NLG applications that have been enhanced by knowledge graph:

- **Commonsense reasoning.** It aims to empower machines to capture the human commonsense from KG during generation. The methods exploit both structural and semantic information of the commonsense KG and perform reasoning over multi-hop relational paths, in order to augment the limited information with chains of evidence for commonsense reasoning. Popular tasks in commonsense reasoning generation include abductive reasoning (e.g., the α NLG task) [7, 57], counterfactual reasoning [56, 57], and entity description generation [21].
- **Dialogue system.** It frequently makes use of KG for the semantics in linked entities and relations [97, 121, 151, 159]. A dialogue may shift focus from one entity to another, breaking

*For brevity, we use “entities” to denote both entities (e.g., prince) and concepts (e.g., musician) throughout the paper.

one discourse into several segments, which can be represented as a linked path connecting the entities and their relations.

- **Creative writing.** This task can be found in both scientific and story-telling domains. Scientific writing aims to explain natural processes and phenomena step by step, so each step can be reflected as a link on KG and the whole explanation is a path [63, 130]. In story generation, the implicit knowledge in KG can facilitate the understanding of storyline and better predict what will happen in the next plot [45, 46, 80].

Compared with separate, independent knowledge triplets, knowledge graph provides comprehensive and rich entity features and relations for models to overcome the influence of the data distribution and enhance its robustness. Therefore, node embedding and relational path have played important roles in various text generation tasks. The corresponding techniques are knowledge graph embedding (KGE) [131] and path-based knowledge graph reasoning [17]. Furthermore, it has been possible to encode multi-hop and high-order relations in KGs using the emerging graph neural network (GNN) [138] and graph-to-sequence (Graph2Seq) frameworks [6].

Definition 4.1 (Knowledge graph (KG)). A knowledge graph (KG) is a directed and multi-relational graph composed of entities and relations which are regarded as nodes and different types of edges. Formally, a KG is defined as $\mathcal{G} = (\mathcal{U}, \mathcal{E}, \mathcal{R})$, where \mathcal{U} is the set of entity nodes and $\mathcal{E} \subseteq \mathcal{U} \times \mathcal{R} \times \mathcal{U}$ is the set of typed edges between nodes in \mathcal{U} with a certain relation in the relation schema \mathcal{R} .

Then given the input/output sequences in the text generation task, a subgraph of the KG which is associated with the sequences can be defined as below.

Definition 4.2 (Sequence-associated K-hop subgraph). A sequence-associated K-hop subgraph is defined as $\mathcal{G}_{sub} = (\mathcal{U}_{sub}, \mathcal{E}_{sub}, \mathcal{R})$, where \mathcal{U}_{sub} is the union of the set of entity nodes mapped through an *entity linking* function $\psi : \mathcal{U} \times \mathcal{X} \rightarrow \mathcal{U}_{sub}$ and their neighbors within K -hops. Similarly, $\mathcal{E}_{sub} \subseteq \mathcal{U}_{sub} \times \mathcal{R} \times \mathcal{U}_{sub}$ is the set of typed edges between nodes in \mathcal{U}_{sub} .

Sequence-associated subgraph provides a graphical form of the task data (i.e., sequences) and thus enables the integration of KGs and the sequences into graph algorithms.

Many methods have been proposed to learn the relationship between KG semantics and input/output sequences. They can be categorized into four methodologies as shown in Figure 6: (M1) incorporate knowledge graph embeddings into language generation; (M2) transfer knowledge into language model with triplet information; (M3) perform reasoning over knowledge graph via path finding strategies; and (M4) improve the graph embeddings with graph neural networks.

4.2.1 M1: Incorporate Knowledge Graph Embeddings into Language Generation. Knowledge graph embedding (KGE) techniques learn node embedding from a KG [131]. KGE aims to capture the semantic relatedness between entity nodes from their connectivity information (i.e., different types of relations) in the KG. The primary idea is to represent entities and relations in a low-dimensional vector space \mathbb{R}^d , where $d \ll |\mathcal{U} \cup \mathcal{R}|$, to reduce data dimensionality while preserving the inherent structure of the KG. TransE [11] is the most widely used KGE technique. In TransE, given a KG edge (u_i, r, u_j) , the relation is seen as a translation vector \mathbf{r} so that the embedded entities \mathbf{u}_i and \mathbf{u}_j can be connected with low translation error, namely $\mathbf{u}_i + \mathbf{r} \approx \mathbf{u}_j$. For example, we have $\overrightarrow{\text{Tokyo}} + \overrightarrow{\text{IsCapitalOf}} \approx \overrightarrow{\text{Japan}}$ for the knowledge edge $(\text{Tokyo}, \text{IsCapitalOf}, \text{Japan})$. As shown in Figure 6(a), a common strategy of incorporating KGE into NLG is to concatenate the original word representations (\mathbf{x}) with the corresponding entity representations (\mathbf{u}) from KGE [151, 159].

4.2.2 M2: Transfer Knowledge into Language Model with Knowledge Triplet Information. The vector spaces of entity embeddings (from KGE) and word embeddings (from pre-trained language models) are usually inconsistent [80]. Beyond a simple concatenation, recent methods

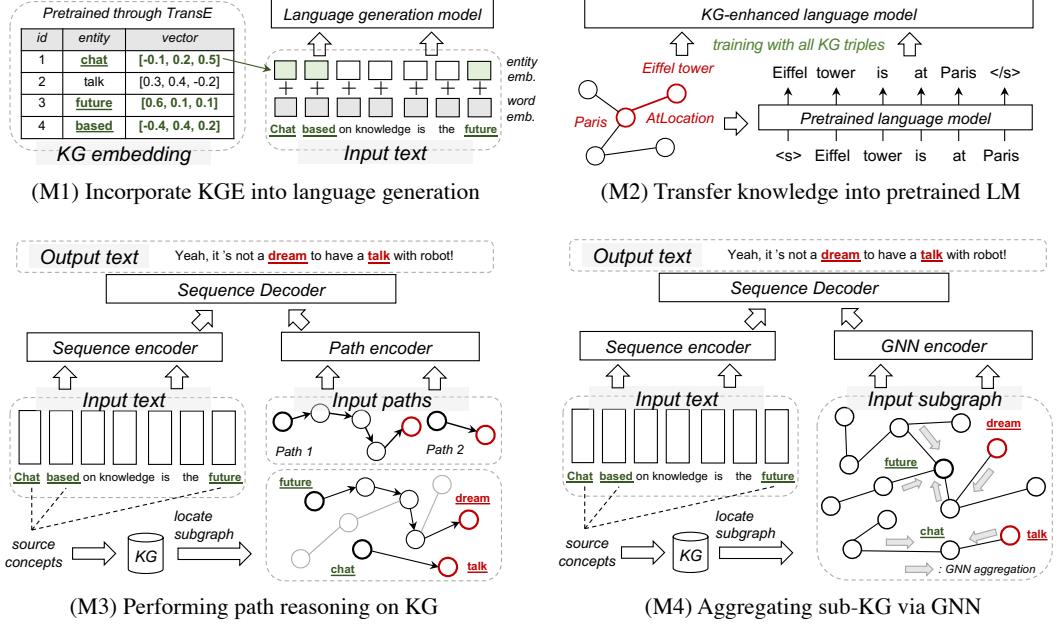


Fig. 6. Four typical methodologies for incorporating KG semantics into text generation.

have explored to fine-tune the language models directly on knowledge graph triplets. Guan et al. transformed the commonsense triplets (in ConceptNet and ATOMIC) into readable sentences using templates, as illustrated in Figure 6(b). And then the language model (e.g., GPT-2) is fine-tuned on the transformed sentences to learn the commonsense knowledge to improve text generation.

4.2.3 M3: Perform Reasoning over Knowledge Graph via Path Finding Strategies. KGE learns node representations from one-hop relations through a certain semantic relatedness (e.g. TransE). However, Xiong et al. argued that an intelligent machine is supposed to be able to conduct explicit reasoning over relational paths to make multiple inter-related decisions rather than merely embedding entities in the KGs [141]. Take the QA task an example. The machine performs reasoning over KGs to handle complex queries that do not have an obvious answer, infer potential answer-related entities, and generate the corresponding answer. So, the challenge lies in identifying a subset of desired entities and mentioning them properly in a response [91]. Because the connected entities usually follow natural conceptual threads, they help generate reasonable and logical answers to keep conversations engaging and meaningful. As shown in Figure 6(c), path-based methods explore various patterns of connections among entity nodes such as meta-paths and meta-graphs. Then they learn from walkable paths on KGs to provide auxiliary guidance for the generation process. The path finding based methods can be mainly divided into two categories: (1) path ranking based methods and (2) reinforcement learning (RL) based path finding methods.

M3.1: Path routing and ranking. Path ranking algorithm (PRA) emerges as a promising method for learning and inferring paths on large KGs [65]. PRA uses random walks to perform multiple bounded depth-first search processes to find relational paths. Coupled with elastic-net based learning [166], PRA picks plausible paths and prunes non-ideal, albeit factually correct KG paths. For example, Tuan et al. proposed a neural conversation model with PRA on dynamic knowledge graphs [121]. In the decoding phase, it selected an output from two networks, a general GRU decoder network and a PRA based multi-hop reasoning network, at each time step. Bauer et

al. ranked and filtered paths to ensure both the information quality and variety via a 3-step scoring strategy: initial node scoring, cumulative node scoring, and path selection [5]. Ji et al. heuristically pruned the noisy edges between entity nodes and proposed a path routing algorithm to propagate the edge probability along multi-hop paths to the entity nodes [56].

M3.2: Reinforcement learning based path finding. Reinforcement learning (RL) based methods make an agent to perform reasoning to find a path in a continuous space. These methods incorporate various criteria in their reward functions of path finding, making the path finding process flexible. Xiong et al. proposed DeepPath, the first work that employed Markov decision process (MDP) and used RL based approaches to find paths in KGs [141]. Leveraging RL based path finding for NLG tasks typically consists of two stages [81, 97]. First, they take a sequence as input, retrieve a starting node u_0 on \mathcal{G} , then perform multi-hop graph reasoning, and finally arrive at a target node u_k that incorporates the knowledge for output sequence generation. Second, they represent the sequence X and selected path $\Phi_k(u_0, u_k)$ through two separate encoders. They decode a sequence with multi-source attentions on the input sequence and selected path. Path-based knowledge graph reasoning converts the graph structure of a KG into a linear path structure that can be easily represented by sequence encoders (e.g., RNN) [30, 97, 121]. For example, Niu et al. encoded selected path and input sequence with two separate RNNs and generated sequence with a general attention-based RNN decoder [97]. To enhance the RL process, Xu et al. proposed six reward functions for training an agent in the reinforcement learning process. For example, the functions looked for accurate arrival at the target node as well as the shortest path between the start and target node, i.e., minimize the length of the selected path $\Phi_k(u_0, u_k)$ [142].

4.2.4 M4: Improve the Graph Embeddings with Graph Neural Networks. The contexts surrounding relevant entities on KGs play an important role in understanding the entities and generating proper text about their interactions [46, 63]. For example, in scientific writing, it is important to consider the neighboring nodes of relevant concepts on a taxonomy and/or the global context of a scientific knowledge graph [63]. However, neither KGE nor relational path could fully represent such information. Graph-based representations aim at aggregating the context/neighboring information on graph data; and recent advances of GNN models demonstrate a promising advancement in graph-based representation learning [138]. In order to improve text generation, graph-to-sequence (Graph2Seq) models encode the structural information of the KG in a neural encoder-decoder architecture [6]. Since then, GNNs have been playing an important role in improving the NLG models. They have been applied to both *encoding* and *decoding* phases.

Learning KG-aware input text representation with GNNs (Encoding). For encoding phase, a general process of leveraging GNNs for incorporating KG is to augment semantics of a word in the input text by combining with the vector of the corresponding entity node vector to the word on the KG [46, 54, 150, 151, 159]. A pre-defined entity linking function $\psi : \mathcal{U} \times \mathcal{X} \rightarrow \mathcal{U}_{sub}$ maps words in the input sequence to entity nodes on the KG. Given an input sequence, all the linked entities and their neighbors within K -hops compose a *sequence-associated K-hop subgraph* \mathcal{G}_{sub} (formally defined in Definition 4.2). For each entity node in \mathcal{G}_{sub} , it uses the KG structure as well as entity and edge features (e.g., semantic description if available) to learn a representation vector \mathbf{u} . Specifically, a GNN model follows a neighborhood aggregation approach that iteratively updates the representation of a node by aggregating information from its neighboring nodes and edges. After k iterations of aggregation, the node representation captures the structural information within its k -hop neighborhood. Formally, the k -th layer of a node $u \in \mathcal{U}_{sub}$ is:

$$\mathbf{u}^{(k)} = \text{COMBINE}_k(\mathbf{u}^{(k-1)}, \text{AGGREGATE}_k(\{(\mathbf{u}_i^{(k-1)}, \mathbf{e}_{ij}^{(k-1)}, \mathbf{u}_j^{(k-1)}) : \forall (u_i, e_{ij}, u_j) \in \mathcal{N}(u)\})). \quad (22)$$

Table 5. Tasks, datasets and KG sources used in different KG-enhanced papers. We also compared the performance of different models before and after incorporating KG into the generation process, in which “w/o KG” performance comes from the best baseline method; “with KG” comes from the KG-enhanced method.

Tasks	Methods	Ref.	Cat.	Dataset Information		Effect of KG			KG source
				Name	#Instance	w/o KG	with KG	ΔBLEU	
Common-sense reasoning	KG-BART	[80]	M4	CommonGen	77,449	28.60	30.90	+2.30	ConceptNet
	CE-PR	[56]	M3	ComVE	30,000	15.70	17.10	+1.60	ConceptNet
	GRF	[57]	M4	αNLG-ART	60,709	9.62	11.62	+2.00	ConceptNet
	MGCN	[21]	M3	EntDesc	110,814	24.90	30.00	+4.30	Self-built KG
Story generation	IE+MSA	[46]	M4	ROCStories (split-1)	98,162	8.25	9.36	+1.11	ConceptNet
	GRF	[57]	M4	ROCStories (split-2)		10.40	11.00	+0.60	ConceptNet
	KEPM	[45]	M2	ROCStories (split-2)	98,162	14.10	14.30	+0.20	ConceptNet & ATOMIC
	MRG	[156]	M3	VisualStory	50,000	3.18	3.23	+0.05	ConceptNet
Scientific writing	GraphWriter	[63]	M4	AGENDA	40,000	12.20	14.30	+1.90	Self-built KG
	PaperRobot	[130]	M4	PaperWriting	27,001	9.20	13.00	+3.80	Self-built KG
Dialogue system	ConceptFlow	[151]	M4	Reddit-10M	3,384K	1.62	2.46	+0.84	ConceptNet
	AKGCM	[81]	M3	EMNLP dialog	43,192	32.45	30.84	-1.61	Self-built KG
	AKGCM	[81]	M3	ICLR dialog	21,569	6.74	6.94	+0.20	Self-built KG
Question answering	MHPGM	[5]	M3	NarrativeQA	46,765	19.79	21.07	+1.28	Self-built KG

Table 6. Qualitative comparison between different KG-enhanced methods.

Methods	Ref.	Method category				Multi-hop info. aggregation	Multi-hop path reasoning	Auxiliary (knowledge related) task(s)
		M1	M2	M3	M4			
THOTH	[92]	✓				✗	✗	✗
CCM	[159]				✓	✗, one-hop	✗	✗
KEPM	[45]		✓			✗	✗	✗
AKGCM	[81]			✓		✗	✓, Markov decision	✓, Path selection
IE+MSA	[46]				✓	✓, by GNN	✗	✗
ConceptFlow	[151]				✓	✓, by GNN	✗	✗
CE-PR	[56]			✓		✗	✓, Path routing	✓, Concept selection
GRF	[57]				✓	✓, by GNN	✓, Path scoring	✓, Link prediction

The sub-graph representation \mathbf{h}_{subG} is learned thorough a READOUT(\cdot) function from all entity node representations (i.e., $\mathbf{h}_{subG} = \text{READOUT}(\{\mathbf{u}^{(k)}, u \in \mathcal{U}_{sub}\})$). Zhou et al. was the first to design such a knowledge graph interpreter to enrich the context representations with neighbouring concepts on ConceptNet using graph attention network (GAT) [159].

Dynamically attending KG representation (Decoding). The sequence decoder uses attention mechanism to find useful semantics from the representation of KG as well as the hidden state of the input text, where the KG’s representation is usually generated by GNNs. Specially, the hidden state is augmented by subgraph representation \mathbf{h}_{subG} , i.e., $\mathbf{s}_0 = \mathbf{h}_n \oplus \mathbf{h}_{subG}$ [6]. Then, the decoder attentively reads the retrieved subgraph to obtain a graph-aware context vector. Then it uses the vector to update the decoding state [46, 57, 80, 151, 159]. It adaptively chooses a generic word or an entity from the retrieved subgraph to generate output words. Because graph-level attention alone might overlook fine-grained knowledge edge information, some recent methods adopted the hierarchical graph attention mechanism [46, 80, 159]. It attentively read the retrieved subgraph \mathcal{G}_{sub} and then attentively read all knowledge edges \mathcal{E}_{sub} involved in \mathcal{G}_{sub} . Ji et al. added a relevance score that reflected the relevancy of the knowledge edge according to the decoding state [57].

4.2.5 Discussion and Analysis of the Methodologies and Methods.

Pros and cons. Knowledge graph embedding (**M1**) was the earliest attempt to embed components of a KG including entities and relations into continuous vector spaces and use them to improve text generation. Those entity and relation embeddings can simply be used to enrich input text representations (e.g., concatenating embeddings), bridging connections between entity words linked from input text in latent space. Because the graph projection and text generation are performed as two separate steps, the embedding vectors from knowledge graph and the hidden states from input text were in two different vector spaces. The model would have to learn to bridge the gap, which might make a negative impact on the performance of text generation.

Fine tuning pre-trained language models on the KG triplets (**M2**) can eliminate the gap between the two vector spaces. Nevertheless, M1 and M2 share two drawbacks. First, they only preserve information of direct (one-hop) relations in a KG, such as pair-wise proximity in M1 and KG triplet in M2, but ignore the indirect (multi-hop) relations of concepts. The indirect relations may provide plausible evidence of complex reasoning for some text generation tasks. Second, from the time KGs were encoded in M1 or M2 methods, the generation models would no longer be able to access the KGs but their continuous representations. Then the models could not support reasoning like commonsense KG reasoning for downstream tasks. Due to these two reasons, M1 and M2 were often used to create basic KG representations upon which the KG path reasoning (**M3**) and GNNs (**M4**) could further enrich the hidden states [151, 159].

The path finding methods of KG reasoning (**M3**) perform multi-hop walks on the KGs beyond one-hop relations. It enables reasoning that is needed in many text generation scenarios such as commonsense reasoning and conversational question answering. At the same time, it provides better interpretability for the entire generation process, because the path selected by the KG reasoning algorithm will be explicitly used for generation. However, the selected paths might not be able to capture the full contexts of the reasoning process due to the limit of number. Besides, reinforcement-learning based path finding uses heuristic rewards to drive the policy search, making the model sensitive to noises and adversarial examples.

The algorithms of GNN and Graph2Seq (**M4**) can effectively aggregate semantic and structural information from multi-hop neighborhoods on KGs, compared to M3 that considers multi-hop paths. Therefore, the wide range of relevant information can be directly embedded into the encoder/decoder hidden states. Meanwhile, M4 enables back propagation for jointly optimizing text encoder and graph encoder. Furthermore, the attention mechanism that has been applied in GNN and Graph2Seq (e.g., graph attention) can explain the model’s output at some extent, though the multi-hop paths from M3 has better interpretability.

M3 and M4 are able to use multi-hop relational information, compared to M1 and M2. However, they have two weak points. First, they have higher complexity than M1 and M2. In M3, the action space of path finding algorithms can be very large due to the large size and sparsity of the knowledge graph. In M4, the decoder has to attentively read both input sequence and knowledge graph. Second, the subgraphs retrieved by M3 and M4 might provide low coverage of useful concepts for generating the output. For example, people use ConceptNet, a widely used commonsense KG, to retrieve the subgraph on three generative commonsense reasoning tasks. The task datasets are ComVE [57], α -NLG [7], and ROCStories [46]. We found 25.1% / 24.2% / 21.1% of concepts in the output could be found on ConceptNet, but only 11.4% / 8.1% / 5.7% of concepts in the output can be found on the retrieved 2-hop sequence-associated subgraph, respectively. It means that a large portion of relevant concepts on the KG are not utilized in the generation process.

Quantitative analysis. Table 5 summarizes tasks, datasets, and KG sources used in existing KG-enhanced works. Three important things should be mentioned. First, all the datasets in the table are public, and we include their links in Table 11. CommonGen [76], ComVE [124] and α -NLG [7]

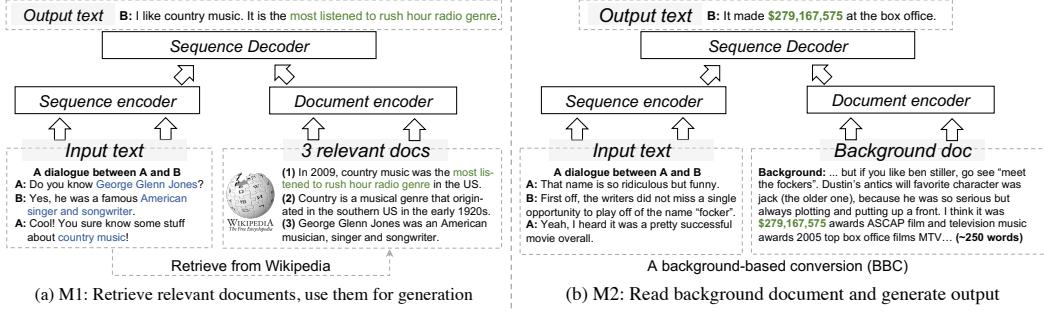


Fig. 7. The left figure demonstrates retrieving relevant documents, then using them for generation; the right figure demonstrate reading background document to conduct conversions.

have a public leaderboard for competition. Second, for KG sources, we observe that eight (57.1%) papers use ConceptNet as external resource, while six (42.9%) papers constructed their own KGs from domain-specific corpus. For example, Koncel et al. created a scientific knowledge graph by applying the SciIE tool (science domain information extraction) [63]. Besides, Zhao et al. compared the performance of models between using ConceptNet and using a self-built KG, and found the model with self-built KG could work better on story generation and review generation tasks [156]. Third, we observed that KG-enhanced NLG methods made the largest improvement on generative commonsense reasoning tasks, in which the average improvement is +2.55% in terms of Δ BLEU, while the average improvement on all different tasks is +1.32%.

Qualitative analysis. Table 6 compares different KG-enhanced methods from three dimensions: multi-hop information aggregation, multi-hop path reasoning, and auxiliary knowledge graph related tasks. M3 is commonly used for multi-hop path reasoning and M4 is used for multi-hop information aggregation, except that CCM [159] only aggregates one-hop neighbors. Besides, the auxiliary KG-related tasks are often used to further help the model learn knowledge from the KG. For example, ablation studies in [56, 57, 81] show that the tasks of path selection, concept selection and link prediction can further boost the generation performance. GRF [57] learns these three abilities at the same time. It achieves the state-of-art performance on three generation tasks.

4.3 NLG enhanced by Grounded Text

Knowledge grounded text refers to textual information that can provide additional knowledge relevant to the input sequence. The textual information may not be found in training corpora or structured databases, but can be obtained from massive textual data from online resources. These online resources include encyclopedia (e.g., Wikipedia), social media (e.g., Twitter), shopping websites (e.g., Amazon reviews). Knowledge grounded text plays an important role in understanding the input sequence and its surrounding contexts. For example, Wikipedia articles may offer textual explanations or background information for the input text. Amazon reviews may contain necessary descriptions and reviews needed to answer a product-related question. Tweets may contain people's comments and summaries towards an event. Therefore, knowledge grounded text is often taken as an important external knowledge source to help with a variety of NLG applications.

Next, we introduce popular NLG applications enhanced by knowledge grounded text:

- **Dialogue system.** Building a fully data-driven dialogue system is difficult since most of the universal knowledge is not presented in the training corpora [42]. The lack of universal knowledge considerably limits the appeal of fully data-driven generation methods, as they are bounded to respond evasively or defectively and seldom include meaningfully factual

contents. To infuse the response with factual information, an intelligent machine is expected to obtain necessary background information to produce appropriate response.

- **Summarization.** Seq2Seq models that purely depend on the input text tend to “lose control” sometimes. For example, 3% of summaries contain less than three words, and 4% of summaries repeat a word for more than 99 times as mentioned in [13]. Furthermore, Seq2Seq models usually focus on copying source words in their exact order, which is often sub-optimal in abstractive summarization. Therefore, leveraging summaries of documents similar as the input document as templates can provide reference for the summarization process [13, 129].
- **Question answering (QA).** It is often difficult to generate proper answers only based on the given question. For example, without knowing any information of an Amazon product, it is hard to deliver satisfactory answer to the user questions such as “*Does the laptop have a long battery life?*” or “*Is this refrigerator frost-free?*” So, the product description and customer reviews can be used as a reference for answering product-related questions [9, 16].

To handle different kinds of relationships between grounded text and input/output sequences, these methods can be categorized into two methodologies as shown in Figure 7: (M1) guiding generation with retrieved information; (M2) modeling background knowledge into response generation.

4.3.1 M1: Guiding Generation with Retrieved Information. Because knowledge grounded text is not presented in the training corpora, an idea is to retrieve relevant textual information (e.g., a review, a relevant document, a summary template) from *external sources* based on the input text and to incorporate the retrieved grounded text into the generation process. This process is similar to designing knowledge acquisition and incorporation of KBs and KGs in text generation tasks. The difference is that ground text is unstructured and noisy. So, researchers design knowledge selection and incorporation methods to address the challenges. Based on the number of stages, we further divide related methods into two categories: retrieve-then-generate (also known as retrieval-augmented generation, short as RAG, in many existing papers [64, 68, 99]) methods (2-stage methods) and retrieve, rerank and rewrite methods (3-stage methods).

M1.1: Retrieval-augmented generation (RAG). RAG follows a two-stage process: retrieval and generation. Specially, as shown in Figure 7(a), a retriever $p(Z|X)$ first returns (usually top-K truncated) distributions over text passages given a query X , and then a generator $p(y_i|X, Z, y_{1:i-1})$ generates a current token based on a context of the previous tokens $y_{1:i-1}$, the original input X and a retrieved passage Z . Methods for retrieving fact or review snippets are various, including matching from a collection of raw text entries indexed by named entities [42]; scoring relevant documents within a large collection by statistical approaches such as BM25 [27], or neural-based retrieval approaches such as dense paragraph retrieval (DPR) [68]. For training the retriever and generator, most of existing work has jointly optimized these two components, without any direct supervision on what document should be retrieve [64, 68]. However, by asking human experts to label what document should be retrieved and adding the retrieval loss (resulting in a multi-task learning setting), the generation performance can be greatly improved [27, 61], though the labelling process is an extremely time-consuming and labor-intensive task.

Ghazvininejad et al. proposed a knowledge grounded neural conversation model (KGNCM), which is the first work to retrieve review snippets from Foursquare and Twitter. Then it incorporates the snippets into dialogue response generation [42]. It uses an end-to-end memory network [116] to generate responses based on the selected review snippets. Lewis et al. introduced a general retrieval-augmented generation (RAG) framework by leveraging a pre-trained neural retriever and generator. It can be easily fine-tuned on downstream tasks, and it has demonstrated state-of-the-art performance on various knowledge intensive NLG tasks [68]. Recently, the fusion-in-decoder

Table 7. Tasks, datasets and evidence sources used in retrieve-then-generate (M1) papers. We also include their document(d)/sentence(s) retrieval space and the number of retrieved document(d)/sentence(s).

Evidence sources	Tasks	Methods	Ref.	Dataset Information		Retrieval space (d/s)	# Retrieved d/s
				Name	#Instance		
Wikipedia	Dialogue system	MemNet SKT	[27] [61]	Wizard of Wikipedia (WoW)	22,311	5.4M/93M	7 7
	Question answering	RAG	[68]	MS-MARCO	267,287	21M/-	10
		BART+DPR RT+C-REALM	[99] [64]	ELI5	274,741	3.2M/- 3.2M/-	- 7
	Argument generation	H&W CANDELA	[53] [52]	ChangeMyView	287,152	5M/- 5M/-	10 10
Online platform (e.g., Amazon)	Dialogue (for business)	AT2T KGNCM	[62] [42]	Amazon books Foursquare	937,032 1M	-/131K -/1.1M	10 10
Gigawords	Summarization	R ³ Sum BiSET	[13] [129]	Gigawords	3.8M	-/3.8M -/3.8M	30 30

Table 8. Qualitative comparison between different grounded text enhanced methods.

Methods	Ref.	Method category			Retrieval supervision	Retriever pre-training	Number of stages
		M1.1	M1.2	M2			
MemNet	[27]	✓			✓, Human annotated labels	✗	2
SKT	[61]	✓			✓, Human annotated labels	✗	2
R ³ Sum	[13]		✓		✓, Pseudo labels	✗	3, with rerank
BiSET	[129]		✓		✓, Pseudo labels	✗	3, with rerank
RefNet	[87]			✓	✗	✗	1, no retrieval
GLKS	[107]			✓	✗	✗	1, no retrieval
RAG	[68]	✓			✗	✓, DPR	2
Kilt	[99]	✓			✗	✓, DPR	2
RT+C-REALM	[64]	✓			✗	✓, REALM	2

methods (i.e., the decoder performs attention over the concatenation of the resulting representations of all retrieved passages [72, 145]) could even outperform RAG as reported in KILT benchmark [99].

M1.2: Retrieve, rerank and rewrite (R³). Different from RAG, a R³-based method is expected to retrieve a most precise reference document that can be directly used for rewriting/editing. R³-based method has proved successful in a number of NLG tasks such as machine translation [44], and summarization [13, 129]. In summarization, Seq2Seq models that purely depend on the input document to generate summaries tend to deteriorate with the accumulation of word generation, e.g., they generate irrelevant and repeated words frequently [13, 129]. Template-based summarization assume the golden summaries of the similar sentences (i.e., templates) can provide a reference point to guide the input sentence summarization process [13, 129]. These templates are often called *soft templates* in order to distinguish from the traditional rule-based templates. Soft template-based summarization typically follows a three-step design: retrieve, rerank, and rewrite. The step of retrieval aims to return a few candidate templates from a summary collection. The reranking identifies the best template from the retrieved candidates. And the rewriting leverages both the source document and template to generate more faithful and informative summaries.

Difference between RAG and R³. Compared with R³-based methods, RAG-based have several differences, including less of emphasis on lightly editing a retrieved item, but on aggregating content from several pieces of retrieved content, as well as learning latent retrieval, and retrieving evidence documents rather than related training pairs.

4.4 M2: Modeling Background Knowledge into Response Generation

Background document, with more global and comprehensive knowledge, has been often used for generating informative responses and ensuring a conversation to not deviate from its topic. Keeping a conversation grounded on a background document is referred as background based conversation (BBC) [9, 90]. Background knowledge plays an important role in human-human conversations. For example, when talking about a movie, people often recall important points (e.g., a scene or review about the movie) and appropriately mention them in the conversation context. Therefore, an intelligent NLG model is expected to find an appropriate background snippet and generate response based on the snippet. As shown in Figure 7(b), the task of BBC is often compared with machine reading comprehension (MRC), in which a span is extracted from the background document as a response to a question [105]. However, since BBC needs to generate natural and fluent responses, the challenge lies in not only locating the right semantic units in the background, but also referring to the right background information at the right time in the right place during the decoding phase.

As MRC models tie together multiple text segments to provide a unified and factual answer, many BBC models use the same idea to connect different pieces of information and find the appropriate background knowledge based on which the next response is to be generated [87, 101]. For instance, Qin et al. proposed an end-to-end conversation model that jointly learned response generation together with on-demand machine reading [101]. The MRC models can effectively encode the input utterance by treating it as a question in a typical QA task (e.g., SQuAD [105]) and encode the background document as the context. Then, they took the utterance-aware background representation as input into decoding phase.

4.4.1 Discussion and Analysis of Different Methods.

Pros and cons. For M1, guiding generation with retrieved information explicitly exposes the role of world knowledge by asking the model to decide what knowledge to retrieve and use during language generation. Since retrieval-augmented generation (RAG) captures knowledge in a interpretable and modular way, it is often used for knowledge-intensive tasks such as long-form QA and argument generation. However, a knowledge retriever is expected to retrieve documents from a large-scale corpus, e.g., the entire Wikipedia, which causes significant computational challenge. Besides, one input often requires retrieved text whose amount is much larger than the input itself (as indicated in Table 7), leading to serious information overwhelming for the generation model.

For M2, background based conversations (BBCs) avoid generating generic responses in a dialogue system and are able to generate more informative responses by exploring related background information. However, existing methods still cannot solve inherent problems effectively, such as tending to break a complete semantic unit and generate shorter responses [87].

Qualitative analysis. Table 7 summarizes tasks, datasets and evidence sources used in existing grounded text enhanced work. Three important things should be mentioned. First, all the datasets in the table are public, and we include their links in Table 11. Second, Wikipedia is the most commonly used evidence source since it is the largest free online encyclopedia. Besides, some online platforms contain plenty of product-related textual information, e.g., product reviews on Amazon, which are often used to build up task/goal oriented dialogue systems for business purpose. Third, the retrieval space of candidate documents are usually larger than 1 million and only 7-10 documents are selected. So, the process of retrieving relevant documents is challenging.

Table 8 compares different grounded text enhanced methods from three dimensions: retrieval supervision, pre-training of the retriever, and number of stages. First, as mentioned above, retrieving relevant documents from a large candidate set is a challenging task. To improve the retrieval accuracy, four (57.1%) papers added the retrieval supervision either by human annotated labels

or pseudo labels, resulting in a multi-task learning setting. Besides, three (42.9%) papers used pre-trained language models to produce document representation for better retrieval. Though existing work has greatly improved the retrieval accuracy, the performance is still far from satisfactory in many text generation tasks [64, 68]. How to learn mutually enhancement between retrieval and generation is still a promising direction in the grounded text enhanced text generation systems.

5 BENCHMARK, TOOLKIT AND LEADERBOARD PERFORMANCE

The development of general evaluation benchmarks for text generation helps to promote the development of research in related fields. Existing text generation benchmarks did not specially focus on choosing the tasks and datasets that have been widely used for knowledge-enhanced text generation. Therefore, we re-screened from the existing four text generation benchmarks, i.e., GLGE [77], GEM [40], Kilt [99], GENIE [60], and determined 9 benchmark datasets for evaluating knowledge-enhanced NLG methods. Here is our criteria for selection:

- We only consider benchmark datasets that have open-access downloading link.
- We focus on diverse text generation tasks, involving various applications.
- We select at most three benchmark datasets for each text generation task.
- We include a mix of internal and external knowledge focused datasets.
- We prefer multi-reference datasets for robust automatic evaluation.

Based on the benchmark selection criteria, we finalize 9 knowledge-centric tasks that covers various NLG tasks, including commonsense reasoning, text summarization, question generation, generative question answering, and dialogue. The data statistics is shown in Table 9. Descriptions and dataset links are listed as follows:

- **Wizard of Wikipedia (WOW):** It is an open-domain dialogue dataset, where two speakers conduct an open-ended conversation that is directly grounded with knowledge retrieved from Wikipedia. (Data link: https://parl.ai/projects/wizard_of_wikipedia/)
- **CommonGen:** It is a generative commonsense reasoning dataset. Given a set of common concepts, the task is to generate a coherent sentence describing an everyday scenario using these concepts. (Data link: <https://inklab.usc.edu/CommonGen/>)
- **α NLG-ART:** It is a generative commonsense reasoning dataset. Given the incomplete observations about the world, the task is to generate a valid hypothesis about the likely explanations to partially observable past and future. (Data link: <http://abductivecommonsense.xyz/>)
- **ComVE:** It is a generative commonsense reasoning dataset. The task is to generate an explanation given a counterfactual statement for sense-making. (Data link: <https://github.com/wangcunxiang/SemEval2020-Task4-Commonsense-Validation-and-Explanation>)
- **ELI5:** It is a dataset for long-form question answering. The task is to produce explanatory multi-sentence answers for diverse questions. Web search results are used as evidence documents to answer questions. (Data link: <https://facebookresearch.github.io/ELI5/>)
- **SQuAD:** It is a dataset for answer-aware question generation. The task is to generate a question asks towards the given answer span based on a given text passage or document. (Data link: <https://github.com/magic282/NQG>)
- **CNN/DailyMail (CNN/DM):** It is a dataset for summarization. Given a news articles, the goal is to produce a summary that represents the most important or relevant information within the original content. (Data link: https://www.tensorflow.org/datasets/catalog/cnn_dailymail)
- **Gigaword:** It is a dataset for summarization. Similar with CNN/DM, the goal is to generate a headline for a news article. (Data link: <https://www.tensorflow.org/datasets/catalog/gigaword>)

Table 9. We choose 9 knowledge-enhanced NLG benchmark datasets. These datasets have been included in four existing general NLG benchmarks (i.e., GLGE [77], GEM [40], Kilt [99], GENIE [60]) or in SemEval tasks.

Tasks	Ref.	Dataset Information				Leader board	In which NLG benchmark	Papers including this dataset
		Name	#Train	#Dev.	#Test			
Dialogue system	[27]	Wizard of Wikipedia	18,430	1,948	1,933	✓*	Kilt	[27, 61, 74]
	[154]	PersonaChat	122,499	14,602	14,056	✗	GLGE	[27, 74]
Question answering	[31]	ELI5	272,634	1,507	600	✓‡	Kilt	[64, 99]
Question generation	[105]	SQuAD	75,722	10,570	11,877	✗	GLGE	[19, 22, 133]
Commonsense reasoning	[76]	CommonGen	67,389	4,018	6,042	✓§	GEM	[32, 80, 126]
	[7]	αNLG-ART	50,481	7,252	2,976	✓¶	GENIE	[7, 57]
	[124]	ComVE	25,596	1,428	2,976	✓¶	SemEval	[56, 57]
Summarization	[109]	CNN/DM	287,226	13,368	11,490	✓**	GLGE	[33, 41, 163]
	[109]	Gigaword	3.8M	189K	1,951	✓††	GLGE	[13, 59, 69]

- **PersonaChat**: It is an open-domain dialogue dataset. It presents the task of making chit-chat more engaging by conditioning on profile information. (Data link: <https://github.com/facebookresearch/ParlAI/tree/master/projects/personachat>)

6 DISCUSSION ON FUTURE DIRECTIONS

Many efforts have been conducted to tackle the problem of knowledge-enhanced text generation and its related applications. To advance the field, there remains several open problems and future directions. Designing more effective ways to represent knowledge and integrate them into the generation process is still the most important trend in knowledge-enhanced NLG systems. From a broader perspective, we provide three directions that make focusing such efforts worthwhile now: (i) incorporating knowledge into visual-language generation tasks, (ii) learning knowledge from broader sources, especially pre-trained language models, (iii) learning knowledge from limited resources, (iv) learning knowledge in a continuous way.

6.1 Incorporate Knowledge into Visual-Language Generation Tasks

Beyond text-to-text generation tasks, recent years have witnessed a growing interest in visual-language (VL) generation tasks, such as describing visual scenes [49], and answering visual-related questions [85]. Although success has been achieved in recent years on VL generation tasks, there is still room for improvement due to the fact that image-based factual descriptions are often not enough to generate high-quality captions or answers [162]. External knowledge can be added in order to generate attractive image/video captions. We observed some pioneer work has attempted to utilize external knowledge to enhance the image/video captioning tasks. For example, Tran et al. proposed to detect a diverse set of visual concepts and generate captions by using an external knowledge base (i.e., Freebase), in recognizing a broad range of entities such as celebrities and landmarks [120]. Zhou et al. used a commonsense knowledge graph (i.e., ConceptNet), to infer a

*https://parl.ai/projects/wizard_of_wikipedia

†<https://nikitacs16.github.io/holl-e-website/>

‡<https://facebookresearch.github.io/ELI5/>

§<https://inklab.usc.edu/CommonGen/leaderboard.html>

¶<https://leaderboard.allenai.org/genie-anlg/submissions/public>

||<https://competitions.codalab.org/competitions/21080#results>

**<https://paperswithcode.com/sota/document-summarization-on-cnn-daily-mail>

††<https://paperswithcode.com/sota/text-summarization-on-gigaword>

set of terms directly or indirectly related to the words that describe the objects found in the scene by the object recognition module [162]. In addition, Mao et al. [85] proposed a neuro-symbolic learner for improving visual-language generation tasks (e.g., visual question answering).

However, existing approaches for knowledge-enhanced visual-language generation tasks still have a lot of space for exploration. Some promising directions for future work include using other knowledge sources, such as retrieving image/text to help solve open-domain visual question answering and image/video captioning tasks; bringing structured knowledge for providing justifications for the captions that they produce, tailoring captions to different audiences and contexts, etc.

6.2 Learning Knowledge from Broader Sources

More research efforts should be spent on learning to discover knowledge more broadly and combine multiple forms of knowledge from different sources to improve the generation process. More knowledge sources can be but not limited to network structure, dictionary and table. For examples, Yu et al. [147] and An et al. [2] augmented the task of scientific papers intention detection and summarization by introducing the citation graph; Yu et al. augmented the rare word representations by retrieving their descriptions from Wiktionary and feed them as additional input to a pre-trained language model [148]. Besides, structured knowledge and unstructured knowledge can play a complementary role in enhancing text generation. To improve knowledge richness, Fu et al. combined both structured (knowledge base) and unstructured knowledge (grounded text) [34].

Leveraging Knowledge from Pre-trained Language Models. Pre-trained language models can learn a substantial amount of in-depth knowledge from data without any access to an external memory, as a parameterized implicit knowledge base [68, 104]. However, as mentioned in [45], directly fine-tuning pre-trained language generation models on the story generation task still suffers from insufficient knowledge by representing the input text thorough a pre-trained encoder, leading to repetition, logic conflicts, and lack of long-range coherence in the generated output sequence. Therefore, discovering knowledge from pre-trained language models can be more flexible, such as knowledge distillation, data augmentation, and using pre-trained models as external knowledge [100]. More efficient methods of obtaining knowledge from pre-trained language models are expected.

6.3 Learning Knowledge from Limited Resources

Most of current NLG research conduct on extensively labelled data to favor model training. However, this is in contrast to many real-world application scenarios, where only a few shots of examples are available for new domains. Limited data resources lead to *limited knowledge* that can be learnt in new domains. For examples, learning topical information of a dialogue occurring under a new domain is difficult since the topic may be rarely discussed before; constructing a syntactic dependency graph of a sequence in a low-resource language is hard since many linguistic features are of great uniqueness. Besides, external knowledge bases are often incomplete and insufficient to cover full entities and relationships due to the human costs of collecting domain-specific knowledge triples. Therefore, quick domain adaptation is an essential task in text generation tasks. One potential route towards addressing these issues is meta-learning, which in the context of NLG means a generation model develops a broad set of skills and pattern recognition abilities at training time, and quickly adapt to a new task given very few examples without retraining the model from scratch. Recently, there has been raising interests in both academia and industry to investigate meta-learning in different NLG tasks. Thus, it is a promising research direction to build efficient meta-learning algorithms that only need a few task-specific fine-tuning to learn the new task quickly. And for knowledge-enhanced text generation, it is of crucial importance to adapt the model quickly on new domains with limited new knowledge (e.g., only a few knowledge triples).

6.4 Learning Knowledge in a Continuous Way

A machine learning is expected to learn continuously, accumulate the knowledge learned in previous tasks, and use it to assist future learning. This research direction is referred as lifelong learning [20]. In the process, the intelligent machine becomes more and more knowledgeable and effective at learning new knowledge. To make an analogy, humans continuously acquire new knowledge and constantly update the knowledge system in the brain. However, existing knowledge-enhanced text generation systems usually do not keep updating knowledge in real time (e.g., knowledge graph expansion). A meaningful exploration of was discussed in [86]. They built a general knowledge learning engine for chatbots to enable them to continuously and interactively learn new knowledge during conversations. Therefore, it is a promising research direction to continuously update knowledge obtained from various information sources, empowering intelligent machines with incoming knowledge and improving the performance on new text generation tasks.

7 CONCLUSIONS

In this survey, we present a comprehensive review of current representative research efforts and trends on knowledge-enhanced text generation, and expect it can facilitate future research. To summarize, this survey aims to answer two questions that commonly appears in knowledge-enhanced text generation: *how to acquire knowledge* and *how to incorporate knowledge to facilitate text generation*. Base on knowledge acquisition, the main content of our survey is divided into three sections according to different sources of knowledge enhancement. Based on knowledge incorporation, we first present general methods of incorporating knowledge into text generation and further discuss a number of specific ideas and technical solutions that incorporate the knowledge to enhance the text generation systems in each section. Besides, we review a variety of text generation applications in each section to help practitioners learn to choose and employ the methods.

ACKNOWLEDGEMENTS

We thank all anonymous reviewers for valuable comments. We also appreciate the suggestions from readers of the pre-print version. We thank Dr. Michael Zeng (Microsoft) and Dr. Nazneen Rajani (Salesforce) for their constructive comments and suggestions. Wenhai Yu and Dr. Meng Jiang's research is supported by National Science Foundation grants IIS-1849816, CCF-1901059, and IIS-2119531. Qingyun Wang and Dr. Heng Ji's research is based upon work supported by Agriculture and Food Research Initiative (AFRI) grant no. 2020-67021-32799/project accession no.1024178 from the USDA National Institute of Food and Agriculture, U.S. DARPA SemaFor Program No. HR001120C0123, DARPA AIDA Program No. FA8750-18-2-0014, and DARPA KAIROS Program No. FA8750-19-2-1004. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copy-right annotation therein.

REFERENCES

- [1] Roe Aharoni and Yoav Goldberg. 2017. Towards String-To-Tree Neural Machine Translation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [2] Chenxin An, Ming Zhong, Yiran Chen, Danqing Wang, Xipeng Qiu, and Xuanjing Huang. 2021. Enhancing scientific papers summarization with citation graph. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference for Learning Representation (ICLR)*.
- [4] Joost Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Sima'an. 2017. Graph Convolutional Encoders for Syntax-aware Neural Machine Translation. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- [5] Lisa Bauer, Yicheng Wang, and Mohit Bansal. 2018. Commonsense for Generative Multi-Hop Question Answering Tasks. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

- [6] Daniel Beck, Gholamreza Haffari, and Trevor Cohn. 2018. Graph-to-Sequence Learning using Gated Graph Neural Networks. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [7] Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Scott Wen-tau Yih, and Yejin Choi. 2020. Abductive commonsense reasoning. In *International Conference for Learning Representation (ICLR)*.
- [8] Bin Bi, Chen Wu, Ming Yan, Wei Wang, Jiangnan Xia, and Chenliang Li. 2019. Incorporating External Knowledge into Machine Reading for Generative Question Answering. In *Conference on Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- [9] Bin Bi, Chen Wu, Ming Yan, Wei Wang, Jiangnan Xia, and Chenliang Li. 2020. Generating Well-Formed Answers by Machine Reading with Stochastic Selector Networks. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [10] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. In *Journal of Machine Learning Research (JMLR)*.
- [11] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [12] Ziqiang Cao, Sujian Li, Yang Liu, Wenjie Li, and Heng Ji. 2015. A novel neural topic model and its supervised extension. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [13] Ziqiang Cao, Wenjie Li, Sujian Li, and Furu Wei. 2018. Retrieve, rerank and rewrite: Soft template based neural summarization. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [14] Ming-Wei Chang, Lev Ratinov, and Dan Roth. 2007. Guiding semi-supervision with constraint-driven learning. In *annual meeting of the association of computational linguistics (ACL)*.
- [15] Kehai Chen, Rui Wang, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2018. Syntax-directed attention for neural machine translation. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [16] Shiqian Chen, Chenliang Li, Feng Ji, Wei Zhou, and Haiqing Chen. 2019. Driven answer generation for product-related questions in e-commerce. In *International Conference on Web Search and Data Mining (WSDM)*.
- [17] Xiaojun Chen, Shengbin Jia, and Yang Xiang. 2020. A review: Knowledge reasoning over knowledge graph. In *Expert Systems with Applications*.
- [18] Xiuyi Chen, Fandong Meng, Peng Li, Feilong Chen, Shuang Xu, Bo Xu, and Jie Zhou. 2020. Bridging the Gap between Prior and Posterior Knowledge Selection for Knowledge-Grounded Dialogue Generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [19] Yu Chen, Lingfei Wu, and Mohammed J Zaki. 2020. Reinforcement learning based graph-to-sequence model for natural question generation. In *International Conference of Learning Representation (ICLR)*.
- [20] Zhiyuan Chen and Bing Liu. 2018. Lifelong machine learning. In *Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool Publishers.
- [21] Liying Cheng, Dekun Wu, Lidong Bing, Yan Zhang, Zhanming Jie, Wei Lu, and Luo Si. 2020. Entdesc: Entity description generation by exploring knowledge graph. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [22] Jaemin Cho, Minjoon Seo, and Hannaneh Hajishirzi. 2019. Mixture Content Selection for Diverse Sequence Generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [23] Sajal Choudhary, Preerna Srivastava, Lyle Ungar, and Joao Sedoc. 2017. Domain aware neural dialog system. In *arXiv preprint arXiv:1708.00897*.
- [24] Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. Plug and play language models: a simple approach to controlled text generation. In *International Conference for Learning Representation (ICLR)*.
- [25] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- [26] P Kingma Diederik, Max Welling, et al. 2014. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [27] Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of wikipedia: Knowledge-powered conversational agents. In *International Conference for Learning Representation (ICLR)*.
- [28] Xiangyu Dong, Wenhao Yu, Chenguang Zhu, and Meng Jiang. 2021. Injecting Entity Types into Entity-Guided Text Generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [29] Mihail Eric and Christopher D Manning. 2017. A Copy-Augmented Sequence-to-Sequence Architecture Gives Good Performance on Task-Oriented Dialogue. In *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- [30] Angela Fan, Claire Gardent, Chloé Braud, and Antoine Bordes. 2019. Using Local Knowledge Graph Construction to Scale Seq2Seq Models to Multi-Document Inputs. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

and International Joint Conference on Natural Language Processing (EMNLP-IJCNLP).

- [31] Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. ELI5: Long Form Question Answering. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [32] Zhihao Fan, Yeyun Gong, Zhongyu Wei, Siyuan Wang, Yameng Huang, Jian Jiao, Xuan-Jing Huang, Nan Duan, and Ruofei Zhang. 2020. An Enhanced Knowledge Injection Model for Commonsense Generation. In *International Conference on Computational Linguistics (COLING)*.
- [33] Xiyan Fu, Jun Wang, Jinghan Zhang, Jinmao Wei, and Zhenglu Yang. 2020. Document Summarization with VHTM: Variational Hierarchical Topic-Aware Mechanism. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [34] Yao Fu and Yansong Feng. 2018. Natural answer generation with heterogeneous memory. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- [35] Kuzman Ganchev, Jennifer Gillenwater, Ben Taskar, et al. 2010. Posterior regularization for structured latent variable models. In *Journal of Machine Learning Research (JMLR)*.
- [36] Ce Gao and Jiangtao Ren. 2019. A topic-driven model for learning to generate diverse sentences. In *Neurocomputing*.
- [37] Cristina Barbacea and Qiaozhu Mei. 2020. Neural Language Generation: Formulation, Methods, and Evaluation. In *arXiv preprint arXiv:2007.15780*.
- [38] Albert Gatt and Emiel Krahmer. 2018. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. In *Journal of Artificial Intelligence Research (JAIR)*.
- [39] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. In *International Conference on Machine Learning (ICML)*.
- [40] Sebastian Gehrmann, Tosin Adewumi, Karmany Aggarwal, Pawan Sasanka Ammanamanchi, Aremu Anuoluwapo, Antoine Bosselut, Khyathi Raghuvi Chandu, Miruna Clinciu, Dipanjan Das, Kaustubh D Dhole, et al. 2021. The gem benchmark: Natural language generation, its evaluation and metrics. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [41] Sebastian Gehrmann, Yuntian Deng, and Alexander M Rush. 2018. Bottom-Up Abstractive Summarization. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [42] Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [43] Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016. Incorporating Copying Mechanism in Sequence-to-Sequence Learning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [44] Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor OK Li. 2018. Search engine guided neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- [45] Jian Guan, Fei Huang, Zhihao Zhao, Xiaoyan Zhu, and Minlie Huang. 2020. A knowledge-enhanced pretraining model for commonsense story generation. In *Transactions of the Association for Computational Linguistics (TACL)*.
- [46] Jian Guan, Yansen Wang, and Minlie Huang. 2019. Story ending generation with incremental encoding and commonsense knowledge. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [47] Dandan Guo, Bo Chen, Ruiying Lu, and Mingyuan Zhou. 2020. Recurrent Hierarchical Topic-Guided RNN for Language Generation. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*.
- [48] Shizhu He, Cao Liu, Kang Liu, and Jun Zhao. 2017. Generating natural answers by incorporating copying and retrieving mechanisms in sequence-to-sequence learning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [49] MD Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratuddin, and Hamid Laga. 2019. A comprehensive survey of deep learning for image captioning. In *ACM Computing Surveys (CSUR)*.
- [50] Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward controlled generation of text. In *International Conference on Machine Learning (ICML)*.
- [51] Zhiting Hu, Zichao Yang, Ruslan Salakhutdinov, Xiaodan Liang, Lianhui Qin, Haoye Dong, and Eric Xing. 2018. Deep Generative Models with Learnable Knowledge Constraints. In *Advances in Neural Information Processing Systems*.
- [52] Xinyu Hua, Zhe Hu, and Lu Wang. 2019. Argument Generation with Retrieval, Planning, and Realization. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [53] Xinyu Hua and Lu Wang. 2018. Neural Argument Generation Augmented with Externally Retrieved Evidence. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [54] Luyang Huang, Lingfei Wu, and Lu Wang. 2020. Knowledge Graph-Augmented Abstractive Summarization with Semantic-Driven Cloze Reward. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [55] Touseef Iqbal and Shaima Qureshi. 2020. The Survey: Text Generation Models in Deep Learning.. In *Journal of King Saud University-Computer and Information Sciences*. Elsevier.
- [56] Haozhe Ji, Pei Ke, Shaohan Huang, Furu Wei, and Minlie Huang. 2020. Generating Commonsense Explanation by Extracting Bridge Concepts from Reasoning Paths. In *Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and International Joint Conference on Natural Language (AAACL-IJCNLP)*.

[57] Haozhe Ji, Pei Ke, Shaohan Huang, Furu Wei, Xiaoyan Zhu, and Minlie Huang. 2020. Language Generation with Multi-Hop Reasoning on Commonsense Knowledge Graph. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

[58] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S Yu. 2020. A survey on knowledge graphs: Representation, acquisition and applications. In *arXiv preprint arXiv:2002.00388*.

[59] Hanqi Jin, Tianming Wang, and Xiaojun Wan. 2020. SemSUM: Semantic Dependency Guided Neural Abstractive Summarization. In *AAAI Conference on Artificial Intelligence (AAAI)*.

[60] Daniel Khashabi, Gabriel Stanovsky, Jonathan Bragg, Nicholas Lourie, Jungo Kasai, Yejin Choi, Noah A Smith, and Daniel S Weld. 2021. Genie: A leaderboard for human-in-the-loop evaluation of text generation. In *arXiv preprint arXiv:2101.06561*.

[61] Byeongchang Kim, Jaewoo Ahn, and Gunhee Kim. 2020. Sequential Latent Knowledge Selection for Knowledge-Grounded Dialogue. In *International Conference for Learning Representation (ICLR)*.

[62] Jihyeok Kim, Seungtaek Choi, Reinald Kim Amplayo, and Seung-won Hwang. 2020. Retrieval-Augmented Controllable Review Generation. In *International Conference on Computational Linguistics (COLING)*.

[63] Rik Koncel-Kedziorski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. Text Generation from Knowledge Graphs with Graph Transformers. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

[64] Kalpesh Krishna, Aurko Roy, and Mohit Iyyer. 2021. Hurdles to Progress in Long-form Question Answering. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

[65] Ni Lao, Tom Mitchell, and William W Cohen. 2011. Random walk inference and learning in a large scale knowledge base. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

[66] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. In *Nature*. Nature Publishing Group.

[67] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

[68] Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Advances in Neural Information Processing Systems (NeurIPS)*.

[69] Chenliang Li, Weiran Xu, Si Li, and Sheng Gao. 2018. Guiding generation for abstractive text summarization based on key information guide network. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

[70] Haoran Li, Junnan Zhu, Jiajun Zhang, Chengqing Zong, and Xiaodong He. 2020. Keywords-guided abstractive sentence summarization. In *AAAI Conference on Artificial Intelligence (AAAI)*.

[71] Jingyuan Li and Xiao Sun. 2018. A Syntactically Constrained Bidirectional-Asynchronous Approach for Emotional Conversation Generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

[72] Wei Li, Xinyan Xiao, Jiachen Liu, Hua Wu, Haifeng Wang, and Junping Du. 2021. Leveraging Graph to Improve Abstractive Multi-Document Summarization. In *Annual Meeting of Association for Computational Linguistics (ACL)*.

[73] Wei Li, Jingjing Xu, Yancheng He, ShengLi Yan, Yunfang Wu, and Xu Sun. 2019. Coherent Comments Generation for Chinese articles with a Graph-to-Sequence Model. In *Annual Meeting of Association Computational Linguistics (ACL)*.

[74] Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. Learning to select knowledge for response generation in dialog systems. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

[75] Kexin Liao, Logan Lebanoff, and Fei Liu. 2018. Abstract Meaning Representation for Multi-Document Summarization. In *International Conference on Computational Linguistics (COLING)*.

[76] Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. 2020. CommonGen: A Constrained Text Generation Challenge for Generative Commonsense Reasoning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP-Findings)*.

[77] Dayiheng Liu, Yu Yan, Yeyun Gong, Weizhen Qi, Hang Zhang, Jian Jiao, Weizhu Chen, Jie Fu, Linjun Shou, Ming Gong, et al. 2021. Glge: A new general language generation evaluation benchmark. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

[78] Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020. K-BERT: Enabling Language Representation with Knowledge Graph.. In *AAAI Conference on Artificial Intelligence (AAAI)*.

[79] Yuanxin Liu, Zheng Lin, Fenglin Liu, Qinyun Dai, and Weiping Wang. 2019. Generating Paraphrase with Topic as Prior Knowledge. In *International Conference on Information and Knowledge Management (CIKM)*.

[80] Ye Liu, Yao Wan, Lifang He, Hao Peng, and Philip S Yu. 2021. KG-BART: Knowledge Graph-Augmented BART for Generative Commonsense Reasoning. In *AAAI Conference on Artificial Intelligence (AAAI)*.

[81] Zhibin Liu, Zheng-Yu Niu, Hua Wu, and Haifeng Wang. 2019. Knowledge aware conversation generation with reasoning on augmented graph. In *Conference on Empirical Methods in Natural Language Processing and International*

Joint Conference on Natural Language Processing (EMNLP-IJCNLP).

[82] Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. 2018. Mem2Seq: Effectively Incorporating Knowledge Bases into End-to-End Task-Oriented Dialog Systems. In *Annual Meeting of Association for Computational Linguistics (ACL)*.

[83] Gideon S Mann and Andrew McCallum. 2007. Simple, robust, scalable semi-supervised learning via expectation regularization. In *International Conference on Machine Learning (ICML)*.

[84] Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Annual Meeting of the Association for Computational Linguistics: System Demonstration (ACL)*.

[85] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B Tenenbaum, and Jiajun Wu. 2019. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. *International Conference for Learning Representation (ICLR)*.

[86] Sahisnu Mazumder, Nianzu Ma, and Bing Liu. 2018. Towards a continuous knowledge learning engine for chatbots. In *arXiv preprint arXiv:1802.06024*.

[87] Chuan Meng, Pengjie Ren, Zhumin Chen, Christof Monz, Jun Ma, and Maarten de Rijke. 2020. RefNet: A reference-aware network for background based conversation. In *AAAI Conference on Artificial Intelligence (AAAI)*.

[88] Tanya Menon and Jeffrey Pfeffer. 2003. Valuing internal vs. external knowledge: Explaining the preference for outsiders. In *Management science*.

[89] Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering discrete latent topics with neural variational inference. In *International Conference on Machine Learning (ICML)*.

[90] Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M Khapra. 2018. Towards Exploiting Background Knowledge for Building Conversation Systems. In *Empirical Methods in Natural Language Processing (EMNLP)*.

[91] Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. Opendialkg: Explainable conversational reasoning with attention-based walks over knowledge graphs. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

[92] Diego Moussallem, Tommaso Soru, and Axel-Cyrille Ngonga Ngomo. 2019. THOTH: Neural Translation and Enrichment of Knowledge Graphs. In *International Semantic Web Conference (ISWC)*.

[93] Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, and Bing Xiang. 2016. Abstractive Text Summarization using Sequence-to-sequence RNNs and Beyond. In *Conference on Computational Natural Language Learning (SIGNL)*.

[94] Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Don't Give Me the Details, Just the Summary! Topic-Aware Convolutional Neural Networks for Extreme Summarization. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

[95] Christina Niklaus, Matthias Cetto, André Freitas, and Siegfried Handschuh. 2018. A Survey on Open Information Extraction. In *International Conference on Computational Linguistics (COLING)*.

[96] Tong Niu and Mohit Bansal. 2018. Polite dialogue generation without parallel data. In *Transactions of the Association for Computational Linguistics (TACL)*.

[97] Zheng-Yu Niu, Hua Wu, Haifeng Wang, et al. 2019. Knowledge Aware Conversation Generation with Explainable Reasoning over Augmented Graphs. In *Conference on Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.

[98] Liangming Pan, Yuxi Xie, Yansong Feng, Tat-Seng Chua, and Min-Yen Kan. 2020. Semantic Graphs for Generating Deep Questions. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

[99] Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vassilis Plachouras, Tim Rocktäschel, et al. 2021. KILT: a benchmark for knowledge intensive language tasks. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

[100] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language Models as Knowledge Bases?. In *Conference on Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.

[101] Lianhui Qin, Michel Galley, Chris Brockett, Xiaodong Liu, Xiang Gao, Bill Dolan, Yejin Choi, and Jianfeng Gao. 2019. Conversing by Reading: Contentful Neural Conversation with On-demand Machine Reading. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

[102] Lianhui Qin, Vered Shwartz, Peter West, Chandra Bhagavatula, Jena Hwang, Ronan Le Bras, Antoine Bosselut, and Yejin Choi. 2020. Backpropagation-based Decoding for Unsupervised Counterfactual and Abductive Reasoning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

[103] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. In *OpenAI Blog*.

[104] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. In *Journal of Machine Learning Research (JMLR)*.

[105] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

[106] Revanth Gangi Reddy, Danish Contractor, Dinesh Raghu, and Sachindra Joshi. 2019. Multi-Level Memory for Task Oriented Dialogs. In *the North American Chapter of the Association for Computational Linguistics (NAACL)*.

[107] Pengjie Ren, Zhumin Chen, Christof Monz, Jun Ma, and Maarten de Rijke. 2020. Thinking Globally, Acting Locally: Distantly Supervised Global-to-Local Knowledge Selection for Background Based Conversation. In *AAAI Conference on Artificial Intelligence (AAAI)*.

[108] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference (ESWC)*.

[109] Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

[110] Rico Sennrich and Barry Haddow. 2016. Linguistic Input Features Improve Neural Machine Translation. In *Conference on Machine Translation (WMT)*.

[111] Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *International Conference for Learning Representations (ICLR)*.

[112] Sifatullah Siddiqi and Aditi Sharan. 2015. Keyword and keyphrase extraction techniques: a literature review. In *International Journal of Computer Applications (IJCA)*. Foundation of Computer Science.

[113] Haoyu Song, Wei-Nan Zhang, Yiming Cui, Dong Wang, and Ting Liu. 2019. Exploiting persona information for diverse generation of conversational responses. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

[114] Zhenqiao Song, Xiaoqing Zheng, Lu Liu, Mu Xu, and Xuan-Jing Huang. 2019. Generating responses with a specific emotion in dialog. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

[115] Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *AAAI Conference on Artificial Intelligence (AAAI)*.

[116] Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Advances in Neural Information Processing Systems (NeurIPS)*.

[117] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*.

[118] Bowen Tan, Lianhui Qin, Eric Xing, and Zhiting Hu. 2020. Summarizing Text on Any Aspects: A Knowledge-Informed Weakly-Supervised Approach. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

[119] Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric Xing, and Zhiting Hu. 2019. Target-Guided Open-Domain Conversation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

[120] Kenneth Tran, Xiaodong He, Lei Zhang, Jian Sun, Cornelia Carapcea, Chris Thrasher, and Chris Buehler. 2016. Rich image captioning in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

[121] Yi-Lin Tuan, Yun-Nung Chen, and Hung-yi Lee. 2019. DyKgChat: Benchmarking Dialogue Generation Grounding on Dynamic Knowledge Graphs. In *Conference on Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.

[122] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*.

[123] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference for Learning Representation (ICLR)*.

[124] Cunxiang Wang, Shuailong Liang, Yili Jin, Yilong Wang, Xiaodan Zhu, and Yue Zhang. 2020. SemEval-2020 Task 4: Commonsense Validation and Explanation. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*.

[125] Hao Wang, Bin Guo, Wei Wu, and Zhiwen Yu. 2020. Towards information-rich, logical text generation with knowledge-enhanced neural models. In *arXiv preprint arXiv:2003.00814*.

[126] Han Wang, Yang Liu, Chenguang Zhu, Linjun Shou, Ming Gong Gong, Yichong Xu, and Michael Zeng. 2021. Retrieval Enhanced Model for Commonsense Generation. In *Annual Meeting of Association for Computational Linguistics (ACL)*.

[127] Hongwei Wang, Fuzheng Zhang, Mengdi Zhang, Jure Leskovec, Miao Zhao, Wenjie Li, and Zhongyuan Wang. 2019. Knowledge-aware graph neural networks with label smoothness regularization for recommender systems. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*.

[128] Jian Wang, Junhao Liu, Wei Bi, Xiaojiang Liu, Kejing He, Ruijing Xu, and Min Yang. 2020. Improving Knowledge-aware Dialogue Generation via Knowledge Base Question Answering. In *Conference on Artificial Intelligence (AAAI)*.

[129] Kai Wang, Xiaojun Quan, and Rui Wang. 2019. BiSET: Bi-directional Selective Encoding with Template for Abstractive Summarization. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

[130] Qingyun Wang, Lifu Huang, Zhiying Jiang, Kevin Knight, Heng Ji, Mohit Bansal, and Yi Luan. 2019. PaperRobot: Incremental Draft Generation of Scientific Ideas. In *Annual Meeting of Association Computational Linguistics (ACL)*.

[131] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. 2017. Knowledge graph embedding: A survey of approaches and applications. In *IEEE Transactions on Knowledge and Data Engineering (TKDE)*.

- [132] Wenlin Wang, Zhe Gan, Hongteng Xu, Ruiyi Zhang, Guoyin Wang, Dinghan Shen, Changyou Chen, and Lawrence Carin. 2019. Topic-Guided Variational Auto-Encoder for Text Generation. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- [133] Zhen Wang, Siwei Rao, Jie Zhang, Zhen Qin, Guangjian Tian, and Jun Wang. 2020. Diversify Question Generation with Continuous Content Selectors and Question Type Modeling. In *Conference on Empirical Methods in Natural Language Processing (EMNLP-Findings)*.
- [134] Xiangpeng Wei, Yue Hu, Luxi Xing, Yipeng Wang, and Li Gao. 2019. Translating with Bilingual Topic Knowledge for Neural Machine Translation. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [135] Chien-Sheng Wu, Richard Socher, and Caiming Xiong. 2019. Global-to-local memory pointer networks for task-oriented dialogue. In *International Conference for Learning Representation (ICLR)*.
- [136] Sixing Wu, Ying Li, Dawei Zhang, Yang Zhou, and Zhonghai Wu. 2020. Diverse and Informative Dialogue Generation with Context-Specific Commonsense Knowledge Awareness. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [137] Sixing Wu, Ying Li, Dawei Zhang, Yang Zhou, and Zhonghai Wu. 2020. TopicKA: Generating Commonsense Knowledge-Aware Dialogue Responses Towards the Recommended Topic Fact. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- [138] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. In *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*.
- [139] Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [140] Wenhan Xiong, Jingfei Du, William Yang Wang, and Veselin Stoyanov. 2020. Pretrained Encyclopedia: Weakly Supervised Knowledge-Pretrained Language Model. In *International Conference of Learning Representation (ICLR)*.
- [141] Wenhan Xiong, Thien Hoang, and William Yang Wang. 2017. DeepPath: A Reinforcement Learning Method for Knowledge Graph Reasoning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [142] Jun Xu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. Conversational graph grounded policy learning for open-domain conversation generation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [143] Minghong Xu, Piji Li, Haoran Yang, Pengjie Ren, Zhaochun Ren, Zhumin Chen, and Jun Ma. 2020. A Neural Topical Expansion Framework for Unstructured Persona-oriented Dialogue Generation. In *European Conference on Artificial Intelligence (ECAI)*.
- [144] Pengcheng Yang, Lei Li, Fuli Luo, Tianyu Liu, and Xu Sun. 2019. Enhancing Topic-to-Essay Generation with External Commonsense Knowledge. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [145] Donghan Yu, Chenguang Zhu, Yuwei Fang, Wenhao Yu, Shuohang Wang, Yichong Xu, Xiang Ren, Yiming Yang, and Michael Zeng. 2022. Kg-fid: Infusing knowledge graph in fusion-in-decoder for open-domain question answering. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [146] Donghan Yu, Chenguang Zhu, Yiming Yang, and Michael Zeng. 2022. Jaket: Joint pre-training of knowledge graph and language understanding. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [147] Wenhao Yu, Mengxia Yu, Tong Zhao, and Meng Jiang. 2020. Identifying referential intention with heterogeneous contexts. In *Proceedings of The Web Conference (WebConf)*.
- [148] Wenhao Yu, Chenguang Zhu, Yuwei Fang, Donghan Yu, Shuohang Wang, Yichong Xu, Michael Zeng, and Meng Jiang. 2021. Dict-BERT: Enhancing Language Model Pre-training with Dictionary. *arXiv preprint arXiv:2110.06490*.
- [149] Wenhao Yu, Chenguang Zhu, Tong Zhao, Zhichun Guo, and Meng Jiang. 2021. Sentence-Permuted Paragraph Generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [150] Qingkai Zeng, Jinfeng Lin, Wenhao Yu, Jane Cleland-Huang, and Meng Jiang. 2021. Enhancing Taxonomy Completion with Concept Generation via Fusing Relational Representations. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*.
- [151] Houyu Zhang, Zhenghao Liu, Chenyan Xiong, and Zhiyuan Liu. 2020. Grounded conversation generation as guided traverses in commonsense knowledge graphs. In *Annual Meeting of the Association for Computational Linguistics*.
- [152] Jian Zhang, Liangyou Li, Andy Way, and Qun Liu. 2016. Topic-Informed Neural Machine Translation. In *International Conference on Computational Linguistics: Technical Papers (COLING)*.
- [153] Jiacheng Zhang, Yang Liu, Huanbo Luan, Jingfang Xu, and Maosong Sun. 2017. Prior Knowledge Integration for Neural Machine Translation using Posterior Regularization. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [154] Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing Dialogue Agents: I have a dog, do you have pets?. In *Annual Meeting of Association Computational Linguistics (ACL)*.
- [155] Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. ERNIE: Enhanced Language Representation with Informative Entities. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

- [156] Liang Zhao, Jingjing Xu, Junyang Lin, Yichang Zhang, Hongxia Yang, and Xu Sun. 2020. Graph-based multi-hop reasoning for long text generation. In *arXiv preprint arXiv:2009.13282*.
- [157] Yinhe Zheng, Rongsheng Zhang, Minlie Huang, and Xiaoxi Mao. 2020. A Pre-Training Based Personalized Dialogue Generation Model with Persona-Sparse Data.. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [158] Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [159] Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Commonsense knowledge aware conversation generation with graph attention. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- [160] Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2017. Neural question generation from text: A preliminary study. In *Conference on Natural Language Processing and Chinese Computing (NLPCC)*.
- [161] Wangchunshu Zhou, Dong-Ho Lee, Ravi Kiran Selvam, Seyeon Lee, Bill Yuchen Lin, and Xiang Ren. 2021. Pre-training text-to-text transformers for concept-centric common sense. *International Conference for Learning Representation*.
- [162] Yimin Zhou, Yiwei Sun, and Vasant Honavar. 2019. Improving image captioning by leveraging knowledge graphs. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*.
- [163] Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang, and Meng Jiang. 2021. Boosting Factual Correctness of Abstractive Summarization with Knowledge Graph. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- [164] Jun Zhu, Ning Chen, and Eric P Xing. 2014. Bayesian inference with posterior regularization and applications to infinite latent SVMs. In *The Journal of Machine Learning Research (JMLR)*.
- [165] Wenya Zhu, Kaixiang Mo, Yu Zhang, Zhangbin Zhu, Xuezheng Peng, and Qiang Yang. 2017. Flexible end-to-end dialogue system for knowledge grounded conversation. In *CoRR, abs/1709.04264*.
- [166] Hui Zou and Trevor Hastie. 2005. Regularization and variable selection via the elastic net. In *Journal of the royal statistical society*. Wiley Online Library.

A APPENDIX

Figure 8 demonstrates the statistics of selected publications in this survey. The left figure shows the paper publishing venues. Most papers were published in top machine learning, artificial intelligence, and natural language processing conferences, such as ACL, EMNLP, AAAI, ICLR, NeurIPS. Besides, many selected papers were published in high-impact journals, such as TNNLS, JMLR, TACL. The right figure shows the paper categories. Among 160 selected papers, 87 papers (“general methods (General)”, “topic”, “keyword”, “knowledge base (KG)”, “knowledge graph (KG)”, “grounded text (Text)”) are directly relevant to the different kinds of knowledge-enhanced text generation methods; 10 papers are relevant to benchmark datasets; 10 papers are related survey papers. Besides, other 43 papers are about basic (pre-trained) generation methods (e.g., Seq2Seq, CopyNet, BART, T5), or necessary background (e.g., TransE, OpenIE, GNN, LDA), or future direction.

Figure 9 summarized different papers according to years, knowledge sources, and methods.

Table 10 lists the leaderboard performance on ten knowledge-enhanced generation benchmarks.

Table 11 lists code links and programming language of representative open-source knowledge-enhanced text generation systems that have been introduced in this survey.

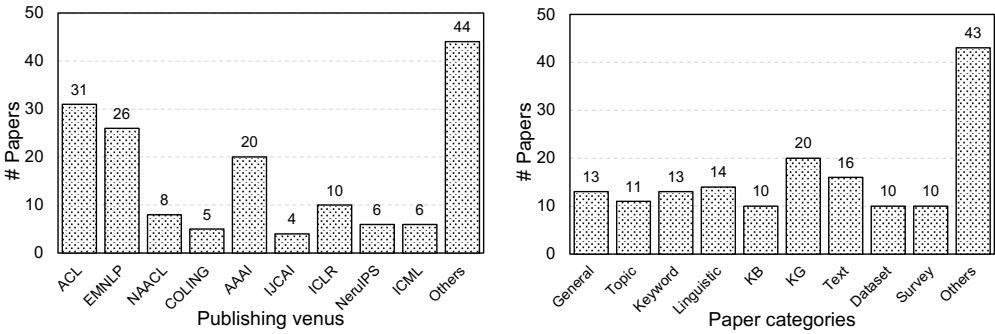


Fig. 8. Paper statistics of selected publications in this survey.

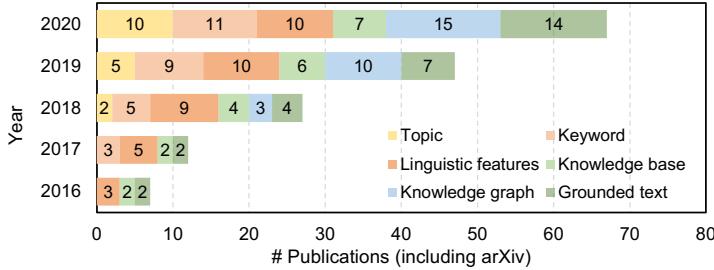


Fig. 9. Knowledge-enhanced text generation has been gaining emerging interests in the recent five years.

A.1 Evaluation Metrics

BLEU-*m* (short as B-*m*): BLEU is a weighted geometric mean of *n*-gram precision scores.

ROUGE-*m* (short as R-*m*): ROUGE measures the overlap of n-grams between the reference and hypothesis; ROUGE-L measures the longest matched words using longest common sub-sequence.

Distinct-*k* (short as D-*k*): Distinct measures the total number of unique *k*-grams normalized by the total number of generated *k*-gram tokens to avoid favoring long sentences.

Table 10. Leaderboard performance on ten knowledge-enhanced generation benchmarks.

(a) Leaderboard performance on two summarization benchmark datasets with different knowledge-enhanced NLG methods. Evaluation metrics are standard n-gram based metrics: ROUGE-2 and ROUGE-L.

Methods	Ref.	Knowledge source	Method category	CNN/DM		Gigaword		
				R-2	R-L	R-2	R-L	
Baseline methods (w/o KG)								
Seq2Seq	[117]			11.81	28.83	11.32	26.42	with attention mechanism
PG	[109]			15.66	33.42	17.63	33.66	w/o coverage mechanism
Knowledge enhanced methods								
VHTM	[33]	Topic	M3	18.05	37.18	-	-	-
SELECTOR	[22]	Keyword	M2	18.31	-	-	-	* Improve generation diversity
FASUM	[163]	OpenKG	-	17.84	37.40	-	-	* Improve factual correctness
KIGN	[69]	Keyword	M1	17.12	35.68	17.93	34.44	-
BottomUp	[41]	Keyword	M2	18.68	38.34	17.61	33.54	-
TGVAE	[132]	Topic	M3	-	-	17.27	33.02	-
HierDualPG	[70]	Keyword	M2	-	-	18.06	34.39	-
R ³ Sum	[13]	Text	M1	-	-	19.03	34.46	-
BiSET	[129]	Text	M1	-	-	19.78	36.87	-
SemSUM	[59]	DepGraph	-	-	-	19.75	36.09	-
ASGARD	[54]	OpenKG	-	20.37	40.48	-	-	-

(b) Leaderboard performance on α NLG-ART dataset. Both B-4 and R-L are commonly used.

Method	Ref.	Source	B-4	R-L
Baseline methods				
Seq2Seq	[117]	-	2.37	22.30
GPT-2	[103]	-	9.80	32.90
Knowledge-enhanced methods				
GPT-COMeT	[7]	KG	9.62	32.88
GRF	[80]	KG	11.62	34.62

(d) Leaderboard performance on CommonGen dataset. SPICE is the primary evaluation metric.

Method	Ref.	Source	B-4	SPICE
Baseline methods				
BART	[67]	-	31.83	27.99
T5	[104]	-	31.96	28.86
Knowledge-enhanced methods				
EKI-BART	[32]	Text	35.95	29.59
KG-BART	[80]	KG	33.87	29.63
RE-T5	[126]	Text	40.87	31.08

(c) Leaderboard performance on ComVE dataset. Both B-4 and R-L are commonly used.

Method	Ref.	Source	B-4	R-L
Baseline methods				
Seq2Seq	[117]	-	6.10	25.80
GPT-2	[103]	-	15.70	36.50
Knowledge-enhanced methods				
CE-PR	[56]	KG	17.10	37.90
GRF	[57]	KG	17.19	38.10

(e) Leaderboard performance on Holl-E (mix-short setting) dataset. R-L are the primary metric.

Method	Ref.	Source	R-L	B-4
Baseline methods				
Seq2Seq	[117]	-	21.48	5.26
BiDAF	[111]	-	35.09	27.44
Knowledge-enhanced methods				
AKGCM	[81]	KG	34.72	30.84
RefNet	[87]	Text	36.17	29.38
GLKS	[107]	Text	39.63	-

(f) Leaderboard performance on Wizard of Wikipedia with seen (S) and unseen (UnS) test set.

Method	Ref.	R-1/R-2 (S)	R-1/R-2 (UnS)
Baseline methods			
Transformer	[122]	17.8/ —	14.0/ —
Knowledge-enhanced methods			
MemNet	[27]	16.9/ —	14.4/ —
PostKS	[74]	18.1/5.3	13.5/2.0
SKT	[61]	19.3/6.8	16.1/4.2
PIPM+KDBTS	[18]	19.9/7.3	17.6/5.4

(h) Leaderboard performance on ELI5 dataset. The Kilt R-L (KRL) is the primary evaluation metric.

Method	Ref.	Source	KRL	R-L
Baseline methods				
T5	[104]	-	0.0	19.1
BART	[67]	-	0.0	20.1
Knowledge-enhanced methods				
RAG	[68]	Text	1.7	17.4
BART+DPR	[99]	Text	1.9	17.4
RT+c-REALM	[61]	Text	2.4	23.2

(g) State-of-the-art performance on SQuAD.

Method	Ref.	Source	B-4
Baseline methods			
Seq2Seq	[117]	-	3.01
Transformer	[122]	-	3.09
Knowledge-enhanced methods			
NQG++	[160]	LF	13.27
SELECTOR	[22]	LF+Keyword	15.87
G2S+BERT	[19]	LF+DepGraph	17.49
G2S+BERT+RL	[19]	LF+DepGraph	18.30

(i) Some state-of-the-art performance on PersonaChat dataset (no leaderboard on this dataset).

Method	Ref.	B-1/B-2	D-1/D-2
Baseline methods			
Seq2Seq	[117]	18.2/9.3	2.6/7.4
Knowledge-enhanced methods			
MemNet(soft)	[27]	17.7/9.1	3.5/9.6
MemNet(hard)	[27]	18.6/9.7	3.7/9.9
PostKS	[74]	19.0/9.8	4.6/13.4
PEE	[143]	23.2/11.5	- / -

Table 11. A list of representative open-source knowledge-enhanced text generation systems.

Task	Ref.	Paper title and open source code/toolkit	Programming language	Venue & Year
Topic-enhanced methods				
Summarization	[94]	Topic-Aware Convolutional Neural Networks for Extreme Summarization --- Code: https://github.com/EdinburghNLP/XSum	PyTorch	EMNLP 2018
	[135]	Friendly Topic Assistant for Transformer Based Abstractive Summarization --- Code: https://github.com/BoChenGroup/TA	-	EMNLP 2020
Keyword-enhanced methods				
Dialogue system	[94]	A Content-Introducing Approach to Generative Short-Text Conversation --- Code: https://github.com/MaZhiyuanBUAA/Seq2BPforDialogueGeneration	Tensorflow	COLING 2016
	[135]	Emotional Chatting Machine: Emotional Conversation Generation with Internal and External Memory --- Code: https://github.com/loader/ECM-tf	PyTorch	AAAI 2018
Summarization	[94]	Coherent Comment Generation with a Graph-to-Sequence Model --- Code: https://github.com/lancopku/Graph-to-seq-comment-generation	PyTorch	ACL 2018
KB-enhanced methods				
Dialogue system	[82]	Mem2Seq: Effectively Incorporating Knowledge Bases into End-to-End Dialog Systems --- Code: https://github.com/HTLCHKUST/Mem2Seq	PyTorch	ACL 2019
	[135]	Global-to-local Memory Pointer Networks for Task-Oriented Dialogue --- Code: https://github.com/jasonwu0731/GLMP	PyTorch	ICLR 2019
	[128]	Improving Knowledge-aware Dialogue Generation via Knowledge Base Question Answering --- Code: https://github.com/siat-nlp/TransDG	PyTorch	AAAI 2020
	[136]	Diverse and Informative Dialogue Generation with Context-Specific Knowledge Awareness --- Code: https://github.com/pku-sixing/ACL2020-ConKADI	Tensorflow	ACL 2020
	[137]	TopicKA: Generating Commonsense Knowledge-Aware Dialogue Responses --- Code: https://github.com/pku-sixing/IJCAI2020-TopicKA	Tensorflow	IJCAI 2020
KG-enhanced methods				
Dialogue system	[159]	Commonsense Knowledge Aware Conversation Generation with Graph Attention --- Code: https://github.com/thu-coai/cmc	Tensorflow	IJCAI 2018
	[121]	DyKgChat: Benchmarking Dialogue Generation Grounding on Dynamic Knowledge Graphs --- Code: https://github.com/Pascalson/DyKGChat	Tensorflow	EMNLP 2019
	[151]	Grounded Conversation Generation as Guided Traverses in Commonsense Knowledge Graphs --- Code: https://github.com/thunlp/ConceptFlow	PyTorch	ACL 2020
Scientific writing	[63]	Text Generation from Knowledge Graphs with Graph Transformers --- Code: https://github.com/rikdz/GraphWriter	PyTorch	NAACL 2019
	[130]	PaperRobot: Incremental Draft Generation of Scientific Ideas --- Code: https://github.com/EagleW/PaperRobot	PyTorch	ACL 2019
Commonsense reasoning & Story generation	[46]	Story Ending Generation with Incremental Encoding and Commonsense Knowledge --- Code: https://github.com/JianGuanTHU/StoryEndGen	Tensorflow	AAAI 2019
	[57]	Language Generation with Multi-Hop Reasoning on Commonsense Knowledge Graph --- Code: https://github.com/cdjhz/multigen	PyTorch	EMNLP 2020
	[80]	KG-BART: Knowledge Graph-Augmented BART for Generative Commonsense Reasoning --- Code: https://github.com/yeliu918/KG-BART	PyTorch	AAAI 2021
	[21]	ENT-DESC: Entity Description Generation by Exploring Knowledge Graph --- Code: https://github.com/LiyingCheng95/EntityDescriptionGeneration	MXNet	EMNLP 2020
Question answering	[5]	Commonsense for Generative Multi-Hop Question Answering Tasks --- Code: https://github.com/yicheng-w/CommonSenseMultiHopQA	Tensorflow	EMNLP 2018
Ground text-enhanced methods				
Dialogue system	[27]	Wizard of Wikipedia: Knowledge-Powered Conversational agents --- Code: https://github.com/facebookresearch/ParlAI	PyTorch	ICLR 2019
	[101]	Conversing by Reading: Contentful Neural Conversation with On-demand Machine Reading --- Code: https://github.com/qkaren/converse_reading_cmr	PyTorch	ACL 2019
	[61]	Sequential Latent Knowledge Selection for Knowledge-Grounded Dialogue --- Code: https://github.com/bckim92/sequential-knowledge-transformer	Tensorflow	ICLR 2020
	[87]	RefNet: A Reference-aware Network for Background Based Conversation --- Code: https://github.com/ChuanMeng/RefNet	Tensorflow	AAAI 2020
Summarization	[129]	BiSET: Bi-directional Selective Encoding with Template for Abstractive Summarization --- Code: https://github.com/InitialBug/BiSET	PyTorch	ACL 2019
Question answering	[68]	Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks --- Code in https://github.com/huggingface/transformers	PyTorch	Neurips 2020
	[99]	KILT: a Benchmark for Knowledge Intensive Language Tasks --- Code: https://github.com/facebookresearch/KILT	PyTorch	NAACL 2021