

# One Token Can Help! Learning Scalable and Pluggable Virtual Tokens for Retrieval-Augmented Large Language Models

Yutao Zhu, Zhaoheng Huang, Zhicheng Dou\*, Ji-Rong Wen

Gaoling School of Artificial Intelligence, Renmin University of China  
yutaozhu94@gmail.com, {huangzh, dou, jrwen}@ruc.edu.cn

## Abstract

Retrieval-augmented generation (RAG) is a promising way to improve large language models (LLMs) for generating more factual, accurate, and up-to-date content. Existing methods either optimize prompts to guide LLMs in leveraging retrieved information or directly fine-tune LLMs to adapt to RAG scenarios. Although fine-tuning can yield better performance, it often compromises the LLMs' general generation capabilities by modifying their parameters. This limitation poses challenges in practical applications, especially when LLMs are already deployed, as parameter adjustments may affect their original functionality. To address this, we propose a novel method that involves learning scalable and pluggable virtual tokens for RAG. By maintaining the LLMs' original parameters and fine-tuning only the embeddings of these pluggable tokens, our approach not only enhances LLMs' performance but also preserves their general generation capabilities. Furthermore, we design several training strategies to improve the scalability, flexibility, and generalizability of our method. Comprehensive experiments across 12 question-answering tasks demonstrate the superiority of our approach.

## Introduction

Large language models (LLMs) have achieved remarkable performance across various natural language processing tasks (Brown et al. 2020; OpenAI 2023; Touvron et al. 2023). Despite their extensive parameters enabling them to learn rich knowledge during pre-training, LLMs may still generate hallucinated, outdated, or inaccurate content, especially in scenarios requiring long-tail knowledge (Ji et al. 2023; Zhang et al. 2023b). To address this problem, retrieval-augmented generation (RAG) has emerged as a pivotal strategy. By explicitly decoupling knowledge retrieval from the backbone LLMs, such architectures have achieved more accurate and reliable content generation and shown particularly enhanced performance on knowledge-intensive tasks such as open-domain question answering (Petroni et al. 2021; Tan et al. 2024; Jin et al. 2024b).

Existing efforts in RAG development can be roughly categorized into two groups (as illustrated in Figure 1). The first group leverages the in-context learning capabilities of

LLMs by incorporating retrieved information into the input along with appropriate prompts (Shi et al. 2023; Ram et al. 2023). This allows for straightforward application to any *off-the-shelf* LLM without tuning its parameters. However, its effectiveness largely depends on the human experience in crafting effective prompts and the LLM's ability to interpret these prompts. The second group focuses on training LLMs to enhance their performance in RAG scenarios. This training might involve either *end-to-end pre-training* (Guu et al. 2020; Borgeaud et al. 2022) or *fine-tuning* (Lin et al. 2023; Wang et al. 2023) for specific tasks. These approaches can often lead to better performance, but they require significant computational resources. Recently, parameter-efficient fine-tuning techniques, such as LoRA (Hu et al. 2022), have been widely studied, significantly reducing training costs. These methods can optimize the LLMs' parameters for RAG, but unfortunately compromise the model's general abilities in non-RAG scenarios, such as commonsense reasoning and in-context learning. All these limitations prevent their application to LLMs already operational in real-world settings.

Therefore, a critical research problem arises: *Is it possible to enhance LLMs' performance under RAG scenarios while preserving their general generation capabilities?* To achieve this, we introduce a novel, lightweight tuning method named **SPRING**, which learns **S**calable and **P**luggable **v**irtual **T**okens for **r**etrieval-augme**N**ted **G**eneration. Our basic idea is to add trainable virtual tokens to help LLMs learn RAG problems. Through fine-tuning, these virtual tokens effectively enhance the LLM's capability to understand retrieved information and its correlation with user inputs. Importantly, as the LLM's original parameters are frozen, its general generation abilities are preserved without any loss. During inference, when retrieval is triggered, these trained virtual tokens can be simply added to the prompt, which includes both the retrieved results and user input, thereby significantly enhancing performance. Moreover, we employ a scalable training approach, allowing the number of virtual tokens to be adjusted according to the needs of the inference scenario. Various training strategies have been implemented to further improve the generalizability of our method, ensuring robustness regardless of the number of the retrieved results.

In experiments, SPRING is trained with the base and instruction fine-tuned versions of Mistral-7b, LLaMA-2-7b, and LLaMA-2-13b models and evaluated on 12 com-

\*Corresponding author.  
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

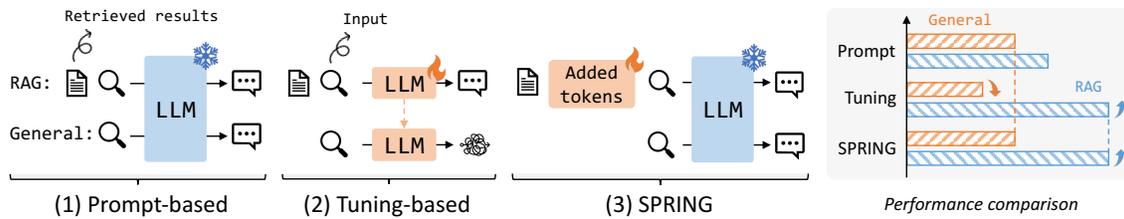


Figure 1: Illustration of existing methods for RAG and our proposed method. Our method can improve LLMs’ performance in RAG scenarios by incorporating trainable virtual tokens, and these tokens can be removed to preserve the general generation abilities in non-RAG scenarios.

monly used QA datasets, covering both in-domain and out-of-domain scenarios. The experimental results demonstrate that SPRING not only effectively improves the RAG performance of LLMs but also successfully preserves their general generation capabilities. Overall, the SPRING method exhibits four main characteristics:

- **Lightweight yet effective.** Instead of updating the full parameters of the LLMs, we opt to freeze the pre-trained models and only learn the embeddings for the added virtual tokens. For example, adding 50 tokens to the Mistral-7b model introduces only 0.2M parameters in total. Despite these minimal parameters, SPRING improves the average EM and F1 scores by more than 43% and 17% across 12 QA datasets, respectively.

- **Scalable.** With our proposed scalable training approach, SPRING can be effective with any number of virtual tokens ( $k \in [1, 50]$  in our experiments). Remarkably, even just one token can substantially improve the LLMs’ performance in RAG scenarios.

- **Pluggable.** Owing to its lightweight design, SPRING can be applied in a plug-and-play manner. When retrieval is triggered, simply adding the virtual tokens can lead to better performance. In non-RAG scenarios, the virtual tokens are not added so the LLMs’ original capabilities can be well preserved. This characteristic is crucial for LLMs that have already been deployed for practical use.

- **Generalizable.** Our robust training strategies ensure that SPRING is adaptable to different retrievers and various numbers of retrieved results. Consequently, there is no need to retrain SPRING with each update to the retrieval system, enhancing its practicality and efficiency.

## Related Work

**Retrieval-Augmented Generation** Compared to standard text generation, retrieval-augmented generation (RAG) incorporates a retrieval module that accesses external knowledge to enhance generation quality (Lewis et al. 2020; Guu et al. 2020; Zhu et al. 2023; Jin et al. 2024a). The mainstream RAG follows a “retrieve-then-read” paradigm, where the retrieval module provides external knowledge as additional context, which is then read by generation models to produce the final output (Shi et al. 2023; Ram et al. 2023; Borgeaud et al. 2022; Lin et al. 2023; Zhu et al. 2024). To optimize the use of external knowledge, some methods focus on crafting effective prompts that guide the utilization

of retrieved information (Shi et al. 2023; Ram et al. 2023). These prompt-based methods are applicable to any LLM without tuning its parameters. However, they depend heavily on skillful prompt writing and the LLMs’ ability to understand instructions. In contrast, other studies attempt to directly train the model to better use the retrieved knowledge. For example, REALM (Guu et al. 2020) and RETRO (Borgeaud et al. 2022) incorporate retrieval in end-to-end retrieval-augmented pre-training. RA-DIT (Lin et al. 2023) employs fine-tuning to enhance LLMs’ retrieval understanding. These tuning-based methods often yield better performance than prompt-based methods by optimizing model parameters for RAG. However, they may compromise the LLMs’ general capabilities, particularly in non-retrieval scenarios. Different from existing methods, we design a new lightweight tuning method for RAG. It is a plug-and-play module that enhances RAG performance using trainable virtual tokens, which can be removed in non-RAG scenarios to preserve the LLMs’ general generation abilities.

**Parameter-Efficient Fine-Tuning** The paradigms of “pre-training then fine-tuning” have demonstrated efficacy across various natural language (Devlin et al. 2019; Raffel et al. 2020; Radford et al. 2019) and vision tasks (He et al. 2020; Chen et al. 2020). The common fine-tuning process involves tuning all parameters of a model, which is computationally intensive, especially for LLMs. To address this, parameter-efficient fine-tuning (PEFT) (Mangrulkar et al. 2022) approaches have been developed. These approaches freeze most of the pre-trained models’ parameters, yet still manage to achieve comparable performance on downstream tasks. PEFT has been widely studied (Wan et al. 2023), and typical methods including adapter-based tuning (Houlsby et al. 2019; Lin, Madotto, and Fung 2020; Chen et al. 2023), low-rank adaptation (LoRA) (Hu et al. 2022; Dettmers et al. 2023), and prompt tuning (Li and Liang 2021; Lester, Al-Rfou, and Constant 2021; Liu et al. 2021b,a). Adapter-based tuning inserts lightweight modules into a model’s existing layers and have been extended to various domains (Gao et al. 2024; Hu et al. 2023; Zhang et al. 2023a). LoRA (Hu et al. 2022) introduces trainable low-rank matrices that adjust the model’s weight updates, achieving promising fine-tuning performance on LLMs (Hu et al. 2023). Prompt tuning incorporates a series of trainable prompt tokens to LLMs. These tokens can be inserted either to the input layer only (Lester, Al-Rfou, and Constant 2021; Liu et al. 2021b)

or to all of the intermediate layers (Li and Liang 2021; Liu et al. 2021a). In this paper, we propose a novel prompt tuning method, SPRING, specifically designed for RAG scenarios. Our method introduces virtual tokens between retrieved results and the input, exploiting the auto-regressive generation paradigm to improve the model’s ability to utilize retrieved information. Additionally, it is designed to be scalable and pluggable, thus broadening its application scope while preserving the original generative capabilities of LLMs.

## Methodology

To take advantage of both the flexibility of prompt-based methods and the efficacy of fine-tuning-based methods, we propose SPRING to learn scalable and pluggable virtual tokens for retrieval-augmented generation (RAG).

### Problem Formulation

Language models are designed to calculate the probability distribution over sequences of natural language texts. Auto-regressive models are commonly used for this through next-token prediction:

$$p_{\text{LM}} = \prod_{i=1}^m p_{\theta}(x_i | x_{<i}), \quad (1)$$

where  $x_{<i}$  denotes the sequence of tokens preceding  $x_i$  at each step, and  $\theta$  represents the parameters of the model. For RAG, a retrieval corpus  $\mathcal{D}$  and a retriever  $M$  are introduced. Then, the generation process is conditioned on both  $x_{<i}$  and the retrieved results  $R = M_{\mathcal{D}}(x_{<i})$  as:

$$p_{\text{RAG}} = \prod_{i=1}^m p_{\theta}(x_i | R; x_{<i}), \quad (2)$$

$$p_{\text{RAG-QA}} = \prod_{i=1}^m p_{\theta}(a_i | R; Q; a_{<i}). \quad (3)$$

Note that here  $x_{<i}$  serves as the *query* for retrieval. In question-answering (QA) tasks,  $x_{<i}$  is usually the question  $Q$ , and the learning objective is to generate the right answer  $A = \{a_i\}_{i=1}^m$ . The retriever can yield multiple passages, which can be concatenated as a long text sequence using proper separator such as “\n\n”. For brevity, this formulation directly concatenates the retrieved results  $R$  with the question  $Q$ , omitting more complex prompt designs. Henceforth, we will use the notations in QA tasks as our evaluation is performed on them.

### Scalable and Pluggable Virtual Tokens for RAG

Our SPRING method, shown in the left side of Figure 2, introduces trainable virtual tokens into the input to optimize LLMs for RAG scenarios. Specifically, following the notation in Equation (3), we add  $n$  trainable tokens  $T = [t_1, t_2, \dots, t_n]$  between the retrieved results  $R$  and the input  $Q$ . The generation process can then be described as:

$$p_{\text{SPRING}} = \prod_{i=1}^m p_{\theta, \delta}(a_i | R; [t_1, t_2, \dots, t_n]; Q; a_{<i}),$$

where  $\delta \in \mathbb{R}^{n \times d}$  represents the added parameters of the trainable tokens (*i.e.*, their embeddings), and  $d$  is the embedding size of the LLM.  $\theta$  denotes the parameters of the backbone LLM, which are *frozen* during training. Given that  $|\delta| \ll |\theta|$ , our method is *highly efficient* for training. For example, with the Mistral-7b model (where  $d = 4,096$ ), when  $n = 50$  tokens are added, we only add and train  $50 \times 4,096 = 0.2\text{M}$  parameters, approximately 0.003% of the full model.

Importantly, we place the virtual tokens  $T$  between the retrieved results  $R$  and the question  $Q$  for two main reasons: (1) In the auto-regressive generation paradigm, positioning the tokens after the retrieved results allows them to attend to this information, thereby aiding the model’s comprehension. (2) Recent studies have indicated that LLMs are particularly sensitive to the end of an input (Liu et al. 2023). By consistently placing these virtual tokens before the question across all test samples, we aim to mitigate any potential adverse effects on the understanding of the question.

**Scalable** In practical developments, LLMs are often constrained by their maximum input lengths, limiting the number of tokens available for retrieval augmentation (especially when the retrieved results are very long). Therefore, it is desired to design a mechanism so that any number of virtual tokens can be used in the inference to improve RAG performance. To achieve this, we propose an optimization strategy working like a “spring” (as shown in Figure 2). Specifically, for a given sample  $\{R, Q, A\}$  with the total number of added tokens  $n$ , we randomly select a number  $k (k \leq n)$  and utilize the *first*  $k$  virtual tokens  $t_{1:k}$  to construct the training example as  $[R; t_1, t_2, \dots, t_k; Q]$ . This method allows for the flexible optimization of any number of virtual tokens. Consequently, the number of virtual tokens incorporated during inference can be arbitrarily selected based on the requirements of the application. The effectiveness of this strategy and its comparison with other methods are further discussed in our experiments.

**Pluggable** Due to its designed structure, our method provides considerable flexibility in application. Practically, if user input is assessed to require external knowledge, our virtual tokens can be simply appended after the retrieval results and then fed, along with the user input, into the LLM for generation. In contrast, if the user input does not necessitate retrieval, it can be processed directly by the LLM. As our approach does not adjust the original parameters of the LLM, it preserves the model’s inherent capabilities. This feature is particularly important for industry or business since existing LLMs may have already been deployed for multiple purposes; our method enhances the retrieval understanding capabilities of these models without compromising their existing functionalities.

**Inference** We illustrate the instructions for using our SPRING in the right side of Figure 2. After training, the embeddings of the added tokens have been optimized for RAG, but these tokens do not correspond to any existing tokens in the vocabulary. To make them easy to use, we can add some special tokens (*e.g.*,  $[r1], \dots, [r50]$ ) to the

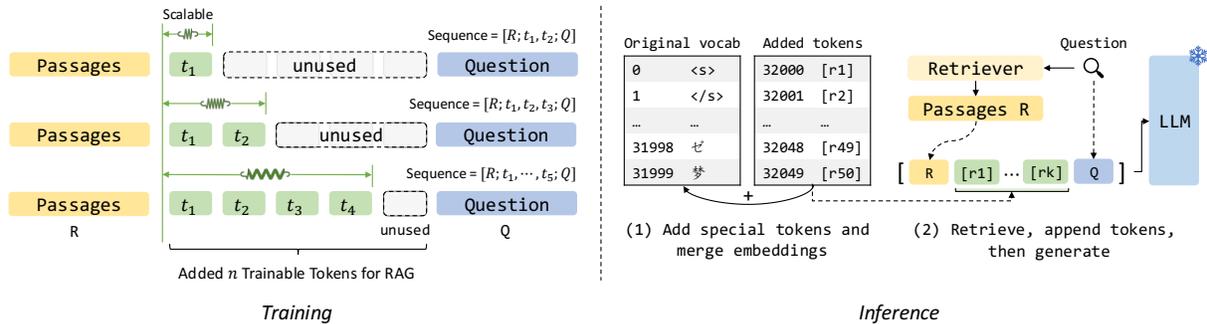


Figure 2: Illustration of SPRING. Only the embeddings of the added  $n$  tokens are trainable during fine-tuning. The added tokens are scalable where any first  $k$  ( $k \leq n$ ) tokens can be used in inference.

vocabulary and initialize their embeddings with the trained embeddings. Then, during inference, after obtaining the retrieved results  $R$ , we can add any number of these special tokens (e.g.,  $[r1] \dots [rk]$ ) after  $R$  and input them with the question  $Q$  to the LLMs for generation.

We refer to our method as SPRING due to its scalable and pluggable nature, making it particularly well-suited for enhancing existing LLMs that have already been deployed. Additionally, it effectively bridges the gap between retrieved results and user input, significantly improving the LLMs’ capabilities in understanding the retrieved external knowledge.

## Experiment

### Datasets and Retrievers

We conduct experiments on twelve commonly used question-answering datasets, including TriviaQA (TQA) (Joshi et al. 2017), Natural Questions (NQ) (Kwiatkowski et al. 2019), HotpotQA (HQA) (Yang et al. 2018), SQuAD 1.1 (Rajpurkar et al. 2016), Web Questions (WebQ) (Berant et al. 2013), 2WikiMultiHopQA (2Wiki) (Ho et al. 2020), CoQA (Reddy, Chen, and Manning 2019), MS MARCO (Nguyen et al. 2016), PopQA (Mallen et al. 2023), Fermi (Kalyan et al. 2021), Musique (Trivedi et al. 2022), and Bamboogle (Press et al. 2023). These datasets are publicly available at HuggingFace or their official websites. To evaluate the generalizability of the methods, we select PopQA, Fermi, Musique, and Bamboogle as held-out datasets. We mix the training set of all remaining datasets for training. For all datasets, we prioritize the use of test sets for evaluation purposes. In cases where the test set is not available, we utilize the development set instead. It is worth noting that, though some datasets have provided golden reference passages for the answer, we do not use them in our experiment but use the passages retrieved from the retrieval sets in both training and inference stages, which aligns with practical applications. Exact match (EM) and F1 score are employed as evaluation metrics.

For the retrieval sets, we follow previous studies (Yoran et al. 2023) and use the combination of Wikipedia and MS MARCO datasets as the retrieval corpus. Wikipedia contains high-quality human knowledge, which is helpful for

many knowledge-intensive tasks. MS MARCO contains a large amount of web pages, which can provide information necessary for curating some natural language questions. We use the datasets that have already been preprocessed into passages and released on HuggingFace.<sup>1</sup> The Wikipedia set has 21M passages, while the MS MARCO set has 8M passages.

We use E5-large (Wang et al. 2022) as the main retriever in our experiments. The impact of other retrievers, *i.e.*, BM25 (Robertson and Zaragoza 2009), BGE-base (Xiao et al. 2023), and E5-base, is studied in our further analysis. Among these retrievers, BM25 is a non-neural sparse retrieval algorithm, while others are neural-based dense retrievers. In general, dense retrievers perform better on several benchmarks (Muennighoff et al. 2023).

### Baseline Methods

We consider both the base and instruction fine-tuned versions of Mistral-7b, LLaMA-2-7b, and LLaMA-2-13b as the backbone models, and compare our SPRING with the following baselines.

- **Concat:** This method directly concatenates the retrieval results and the question for evaluation.
- **Prompt:** This method uses a manually-crafted prompt to indicate the use of retrieval information.
- **Prefix-tuning** (Li and Liang 2021): This method uses prefix-tuning to fine-tune the backbone models. To make a fair comparison with our method, we add 50 prefix tokens for training.
- **LoRA** (Hu et al. 2022): This method uses LoRA to fine-tune the backbone models. We use the hyperparameters suggested by the LLaMA’s official guidance.<sup>2</sup> To further validates the effectiveness of our SPRING on models that have already been optimized for RAG, we also train our method based on the LoRA checkpoint, and denote this variant as SPRING<sup>+</sup>.

<sup>1</sup>Wikipedia passages: <https://huggingface.co/datasets/Tevatron/wikipedia-nq-corpus>. MS MARCO passages: <https://huggingface.co/datasets/Tevatron/msmarco-passages-corpus>.

<sup>2</sup>LLaMA Recipes, [https://github.com/meta-llama/llama-recipes/blob/main/src/llama\\_recipes/configs/peft.py](https://github.com/meta-llama/llama-recipes/blob/main/src/llama_recipes/configs/peft.py)

Dataset	Metric	with Retrieval						without Retrieval					
		Concat	Prompt	Prefix	LoRA	SPRING	SPRING <sup>+</sup>	Concat	Prompt	Prefix	LoRA	SPRING	SPRING <sup>+</sup>
<i>Tuning Parameters</i>		0	0	0.2M	4M	0.2M	4.2M	0	0	0.2M	4M	0.2M	4.2M
Trivia QA	EM	0.00	57.79	11.74	62.76	<b>65.71</b>	63.89	0.01	39.90	0.00	0.03	<b>46.56</b>	43.37
	F1	65.60	80.33	59.97	85.44	85.26	<b>85.94</b>	63.96	69.72	28.30	34.91	74.48	<b>74.89</b>
NQ	EM	0.00	28.99	13.04	47.95	42.35	<b>49.78</b>	0.00	13.36	0.00	0.00	18.80	<b>25.54</b>
	F1	41.77	58.72	38.22	74.15	70.73	<b>75.22</b>	43.74	48.63	17.74	29.10	55.75	<b>60.82</b>
HQA	EM	0.00	26.36	5.79	39.95	35.26	<b>41.14</b>	0.00	17.07	0.00	0.03	20.15	<b>23.38</b>
	F1	44.91	56.15	42.59	68.93	65.44	<b>69.95</b>	47.54	49.50	17.45	27.57	54.79	<b>57.99</b>
SQuAD	EM	0.00	23.92	7.19	35.71	33.67	<b>35.98</b>	0.00	8.61	0.00	0.00	12.71	<b>13.90</b>
	F1	43.05	57.66	39.75	68.05	66.99	<b>68.41</b>	43.32	46.81	21.88	27.51	53.58	<b>54.93</b>
WebQ	EM	0.00	17.53	4.44	43.65	31.84	<b>48.10</b>	0.00	14.79	0.00	0.00	24.95	<b>28.81</b>
	F1	37.46	52.10	31.55	71.99	64.78	<b>73.88</b>	44.34	50.60	20.14	32.36	59.83	<b>62.95</b>
2Wiki	EM	0.00	22.64	4.38	35.93	31.80	<b>37.12</b>	0.00	23.45	0.00	0.01	24.62	<b>28.60</b>
	F1	47.77	55.58	41.82	63.85	62.03	<b>64.60</b>	52.83	53.55	21.14	37.09	56.83	<b>59.12</b>
CoQA	EM	0.00	8.20	1.56	12.89	13.28	<b>13.87</b>	0.00	8.59	0.00	0.00	9.96	<b>12.50</b>
	F1	27.98	36.72	20.02	41.19	<b>42.41</b>	42.04	32.97	36.58	13.15	18.99	39.96	<b>41.36</b>
MS MARCO	EM	0.00	5.73	0.60	8.13	6.57	<b>8.27</b>	0.00	<b>2.56</b>	0.00	0.01	2.09	<b>3.24</b>
	F1	56.44	53.81	50.56	54.81	53.48	<b>55.90</b>	49.50	47.75	47.44	52.44	<b>51.41</b>	49.84
PopQA*	EM	0.00	39.79	10.02	47.15	<b>48.71</b>	46.98	0.00	16.05	0.00	0.00	<b>20.25</b>	18.70
	F1	56.49	68.26	44.54	73.12	<b>73.90</b>	73.29	53.61	54.85	20.39	25.09	<b>58.32</b>	58.05
Fermi*	EM	0.00	0.06	0.00	0.12	0.18	<b>0.24</b>	0.00	0.06	0.06	0.00	0.06	0.19
	F1	21.66	18.16	12.65	29.63	<b>31.15</b>	30.67	25.51	17.84	20.33	25.42	29.28	30.37
Musique*	EM	4.01	4.05	0.04	10.86	8.80	<b>12.58</b>	0.00	1.93	0.66	0.17	3.64	<b>4.43</b>
	F1	38.79	38.64	20.91	52.65	48.74	<b>53.65</b>	42.30	36.18	35.22	45.31	44.93	<b>47.70</b>
Bamboogle*	EM	12.80	12.80	0.00	24.22	22.66	<b>28.13</b>	0.00	4.69	3.20	0.00	12.00	<b>12.80</b>
	F1	45.81	48.13	13.67	59.05	56.95	<b>62.00</b>	42.61	42.24	36.22	45.14	<b>47.69</b>	46.89
Average	EM	1.40	19.80	4.90	30.78	28.40	<b>32.17</b>	0.00	12.59	0.33	0.02	16.32	<b>17.96</b>
	F1	43.98	51.30	34.69	61.91	60.15	<b>62.96</b>	45.19	46.19	24.95	33.41	52.24	<b>53.74</b>

Table 1: Evaluation results of different methods on twelve QA datasets. The retriever is E5-large model, and the number of retrieved passages is set as three. The number of virtual tokens used in SPRING is set as 50. \*PopQA, Fermi, Musique, and Bamboogle are invisible during training. “Prefix” stands for prefix-tuning, and “SPRING<sup>+</sup>” is trained based on the LoRA’s checkpoint. The best results are in **bold**.

## Implementation Details

We use PyTorch (Paszke et al. 2019) and Huggingface Accelerate library to implement all methods. The learning rate is set as 1e-4 with a warm-up ratio of 0.1. All methods are trained for three epochs, with a training batch size of 256. We use eight NVIDIA A800 GPUs for training. Training our SPRING for Mistral-7b models consumes around 2.2 hours per epoch. The embeddings of the virtual tokens are initialized by the embeddings of the prompt: “According to the previous relevant passages, please answer the following question. Only return the answer without any other words.” Following the settings of prefix-tuning, if the number of tokens required exceeds those available in the prompt, the prompt is repeated to complete the initialization; if fewer are needed, the prompt is truncated accordingly. Additionally, we experiment with random initialization of tokens but observe that its performance is slightly worse than that achieved through prompt-based initialization. Our code is available at <https://github.com/DaoD/SPRING>.

## Experimental Results

We fine-tune the prefix-tuning, LoRA, and SPRING methods on RAG tasks, and then evaluate their performance

in scenarios both with (RAG) and without (non-RAG) retrieval. For SPRING, we use  $k = 50$  virtual tokens for inference by default, and the impact of token quantity  $k$  is discussed in later. The experimental results are shown in Table 1. To save space, we only show the results based on Mistral-7b-instruct.

We can observe: (1) It is evident that SPRING significantly improves the RAG performance of the original LLM with manually-crafted prompts (the average EM and F1 scores are improved by 43.4% and 17.3%, respectively). It outperforms LoRA on certain datasets, such as TriviaQA and CoQA. Given that SPRING involves only 0.2M trainable parameters, these results demonstrate its remarkable efficiency and effectiveness. (2) While LoRA achieves slightly better performance on some datasets, it adjusts the LLMs’ original parameters, which adversely impact their performance in non-RAG scenarios—a significant drop has been observed, even far worse than the original models. This challenge also extends to other general generation tasks, which will be discussed in the next section. (3) In non-RAG evaluation, only SPRING and SPRING<sup>+</sup> demonstrate better performance than the Prompt method. This indicates that even in the absence of retrieved results, adding virtual tokens is

Dataset	$n$ -shot	LoRA	SPRING	Diff
BoolQ	0	79.30	<b>82.97</b>	3.67
CommonsenseQA	0	55.45	<b>63.80</b>	8.35
CommonsenseQA	4	59.87	<b>67.07</b>	7.20
GSM8K	8	17.33	<b>31.89</b>	14.56
MMLU	0	51.30	<b>53.62</b>	2.32
MMLU	5	48.76	<b>54.96</b>	6.20

Table 2: Performance comparison on other datasets.

still beneficial. We speculate that beyond simply utilizing retrieved results, virtual tokens can help the LLM understand the task goal and format (e.g., the task is question-answering rather than text continuation). (4) Based on the LoRA’s checkpoint, SPRING+ achieves the best performance on most datasets. Additionally, all backbone models show improvements with SPRING. These findings verify the versatility and flexibility of our approach, confirming its suitability for enhancing various LLMs in RAG scenarios. (5) Using manually-crafted prompts is effective for improving LLMs’ performance on RAG tasks. However, this improvement is limited as no training is involved. (6) SPRING achieves robust performance on the held-out datasets, validating the good generalizability of our method. (7) Interestingly, prefix-tuning cannot perform well for RAG, highlighting that the insertion position of the virtual tokens in SPRING is both reasonable and effective.

### Further Analysis

We further conduct a series of experiments to investigate the impact of different settings in SPRING. All the following experiments are conducted based on fine-tuning the `Mistral-7b-instruct` model.

**Performance on Other Tasks** To examine the impact of different fine-tuning methods on the inherent capabilities of LLMs, we evaluate the performance of models fine-tuned by LoRA and SPRING on several other (non-RAG) tasks. These tasks are commonly used to evaluate LLMs’ reasoning, mathematical abilities, and world knowledge, including BoolQ (Clark et al. 2019), CommonsenseQA (Talmor et al. 2019), GSM8K (Cobbe et al. 2021), and MMLU (Hendrycks et al. 2021). The experimental results are shown in Table 2.<sup>3</sup> From the results, we can observe: (1) Thanks to the plug-and-play design of our method, SPRING can revert to to the original LLMs by not using virtual tokens. Therefore, it successfully preserves the original capabilities of the LLMs. In contrast, LoRA, which adjusts the model’s parameters for RAG tasks, inevitably compromises the model’s performance on other tasks. (2) A noticeable decline is observed in the few-shot evaluation, reflecting a decrease in the in-context learning abilities of LLMs. This decline may stem from the fact that RAG fine-tuning does not incorporate in-context learning capabilities. Besides, fine-tuning for RAG

<sup>3</sup>We notice that our results are quite different from those officially reported, which we attribute to the impact of different prompts.

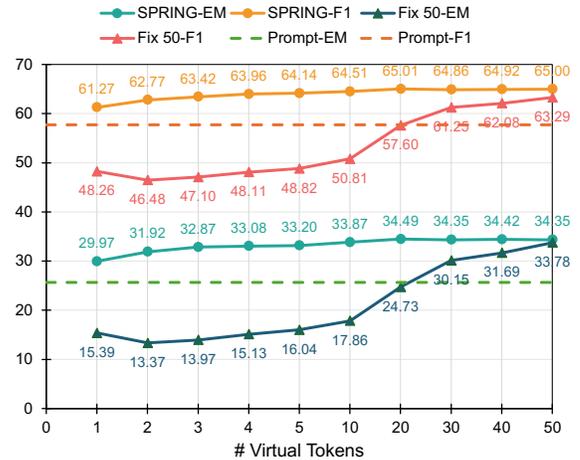


Figure 3: Average performance on nine QA datasets with various numbers of virtual tokens.

Retriever	Prompt		SPRING	
	EM	F1	EM	F1
BM25	21.23	54.94	30.94	62.73
BGE-base	23.07	56.12	31.81	63.46
E5-base	24.38	56.84	33.34	64.49
E5-large	25.66	57.70	<b>34.35</b>	<b>65.00</b>
Average	23.58	56.40	32.61	63.92
Variance	2.69	1.02	1.75	0.78

Table 3: Average performance on nine QA datasets with different retrievers.

tasks may lead the model to overfit to specific task formats (prompts), thereby impairing its general generation abilities.

**Impact of Token Quantity** In SPRING, we design a scalable training approach that enables to use arbitrary numbers of virtual tokens in inference. To validate its effectiveness, we test the performance of our method with various numbers of virtual tokens and compare it with a variant model trained with a fixed number of tokens ( $k = 50$ ). The experimental results are illustrated in Figure 3. In general, we can observe that the performance of SPRING increases with more virtual tokens used. Surprisingly, SPRING can significantly enhance LLMs’ performance in RAG scenarios with just a single token, which is very encouraging.<sup>4</sup> In comparison, training with a fixed number of tokens limits the flexibility of SPRING, making it can only be used with the same number of tokens in inference (i.e.,  $k = 50$ ).

**Effects of Different Retrievers** In our experiments, SPRING is fine-tuned using passages retrieved by E5-large. To investigate its effectiveness with other retrievers, we conduct an experiment by testing its performance with passages retrieved by BM25, BGE-base, and

<sup>4</sup>This varies across different LLMs.

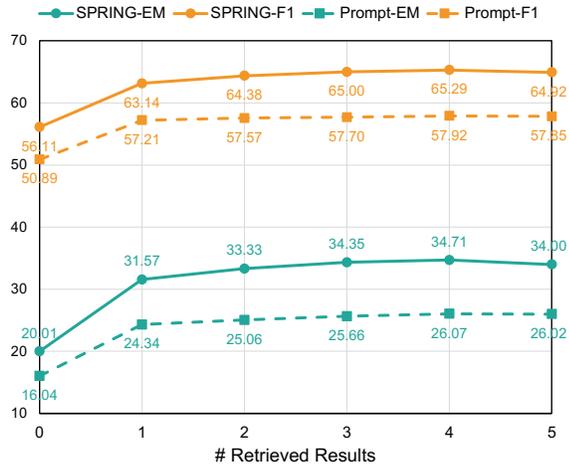


Figure 4: Average performance on nine QA datasets with different number of retrieved passages.

E5-large. The results are presented in Table 3. First, SPRING achieves consistent improvement over the original model using manually crafted prompt, thereby confirming the generalizability of our approach. Second, compared to the original model, the performance gap (variance) among different retrievers becomes smaller, highlighting SPRING’s robustness to variations in retrievers. Finally, even fine-tuned with a superior retriever (*i.e.*, E5-large), SPRING maintains strong performance well with less effective retrievers (such as BM25). This indicates that our method can effectively adapt to varying quality of retrieved results. Hence, there is no necessity to retrain the virtual tokens with each update of retrievers in practical applications, significantly enhancing its applicability.

**Influence of Retrieved Passages** During the fine-tuning of SPRING, we construct training samples by randomly selecting the top- $m$  ( $m \in [1, 5]$ ) retrieved passages. This aims to enhance SPRING’s adaptability by ensuring it can operate effectively with varying numbers of retrieved passages in real-world scenarios. To evaluate the effect of this training strategy, we test the SPRING’s performance across a range from zero to five passages. Figure 4 illustrates the results. We can find that SPRING’s performance gradually improves as more retrieved passages are used ( $m = 0 \rightarrow 4$ ), suggesting that more retrieved passages contribute valuable knowledge for question answering. However, the performance peaks at four passages and declines when more passages are added. This decrease could be attributed to noise accumulation within the retrieved knowledge, a phenomenon also reported in recent studies (Yoran et al. 2023). Despite this, the use of retrieved passages still results in performance gains compared to scenarios without retrieval ( $m = 0$ ), highlighting again the benefits of RAG.

**Cross-Dataset Generalizability** Inspired by previous studies in multi-task learning (Raffel et al. 2020; Khashabi et al. 2020), we mix eight QA datasets for training as they

Training $\rightarrow$	TQA		NQ		Mix	
Test $\downarrow$	EM	F1	EM	F1	EM	F1
TQA	62.80	84.65	65.51	84.73	<b>65.71</b>	<b>85.26</b>
NQ	32.19	64.98	<b>45.26</b>	<b>72.68</b>	42.35	70.73
HQA	22.97	56.95	28.11	59.45	<b>35.26</b>	<b>65.44</b>
SQuAD	25.24	61.39	30.81	64.32	<b>33.67</b>	<b>66.99</b>
WebQ	26.03	62.40	<b>34.13</b>	<b>67.30</b>	31.84	64.78
2Wiki	17.32	54.41	25.43	58.21	<b>31.80</b>	<b>62.03</b>
CoQA	6.05	36.61	7.62	38.54	<b>13.28</b>	<b>42.41</b>
MARCO	3.26	31.88	3.88	33.33	<b>6.57</b>	<b>53.48</b>
PopQA	44.07	72.50	48.28	74.07	<b>48.71</b>	<b>73.90</b>
Average	26.66	58.42	32.11	61.40	<b>34.35</b>	<b>65.00</b>

Table 4: Performance comparison between training on a specific dataset or a mixture of all datasets.

require similar LLM capabilities (*e.g.*, reasoning). To study the impact of this strategy, we conduct experiments by training SPRING on each dataset individually and then testing its performance on the others. Table 4 shows partial results. As indicated, training on a mixed dataset generally enhances performance on most datasets, thereby validating the benefits of multi-task learning. While training on a single dataset, such as NQ, may yield superior results on its specific test set, such improvements often fail to generalize to other datasets. Notably, training solely on NQ may negatively impact performance on MS MARCO, where the original LLM using a prompt could outperform it. These findings inspire us to carefully consider the interaction between different datasets when applying our method in future applications.

## Conclusion

In this paper, we introduced scalable and pluggable virtual tokens for retrieval-augmented large language models. Our method, SPRING, serves as a parameter-efficient fine-tuning approach that significantly enhances RAG performance with the addition of only 0.2M trainable parameters. More importantly, the plug-and-play nature of our approach successfully preserves the performance of LLMs on non-RAG tasks, while its scalable training strategy broadens the method’s applicational flexibility. Through extensive experiments across various datasets, we have demonstrated the effectiveness, generalizability, flexibility, and high efficiency of our method. We believe that our research will foster further integration of information retrieval and LLMs, and advance the development of other parameter-efficient fine-tuning technologies for LLMs.

## Acknowledgments

This work was supported by National Natural Science Foundation of China (Grant No. 62402497 and 62272467), Beijing Natural Science Foundation L233008, and Beijing Municipal Science and Technology Project No. Z231100010323009. The work was partially done at the Engineering Research Center of Next-Generation Intelligent Search and Recommendation, MOE.

## References

- Berant, J.; Chou, A.; Frostig, R.; and Liang, P. 2013. Semantic Parsing on Freebase from Question-Answer Pairs. In *EMNLP*, 1533–1544.
- Borgeaud, S.; Mensch, A.; Hoffmann, J.; Cai, T.; Rutherford, E.; Millican, K.; van den Driessche, G.; Lespiau, J.; Damoc, B.; Clark, A.; de Las Casas, D.; Guy, A.; Menick, J.; Ring, R.; Hennigan, T.; Huang, S.; Maggiore, L.; Jones, C.; Cassirer, A.; Brock, A.; Paganini, M.; Irving, G.; Vinyals, O.; Osindero, S.; Simonyan, K.; Rae, J. W.; Elsen, E.; and Sifre, L. 2022. Improving Language Models by Retrieving from Trillions of Tokens. In *ICML*, 2206–2240.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. In *NeurIPS*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. E. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *ICML*, 1597–1607.
- Chen, Z.; Duan, Y.; Wang, W.; He, J.; Lu, T.; Dai, J.; and Qiao, Y. 2023. Vision Transformer Adapter for Dense Predictions. In *ICLR*.
- Clark, C.; Lee, K.; Chang, M.; Kwiatkowski, T.; Collins, M.; and Toutanova, K. 2019. BoolQ: Exploring the Surprising Difficulty of Natural Yes/No Questions. In *NAACL-HLT*, 2924–2936.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training Verifiers to Solve Math Word Problems. *CoRR*.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. In *NeurIPS*.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*, 4171–4186.
- Gao, P.; Geng, S.; Zhang, R.; Ma, T.; Fang, R.; Zhang, Y.; Li, H.; and Qiao, Y. 2024. CLIP-Adapter: Better Vision-Language Models with Feature Adapters. *Int. J. Comput. Vis.*, (2): 581–595.
- Guu, K.; Lee, K.; Tung, Z.; Pasupat, P.; and Chang, M. 2020. Retrieval Augmented Language Model Pre-Training. In *ICML*, 3929–3938.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. B. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *CVPR*, 9726–9735.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multitask Language Understanding. In *ICLR*.
- Ho, X.; Nguyen, A. D.; Sugawara, S.; and Aizawa, A. 2020. Constructing A Multi-hop QA Dataset for Comprehensive Evaluation of Reasoning Steps. In *COLING*, 6609–6625.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; de Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-Efficient Transfer Learning for NLP. In *ICML*, 2790–2799.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*.
- Hu, Z.; Wang, L.; Lan, Y.; Xu, W.; Lim, E.; Bing, L.; Xu, X.; Poria, S.; and Lee, R. K. 2023. LLM-Adapters: An Adapter Family for Parameter-Efficient Fine-Tuning of Large Language Models. In *EMNLP*, 5254–5276.
- Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.; Madotto, A.; and Fung, P. 2023. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.*, (12): 248:1–248:38.
- Jin, J.; Zhu, Y.; Yang, X.; Zhang, C.; and Dou, Z. 2024a. FlashRAG: A Modular Toolkit for Efficient Retrieval-Augmented Generation Research. *CoRR*.
- Jin, J.; Zhu, Y.; Zhou, Y.; and Dou, Z. 2024b. BIDER: Bridging Knowledge Inconsistency for Efficient Retrieval-Augmented LLMs via Key Supporting Evidence. *CoRR*.
- Joshi, M.; Choi, E.; Weld, D. S.; and Zettlemoyer, L. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. In *ACL*, 1601–1611.
- Kalyan, A.; Kumar, A.; Chandrasekaran, A.; Sabharwal, A.; and Clark, P. 2021. How much coffee was consumed during EMNLP 2019? Fermi Problems: A New Reasoning Challenge for AI. In *EMNLP*, 7318–7328.
- Khashabi, D.; Min, S.; Khot, T.; Sabharwal, A.; Tafjord, O.; Clark, P.; and Hajishirzi, H. 2020. UnifiedQA: Crossing Format Boundaries With a Single QA System. In *Findings of EMNLP*, 1896–1907.
- Kwiatkowski, T.; Palomaki, J.; Redfield, O.; Collins, M.; Parikh, A. P.; Alberti, C.; Epstein, D.; Polosukhin, I.; Devlin, J.; Lee, K.; Toutanova, K.; Jones, L.; Kelcey, M.; Chang, M.; Dai, A. M.; Uszkoreit, J.; Le, Q.; and Petrov, S. 2019. Natural Questions: a Benchmark for Question Answering Research. *Trans. Assoc. Comput. Linguistics*, 452–466.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In *EMNLP*, 3045–3059.
- Lewis, P. S. H.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.; Rocktäschel, T.; Riedel, S.; and Kiela, D. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *NeurIPS*.
- Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. In *ACL*, 4582–4597.
- Lin, X. V.; Chen, X.; Chen, M.; Shi, W.; Lomeli, M.; James, R.; Rodriguez, P.; Kahn, J.; Szilvasy, G.; Lewis, M.; Zettlemoyer, L.; and Yih, S. 2023. RA-DIT: Retrieval-Augmented Dual Instruction Tuning. *CoRR*.
- Lin, Z.; Madotto, A.; and Fung, P. 2020. Exploring Versatile Generative Language Model Via Parameter-Efficient Transfer Learning. In *Findings of EMNLP*, 441–459.

- Liu, N. F.; Lin, K.; Hewitt, J.; Paranjape, A.; Bevilacqua, M.; Petroni, F.; and Liang, P. 2023. Lost in the Middle: How Language Models Use Long Contexts. *CoRR*.
- Liu, X.; Ji, K.; Fu, Y.; Du, Z.; Yang, Z.; and Tang, J. 2021a. P-Tuning v2: Prompt Tuning Can Be Comparable to Fine-tuning Universally Across Scales and Tasks. *CoRR*.
- Liu, X.; Zheng, Y.; Du, Z.; Ding, M.; Qian, Y.; Yang, Z.; and Tang, J. 2021b. GPT Understands, Too. *CoRR*.
- Mallen, A.; Asai, A.; Zhong, V.; Das, R.; Khashabi, D.; and Hajishirzi, H. 2023. When Not to Trust Language Models: Investigating Effectiveness of Parametric and Non-Parametric Memories. In *ACL*, 9802–9822.
- Mangrulkar, S.; Gugger, S.; Debut, L.; Belkada, Y.; Paul, S.; and Bossan, B. 2022. PEFT: State-of-the-art Parameter-Efficient Fine-Tuning methods.
- Muennighoff, N.; Tazi, N.; Magne, L.; and Reimers, N. 2023. MTEB: Massive Text Embedding Benchmark. In *EACL*, 2006–2029.
- Nguyen, T.; Rosenberg, M.; Song, X.; Gao, J.; Tiwary, S.; Majumder, R.; and Deng, L. 2016. MS MARCO: A Human Generated MACHine Reading COMprehension Dataset. In *NeurIPS*.
- OpenAI. 2023. GPT-4 Technical Report. *CoRR*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E. Z.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*, 8024–8035.
- Petroni, F.; Piktus, A.; Fan, A.; Lewis, P. S. H.; Yazdani, M.; Cao, N. D.; Thorne, J.; Jernite, Y.; Karpukhin, V.; Mail-lard, J.; Plachouras, V.; Rocktäschel, T.; and Riedel, S. 2021. KILT: a Benchmark for Knowledge Intensive Language Tasks. In *NAACL-HLT*, 2523–2544.
- Press, O.; Zhang, M.; Min, S.; Schmidt, L.; Smith, N. A.; and Lewis, M. 2023. Measuring and Narrowing the Compositional Gap in Language Models. In *Findings of EMNLP*, 5687–5711.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language Models are Unsupervised Multitask Learners.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.*, 140:1–140:67.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. SQuAD: 100, 000+ Questions for Machine Comprehension of Text. In *EMNLP*, 2383–2392.
- Ram, O.; Levine, Y.; Dalmedigos, I.; Muhlgay, D.; Shashua, A.; Leyton-Brown, K.; and Shoham, Y. 2023. In-Context Retrieval-Augmented Language Models. *CoRR*.
- Reddy, S.; Chen, D.; and Manning, C. D. 2019. CoQA: A Conversational Question Answering Challenge. *Trans. Assoc. Comput. Linguistics*, 249–266.
- Robertson, S. E.; and Zaragoza, H. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Found. Trends Inf. Retr.*, (4): 333–389.
- Shi, W.; Min, S.; Yasunaga, M.; Seo, M.; James, R.; Lewis, M.; Zettlemoyer, L.; and Yih, W. 2023. REPLUG: Retrieval-Augmented Black-Box Language Models. *CoRR*.
- Talmor, A.; Herzig, J.; Lourie, N.; and Berant, J. 2019. CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. In *NAACL-HLT*, 4149–4158.
- Tan, J.; Dou, Z.; Zhu, Y.; Guo, P.; Fang, K.; and Wen, J. 2024. Small Models, Big Insights: Leveraging Slim Proxy Models To Decide When and What to Retrieve for LLMs. *CoRR*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. *CoRR*.
- Trivedi, H.; Balasubramanian, N.; Khot, T.; and Sabharwal, A. 2022. MuSiQue: Multihop Questions via Single-hop Question Composition. *Trans. Assoc. Comput. Linguistics*, 539–554.
- Wan, Z.; Wang, X.; Liu, C.; Alam, S.; Zheng, Y.; Liu, J.; Qu, Z.; Yan, S.; Zhu, Y.; Zhang, Q.; Chowdhury, M.; and Zhang, M. 2023. Efficient Large Language Models: A Survey. *CoRR*.
- Wang, L.; Yang, N.; Huang, X.; Jiao, B.; Yang, L.; Jiang, D.; Majumder, R.; and Wei, F. 2022. Text Embeddings by Weakly-Supervised Contrastive Pre-training. *CoRR*.
- Wang, Z.; Araki, J.; Jiang, Z.; Parvez, M. R.; and Neubig, G. 2023. Learning to Filter Context for Retrieval-Augmented Generation. *CoRR*.
- Xiao, S.; Liu, Z.; Zhang, P.; and Muennighof, N. 2023. C-Pack: Packaged Resources To Advance General Chinese Embedding. *CoRR*.
- Yang, Z.; Qi, P.; Zhang, S.; Bengio, Y.; Cohen, W. W.; Salakhutdinov, R.; and Manning, C. D. 2018. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *EMNLP*, 2369–2380.
- Yoran, O.; Wolfson, T.; Ram, O.; and Berant, J. 2023. Making Retrieval-Augmented Language Models Robust to Irrelevant Context. *CoRR*.
- Zhang, R.; Han, J.; Zhou, A.; Hu, X.; Yan, S.; Lu, P.; Li, H.; Gao, P.; and Qiao, Y. 2023a. LLaMA-Adapter: Efficient Fine-tuning of Language Models with Zero-init Attention. *CoRR*.
- Zhang, Y.; Li, Y.; Cui, L.; Cai, D.; Liu, L.; Fu, T.; Huang, X.; Zhao, E.; Zhang, Y.; Chen, Y.; Wang, L.; Luu, A. T.; Bi, W.; Shi, F.; and Shi, S. 2023b. Siren’s Song in the AI Ocean: A Survey on Hallucination in Large Language Models. *CoRR*.
- Zhu, Y.; Yuan, H.; Wang, S.; Liu, J.; Liu, W.; Deng, C.; Dou, Z.; and Wen, J. 2023. Large Language Models for Information Retrieval: A Survey. *CoRR*.
- Zhu, Y.; Zhang, P.; Zhang, C.; Chen, Y.; Xie, B.; Dou, Z.; Liu, Z.; and Wen, J. 2024. INTERS: Unlocking the Power of Large Language Models in Search with Instruction Tuning. *CoRR*.