

# Weakly Supervised Pre-Training for Multi-Hop Retriever

Yeon Seonwoo<sup>†</sup>, Sang-Woo Lee<sup>‡§</sup>, Ji-Hoon Kim<sup>‡§</sup>, Jung-Woo Ha<sup>‡§</sup>, Alice Oh<sup>†</sup>

<sup>†</sup>KAIST

<sup>‡</sup>NAVER AI Lab, <sup>§</sup>NAVER Clova

yeon.seonwoo@kaist.ac.kr

{sang.woo.lee, genesis.kim, jungwoo.ha}@navercorp.com

alice.oh@kaist.edu

## Abstract

In multi-hop QA, answering complex questions entails iterative document retrieval for finding the missing entity of the question. The main steps of this process are sub-question detection, document retrieval for the sub-question, and generation of a new query for the final document retrieval. However, building a dataset that contains complex questions with sub-questions and their corresponding documents requires costly human annotation. To address the issue, we propose a new method for weakly supervised multi-hop retriever pre-training without human efforts. Our method includes 1) a pre-training task for generating vector representations of complex questions, 2) a scalable data generation method that produces the nested structure of question and sub-question as weak supervision for pre-training, and 3) a pre-training model structure based on dense encoders. We conduct experiments to compare the performance of our pre-trained retriever with several state-of-the-art models on end-to-end multi-hop QA as well as document retrieval. The experimental results show that our pre-trained retriever is effective and also robust on limited data and computational resources.

## 1 Introduction

Multi-hop QA is the task of answering complex questions that requires reasoning across multiple documents (Nogueira and Cho, 2017; Nie et al., 2019; Sun et al., 2019; Fang et al., 2020; Zhao et al., 2020). The core components of multi-hop reasoning are identifying the missing entity in the question and generating a new query with the missing entity. Figure 1 shows an example of the reasoning process in multi-hop QA. In the example, the missing entity, which we call bridge entity, of the question is “*Jupiter*”. To answer the question, the correct document for the sub-question “the largest

planet in the Solar System,” should be retrieved. Supervised training of the multi-hop QA models for these intermediate reasoning steps requires a dataset of complex questions, sub-questions, and their corresponding documents. However, building such dataset requires costly human annotation and cannot be done at scale (Min et al., 2019; Wolfson et al., 2020).

When there is limited annotated supervision signal, weakly supervised pre-training can be a solution (Devlin et al., 2019; Liu et al., 2019) which has shown effectiveness in open-domain QA (Lee et al., 2019; Guu et al., 2020). Unlike open-domain QA, it is not trivial to apply a pre-training method to multi-hop QA due to the complexity in generating weak supervision data. In open-domain QA, weak supervision is generated by selecting a document from a corpus and extracting a sentence from the document. The sentence becomes a pseudo question, and the document becomes a pseudo supporting document to be predicted by retrievers. This two-step process for weak supervision cannot be directly applied in multi-hop QA since each multi-hop question refers to multiple documents.

In this paper, we propose a novel weakly supervised pre-training method for multi-hop retriever, **LOUVRE** (Learning frOm mUlti-hop Variation of document RElations). Our method contains three core elements: 1) a pre-training task, 2) a scalable method to generate pre-training data with weak supervision, and 3) a model to pre-train a retriever for multi-hop QA. Specifically, we define a task for pre-training, “Next Document Prediction” (NDP), which is to retrieve documents for sub-questions. We then propose “Bridge Entity Re-Phrasing” to generate the pre-training data. “Bridge Entity Re-Phrasing” generates complex questions that contain sub-questions of the bridge entities and their corresponding documents. To generate a complex question using a bridge entity without human effort, we




Figure 1: An example of chain reasoning in multi-hop QA. Answering the question requires finding “the largest planet in the Solar System” which is the bridge entity “Jupiter.” With the retrieved document A, the question has enough information to retrieve the correct answer in document B.

use two documents connected by Wikipedia hyperlinks. The hyperlinked entity becomes the bridge entity, and the introductory phrase of the entity becomes the sub-question in the complex question. This approach enables our weak supervision data generation to be scalable, as shown in Figure 2. We use a dense retriever consisting of a question encoder and a document encoder for the pre-trained model structure (Karpukhin et al., 2020; Xiong et al., 2021). The two encoders calculate vector representations of questions and documents. Document retrieval is performed by comparing the vectors with MIPS (maximum inner product search).

Pre-training multi-hop retriever with our weak supervision method brings three benefits: significant performance improvement, robustness on few-shot settings, and computational efficiency. We evaluate our weakly supervised pre-trained retriever with two types of experiments on HOTPOTQA dataset: supporting documents prediction and end-to-end multi-hop QA. In both experiments, LOUVRE outperforms previous multi-hop retrievers. Also, we fine-tune LOUVRE on 1% of training data and show that the performance of LOUVRE is comparable to the baselines. We evaluate the performance of LOUVRE according to the computational efficiency. The results show that LOUVRE requires less inference time than baselines.

Contributions of this paper are as follows: 1) we propose a novel scalable weakly supervised pre-training method for multi-hop retrievers, 2) we provide the implementation of LOUVRE and the pre-trained checkpoint publicly available<sup>1</sup>, 3) we show the effect of our pre-training method in multi-hop QA with various experimental results.

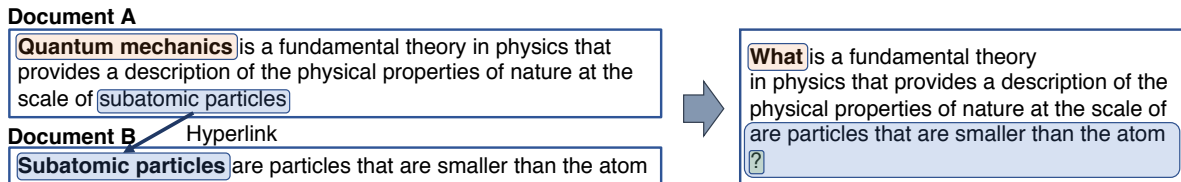
## 2 Related Work

**Distant Supervision in Open-Domain QA:** Many open-domain QA datasets only provide

question-answer pairs; some also provide weakly annotated supporting documents, but they are predicted by simple heuristics (Joshi et al., 2017; Berant et al., 2013). Document retrieval has suffered from insufficient strong supervision issues. Hence, document retrieval has suffered from lack of strong supervision. To resolve this issue, Karpukhin et al. (2020) use a document retrieved by TF-IDF as the supporting document of the given question. Weak supervision is also an effective method in the distant supervision setting of open-domain QA. Lee et al. (2019) use ICT (inverse cloze task) to generate pseudo question-document pairs and pre-train their retriever. They select documents from Wikipedia and extract sentences from the documents. The selected sentence-document pairs become pseudo-question-document pairs. Guu et al. (2020) propose a pre-training method for a language model that uses knowledge retriever (document retriever). They train the knowledge retriever only with the language modeling loss without using any supervision signal of supporting documents. Although pre-training methods show effectiveness in open-domain QA, they are limited to single-hop retrievers.

**Multi-Hop QA:** To overcome the lack of supervision signal in multi-hop QA, weak supervision methods have been proposed. Qi et al. (2019) propose a sub-question generation method. They use heuristically generated pseudo-questions as supervision for the question generation model. Perez et al. (2020) generate weak supervision for question decomposition by mapping a complex question to multiple single-hop questions in existing QA datasets. They use complex questions in HOTPOTQA (Yang et al., 2018) and single-hop questions in SQuAD 2.0 (Rajpurkar et al., 2018). Another method to train multi-hop QA models without human annotated datasets is by taking two simple questions and generating a complex question (Pan et al., 2020). They generate complex ques-

<sup>1</sup><https://github.com/yeonsw/LOUVRE>



**Question:** What is a fundamental theory in physics that provides a description of the physical properties of nature at the scale of are particles that are smaller than the atom ? **Answer:** Quantum mechanics

**Supporting documents in order:** Document B -> Document A

Figure 2: Proposed pre-training data generation process. Two documents connected by Wikipedia hyperlink are selected. In “Bridge Entity Re-Phrasing” process, document B which describes the entity “subatomic particles” is used to re-phrase the entity in document A. After replacing the answer entity, “Quantum mechanics”, the complex question and its corresponding supporting documents are generated.

tions with GPT-2 (Radford et al., 2019) fine-tuned on SQuAD1.1. Our work improves upon previous research by providing a more general method that leverages a large open corpus with retriever pre-training.

### 3 Method

We propose an effective and scalable pre-training method that provides weak supervision of the complex questions with sub-questions and their corresponding supporting documents.

#### 3.1 Next Document Prediction

We propose the “Next Document Prediction” (NDP) task for pre-training. NDP refers to the process of recurrent document retrieval used in (Qi et al., 2019; Asai et al., 2020; Xiong et al., 2021). We apply the common definitions in the existing studies to our “Next Document Prediction” task. We define NDP as the task that predicts documents in the reasoning sequence  $[d_1, \dots, d_n]$  recurrently as follows:

$$d_k = \text{Retriever}(q, D_{k-1}), \quad (1)$$

where  $q$  is a question,  $d_k$  is a predicted document at step  $k$ , and  $D_{k-1}$  is a set of documents retrieved in the previous steps,  $\{d_1, \dots, d_{k-1}\}$ .

#### 3.2 Bridge Entity Re-Phrasing

Our pre-training requires a dataset of questions, sub-questions, and their corresponding reasoning chains (i.e., a sequence of documents). We propose “Bridge Entity Re-Phrasing” for generating this pre-training dataset. “Bridge Entity Re-Phrasing” takes two steps: entity selection and re-phrasing. Figure 2 provides an overview of our data generation process. We provide the detailed description of

“Bridge Entity Re-Phrasing” in the following paragraphs.

The “Bridge Entity Re-Phrasing” process requires informative entities and the description of the entity. We assume that an entity with a Wikipedia hyperlink is an informative entity. Also, hyperlink entities often have Wikipedia articles describing the entities. The hyperlink entity becomes the bridge entity. In Figure 2, document A and document B are connected with the bridge entity, “subatomic particles”. We re-phrase the selected entity with the first line of the document. In Figure 2, “subatomic particles” is re-phrased with the first line in document B. When the bridge entity appears in the question, multi-hop retrievers easily find the bridge document using only the word. To prevent this issue, we remove the bridge entity from document B. The generated document becomes the document to be used for question-answer pair generation.

Generating a question-answer pair from a single document has been studied by pre-training research in open-domain QA (Lee et al., 2019; Guu et al., 2020). We extend their work to generate questions, reasoning chains, and answers. We randomly select an entity from the merged document and replace the entity with the word “what”. In Figure 2, “Quantum mechanics” is the entity word and replaced with “what.” The new sentence becomes a pseudo-question, and the replaced entity becomes the answer. Since document B contains the bridge entity, the pseudo-question reasoning chain becomes [document B, document A].

#### 3.3 Model Architecture

Model structure for our pre-training is subject to two requirements: general model structure and recurrent retrieval. We use multi-hop dense retriever

(Xiong et al., 2021) which meets the two requirements and is based on the DPR (dense passage retriever) (Karpukhin et al., 2020). DPR consists of a question encoder  $E_Q$  and a document encoder  $E_D$ , both of which are based on RoBERTa-base. Documents are retrieved by MIPS (maximum inner product search) with similarity between the question vectors and the document vectors as follows:

$$\text{sim}(q, d) = E_Q(q)^\top E_D(d). \quad (2)$$

MDR retrieves documents recurrently by taking the previously retrieved documents as input. MDR concatenates the question  $q$  and the retrieved documents  $\{d_1, \dots, d_{k-1}\}$  and calculates a question vector for  $k$ -th step as follows:

$$d_k = \text{argmax}_d (E_Q(q, d_1, \dots, d_{k-1})^\top E_D(d)), \quad (3)$$

where  $d$  is a document in the corpus.

We train the dense encoder to assign the highest probability for the ground truth document among the documents in the huge corpus. The loss function for our pre-training is as follows:

$$L_{\text{NBP}}(q_k, d_k) = -\log \frac{e^{\text{sim}(q_k, d_k)}}{e^{\text{sim}(q_k, d_k)} + \sum_{d \in \text{neg}(d_k)} e^{\text{sim}(q_k, d)}}, \quad (4)$$

where  $q_k$  is a concatenation of  $q$  and  $D_{k-1}$ , and  $\text{neg}(d_k)$  is a set of documents excluding  $d_k$ . Since computing the softmax over the whole corpus is computationally expensive, we use in-batch negatives for  $\text{neg}(d_k)$  (Karpukhin et al., 2020).

## 4 Experimental Setup

### 4.1 Pre-Training Details

We generate our pre-training data from 5,233,329 Wikipedia articles provided by Yang et al. (2018). We select all sentences that contain at least one hyperlinked entity to generate pseudo questions and randomly select “answer” entities from the sentences<sup>2</sup>. Our data generation process builds 13.9 million question-document-answer triples. We pre-train our dense retriever with a batch size of 256 for 200K+ steps. We use Adam with a warm-up ratio of 0.1 and set the learning rate to  $2 \times 10^{-5}$ . We use a machine with eight V100 (32G) GPUs.

<sup>2</sup>We use spaCy for entity recognition

	TF-IDF	Wiki	RR	Eff
LOUVRE				
- eff	✓	✓		✓
- Wiki		✓		
- reranking			✓	
- reranking-Wiki		✓	✓	

Table 1: The five variations of LOUVRE. Each column represents sparse retrieval (TF-IDF), Wikipedia hyperlinks (Wiki), reranking (RR), and efficient fine-tuning (Eff). “Eff” represents whether the model uses efficient hyper-parameter setting, a batch size of 32 and a number of epochs of 15.

### 4.2 Fine-Tuning Details

We use TF-IDF negatives in addition to in-batch negatives for fine-tuning as in Karpukhin et al. (2020); Xiong et al. (2021). We set the number of TF-IDF negatives to 2. We use the Adam optimizer with a warm-up rate of 0.1 and set the learning rate to  $2 \times 10^{-5}$ . We set the batch size to 32, the number of epochs to 15. To achieve better performance, we additionally fine-tune our model with another hyper-parameter setting: a batch size of 150 and a longer training time of 50 epochs.

### 4.3 Tasks

**Supporting Document Prediction:** In this task, retrievers and rerankers predict supporting documents for each question in HOTPOTQA dataset (Yang et al., 2018). The models predict possible combinations of supporting documents. Formally, when a question,  $q_i$ , has been taken as input of the model, the models yield a ranked list of K-sets of documents,  $L_i = [\{d_a, d_b\}]_{n=1}^K$ . Each  $\{d_a, d_b\}$  is a pair of candidate supporting documents. In HOTPOTQA, the number of supporting documents is fixed to 2. We use the 5 million Wikipedia articles as the knowledge source.

**End-to-End Multi-Hop QA:** We evaluate the supporting facts prediction performance and the answer prediction performance of LOUVRE on HOTPOTQA full wiki setting (Yang et al., 2018).

### 4.4 Multi-Hop Retrieval Strategy

We propose five variations of LOUVRE based on existing multi-hop retrieval strategies. Multi-hop document retrievers leverage three strategies for performance improvement and computational efficiency: sparse retrieval methods such as TF-IDF (Nie et al., 2019), Wikipedia hyperlinks (Asai et al.,



2020), and reranking (Xiong et al., 2021). Sparse retrieval methods select a small number of candidate documents relevant to the given question and are used to narrow down the search space of dense retrievers. We use TF-IDF and keyword matching as Nie et al. (2019) to retrieve 200 candidate documents. Existing multi-hop retrievers select reasoning paths (document chains) from documents connected with Wikipedia hyperlinks. We iteratively select the next-hop documents from the documents connected with the previously retrieved documents. Rerankers take the candidate reasoning paths (pairs of documents) from the retriever and predict the most probable reasoning path. We use the reranker proposed by Xiong et al. (2021). Table 1 shows the detailed information of these five variations.

#### 4.5 Baselines

We compare our model to two types of multi-hop retrievers with and without reranking. For retrievers, we use TF-IDF, DPR (Karpukhin et al., 2020), Cognitive Graph (Ding et al., 2019), GOLDEN Retriever (Qi et al., 2019), and MDR (Xiong et al., 2021). For rerankers, we use SemanticRetrievalMRS (Nie et al., 2019), PathRetriever (Asai et al., 2020), MDR-reranking (Xiong et al., 2021), and HopRetriever (Li et al., 2020). We describe detailed experimental settings of each baseline in Appendix A.

#### 4.6 Metric

We use five evaluation metrics: EM, F1, R@K, PathR@K, and AR@K. EM and F1 measure answer prediction and supporting fact prediction performance of multi-hop QA models (Yang et al., 2018). In addition to R@K, which measures the performance of supporting document prediction, we use another metric PathR@K to evaluate how well the retriever predicts the entire set of supporting documents. Since the readers predict answers by reading each path, PathR@K is a more appropriate estimate of answer prediction. The definitions of R@K and PathR@K are:

$$\begin{aligned} \text{R@K} &= \mathbb{1}(G \subseteq D) \\ \text{PathR@K} &= \mathbb{1}((G = P_1) \vee \dots \vee (G = P_K)), \end{aligned} \quad (5)$$

where  $G = \{g_i, g_j\}$  is the set of ground truth supporting documents,  $D$  is the set of retrieved documents, and  $P_i = \{d_a, d_b\}$  is a reasoning path ranked at  $i$ .

In our experimental setting,  $D$  is set to all documents in  $\bigcup_{i=1}^{K/2} P_i$ . AR@K measures the percentage of predictions that at least one passage in the top  $K$  predicted paths contains the answer text.

## 5 Results & Discussion

LOUVRE overcomes limited supervision in multi-hop QA with our weakly supervision data. Training with additional data brings progress in three ways: overall retrieval performance improvement, 2) robustness on a few-shot setting, and 3) overall improvement in the end-to-end multi-hop QA. We verify these improvements on the two tasks described in section 4: supporting document prediction and end-to-end multi-hop QA.

### 5.1 Supporting Document Prediction

In this experiment, we demonstrate the efficacy of LOUVRE with document retrieval experiments. First, we show the performance gain that comes from using our pre-trained model. Then, we show that the result becomes more significant in few-shot settings.

**Effect of Our Pre-Training:** We compare LOUVRE with MDR which is a multi-hop retriever fine-tuned on RoBERTa-base. We use the same fine-tuning method as MDR but initialize the parameters with LOUVRE. Table 2 shows the results. LOUVRE achieves 1.1% absolute performance improvement than when using RoBERTa (65.9); PathR@1 of LOUVRE is 67.0. Also, LOUVRE outperforms MDR in other evaluation metrics. In reranking experiments, we use the same reranking model as MDR-rerank. The only difference between LOUVRE-rerank and MDR-rerank is the parameter initialization method in the fine-tuning step same as the retriever experiment. These results show that our pre-training method is effective even after reranking; PathR@1 of LOUVRE-rerank is 83.2, and PathR@1 of MDR-rerank is 81.2.

**Weak Supervision and Training Time:** LOUVRE’s pre-training method uses additional training with the multi-hop weak supervision dataset and results in the performance improvement shown above. To verify that the performance gap between RoBERTa and LOUVRE is not from the additional training time that LOUVRE uses in pre-training, we train RoBERTa with much a longer training time, 50 epochs, and compare with LOUVRE.

In Figure 3, we show the performance of RoBERTa fine-tuned for  $\{2, 5, 15, 50\}$  epochs and

		Wiki link	PathR@1	R@10	R@20	AR@1
Retriever	TF-IDF		9.8	27.6	35.1	43.4
	DPR		25.2	45.4	52.1	-
	Cognitive Graph		57.8	-	-	76.0
	GoldEnRetriever		-	75.4	-	-
	MDR		65.9	77.5	80.2	75.4
	LOUVRE-eff (1%)	✓	53.5	75.5	80.0	72.3
	LOUVRE-eff	✓	65.3	80.4	<b>83.0</b>	<b>78.9</b>
	LOUVRE		67.0	77.8	80.3	76.3
	LOUVRE-Wiki	✓	<b>69.5</b>	<b>80.7</b>	82.4	77.8
	Reranker	SemanticRetrievalMRS	✓	63.9	81.7	82.1
PathRetriever		✓	75.7	82.4	-	87.5
MDR-rerank			81.2	86.4	86.6	88.2
HopRetriever (w/o Wiki link)			66.2	78.8	-	76.3
HopRetriever		✓	82.5	88.6	-	86.8
LOUVRE-rerank			83.2	89.1	89.7	<b>90.0</b>
LOUVRE-rerank-Wiki		✓	<b>83.5</b>	<b>89.5</b>	<b>90.1</b>	89.8

Table 2: Performance of retrievers. LOUVRE outperforms baseline retrievers and rerankers. Wiki link denotes whether the model uses Wikipedia hyperlinks.




Figure 3: Comparison of LOUVRE-eff fine-tuned for 15 epochs with RoBERTa fine-tuned for 2 to 50 epochs. The performance of RoBERTa flattens at 15 epochs.

the performance of LOUVRE fine-tuned for 15 epochs. The performance of RoBERTa stabilizes at approximately 80% in terms of PathR@20 after 15 epochs. This result shows that the main factor of the performance improvement from our pre-training method is not merely from longer training time but the unique information for multi-hop retrievers provided by our weak supervision.

**Retrieval Performance of Variations of LOUVRE-eff:** Table 2 shows the effect of using LOUVRE. Similar results are observed in LOUVRE-eff. Table 3 shows the same experiments as Table 2 but in efficient fine-tuning setting, a small batch size and a short train time. We compare LOUVRE-eff with the retrieval performance of

	R		PathR	
	@10	@20	@8	@20
LOUVRE-eff	<b>80.4</b>	<b>83.0</b>	<b>81.9</b>	<b>83.5</b>
- 1% train data	75.5	80.0	77.1	81.5
- zeroshot	44.8	55.8	48.3	59.0
- w/o pre-training	75.7	78.5	76.9	79.8
DPR-zeroshot	39.0	51.6	42.8	56.6

Table 3: Supporting document prediction performance of retrievers.

LOUVRE-eff without our pre-training method, which is fine-tuned on RoBERTa. Applying our pre-training method increases the retrieval performance by 4.7% point (R@10); R@10 of LOUVRE-eff is 80.4 and R@10 of LOUVRE-eff without pre-training is 75.7. Taking the results in Table 2 (the performance gain from our method with a big batch size/train epochs is 1.1) and the results in Table 3, we see the performance gain increases as there is more limitation on computation time.

**Robustness on Few-shot Settings:** Pre-training alleviates the model’s drastic performance drop when the number of training data is insufficient. We demonstrate the robustness of LOUVRE on few-shot settings with different sizes of train data. We




Figure 4: Retrieval performance of LOUVRE, LOUVRE-eff, and RoBERTa. The only difference between each model in the comparison models, (LOUVRE and MDR) and (LOUVRE-eff and RoBERTa), is the pre-trained model being used for parameter initialization.

fine-tune LOUVRE and MDR on a small portion of the HOTPOTQA train set within 0.1% to 100%. Figure 4 shows that the performance gap between LOUVRE and MDR increases as the size of train data decreases. When we use 0.1% of HOTPOTQA train data, almost 30% of LOUVRE’s predictions contain correct supporting documents; the performance of MDR with the same amount of train data is close to 0. We conduct the same experiment with LOUVRE-Wiki and verify that using Wikipedia hyperlinks improves the robustness on a few-shot setting by 10.5% point in terms of PathR@5 when there is only 0.1% train data.

We conduct the same experiment with a small batch size of 32 and 10 epochs. Figure 4 illustrates the retrieval performance of LOUVRE-eff and LOUVRE-eff without our pre-training (RoBERTa) depending on the proportion of the data used for fine-tuning. It is worth noting that LOUVRE-eff fine-tuned with 10% data outperforms RoBERTa with 100% and shows little performance degradation compared to fine-tuning with 100%. We report the detailed results of LOUVRE-eff (1%) in Table 2. Table 2 shows that LOUVRE-eff (1%) achieves comparable performance to MDR trained on full data with a larger batch size and a longer train

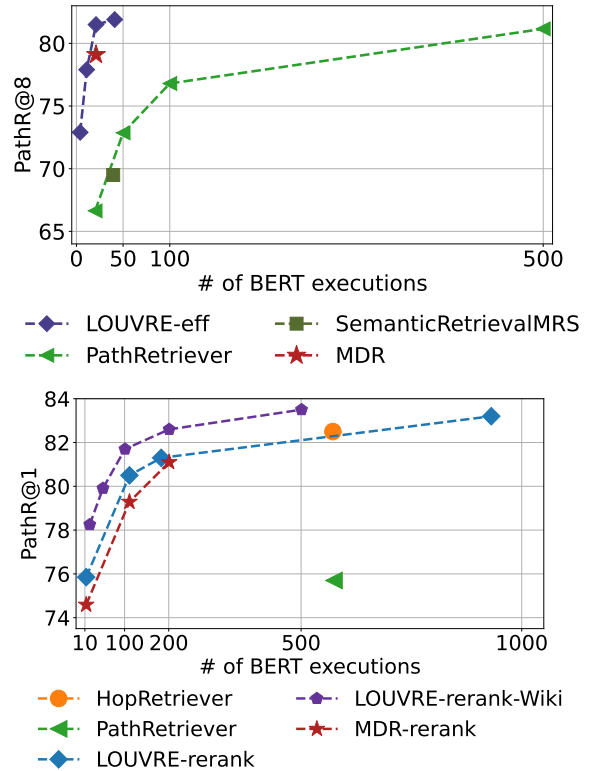


Figure 5: The retrieval performance of baselines and their computational efficiency. The beam sizes used for each model are as follows: [3, 10, 20, 40] for LOUVRE-eff, [3, 10, 13, 30] for LOUVRE-rerank, [10, 10, 50, 100, 100] for LOUVRE-rerank-wiki, and [3, 10, 100] for MDR-rerank. We use the input size of the reranker of each model as follows: [9, 100, 169, 900] for LOUVRE-rerank, [10, 40, 50, 100, 400] for LOUVRE-rerank-wiki, and [9, 100, 100] for MDR. We evaluate PathRetriever with different number of initial n-documents retrieved by TF-IDF;  $n \in \{20, 50, 100, 500\}$ .

time; R@10 of LOUVRE-eff (1%) is 2.0% point lower than R@10 of MDR, which is 97% of the performance of MDR.

Furthermore, we evaluate the zero-shot performance of LOUVRE-eff and compare to DPR not fine-tuned on HOTPOTQA. To adapt DPR to the multi-hop retrieval task, we replace encoders in LOUVRE-eff with DPR encoders. We report this result in Table 3. In Table 3, LOUVRE-zeroshot achieves higher performance (R@10: 44.8 and R@20: 55.8) than DPR-zeroshot (R@10: 39.0 and R@20: 51.6).

**Computational Efficiency:** We compare the inference time of baselines and LOUVRE, with the number of BERT-base executions needed for each question. We exclude the inference time for document indexing which can be done a pri-

	Answer		Support		Joint	
	EM	F1	EM	F1	EM	F1
SemanticRetrievalMRS	45.32	57.34	38.67	70.83	25.14	47.60
Transformer-XH	51.60	64.07	40.91	71.42	26.14	51.29
PathRetriever	60.04	72.96	49.08	76.41	35.35	61.18
MDR	62.28	75.29	57.46	80.86	41.78	66.55
<b>LOUVRE</b>	<b>62.90</b>	<b>75.82</b>	<b>57.71</b>	<b>81.26</b>	<b>42.18</b>	<b>67.08</b>

Table 4: End-to-end multi-hop QA performance of models on HOTPOTQA test set.

Retriever	#BERT	Joint F1
Reader: BERT		
- PathRetriever-eff	50 <sup>+</sup>	56.85
- LOUVRE-eff	50 <sup>+</sup>	<b>60.23</b>
Reader: ELECTRA		
- MDR	450 <sup>+</sup>	66.55
- LOUVRE-Wiki	450 <sup>+</sup>	<b>66.87</b>

Table 5: The end-to-end multi-hop performance of LOUVRE-eff and LOUVRE-Wiki on HOTPOTQA test set with the same inference speed as baselines.

ori. The number of BERT executions for each baseline is derived from each paper and its implementation. We measure the inference time of LOUVRE, PathRetriever, and MDR in various hyper-parameter settings by adjusting the number of beam size and the number of documents retrieved by the sparse retriever, TF-IDF. The number of BERT executions of MDR, LOUVRE-eff, and LOUVRE-eff with a beam size of  $b$  is calculated as follows:  $\#BERT = 1(\text{question encoding}) + b(\text{question-passage encoding})$ . The inference time of MDR-rerank, LOUVRE-rerank, and LOUVRE-rerank-wiki involves another factor, the input size of the reranker. The inference time of these reranking models with a beam size of  $b$  and a input size of  $r$  becomes  $\#BERT = 1 + b + r$ . For PathRetriever, we vary the number of documents retrieved by the sparse retriever, TF-IDF. Figure 5 illustrates that LOUVRE is more effective and efficient than the baselines because it yields better retrieval performance with a much smaller number of BERT executions.

## 5.2 End-to-End Multi-Hop QA

In this section, we demonstrate that the end-to-end multi-hop QA pipeline using LOUVRE retains the three outcomes of LOUVRE: overall performance

improvement, robustness on a few-shot setting, and the fast inference speed. We use multi-hop QA pipelines of MDR and PathRetriever. All the components of the multi-hop QA pipelines except the retriever are fixed. We plug in each baseline retriever and LOUVRE to the pipeline and evaluate the end-to-end performance of each model.

**End-to-End Performance:** Table 4 shows the end-to-end multi-hop QA performance of baselines and LOUVRE. In this experiment, we replace the retriever of MDR’s pipeline with LOUVRE-rerank. We set the beam size to 30 and the input size of the reranker to 900. LOUVRE outperforms baselines with a Joint F1 of 67.08. Table 5 shows the performance of LOUVRE-Wiki using the same inference time as MDR. In this experiment, we use a beam size of 100 and the reranker’s input size of 350. MDR uses a beam size of 200 and the reranker’s input size of 250. LOUVRE-Wiki outperforms MDR by 0.32% point in terms of Joint F1. We conduct the same experiment with LOUVRE-eff and PathRetriever. Table 5 shows the results. LOUVRE-eff in Table 5 represents the end-to-end pipeline of PathRetriever-eff using LOUVRE-eff as the initial candidate document retriever not TF-IDF. We set the number of initial candidate documents of LOUVRE-eff and PathRetriever-eff to 50. LOUVRE-eff outperforms PathRetriever-eff by 3.38% point without any loss of computational efficiency. We provide the detailed experimental results of LOUVRE-eff in Appendix B.

**End-to-End Performance of Variations of LOUVRE-eff:** Table 6 shows the end-to-end performance of each model on different inference time, size of train data, and pre-training. We evaluate two types of LOUVRE-eff. LOUVRE-eff (RR) represents the same pipeline used in Table 5. LOUVRE-eff (w/o RR) represents the pipeline that consists of LOUVRE-eff and the reader. Comparison between PathRetriever(500<sup>+</sup>)



	RR	#BERT	Joint F1
PathRetriever	✓	500+	59.5
- fast inference	✓	100+	58.6
- fast inference	✓	50+	56.7
LOUVRE-eff	✓	100+	60.5
- fast inference	✓	50+	60.1
LOUVRE-eff		41	57.9
- 1% train set		41	57.1
- w/o pre-training		41	54.3
- fast inference		5	53.3

Table 6: End-to-end multi-hop QA performances of baseline models and LOUVRE on HOTPOTQA dev set. #BERT denotes the number of BERT-base executions in supporting document prediction and the number of documents that the reader takes as input. We exclude the inference time of the supporting sentence selector. We use the same reader model as PathRetriever.

and LOUVRE-eff (100+) shows that applying LOUVRE-eff to PathRetriever increases Joint F1 by 1 with 5 times faster inference speed. We conduct the same experiment by reducing the number of documents retrieved by LOUVRE-eff and achieve 0.6% point higher Joint F1 than PathRetriever with 10 times faster inference speed.

We conduct ablation studies with two factors of retrievers: 1) size of train data and 2) pre-training. We fix the reader with BERT-wwm fully fine-tuned on HOTPOTQA train set. Table 6 shows the results. When we train LOUVRE-eff with only 1% of the train set, the end-to-end performance drops by 0.8% point. However, using RoBERTa with the whole train set decreases the performance by 3.6% point. This result indicates that our pre-training methods bring more significant improvement to the end-to-end multi-hop QA pipeline when the size of the train data is small.

Decreasing the search space of multi-hop retrievers increases the retriever’s computational efficiency but results in a significant performance drop. We demonstrate the robustness of LOUVRE-eff when the computation time is limited. We decrease the beam size of LOUVRE-eff to 2 and the number of output paths to 2; the total number of BERT executions of this model is #BERT = 1 + 2(beam size) + 2(input size). Table 6 shows that LOUVRE-eff achieves 89% of PathRetriever’s performance with almost 100 times faster infer-

ence speed of PathRetriever. We adjust the inference speed of PathRetriever and compare it with LOUVRE-eff (w/o RR). LOUVRE-eff (w/o RR) outperforms PathRetriever (50+) by 1.2% point with less computation time. LOUVRE-eff (w/o RR) trained only with 1% train data even outperforms PathRetriever (50+).

## 6 Conclusion

Answering complex questions includes reasoning across multiple documents. Recent studies have found that reasoning requires learning sub-question detection and relevant document retrieval to predict a correct answer with supporting facts. However, building such datasets requires costly human annotation and has limited scalability. To address this issue, we proposed a weakly supervised pre-training method for multi-hop retriever, LOUVRE. Our pre-training method contains three elements: “Next Document Prediction” task, “Bridge Entity Re-Phrasing”, and a model. We demonstrated the efficacy of LOUVRE and its robustness on few-shot settings with extensive experiments on supporting document retrieval task and end-to-end multi-hop QA task. We also showed that our method performs very well at a much lower inference cost.

## Acknowledgements

This work was partly supported by NAVER Corp. and Institute for Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korean government(MSIT) (No. 2017-0-01780, The technology development for event recognition/relational reasoning and learning knowledge based system for video understanding).

## References

- Akari Asai, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caiming Xiong. 2020. Learning to retrieve reasoning paths over wikipedia graph for question answering. In *ICLR*.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *EMNLP*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Ming Ding, Chang Zhou, Qibin Chen, Hongxia Yang, and Jie Tang. 2019. Cognitive graph for multi-hop reading comprehension at scale. In *ACL*.

- Yuwei Fang, Siqi Sun, Zhe Gan, Rohit Pillai, Shuhang Wang, and Jingjing Liu. 2020. Hierarchical graph network for multi-hop question answering. In *EMNLP*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Papat, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *ICML*.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *ACL*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *EMNLP*.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. Latent retrieval for weakly supervised open domain question answering. In *ACL*.
- Shaobo Li, Xiaoguang Li, Lifeng Shang, Xin Jiang, Qun Liu, Chengjie Sun, Zhenzhou Ji, and Bingquan Liu. 2020. Hopretriever: Retrieve hops over wikipedia to answer complex questions. *arXiv*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv*.
- Sewon Min, Victor Zhong, Luke Zettlemoyer, and Hananeh Hajishirzi. 2019. Multi-hop reading comprehension through question decomposition and rescoring. In *ACL*.
- Yixin Nie, Songhe Wang, and Mohit Bansal. 2019. Revealing the importance of semantic retrieval for machine reading at scale. In *EMNLP-IJCNLP*.
- Rodrigo Nogueira and Kyunghyun Cho. 2017. Task-oriented query reformulation with reinforcement learning. In *EMNLP*.
- Liangming Pan, Wenhui Chen, Wenhan Xiong, Min-Yen Kan, and William Yang Wang. 2020. Unsupervised multi-hop question answering by question generation. *arXiv*.
- Ethan Perez, Patrick Lewis, Wen-tau Yih, Kyunghyun Cho, and Douwe Kiela. 2020. Unsupervised question decomposition for question answering. In *EMNLP*.
- Peng Qi, Xiaowen Lin, Leo Mehr, Zijian Wang, and Christopher D Manning. 2019. Answering complex open-domain questions through iterative query generation. In *EMNLP-IJCNLP*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don't know: Unanswerable questions for squad. In *ACL*.
- Haitian Sun, Tania Bedrax-Weiss, and William Cohen. 2019. Pullnet: Open domain question answering with iterative retrieval on knowledge bases and text. In *EMNLP-IJCNLP*.
- Tomer Wolfson, Mor Geva, Ankit Gupta, Yoav Goldberg, Matt Gardner, Daniel Deutch, and Jonathan Berant. 2020. Break it down: A question understanding benchmark. *TACL*.
- Wenhan Xiong, Xiang Lorraine Li, Srinu Iyer, Jingfei Du, Patrick Lewis, William Yang Wang, Yashar Mehdad, Wen-tau Yih, Sebastian Riedel, Douwe Kiela, et al. 2021. Answering complex open-domain questions with multi-hop dense retrieval. In *ICLR*.
- 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 *EMNLP*.
- Chen Zhao, Chenyan Xiong, Corby Rosset, Xia Song, Paul Bennett, and Saurabh Tiwary. 2020. Transformer-xh: Multi-evidence reasoning with extra hop attention. In *ICLR*.

## A Appendix: Baselines

**Retrievers:** Retrievers encode questions and documents independently, and search documents by comparing the question vector and the document vectors. Retrievers are computationally efficient compared to rerankers since the document vectors can be indexed before the questions are given. In **TF-IDF**, to get the list of supporting documents from the retrieved documents, we rearrange the ranked documents  $[d_1, d_2, d_3, d_4, \dots]$  to  $[\{d_1, d_2\}, \{d_3, d_4\}, \dots]$ . **DPR** (Karpukhin et al., 2020) is a single-hop document retrieval model. We compare LOUVRE to the performance of DPR fine-tuned on HOTPOTQA reported from Xiong et al. (2021). **Cognitive Graph** (Ding et al., 2019) and **GOLDEN Retriever** (Qi et al., 2019) are multi-hop document retrieval models. We report the performance of Cognitive Graph reported from Asai et al. (2020) and the performance of GOLDEN Retriever reported from Qi et al. (2019). We fine-tune LOUVRE with the same training method proposed in **MDR** (Xiong et al., 2021); thus, the performance difference between LOUVRE and MDR only results from our pre-training method. We report the performance of MDR from Xiong et al. (2021).

**Rerankers:** Rerankers reorder reasoning paths/documents predicted by retrievers. Rerankers

	RR	#BERT	Answer		Support		Joint	
			EM	F1	EM	F1	EM	F1
PathRetriever	✓	500+	60.2	72.8	46.3	73.7	33.8	59.5
	✓	100+	59.7	72.6	44.8	72.6	32.7	58.6
	✓	50+	58.3	71.1	42.7	70.8	31.0	56.7
LOUVRE-eff - fast inference	✓	100+	60.4	73.4	48.0	74.7	35.1	60.5
	✓	50+	60.2	72.9	47.6	74.5	34.8	60.1
LOUVRE-eff		41	59.3	71.5	44.9	71.7	33.4	57.9
- 1% train data		41	58.7	71.0	44.3	70.7	32.6	57.1
- w/o pre-training		41	57.9	70.1	42.4	69.0	31.5	54.3
- fast inference		5	54.3	66.3	41.9	67.2	30.5	53.3

Table 7: End-to-end multi-hop QA performances of baseline models and LOUVRE on HOTPOTQA dev set.

	#BERT	Answer		Support		Joint	
		EM	F1	EM	F1	EM	F1
SemanticRetrievalMRS	39.4/-	45.32	57.34	38.67	70.83	25.14	47.60
Transformer-XH	100/-	51.60	64.07	40.91	71.42	26.14	51.29
PathRetriever	50+	58.21	70.86	42.91	71.30	30.95	56.85
LOUVRE-eff	50+	<b>59.79</b>	<b>72.65</b>	<b>47.95</b>	<b>74.89</b>	<b>34.54</b>	<b>60.23</b>

Table 8: Evaluation results on HOTPOTQA test set. We retrieve the test results of PathRetriever (50+) model and our model from the HOTPOTQA leaderboard.

calculate score of each paths/documents by jointly encode each document with the given question.

As a result, reranking takes a huge portion of computation time of the end-to-end multi-hop QA pipeline. **SemanticRetrievalMRS** (Nie et al., 2019) propose the document reranking model that takes output of sparse retrievers such as TF-IDF. Since the model outputs documents not a list of supporting documents, we use the same document rearranging method as TF-IDF above. **PathRetriever** (Asai et al., 2020) and **HopRetriever** (Li et al., 2020) are reasoning path prediction models. These models use TF-IDF and BERT to retrieve and rerank the candidate documents. They use Wikipedia hyperlinks for candidate documents selection as described in section 4.4 and beam search with size 8 to rank each predicted supporting documents. **MDR** (Xiong et al., 2021) provides a reranking model as well as their retriever. We report the performance of MDR-reranking from Xiong et al. (2021).

## B Appendix: End-to-End Performance of LOUVRE-eff

Table 7 shows the additional results of Table 6. We evaluate LOUVRE-eff on other evaluation metrics used in the HOTPOTQA benchmark and verify the efficacy of LOUVRE-eff. We report the detailed results of Table 5 in Table 8. The results show the end-to-end performance of LOUVRE-eff and PathRetriever on the HOTPOTQA test set.