

Multi-Hop Question Generation via Dual-Perspective Keyword Guidance

Maodong Li¹, Longyin Zhang², Fang Kong^{1*}

¹School of Computer Science and Technology, Soochow University, China

²Institute for Infocomm Research, A*STAR, Singapore

{20234227022@stu, kongfang@}suda.edu.cn zhang_longyin@i2r.a-star.edu.sg

Abstract

Multi-hop question generation (MQG) aims to generate questions that require synthesizing multiple information snippets from documents to derive target answers. The primary challenge lies in effectively pinpointing crucial information snippets related to question-answer (QA) pairs, typically relying on keywords. However, existing works fail to fully utilize the guiding potential of keywords and neglect to differentiate the distinct roles of question-specific and document-specific keywords. To address this, we define dual-perspective keywords—question and document keywords—and propose a **Dual-Perspective Keyword-Guided (DPKG)** framework, which seamlessly integrates keywords into the multi-hop question generation process. We argue that question keywords capture the questioner’s intent, whereas document keywords reflect the content related to the QA pair. Functionally, question and document keywords work together to pinpoint essential information snippets in the document, with question keywords required to appear in the generated question. The DPKG framework consists of an expanded transformer encoder and two answer-aware transformer decoders for keyword and question generation, respectively. Extensive experiments demonstrate the effectiveness of our work, showcasing its promising performance and underscoring its significant value in the MQG task.

1 Introduction

Question generation (QG) is the task of generating questions based on given documents and target answers (Zhang et al., 2021; Ushio et al., 2023). It is a vital area in natural language processing (NLP), with considerable potential for applications in automated question answering, reading comprehension, and related fields (Liang et al., 2023; Guo et al.,

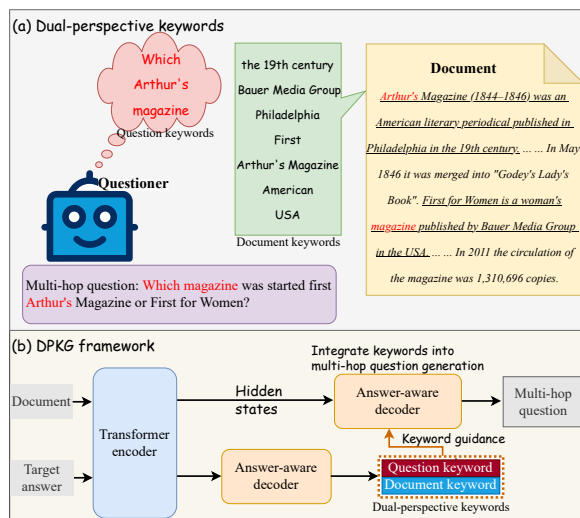


Figure 1: (a) shows the definition of dual-perspective keywords: question keywords and document keywords. The red highlights the question keywords, while the underlined text indicates crucial information snippets related to question-answer pairs. (b) illustrates the proposed DPKG framework.

2024). Compared to the QG task, multi-hop question generation (MQG) aims to generate questions that require synthesizing multiple information snippets from documents to derive target answers, making it more representative of real-world scenarios (Fei et al., 2021, 2022a; Chen et al., 2023).

The primary challenge in MQG is effectively pinpointing crucial information snippets related to question-answer (QA) pairs in documents, a task that many studies have approached using keywords (Xia et al., 2023; Ding et al., 2024; Liu et al., 2024a; Fazili et al., 2024). However, current research does not fully leverage the guiding potential of keywords. Some studies focus exclusively on document-specific keywords (Su et al., 2020; Cao and Wang, 2021), while others, including more recent, higher-performing models (Fei et al., 2022a,b) consider constraints on question-specific keywords during question generation. According to our un-

*Corresponding author

derstanding, as illustrated in Figure 1(a), question keywords reflect the questioner’s intent, whereas document keywords reflect the content related to the QA pair in documents. Functionally, question and document keywords work together to pinpoint essential information snippets in the document, with question keywords required to appear in the generated question. However, existing studies overlook the distinction between the two keyword types, which poses a challenge to further improving performance (Fei et al., 2022a; Xia et al., 2023; Ding et al., 2024).

To break through the above limitation, we define dual-perspective keywords: question keywords, which capture the intent of the questioner, and document keywords, which reflect the content related to the QA pair in documents (Figure 1(a)). Based on this, we further propose a **Dual-Perspective Keyword-Guided (DPKG)** framework, which seamlessly integrates keywords into the multi-hop question generation process (Figure 1(b)). The framework clarifies essential information snippets by differentiating the two categories of keywords and ensures that specific keywords appear in the generated question, aligning closely with natural human intuition. More specifically, the DPKG framework consists of three main components: (1) an expanded transformer encoder that processes the document and target answer; (2) a decoder that generates dual-perspective keywords; and (3) another decoder that generates multi-hop questions by leveraging crucial information snippets pinpointed by keywords. This framework tackles MQG challenges and aligns with practical requirements. Our contributions are as follows:

- We define dual-perspective keywords: question keywords, which capture the questioner’s intent, and document keywords, which reflect the content related to QA pairs in documents.
- We propose a DPKG framework that seamlessly integrates keywords into the multi-hop question generation process.
- Extensive experiments demonstrate the effectiveness of our work, showcasing its promising performance and underscoring its significant value in the MQG task ¹.

¹<https://github.com/imaodong/DPKG>

2 Related Work

Question generation (QG), also known as shallow QG, involves generating questions based on given documents and target answers. Early approaches to QG were rule-based (Heilman and Smith, 2010; Chali and Hasan, 2012; Mazidi and Nielsen, 2014), but their performance was limited. As the field advanced, attention-based sequence-to-sequence architectures gained popularity (Du et al., 2017; Zhou et al., 2018; Sun et al., 2018; Song et al., 2018). Some studies emphasized feature-rich encoders to capture more contextual information (Zhou et al., 2018; Sun et al., 2018), and answer-aware features were introduced to extract additional information (Sun et al., 2018; Song et al., 2018). To address the challenge posed by long passages, a maxout pointer mechanism was developed (Zhao et al., 2018). Moreover, reinforcement learning (RL) has been applied in a graph-to-sequence model (G2S) for question generation (Chen et al., 2020, 2023). All of these studies focus on shallow QG, with the SQuAD dataset (Rajpurkar et al., 2016) being the most commonly used benchmark.

Multi-hop question generation (MQG) focuses on generating questions that synthesize multiple information snippets from documents to derive target answers. Compared to shallow QG, MQG is more demanding and has garnered significant attention in recent years (Su et al., 2020; Fei et al., 2021, 2022a; Chen et al., 2023; Ding et al., 2024). In the early stages of MQG, graph neural networks (GNNs) were integrated into multi-hop reasoning steps to facilitate question generation (Su et al., 2020; Pan et al., 2020). However, with the evolution of transformer-based architectures, transformers have been demonstrated to outperform GNNs in this task (Vaswani, 2017; Su et al., 2022). To further enhance the performance of pre-trained language models (PLMs) in multi-hop question generation, researchers have incorporated mechanisms such as additional question-answer modules and multi-level content planning (Su et al., 2022; Xia et al., 2023). Numerous studies have also explored auxiliary tasks, proposing methods like order-agnostic learning and sequential rewriting to improve question quality (Murakhovs’ka et al., 2022; Kim et al., 2024; Hwang et al., 2024). More recently, question-specific keywords have been utilized during decoding to enhance high-quality generation (Fei et al., 2022a), and Ding et al. (2024) supported multi-hop question generation by identi-

fying salient, related keywords or sentences. However, they do not fully harness the guiding potential of keywords, which limits improvements in generation performance.

3 Methodology

3.1 Task Formulation

Assume a MQG dataset $U = \{D, A, Q\}_{N}^{i=1}$, where D^i denotes contextual documents, A^i represents a target answer, and N is the number of samples. The task is to generate a multi-hop question Q^i based on D^i and A^i . We extend the dataset to $U = \{D, A, Q, K\}_{N}^{i=1}$, where K^i is a keyword sequence consisting of two sets: $K_{j=1:M_i^q}^q$ and $K_{j=1:M_i^d}^d$, representing question keywords and document keywords annotated with SpaCy (see §4.1 for details). Here, M_i^q and M_i^d indicate the number of question keywords and document keywords, respectively. The updated task now includes two primary objectives: (1) generating keywords, $K^i = f_{\theta_1}(D^i, A^i)$; and (2) generating multi-hop question, $Q^i = f_{\theta_2}(D^i, A^i, K^i)$, where θ_1 and θ_2 are model parameters.

3.2 DPKG Framework

In this subsection, we introduce the DPKG framework in detail, which seamlessly integrates keywords into the multi-hop question generation process. As illustrated in Figure 2, the DPKG framework consists of an expanded transformer encoder and two answer-aware transformer decoders, designated for keyword generation and question generation, respectively. Additionally, inspired by the success of previous prompt-based approaches, we adopt a similar method for this task. We prepend an additional prefix defined as follows: $D^i = \text{"Document: " + } D^i$, $A^i = \text{"Answer: " + } A^i$, and $Q^i = \text{"Question: " + } Q^i$.

As illustrated in Figure 2(a), this is an expansion of the transformer encoder (Vaswani, 2017), which encodes three components and produces the corresponding hidden states:

$$H_{doc}^i, H_{ans}^i, H_{da}^i = Encoder_a([D^i, A^i, D^i; A^i]) \quad (1)$$

where H_{doc}^i and H_{ans}^i represent document states and answer states, respectively. H_{da}^i denotes document-answer states, and $[\cdot; \cdot]$ denotes concatenation. Notably, we modify the internal structure of the transformer encoder to ensure that, when

encoding the three input components simultaneously, each part remains independent, except for the shared self-attention mechanism. We encode the document and answer separately to facilitate the subsequent computation of answer-aware states. In Figure 2, (b) and (c) are identical, except (c) incorporates an additional vanilla BART encoder (Lewis et al., 2020) to encode keywords. Although structurally similar, they are independent because of their distinct tasks. Specifically, (b) generates keywords based on the document and answer, while (c) leverages the keywords to pinpoint essential information snippets related to QA pairs in the document and generate multi-hop questions. Formalizing these processes:

$$K^i = Decoder_b(H_{doc}^i, H_{ans}^i, H_{da}^i) \quad (2)$$

$$Q^i = Decoder_c(H_{doc}^i, H_{ans}^i, H_{da}^i, E(K^i)) \quad (3)$$

where $E(\cdot)$ denotes the vanilla BART encoder and K^i represents the i -th keyword sequence. Additionally, we have modified the transformer decoder by incorporating an answer-aware attention mechanism to enhance answer-awareness in the generation of both keywords and questions.

3.3 Answer-aware Attention Mechanism and Fusion Module

It is worth noting that (b) and (c) in Figure 2 are similar processes to each other, with (b) used to generate keywords and (c) used to generate questions. Here, we use (b) as an illustrative example. Assume that in the previous step, we obtain the keyword representation, denoted as H_k^{t-1} . The answer-aware attention mechanism (Wang et al., 2024) is formalized as follows:

$$H_a = softmax\left(\frac{H_k^{t-1} H_{doc}^T}{\sqrt{d}} \odot K_{weight}\right) H_{doc} \quad (4)$$

$$K_{weight} = MeanPooling\left(\frac{H_{doc} H_{ans}^T}{\sqrt{d}}\right) \quad (5)$$

where d represents the hidden size and \odot denotes element-wise multiplication. We argue that K_{weight} encapsulates the relationship between the document and the answer. By applying the softmax function, we quantify these interconnections. Finally, we take the dot product with H_{doc} and derive the answer-aware states, denoted as H_a . Correspondingly, in the cross-attention mechanism (Vaswani, 2017), we also obtain hidden states, denoted by H_h . Next, we apply the fusion module to

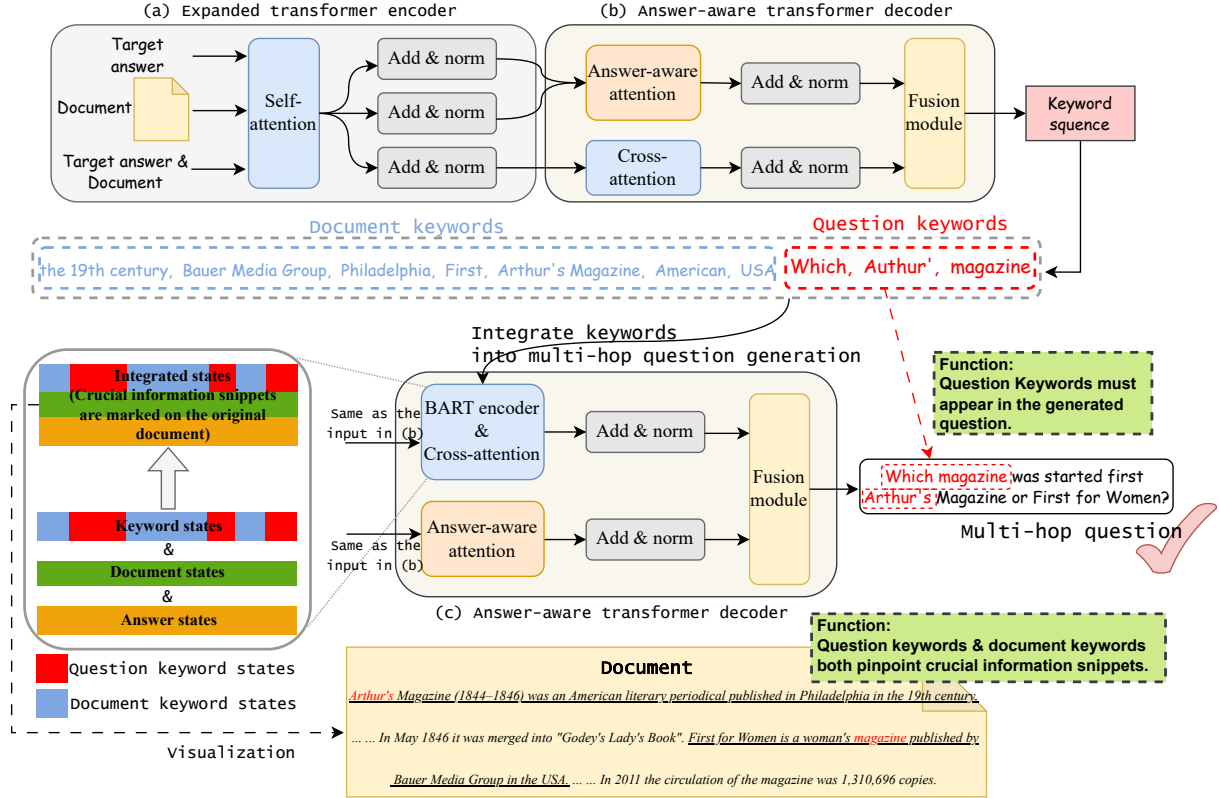


Figure 2: Overall architecture of DPKG framework. In (c), both modules receive the same inputs as in (b), originating from (a). To simplify the illustration and highlight key points, some Transformer details—such as the Feed-Forward sub-layer—are omitted. The encoder includes three Feed-Forward and Add&Norm sub-layers, and the decoder adds a Fusion module.

combine the states of H_a and H_h , producing the keyword representation H_k^t in the following step. The fusion module operates as follows:

$$H_k^t = gate \odot H_a + (1 - gate) \odot H_h \quad (6)$$

$$gate = sigmoid([H_a; H_h]) \quad (7)$$

3.4 Keyword Guidance

The keyword guidance concept of DPKG integrates keywords into the multi-hop question generation process by pinpointing crucial information snippets related to QA pairs in the documents and specifying the question keywords that must appear in the generated questions. Our keyword guidance approach is grounded in two primary principles: First, after generating keywords, they are encoded by the BART encoder and aggregated with H_{da}^i to form new states, seamlessly incorporating keywords into the question generation process. This integration enables the identification of crucial information snippets and ensures the inclusion of specific keywords in the generated question, as shown in Figure 2. Second, keyword generation and question generation are jointly trained using a shared, expanded

transformer encoder, which allows for deep integration of representations from both tasks.

There are two modes when using dual-perspective keywords in the DPKG framework: (1) In hard mode, special prefixes `<qes>` or `<doc>` are added to each keyword to indicate whether it is a question or document keyword. (2) In soft mode, no special prefixes are added; instead, the framework dynamically identifies each keyword’s role (question or document) during question generation. This means that, in hard mode, the recognition of the two types of keywords occurs during the keyword generation process, whereas in soft mode, the identification happens during question generation. The soft mode has already achieved impressive performance in other areas of NLP (Liu et al., 2024a). The keywords in this paper are presented sequentially. In hard mode, the keyword sequence is defined as $K^i = \langle qes \rangle K_{j=1:M_i^q}^q \langle doc \rangle K_{j=1:M_i^d}^d$; while in soft mode, it is defined as $K^i = K_{j=1:M_i^q}^q K_{j=1:M_i^d}^d$, where the framework dynamically identifies each keyword’s role (question or document) during

multi-hop question generation.

3.5 Training Loss

We present a joint training approach for keyword and multi-hop question generation. During training, keyword and question generation are trained in a pipeline fashion, where ground-truth keywords are used for question generation. In the inference phase, keywords are first generated and then used for question generation, creating a gap between training and inference. To bridge this gap, we introduce a novel loss function, defined as follows:

$$\mathcal{L}_3 = \|F_k - F_g\|_2 \quad (8)$$

$$F_k = \text{AvgPooling}(H_k^t) \quad (9)$$

$$F_g = \text{AvgPooling}(E(g)) \quad (10)$$

where g denotes the ground-truth keywords. This implies that, during joint training, we aim for the ground-truth keyword representations used in multi-hop question generation to closely resemble the generated keyword representations, thereby minimizing this gap. In joint training, the loss for keyword generation is denoted as \mathcal{L}_1 , while the loss for question generation is denoted as \mathcal{L}_2 . Both losses are computed using cross-entropy. The final loss is expressed as follows:

$$\mathcal{L} = \beta_1 \mathcal{L}_1 + \beta_2 \mathcal{L}_2 + \beta_3 \mathcal{L}_3 \quad (11)$$

where β_1 , β_2 and β_3 are hyperparameters ².

4 Experiments

4.1 Dataset

We conduct experiments on a common MQG dataset, HotpotQA (Yang et al., 2018), which consists of two settings: supporting fact (SF) and full document context (Full). In the SF setting, only the sentences containing supporting facts for the answers are provided, whereas in the Full setting, several paragraphs from different documents are presented, which offers a more challenging evaluation of the model’s capabilities. Since the original test set is not publicly available, we use the split version provided by (Fei et al., 2022a). We annotate question keywords and document keywords in the HotpotQA dataset using SpaCy ³. During this process, we focus exclusively on sentences relevant to both answers and questions to ensure

²Experimental details are in Appendix A.

³<https://spacy.io/>. We utilize the en_core_web_sm model.




Figure 3: (a) shows our data annotation process; (b) illustrates the functions of question keywords and document keywords. K^q and K^d denote question keywords and document keywords, respectively.

	#Sam.	#Qes.			#Doc.		
		Min.	Max.	Avg.	Min.	Max.	Avg.
Train	89947	1	38	6	0	90	9
Dev	500	1	21	6	1	36	9
Test	7405	1	20	5	0	55	10

Table 1: Statistics of HotpotQA, where "Sam." denotes "Sample," "Qes." denotes "Question keyword," and "Doc." denotes "Document keyword."

that extracted keywords are directly related to them. Entities or phrases, drawn from both relevant sentences and questions, serve as keywords. Figure 3(a) illustrates our process of annotating keywords in detail. Notably, we observe that interrogative words (such as "what," "how," and so on) play a crucial role in the questioning process; therefore, we treat them as question keywords. Through this process, we effectively construct both question and document keywords, ensuring that both types align with the description provided. The details of the keyword statistics are provided in Table 1.

However, in Figure 3(a), it is found that a portion is marked as "Not utilized," which means that we do not adopt this portion’s keywords. This is because, after questioners view the answer and the document, they will pinpoint the crucial answer-related information snippets in the document while selecting the keywords from those snippets that are relevant to the question (i.e., keywords that must appear in the generated question). Hence, both question and document keywords should be included in the document. This is an intuitive explanation, and empirical explanations are offered in Appendix B.

4.2 Baselines and Evaluation Metrics

In our experiments, we compare several advanced models from recent years: **BART** (Lewis et al., 2020), an encoder-decoder pre-trained model designed for natural language generation; **MulQG**

Model	Support Fact Sentences (SF)			Full Document Context (Full)		
	B4	MTR	R-L	B4	MTR	R-L
BART (Lewis et al., 2020)	20.39	23.46	37.23	16.77	20.07	33.69
MulQG (Su et al., 2020)	-	-	-	15.20	20.51	35.30
CQG (Fei et al., 2022a)	25.09	27.45	41.83	21.46	24.97	39.61
MixQG (Murakhovs'ka et al., 2022)	25.45	26.36	43.21	22.13	23.78	41.21
QA4QG-large (Su et al., 2022)	25.70	27.44	<u>46.48</u>	21.21	<u>25.53</u>	42.44
E2EQR (Hwang et al., 2024)	21.73	27.47	41.34	-	-	-
SGCM (Ding et al., 2024)	26.16	28.51	44.06	22.61	26.04	40.61
(Ours) TS-BART	25.59	27.21	44.68	19.89	22.28	41.33
(Ours) DPKG _{hard}	26.80	<u>27.87</u>	46.50	<u>22.74</u>	24.90	43.29
(Ours) DPKG _{soft}	<u>26.19</u>	28.51	46.36	23.33	25.21	<u>43.18</u>

Table 2: Performance comparison between DPKG and baselines. **Bold** text highlights the best results, while underlined text indicates the second-best. "B4" denotes "BLEU-4," "MTR" denotes "METEOR," and "R-L" denotes "ROUGE-L."

(Su et al., 2020), which incorporates graph neural networks into multi-hop reasoning to facilitate question generation; **DP-Graph** (Pan et al., 2020), a generative model based on Graph2Seq to generate questions; **CQG** (Fei et al., 2022a), a model that controls key entities during decoding to ensure high-quality generation; **MixQG** (Murakhovs'ka et al., 2022), a generative model trained on diverse question-answering datasets with various answer types; **QA4QG** (Su et al., 2022), an enhanced version of BART augmented with a QA module to further constrain generation; **E2EQR** (Hwang et al., 2024), an end-to-end question rewriting model that increases question complexity through sequential rewriting; **SGCM** (Ding et al., 2024), which supports question generation by identifying salient and relevant sentences. In addition, we introduce **TS-BART** (Two-stage-BART). In the first stage, BART is fine-tuned to generate dual-perspective keywords; in the second stage, these keywords are integrated into the same BART to generate multi-hop questions. TS-BART is designed to further validate the effectiveness of both the dual-perspective keywords and the overall DPKG framework.

In line with previous studies, we primarily evaluate the DPKG framework using BLEU-4 (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE-L (Lin, 2004), and BERTScore (Zhang et al., 2020).

4.3 Results and Analysis

Our comparison between the DPKG framework and advanced MQG models is presented in Table 2. In the SF setting, DPKG_{hard} achieves the highest BLEU-4 score of 26.80 and ROUGE-L score of 46.50, while DPKG_{soft} attains the best METEOR score of 28.51. Both DPKG_{hard} and DPKG_{soft} achieve results comparable to current

leading levels. In the Full setting, DPKG_{hard} achieves the highest ROUGE-L score of 43.29, while DPKG_{soft} achieves the highest BLEU-4 score of 23.33. DPKG_{hard} and DPKG_{soft} secure the best or second-best performance on BLEU-4 and ROUGE-L, although their performance on METEOR is slightly worse. Across the six evaluation metrics, DPKG_{hard} and DPKG_{soft} rank first or second in most cases, especially in BLEU-4 and ROUGE-L. These results demonstrate that the DPKG framework excels at local matching and capturing long-range semantic dependencies, as reflected in its higher BLEU-4 and ROUGE-L scores.

On the whole, the DPKG framework demonstrates promising performance compared to recent advanced models, achieving state-of-the-art (SOTA) results in most cases. Notably, it outperforms the representative model CQG, which relies on question keywords during the generation process, thereby highlighting the positive impact of dual-perspective keywords in guiding multi-hop question generation. Although TS-BART demonstrates competitive results in the SF setting, its performance in the Full setting is relatively weaker. Nevertheless, it still outperforms BART by a large margin, further demonstrating the advantages of dual-perspective keywords and the overall effectiveness of the DPKG framework. These results underscore the feasibility and value of our dual-perspective keywords and the DPKG framework for the MQG task. It is worth noting that in the SF setting, DPKG_{hard} slightly outperforms DPKG_{soft}, while in the Full setting, DPKG_{soft} slightly outperforms DPKG_{hard}. This suggests that the role of keyword-guided question generation differs between the two modes, and we will analyze the effects of different keywords on question generation performance in detail in the next subsection.

Model	w/ generated keywords (w/ gk)					w/ ground-truth keywords (w/ gt)				
	KP	B4	MTR	R-L	BS	B4	MTR	R-L	BS	
Support Fact Sentences (SF)										
DPKG _{hard}	88.13	26.80	27.87	46.50	54.14	39.08	35.80	62.61	67.34	
DPKG _{soft}	88.86	26.19	28.51	46.36	54.05	39.22	34.58	61.98	66.51	
DPKG_D	87.05	<u>23.95</u>	25.20	<u>45.62</u>	<u>52.88</u>	27.04	27.33	47.88	<u>55.30</u>	
DPKG_Q	79.36	25.77	27.24	45.16	53.24	38.24	36.41	57.57	65.03	
Full Document Context (Full)										
DPKG _{hard}	85.39	22.74	24.90	43.29	50.68	37.67	33.28	60.63	65.35	
DPKG _{soft}	86.74	23.33	25.21	<u>43.18</u>	<u>50.26</u>	<u>35.90</u>	32.06	<u>59.68</u>	<u>64.06</u>	
DPKG_D	84.53	18.62	20.83	41.52	48.19	20.20	21.91	42.97	49.88	
DPKG_Q	77.85	21.52	23.92	42.53	49.90	34.24	<u>32.64</u>	55.61	62.95	

Table 3: Effect of different keywords on question generation performance. "BS" denotes "BERTScore."

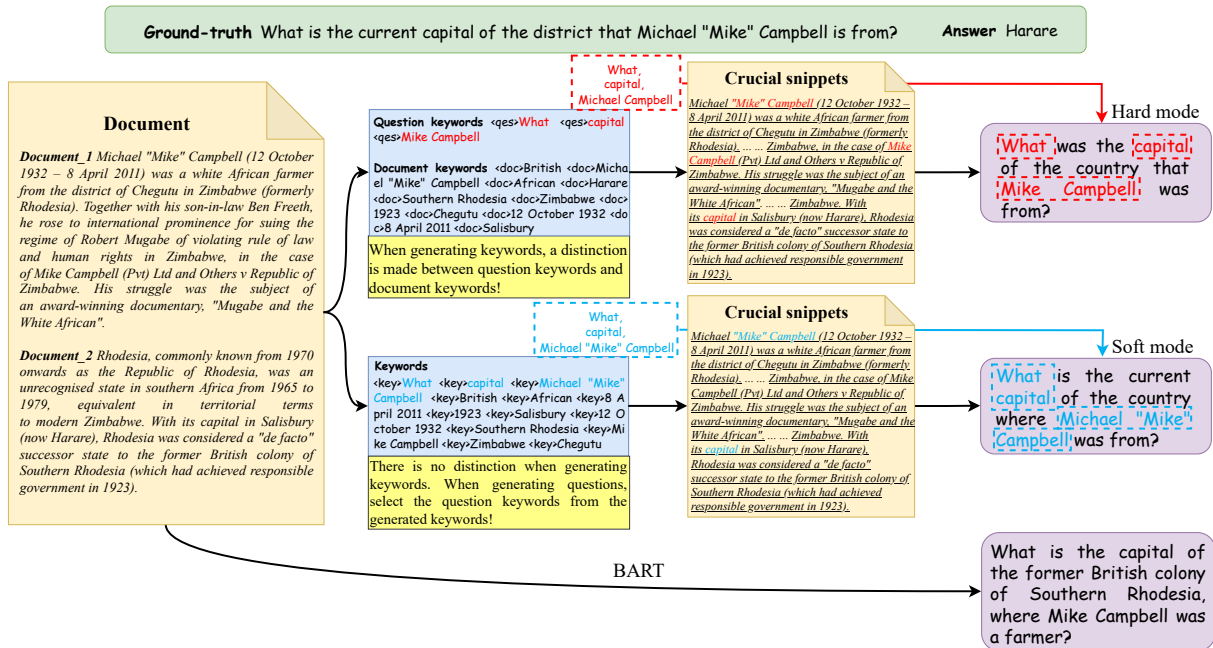


Figure 4: An example of multi-hop question generation is provided, where the underlined text highlights crucial information snippets related to the QA pair pinpointed by dual-perspective keywords. The **red** denotes the question keywords in hard mode, while the **blue** represents the question keywords in soft mode.

4.4 Effect of Different Keywords on Results

In this subsection, we examine the effects of different keywords on multi-hop question generation performance. In addition to DPKG_{hard} and DPKG_{soft}, we introduce two new model variants: DPKG_D, which uses only document keywords in the DPKG framework, and DPKG_Q, which uses only question keywords. Table 3 shows the performance of these model variants. "KP" in the table denotes keyword generation performance, evaluated using BERTScore. We evaluate DPKG's question generation performance using self-generated keywords ("w/ generated keywords," abbreviated as "w/ gk") and ground-truth keywords ("w/ ground-truth keywords," abbreviated as "w/ gt"). This al-

lows us to assess its ability to generate questions in practice and evaluate the effectiveness of keyword-guided question generation.

As shown in Table 3, the question generation performance of DPKG_Q significantly surpasses that of DPKG_D in both the SF and Full settings, despite its slightly lower keyword generation performance. Furthermore, when ground-truth keywords are used, DPKG_Q still outperforms DPKG_D. This reveals that question keywords tend to capture richer information than document keywords in question generation. The keyword generation performance of DPKG_D and DPKG_Q is lower than that of DPKG_{hard} and DPKG_{soft}, indicating that considering both types of keywords improves keyword generation performance. While DPKG_Q has

a relatively lower keyword generation performance, its question generation performance remains competitive. However, this does not imply that considering only question keywords is better than considering both, as even with ground-truth keywords, DPKG_{hard} and DPKG_{soft} consistently outperform DPKG_Q . DPKG_{hard} and DPKG_{soft} rank first or second in nearly all cases, demonstrating that guidance from dual-perspective keywords significantly enhances question generation performance.

With generated keywords, DPKG_{hard} generally outperforms DPKG_{soft} in the SF setting, except for its marginally worse performance on METEOR. Conversely, DPKG_{soft} performs particularly well in the Full setting, especially on BLEU-4 and METEOR. In terms of keyword generation performance, DPKG_{soft} slightly exceeds DPKG_{hard} , particularly in the Full setting. This suggests that DPKG_{soft} excels with generated questions mainly due to its superior keyword generation performance. However, DPKG_{hard} still achieves the best performance in the SF setting, despite its slightly lower keyword generation performance. With ground-truth keywords, DPKG_{hard} consistently and significantly outperforms DPKG_{soft} . These results highlight that dual-perspective keywords offer more powerful guidance in hard mode than in soft mode.

Moreover, the superior performance of DPKG_{hard} with ground-truth keywords, especially in the Full setting, highlights the ability of dual-perspective keywords to effectively pinpoint crucial information snippets in long documents with abundant noise. Hard mode incorporates special prefixes to indicate each keyword’s role during keyword generation, making the process more challenging. This explains why keyword generation performance is slightly worse in hard mode compared to soft mode. In contrast, soft mode requires identifying each keyword’s role and generating questions during question generation, which provides weaker guidance than hard mode. These analyses demonstrate the importance of further improving keyword generation performance in hard mode to enhance question generation performance. We also investigate the performance of large language models (LLMs) on multi-hop question generation, as shown in Appendix C.

4.5 Ablation Study

In this subsection, we conduct an ablation study to verify the effectiveness of each mechanism in

Model	Support Fact Sentences (SF)			
	B4	MTR	R-L	BS
DPKG_{hard}	26.80	27.87	46.50	54.14
w/o \mathcal{L}_3	24.74	27.83	45.08	52.93
w/o AA	24.83	27.80	45.09	52.94
DPKG_{soft}	26.19	28.51	46.36	54.05
w/o \mathcal{L}_3	24.59	27.56	44.93	52.68
w/o AA	24.84	27.28	44.82	52.65
Full Document Context (Full)				
DPKG_{hard}	22.74	24.90	43.29	50.68
w/o \mathcal{L}_3	21.16	23.90	42.01	49.61
w/o AA	21.70	24.31	42.15	49.57
DPKG_{soft}	23.33	25.21	43.18	50.26
w/o \mathcal{L}_3	21.42	24.04	41.77	49.20
w/o AA	22.41	24.89	42.59	50.05

Table 4: Ablation study.

the DPKG framework, as shown in Table 4. We examine the following mechanisms in the ablation experiments: (1) without the \mathcal{L}_3 loss (**w/o \mathcal{L}_3**); and (2) without the answer-aware attention mechanism (**w/o AA**). From the results in Table 4, we observe that both \mathcal{L}_3 and the answer-aware attention mechanism significantly contribute to the performance of DPKG_{hard} and DPKG_{soft} in both the SF and Full settings. These findings demonstrate that \mathcal{L}_3 effectively bridges the gap discussed in §3.5, while the answer-aware attention mechanism, as detailed in §3.3, enhances the quality of the generated questions. The synergistic interaction of DPKG’s mechanisms further elevates overall performance.

4.6 Case Study

To present the quality of multi-hop question generation across different models and to emphasize the role of keywords in guiding this process, we conduct a case study. We focus on DPKG_{hard} , DPKG_{soft} , and BART, with BART being unguided by keywords. The case is presented in Figure 4.

It is clear that BART, lacking keyword guidance, generates questions of lower quality. In contrast, both DPKG_{hard} and DPKG_{soft} benefit from keyword guidance and produce higher-quality questions. Specifically, in hard mode, question keywords directly specify terms that should appear in the question, while both question and document keywords work together to pinpoint essential information snippets related to the QA pair in the document, resulting in more concise and clearer questions. In soft mode, DPKG dynamically identifies each keyword’s role during question generation, leading to more flexible questions. While both DPKG_{hard} and DPKG_{soft} have their unique

characteristics, they pinpoint the same crucial information snippets in the document. Overall, compared with the model without keyword guidance, the DPKG framework generates more understandable and answerable questions, demonstrating the effectiveness of both the DPKG framework and dual-perspective keywords. We also conduct extended analytical experiments; see Appendix D for details.

5 Conclusion

In this paper, we define dual-perspective keywords: question keywords, which capture the questioner’s intent, and document keywords, which reflect the content related to QA pairs in documents. Furthermore, we propose a DPKG framework that seamlessly integrates keywords into the multi-hop question generation process. We analyze the role of different keywords in guiding question generation and demonstrate that dual-perspective keywords are effective, especially in hard mode, providing more powerful guidance. Extensive experiments demonstrate the effectiveness of our work, showcasing its promising performance and underscoring its significant value in the MQG task.

Limitations

We categorize keywords into question keywords and document keywords and propose a DPKG framework to seamlessly integrate them into the multi-hop question generation process. Although extensive experiments have demonstrated the effectiveness of our work, several limitations remain, such as the limited performance of the proposed keywords in guiding large language models (LLMs) for multi-hop question generation (Appendix C). This is because multi-hop question generation requires not only a better understanding of semantics but also alignment with the answers. LLMs are often prone to hallucinations, and recent studies have shown that training dedicated PLMs can achieve better results than using large models (Ushio et al., 2023; Liu et al., 2024b). Eliminating hallucinations when prompting LLMs to generate multi-hop questions will be an interesting challenge. Additionally, the focus of our study is on MQG, also known as deep question generation (DQG) (Pan et al., 2020; Fei et al., 2022a; Xia et al., 2023). Suitable datasets for this research are limited, highlighting the need to develop new, appropriate datasets, which is highly meaningful.

Acknowledgement

This work was supported by the Project 62276178 under the National Natural Science Foundation of China, the Key Project 23KJA520012 under the Natural Science Foundation of Jiangsu Higher Education Institutions, the project 22YJCZH091 of Humanities and Social Science Fund of Ministry of Education and the Priority Academic Program Development of Jiangsu Higher Education Institutions.

References

- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Shuyang Cao and Lu Wang. 2021. [Controllable open-ended question generation with a new question type ontology](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6424–6439, Online. Association for Computational Linguistics.
- Yllias Chali and Sadid A. Hasan. 2012. [Towards automatic topical question generation](#). In *Proceedings of COLING 2012*, pages 475–492, Mumbai, India. The COLING 2012 Organizing Committee.
- Yu Chen, Lingfei Wu, and Mohammed J. Zaki. 2020. [Reinforcement learning based graph-to-sequence model for natural question generation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Yu Chen, Lingfei Wu, and Mohammed J Zaki. 2023. [Toward subgraph-guided knowledge graph question generation with graph neural networks](#). *IEEE Transactions on Neural Networks and Learning Systems*.
- Chuyao Ding, Yu Hong, and Jianmin Yao. 2024. [SGCM: Saliency-guided context modeling for question generation](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 14755–14762, Torino, Italia. ELRA and ICCL.
- Xinya Du, Junru Shao, and Claire Cardie. 2017. [Learning to ask: Neural question generation for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1342–1352, Vancouver, Canada. Association for Computational Linguistics.

- Barah Fazili, Koustava Goswami, Natwar Modani, and Inderjeet Nair. 2024. Gensco: Can question decomposition based passage alignment improve question answering? *arXiv preprint arXiv:2407.10245*.
- Zichu Fei, Qi Zhang, Tao Gui, Di Liang, Sirui Wang, Wei Wu, and Xuanjing Huang. 2022a. **CQG: A simple and effective controlled generation framework for multi-hop question generation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6896–6906, Dublin, Ireland. Association for Computational Linguistics.
- Zichu Fei, Qi Zhang, and Yaqian Zhou. 2021. **Iterative GNN-based decoder for question generation**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2573–2582, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zichu Fei, Xin Zhou, Tao Gui, Qi Zhang, and Xuanjing Huang. 2022b. **LFKQG: A controlled generation framework with local fine-tuning for question generation over knowledge bases**. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6575–6585, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Shasha Guo, Lizi Liao, Cuiping Li, and Tat-Seng Chua. 2024. **A survey on neural question generation: Methods, applications, and prospects**. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, pages 8038–8047. International Joint Conferences on Artificial Intelligence Organization. Survey Track.
- Michael Heilman and Noah A. Smith. 2010. **Good question! statistical ranking for question generation**. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 609–617, Los Angeles, California. Association for Computational Linguistics.
- Seonjeong Hwang, Yunsu Kim, and Gary Geunbae Lee. 2024. **Explainable multi-hop question generation: An end-to-end approach without intermediate question labeling**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy*, pages 6855–6866. ELRA and ICCL.
- Kyungho Kim, Seongmin Park, Junseo Lee, and Jihwa Lee. 2024. **Non-essential is NEcessary: Order-agnostic multi-hop question generation**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 12300–12306, Torino, Italia. ELRA and ICCL.
- 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 *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yuanyuan Liang, Jianing Wang, Hanlun Zhu, Lei Wang, Weining Qian, and Yunshi Lan. 2023. **Prompting large language models with chain-of-thought for few-shot knowledge base question generation**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4329–4343, Singapore. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Jintao Liu, Ruixue Ding, Linhao Zhang, Pengjun Xie, and Fie Huang. 2024a. **Cofe-rag: A comprehensive full-chain evaluation framework for retrieval-augmented generation with enhanced data diversity**. *Preprint*, arXiv:2410.12248.
- Naiming Liu, Zichao Wang, and Richard Baraniuk. 2024b. **Synthetic context generation for question generation**. *Preprint*, arXiv:2406.13188.
- Karen Mazidi and Rodney D. Nielsen. 2014. **Linguistic considerations in automatic question generation**. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 321–326, Baltimore, Maryland. Association for Computational Linguistics.
- Lidiya Murakhovs’ka, Chien-Sheng Wu, Philippe Laban, Tong Niu, Wenhao Liu, and Caiming Xiong. 2022. **MixQG: Neural question generation with mixed answer types**. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1486–1497, Seattle, United States. Association for Computational Linguistics.
- Liangming Pan, Yuxi Xie, Yansong Feng, Tat-Seng Chua, and Min-Yen Kan. 2020. **Semantic graphs for generating deep questions**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1463–1475, Online. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. **SQuAD: 100,000+ questions for machine comprehension of text**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.

- Lin Feng Song, Zhiguo Wang, Wael Hamza, Yue Zhang, and Daniel Gildea. 2018. [Leveraging context information for natural question generation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 569–574, New Orleans, Louisiana. Association for Computational Linguistics.
- Dan Su, Peng Xu, and Pascale Fung. 2022. Qa4qg: using question answering to constrain multi-hop question generation. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8232–8236. IEEE.
- Dan Su, Yan Xu, Wenliang Dai, Ziwei Ji, Tiezheng Yu, and Pascale Fung. 2020. [Multi-hop question generation with graph convolutional network](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4636–4647, Online. Association for Computational Linguistics.
- Xingwu Sun, Jing Liu, Yajuan Lyu, Wei He, Yanjun Ma, and Shi Wang. 2018. [Answer-focused and position-aware neural question generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3930–3939, Brussels, Belgium. Association for Computational Linguistics.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. [MuSiQue: Multi-hop questions via single-hop question composition](#). *Transactions of the Association for Computational Linguistics*, 10:539–554.
- Asahi Ushio, Fernando Alva-Manchego, and Jose Camacho-Collados. 2023. [A practical toolkit for multilingual question and answer generation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 86–94, Toronto, Canada. Association for Computational Linguistics.
- A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Jian Wang, Dongding Lin, and Wenjie Li. 2024. [A target-driven planning approach for goal-directed dialog systems](#). *IEEE Transactions on Neural Networks and Learning Systems*, 35(8):10475–10487.
- Zehua Xia, Qi Gou, Bowen Yu, Haiyang Yu, Fei Huang, Yongbin Li, and Cam-Tu Nguyen. 2023. [Improving question generation with multi-level content planning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 800–814. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Ruqing Zhang, Jiafeng Guo, Lu Chen, Yixing Fan, and Xueqi Cheng. 2021. [A review on question generation from natural language text](#). *ACM Trans. Inf. Syst.*, 40(1).
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Yao Zhao, Xiaochuan Ni, Yuanyuan Ding, and Qifa Ke. 2018. [Paragraph-level neural question generation with maxout pointer and gated self-attention networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3901–3910, Brussels, Belgium. Association for Computational Linguistics.
- Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2018. Neural question generation from text: A preliminary study. In *Natural Language Processing and Chinese Computing*, pages 662–671, Cham. Springer International Publishing.

	#Sam.	#Qes.			#Doc.		
		Min.	Max.	Avg.	Min.	Max.	Avg.
Train	89947	1	52	9	0	96	13
Dev	500	2	45	8	2	43	13
Test	7405	1	24	8	0	63	14

Table 5: Statistics of HotpotQA (version 2), where "Sam." denotes "Sample," "Qes." denotes "Question keyword," and "Doc." denotes "Document keyword."

A Experimental Details

We employ AdamW as the optimizer, incorporating warmup and gradient clipping techniques. The batch size is set to 8, with a learning rate of 1×10^{-5} . Training is conducted for 50 epochs with an early stopping strategy. The hyperparameters β_1 , β_2 , and β_3 are all set to 1.0, and the beam size is 8. All experiments are performed on a GeForce RTX 3090 GPU.

In our DPKG framework, the encoder and decoder are modified based on the transformer architecture, with most of their parameters being compatible with BART’s pre-trained model. Therefore, during training, we load the parameters compatible with BART’s pre-trained model⁴, initialize the incompatible parameters with a mean of 0 and a variance of 0.02, and then fine-tune all parameters. Additionally, it is worth noting that the hyperparameter settings not discussed in this paper are identical to those of BART-base, including the number of encoder and decoder layers, among others.

B Exploring Different Data Annotation Methods

In this paper, we use the method shown in Figure 3(a) to annotate question keywords and document keywords. However, a question arises: why is there a "Not utilized" portion? While we provide an intuitive explanation in the main text, this subsection offers an empirical explanation.

We consider the "Not utilized" portion in Figure 3(a) as the question keywords. In this way, question keywords are extracted from the question, and document keywords are extracted from the relevant sentences. Question keywords and document keywords are independent, with no overlap. Additionally, interrogative words (such as "what," "how," and so on) are treated as question keywords. The statistics for version 2 annotations are shown in Table 5. We observe that the number of question

keywords and document keywords in this version has increased significantly compared to those in Table 1. We use the annotated dataset here and adopt the hard mode of DPKG, with the new model called DPKG_{v2}. The performance comparison with DPKG_{hard} and DPKG_{soft} is shown in Table 6.

From the results in Table 6, it is worth noting that the multi-hop question generation performance of DPKG_{v2} with generated keywords is significantly worse than that of DPKG_{hard} and DPKG_{soft}, especially in the Full setting. The performance of keyword generation remains poor. These results indicate that increasing the number of keywords hinders the performance of keyword generation and further inhibits the improvement of question generation performance. DPKG_{v2} performs significantly better under ground-truth keywords, outperforming both DPKG_{hard} and DPKG_{soft}. This is intuitive, as both the question and document keywords in this annotation method provide more detailed information. For example, when generating the question, all terms related to it are included (with the "Not utilized" portion treated as question keywords). In summary, while this annotation method provides more information, its practical performance is often unsatisfactory. The data annotation method used in the main text strikes a trade-off: it achieves SOTA performance in practice while also providing valuable information. This empirically explains why the "Not utilized" portion is included in the annotated approach of our main text.

C Results of Large Language Models on Question Generation

C.1 Results of Prompting Large Language Models

In this subsection, we examine the performance of large language models (LLMs) prompted by keywords on multi-hop question generation. Specifically, to explore the guiding role of dual-perspective keywords, we conduct experiments using ground-truth keywords. Our primary focus is the Qwen-Plus model⁵, which performs well in English and understands prompt instructions. We evaluate its question generation capabilities in both zero-shot and one-shot settings, using the prompt template illustrated in Figure 5. The experimental results are presented in Table 7, where "w/o keyword guidance" means no keyword guidance, and

⁴<https://huggingface.co/facebook/bart-base>

⁵<https://bailian.console.aliyun.com>

	w/ generated keywords (w/ gk)					w/ ground-truth keywords (w/ gt)			
	Support Fact Sentences (SF)								
Model	KP	B4	MTR	R-L	BS	B4	MTR	R-L	BS
DPKG _{hard}	<u>88.13</u>	26.80	<u>27.87</u>	46.50	54.14	39.08	<u>35.80</u>	<u>62.61</u>	<u>67.34</u>
DPKG _{soft}	88.86	<u>26.19</u>	28.51	<u>46.36</u>	<u>54.05</u>	<u>39.22</u>	34.58	<u>61.98</u>	<u>66.51</u>
DPKG _{v2}	86.38	<u>25.79</u>	26.95	45.94	53.34	52.34	42.73	76.09	78.31
	Full Document Context (Full)								
DPKG _{hard}	<u>85.39</u>	<u>22.74</u>	<u>24.90</u>	43.29	50.68	<u>37.67</u>	33.28	<u>60.63</u>	<u>65.35</u>
DPKG _{soft}	86.74	23.33	25.21	<u>43.18</u>	<u>50.26</u>	35.90	32.06	59.68	64.06
DPKG _{v2}	83.48	18.40	21.05	41.53	47.98	51.77	42.45	75.44	77.73

Table 6: Performance comparison. **Bold** text highlights the best results, while underlined text indicates the second-best. "B4" denotes "BLEU-4," "MTR" denotes "METEOR," "R-L" denotes "ROUGE-L," "BS" denotes "BERTScore" and "KP" stands for "keyword generation performance."

<u>Prompt (w/ keyword guidance)</u>
<p>You are very good at multi-hop question generation. Multi-hop question generation (MQG) aims to generate questions that require synthesizing multiple information snippets in documents to derive answers.</p> <p>Document: {}; Answer: {}; Question keywords: {}; Document keywords: {}.</p> <p>Important Note:</p> <p>(1) It is known that question keywords must appear in the generated question, and both question keywords and document keywords can help identify snippets related to question-answer pairs.</p> <p>(2) You only need to output the generated question without any redundant information.</p> <p>Task: You need to use these snippets to generate a multi-hop question that matches the answer.</p>
<u>Prompt (w/o keyword guidance)</u>
<p>You are very good at multi-hop question generation. Multi-hop question generation (MQG) aims to generate questions that require synthesizing multiple information snippets in documents to derive answers.</p> <p>Document: {}; Answer: {}.</p> <p>Important Note: You only need to output the generated question without any redundant information.</p> <p>Task: You need to use these snippets to generate a multi-hop question that matches the answer.</p>

Figure 5: Prompt template.

"w/ keyword guidance" refers to dual-perspective keyword guidance.

According to the results in Table 7, we find that under the guidance of dual-perspective keywords, the quality of LLM-generated questions has greatly improved, and the performance of Qwen-Plus in

the one-shot setting has significantly outperformed that in the zero-shot setting. These results show that keyword guidance improves question generation performance and demonstrate the effectiveness of the dual-perspective keywords proposed in this paper. However, the performance of Qwen-Plus in

Model	Support Fact Sentences (SF)			
	B4	MTR	R-L	BS
w/o keyword guidance				
Qwen-plus				
w/ zero-shot	13.60	19.56	33.88	41.92
w/ one-shot	14.97	20.68	34.42	42.87
w/ keyword guidance				
Qwen-plus				
w/ zero-shot	20.35	25.44	44.28	52.07
w/ one-shot	22.81	27.57	46.46	54.47
DPKG _{hard}	<u>39.08</u>	35.80	62.61	67.34
DPKG _{soft}	39.22	<u>34.58</u>	<u>61.98</u>	<u>66.51</u>
Full Document Context (Full)				
w/o keyword guidance				
Qwen-plus				
w/ zero-shot	10.62	16.82	30.34	38.11
w/ one-shot	11.17	17.37	30.39	38.18
w/ keyword guidance				
Qwen-plus				
w/ zero-shot	18.77	24.14	42.71	50.52
w/ one-shot	20.91	26.18	44.88	52.70
DPKG _{hard}	37.67	33.28	60.63	65.35
DPKG _{soft}	<u>35.90</u>	<u>32.06</u>	<u>59.68</u>	<u>64.06</u>

Table 7: Performance comparison between the DPKG framework and prompting LLMs on multi-hop question generation.

all cases still lags behind that of the DPKG framework proposed in this paper. This is because multi-hop question generation requires not only a better understanding of semantics but also alignment with the answers. LLMs are often prone to hallucinations. Although hallucinations are somewhat mitigated under the guidance of keywords, they remain prevalent. Some studies have shown that training dedicated PLMs can achieve better results than prompting large models (Ushio et al., 2023; Liu et al., 2024b). Overall, we find that the dual-perspective keywords presented in this paper effectively prompt LLMs, though they are less effective than the dedicated DPKG framework. These results highlight the urgent need for a framework that effectively prompts LLMs using keywords and eliminates hallucinations when generating multi-hop questions.

C.2 Results of Fine-tuned Large Language Models

In this subsection, we evaluate the performance of fine-tuned LLMs for multi-hop question generation.

Model	B4	MTR	R-L	BS
	Support Fact Sentences (SF)			
Llama-3.1-8B (Lora)	25.29	26.68	46.73	54.63
DPKG _{hard}	26.80	<u>27.87</u>	<u>46.50</u>	<u>54.14</u>
DPKG _{soft}	<u>26.19</u>	28.51	46.36	54.05
Full Document Context (Full)				
Llama-3.1-8B (Lora)	22.49	24.67	44.38	51.98
DPKG _{hard}	22.74	24.90	43.29	50.68
DPKG _{soft}	23.33	25.21	43.18	50.26

Table 8: Performance comparison between the DPKG framework and fine-tuned LLMs on multi-hop question generation.

Specifically, we use Llama-3.1-8B⁶ (abbreviated as Llama), fine-tuned via the LoRA method. The results of the experiment are displayed in Table 8.

It can be seen from Table 8 that in both the SF and Full settings, Llama’s performance on BLEU-4 and METEOR is inferior to that of the DPKG framework, while it performs better on ROUGE-L and BERTScore. Although Llama has a large number of parameters, it does not demonstrate stronger performance compared to the DPKG framework. On the one hand, this highlights the difficulty of the MQG task; on the other, it underscores the effectiveness of the dual-perspective keywords and the DPKG framework. These findings again suggest the challenges large models face in generating multi-hop questions — a problem that remains to be addressed in future research.

D Extended Analysis

D.1 Additional Case Analysis

We provide a more detailed analysis here. We have previously examined the specific case of keyword-guided multi-hop question generation (see section 4.6). In this section, we present another concrete example.

As you can see from Table 9, DPKG_{hard} and DPKG_{soft} pinpoint the same essential information snippets. While they identify different question keywords, the resulting questions are multi-hop and contain a wealth of information that can be understood and answered. The table contains rich information, and we use different colors to highlight it. We observe that the same words may be assigned different roles in different modes, which affects the resulting question. For example, "Trig-

⁶<https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>

Document	<u>Dana Brunetti</u> (born June 11, 1973) is an <u>American film producer</u> and social networking entrepreneur. He is the <u>president</u> of <u>Kevin Spacey's production company</u> , <u>Trigger Street Productions</u> . In 2016, he became the presidents of Relativity Media with Kevin Spacey taking on a chairman position. <u>Kevin Spacey Fowler, KBE</u> (born July 26, 1959) is an <u>American actor, film director, producer</u> , screenwriter, and singer. He began his career as a stage <u>actor</u> during the 1980s before obtaining supporting roles in <u>film</u> and television. He gained critical acclaim in the early 1990s that culminated in his first Academy Award for Best Supporting Actor for the neo-noir crime thriller "The Usual Suspects" (1995), and an Academy Award for Best Actor for midlife crisis-themed drama " <u>American Beauty</u> " (1999).
DPKG _{hard}	When was the <u>American actor, film director, producer</u> , screenwriter, and singer <u>born</u> who's <u>production company</u> is <u>Trigger Street Productions</u> ?
Document	<u>Dana Brunetti</u> (born June 11, 1973) is an <u>American film producer</u> and social networking entrepreneur. He is the <u>president</u> of <u>Kevin Spacey's production company</u> , <u>Trigger Street Productions</u> . In 2016, he became the presidents of Relativity Media with Kevin Spacey taking on a chairman position. <u>Kevin Spacey Fowler, KBE</u> (born July 26, 1959) is an <u>American actor, film director, producer</u> , screenwriter, and singer. He began his career as a stage <u>actor</u> during the 1980s before obtaining supporting roles in <u>film</u> and television. He gained critical acclaim in the early 1990s that culminated in his first Academy Award for Best Supporting Actor for the neo-noir crime thriller "The Usual Suspects" (1995), and an Academy Award for Best Actor for midlife crisis-themed drama " <u>American Beauty</u> " (1999).
DPKG _{soft}	<u>Dana Brunetti</u> is the <u>president</u> of a <u>production company</u> owned by an <u>American actor, film director, producer</u> , screenwriter, and singer, <u>born</u> on what date?
BART	Dana Brunetti is the president of the production company of an American actor born on what date?
Ground-truth	When was the American actor, film director which Dana Brunetti is the president of his company born

Table 9: Another case: The **red** represents the question keywords in hard mode, while the **blue** represents those in soft mode. The **purple** indicates the document keywords, and the underlined text highlights essential information snippets.

ger Street Productions" is identified as a question keyword in hard mode and as a document keyword in soft mode. There are many more cases like this, which demonstrate that hard mode and soft mode have their own characteristics, as we discussed in section 4.6. In addition, we find that the questions generated by BART are more straightforward, contain less multi-hop information, and are not as rich as DPKG_{hard} and DPKG_{soft}. This is mainly because BART does not have keyword guidance, does not know what is essential, and can simply generate questions through the mappings it learns during training.

In short, through the analysis of the above cases, we gain a more intuitive understanding of the necessity of keyword-guided multi-hop question genera-

tion, especially the two types of keywords proposed in this paper. The two modes of the framework proposed here each have their own characteristics, revealing their significant value in the MQG task.

D.2 Results of Generated Questions on QA

To further verify the validity of the generated questions, we evaluate them using a question-answering (QA) system. Specifically, we use Qwen-plus⁷ as our QA system, to which we input the generated questions and obtain answers. The evaluation metrics we adopt are Exact Match (EM) and F1 score (F1). The results of the experiment are shown in Table 10. From the experimental results, we can see that even Ground-truth does not achieve strong per-

⁷<https://bailian.console.aliyun.com>

Model	SF		Full	
	EM	F1	EM	F1
BART	54.71	70.43	54.86	69.28
DPKG _{hard}	58.76	73.65	<u>56.31</u>	71.05
DPKG _{soft}	<u>57.19</u>	<u>72.21</u>	56.79	<u>70.99</u>
Ground-truth	63.24	79.02	61.63	77.29

Table 10: Results of generated questions on the QA system. "Ground-truth" indicates the use of gold-standard questions.

Model	Complexity	Relevance	Fluency	Average
BART	3.17	3.28	3.24	3.23
TS-BART	3.68	3.75	3.62	3.68
DPKG _{hard}	4.54	<u>4.55</u>	<u>4.50</u>	<u>4.53</u>
DPKG _{soft}	<u>4.53</u>	4.57	4.56	4.55
Ground-truth	4.94	4.97	4.92	4.94

Table 11: Human evaluation results. "Average" denotes the average of complexity, relevance, and fluency. "Ground-truth" indicates the use of gold-standard questions.

formance, which may be related to the limitations of the QA system itself. Overall, the DPKG framework achieves competitive results, significantly outperforming BART, though still lagging behind the Ground-truth. These results further validate the effectiveness of the DPKG framework from the perspective of question answering.

D.3 Human Evaluation

While automatic evaluation offers valuable insights into model performance, n-gram-based metrics have inherent limitations. Therefore, we conduct a human evaluation by randomly selecting 50 questions from the generated set in the Full setting. Based on experimental results, we evaluate the following representative models: (1) **DPKG_{hard}**; (2) **DPKG_{soft}**; (3) **TS-BART**; and (4) **BART**. Three crowdsourced workers rate each sample on a scale from 1 (poor) to 5 (good) across three key dimensions: (i) complexity, which assesses whether the question requires reasoning across multiple snippets in the document; (ii) relevance, which determines if the question is answerable and contextually appropriate; and (iii) fluency, which evaluates whether the question adheres to grammatical rules and maintains coherent logic. We summarize and average the workers' ratings, presenting the results in Table 11.

According to the results displayed in Table 11, the DPKG framework has achieved very promis-

ing results compared with BART and TS-BART. Specifically, BART is a vanilla model that does not use dual-perspective keywords, while TS-BART first uses vanilla BART to generate the keywords, and then reuses it to generate questions guided by those keywords. TS-BART outperforms BART in both automatic and human evaluations, fully demonstrating the guiding role of dual-perspective keywords. The DPKG framework outperforms TS-BART, further confirming its overall effectiveness. DPKG_{hard} performs slightly worse than DPKG_{soft}. DPKG_{hard} and DPKG_{soft} perform comparably in terms of Complexity and Relevance, while DPKG_{soft} shows slightly better performance in Fluency. These results indicate that the questions generated by DPKG_{soft} are more fluent in the Full setting, which is consistent with the automatic evaluation results. All in all, the human evaluation results further demonstrate the effectiveness of the dual-perspective keywords and the DPKG framework.

E Generalization to Other Datasets

To further verify the effectiveness of the DPKG framework, we also use another common dataset, MusiQue-2hop (Trivedi et al., 2022), and follow the same keyword annotation procedure for questions and documents as in HotpotQA. Notably, since the test set of MusiQue-2hop is not publicly available, we construct a custom split. Table 12 provides detailed keyword statistics for MusiQue-2hop. We conduct two types of experiments: comparing our results with current advanced models, and verifying the generated questions using a QA system. For the comparative evaluation, our baselines include—in addition to CQG, MulQG, and E2EQR—DP-Graph (Pan et al., 2020), a generative Graph2Seq-based question generation model. For the QA evaluation, we use Qwen-plus as our QA system, into which we input the generated questions to obtain answers. Table 13 shows the results of the comparative evaluation, while Table 14 presents the QA evaluation results.

Table 13 shows that the DPKG framework achieves the best performance compared to other models, indicating its effectiveness. More specifically, DPKG_{soft} performs slightly better than DPKG_{hard}, consistent with our earlier conclusion. Table 14 shows that Ground-truth does not perform strongly, which may be due to the limitations of the QA system. Although DPKG_{soft} performs slightly

	#Sam.	#Qes.			#Doc.		
		Min.	Max.	Avg.	Min.	Max.	Avg.
Train	13124	1	12	3	3	170	24
Dev	1252	1	11	3	4	134	24
Test	1252	1	9	3	4	134	24

Table 12: Statistics of MusiQue-2hop, where "Sam." denotes "Sample," "Qes." denotes "Question keyword," and "Doc." denotes "Document keyword."

Model	B4	MTR	R-L
DP-Graph (Pan et al., 2020)	5.14	10.80	28.69
CQG (Fei et al., 2022a)	9.64	16.20	33.98
MulQG (Su et al., 2020)	9.56	15.41	37.18
E2EQR (Hwang et al., 2024)	20.33	25.64	44.01
(Ours) DPKG _{hard}	28.33	27.77	51.93
(Ours) DPKG _{soft}	29.03	28.14	52.38

Table 13: Performance comparison between DPKG and baselines on MusiQue-2hop.

Model	EM	F1
DPKG _{hard}	<u>41.45</u>	50.95
DPKG _{soft}	41.53	<u>50.93</u>
Ground-truth	58.79	71.42

Table 14: Results of generated questions evaluated by the QA system on MusiQue-2hop. "Ground-truth" indicates the use of gold-standard questions.

better than DPKG_{hard} on n-gram-based metrics, DPKG_{hard} slightly outperforms DPKG_{soft} on the QA system. This suggests that the QA system itself has some level of semantic understanding. This also highlights the limitations of n-gram-based metrics and the necessity of using the QA system to evaluate the generated questions. To sum up, the experiments on the MusiQue-2hop further demonstrate the effectiveness and generalizability of the DPKG framework.