

Imagination Augmented Generation: Learning to Imagine Richer Context for Question Answering over Large Language Models

Anonymous ACL submission

Abstract

Retrieval-Augmented-Generation and *Generation-Augmented-Generation* have been proposed to enhance the knowledge required for question answering over Large Language Models (LLMs). However, the former relies on external resources, and both require incorporating explicit documents into the context, which increases execution costs and susceptibility to noise data. Recent works indicate that LLMs have modeled rich knowledge, albeit not effectively triggered or awakened. Inspired by this, we propose a novel knowledge-augmented framework, **Imagination-Augmented-Generation (IAG)**, which simulates the human capacity to compensate for knowledge deficits while answering questions solely through imagination, thereby awakening relevant knowledge in LLMs without relying on external resources. Guided by IAG, we propose an **imagine richer context** method for **question answering (IMcQA)**. IMcQA consists of two modules: *explicit imagination*, which generates a short dummy document by learning from long context compression, and *implicit imagination*, which creates flexible adapters by distilling from a teacher model with a long context. Experimental results on three datasets demonstrate that IMcQA exhibits significant advantages in both open-domain and closed-book settings, as well as in out-of-distribution generalization¹.

1 Introduction

Knowledge-intensive tasks like question answering (QA) necessitate utilizing extensive world and domain knowledge (Berant et al., 2013; Joshi et al., 2017; Kwiatkowski et al., 2019). Nowadays, Large Language Models (LLMs) have displayed notable competencies in almost every task and industry (Liu et al., 2023b). However, LLMs lack the sufficient capability to independently handle

¹Our code will be available at <https://anonymous.4open.science/r/IMcQA>

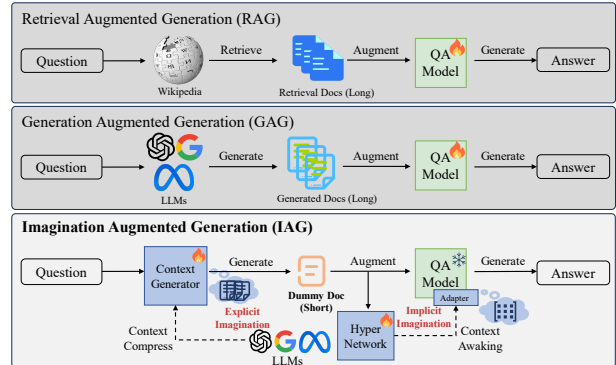


Figure 1: Compared with RAG and GAG, the proposed IAG eschews external resources, generates a shorter context (explicitly imagination) and creates flexible adapters (implicitly imagination) for each question.

knowledge-intensive tasks (Frisoni et al., 2024) and usually generate hallucinations (Zhao et al., 2023).

In recent years, to mitigate hallucinations in LLMs and improve performance in knowledge-intensive tasks such as QA, researchers have proposed numerous knowledge-augmented methods for LLMs. These methods primarily fall into two categories: *Retrieval-Augmented-Generation (RAG)* (Guu et al., 2020) and *Generation-Augmented-Generation (GAG)* (Kim et al., 2024). RAG (Top part of Figure 1) retrieves related documents from external resources (e.g., Wikipedia) and then inputs both the retrieved documents and the question into LLMs (Izacard and Grave, 2021). GAG (Middle part of Figure 1) leverages external LLMs such as ChatGPT (Ouyang et al., 2022) to generate relevant documents (Yu et al., 2023).

However, these methods have the following disadvantages²: 1) **Dependence on external resources**, RAG relies on external domain knowledge resources (Ke et al., 2024), while GAG depends on a more powerful external LLM as a knowledge generator. This dependency hinders their broader application. 2) **Increased execution costs**, The

²A more intuitive comparison can be seen in A.1.

required computing resources and inference time continue to increase as the number of documents increases (Xu et al., 2023a). For example, the typical RAG method FiD (Izacard and Grave, 2021) needs to handle over 12K tokens for retrieving 100 documents, leading to more than a 100-fold increase in prompt length and over 100²-fold increase in inference time (Liu et al., 2023a). The GAG method (Yu et al., 2023) requires additional financial costs (e.g., API calls). 3) **Susceptibility to noise data**, both methods are prone to the inclusion of irrelevant or noisy data (Li et al., 2023; Shaier et al., 2024). For example, Gao et al. (2024) indicates that noise in the documents negatively affects the performance. Therefore, there is an urgent need to explore new knowledge-augmented methods.

In fact, LLMs inherently contain rich knowledge and possess significant potential for resolving knowledge-intensive tasks (Bhagavatula et al., 2020). Enhancing the performance of specific tasks can be achieved by better activating and awakening relevant knowledge or expanding memory capacity without relying on external resources. For example, simply repeating the question twice (Xu et al., 2023b), just reviewing and consolidating knowledge by appending a straightforward prompt “As far as I know” (Yao et al., 2023), and using visual-language models to imagine images (Tang et al., 2023), all these approaches enhance the performance of LLMs on downstream tasks.

Therefore, inspired by the above phenomena and aiming to address the challenges faced by RAG and GAG, we propose a novel knowledge-augmented framework, *Imagination-Augmented-Generation* (IAG). **IAG simulates the human capacity to compensate for knowledge deficits solely through imagination in QA.** For resolving knowledge-intensive tasks, IAG utilizes solely LLMs to generate a shorter context and create a more flexible adapter (Bottom part of Figure 1). Through a knowledge distillation framework, IAG achieves knowledge enhancement through imagination. The fundamental idea is to **permit student models, which are deficient in rich contextual information (e.g., without abundant retrieved documents), to imitate teacher models that are endowed with such information.**

Within the framework of IAG, we introduce an **imagine richer context** method for question answering (IMcQA). To sufficiently utilize the inherent knowledge of LLMs, we design two main

modules to activate the various potential knowledge modeling in LLMs and obtain a richer context. Specifically, the **explicit imagination** module first uses symbol distillation to obtain the compressed context and then guides the context generator to generate a short dummy document. Subsequently, within the knowledge distillation framework, the **implicit imagination** module utilizes the hypernetwork to generate adapters to awaken the question-aware ability of LLMs. Unlike the LoRA (Hu et al., 2021) stores task knowledge and ability in modules, the hypernetwork learns to generate a question-aware knowledge module.

We evaluate the proposed IMcQA on various LLMs, including T5 (Roberts et al., 2020a) and Llama2 (Touvron et al., 2023). The experimental results across three QA datasets indicate that the proposed method yields performance gains while reducing computational expenses and time. Notably, it even outperforms baseline methods that retrieve and generate knowledge under the same document settings. In conclusion, the contributions of this paper are summarized as follows:

- We propose a new knowledge augmentation framework IAG to fully leverage the LLMs’ intrinsic knowledge more efficiently without relying on external resources.
- We propose a novel QA method IMcQA that employs two main modules (explicit imagination and implicit imagination) to utilize better the knowledge stored in the LLMs and awaken a richer context in QA.
- Experimental results indicate that the proposed method successfully activates the relevant internal knowledge of LLMs. IMcQA exhibits significant advantages in both open-domain and closed-book settings, as well as in out-of-distribution generalizations.

2 Related Work

This paper mainly utilizes context compression, hypernetworks and knowledge distillation to achieve knowledge enhancement. The following will elucidate pertinent research across four facets.

Knowledge Enhancement has usually been adopted to alleviate the issue of insufficient knowledge in LLMs. There are two main methods: RAG (Sun et al., 2019; Wang et al., 2024) and GAG (Abdallah and Jatowt, 2023). The typical RAG method

FiD (Izacard and Grave, 2021) retrieves documents from Wikipedia to answer questions. LLMs serving as a knowledge base have been the focus of numerous studies that advocate the extraction of knowledge from such models (e.g., GPT-3). For instance, Yu et al. (2023) generates 10 documents for each question. However, RAG requires external resources, and both RAG and GAG need verbose long contexts. Recently, methods have been developed to enhance LLMs’ abilities by simulating human imagination of visual information using existing visual-language models (Tang et al., 2023; Akter et al., 2024). We prefer self-imagination to augment knowledge and aim to leverage the parameterized knowledge within the models (Xu et al., 2023b; Kazemnejad et al., 2023). Our proposed method not only eliminates the need for external resources but also improves the efficiency of activating internal knowledge within LLMs.

Context Compression has often been used to improve the efficiency of LLMs in processing long contexts. Recent studies (Mu et al., 2023) propose that long contexts be condensed into summary vectors (soft prompts) to ensure their effective utilization by LLMs. Simultaneously, some studies (Jiang et al., 2023; Pan et al., 2024) suggest utilizing information redundancy and entropy in lengthy texts to compress contexts (Li et al., 2023). Unlike these approaches, this paper aims to enhance the long-context modeling ability of LLMs. By developing a context generator that creates compressed contexts, the QA model operating on short contexts can achieve a rich contextual understanding similar to models designed for longer contexts.

Knowledge Distillation is a technique where a smaller model learns to mimic the predictions of a larger model, aiming to retain performance while reducing computational resources (Hinton et al., 2015). Recent studies (West et al., 2022) present symbolic knowledge distillation, a process that facilitates knowledge transfer from a teacher model via extracting training data to subsequently train a student model (Wang et al., 2023b; Ranaldi and Freitas, 2024). In this paper, the process of obtaining compressed context during the explicit imagination resembles a form of symbolic distillation. Regarding training, our emphasis lies in distilling the long-context modeling abilities of LLMs.

Hypernetworks is designed to reduce the number of parameters (Ha et al., 2016), i.e., a small neural network generates parameters for another big

neural network. Recent studies (Phang et al., 2022; Ivison et al., 2023) have explored the enhancement of model performance in zero/few-shot settings through meta-learning involving hypernetworks. We utilize hypernetworks to acquire implicit imagine capabilities by dynamically generating LoRA for efficiency and generalization.

3 Method

In this section, we introduce the detailed method of IMcQA to activate LLMs’ intrinsic knowledge and obtain a richer context for QA. The fundamental premise underlying this method is that QA with a richer context yields greater performance (e.g., RAG with retrieved documents). Consequently, diverse methods are employed for questions lacking in richer contexts to activate knowledge within LLMs to replicate comparable effects to those achieved with richer contexts.

Specifically, IMcQA comprises two main modules. Explicit imagination with long context compression learns to imagine a short dummy document (§ 3.2). And implicit imagination with the hypernetwork models’ hidden knowledge that learns a shared knowledge feature projection across questions (§ 3.3). The hypernetwork is trained to generate lightweight LoRA modules, aiming to align the question and the internal knowledge. Besides, there is long context distillation in training, which learns the teacher’s rich representations to compensate for missing knowledge in imagination (§ 3.4).

3.1 Formulation

The formulation of our task follows RAG for QA (Guu et al., 2020). Let \mathcal{V}^* denote the infinite set, encompassing all potential strings over the tokens in vocabulary \mathcal{V} , and this includes the empty string. An instance within a QA dataset is defined as a triplet (q, a, c) comprising question q , answer a , and context c , where $q, a, c \in \mathcal{V}^*$. Conventionally, the context c is drawn from the knowledge corpus \mathcal{Z} , like Wikipedia, whereby $\mathcal{Z} \subset \mathcal{V}^*$. Additional background details are available in B.1.

3.2 Explicit Imagination with Context Generator

To obtain a context c for each question, we utilize LLMs to generate a short dummy document, which can mitigate *knowledge corpus error* (Lee et al., 2023) by considering potentially useful contexts. We employ the "teacher-student" framework. Dur-

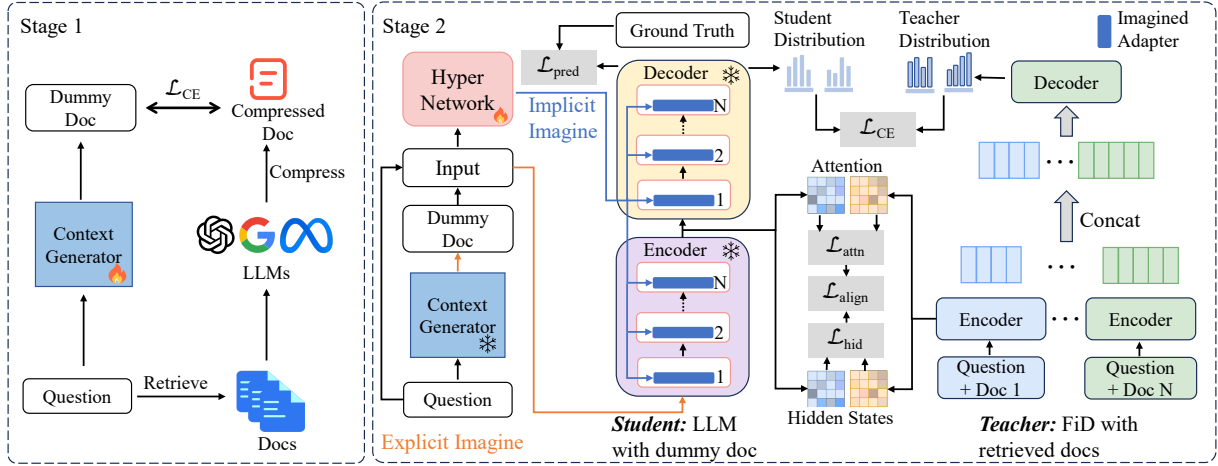


Figure 2: Overview of IMcQA method. In the inference phase, for each question, the explicit imagination generates a short dummy document and the implicit imagination with hypernetwork creates a specific LoRA module. During training, there are two stages: the first stage is the **pre-training** of the **context generator** (§ 3.2), aiming at its ability to imagine a short dummy document based on the question, and the second stage is the **hypernetwork fine-tuning** (§ 3.3) using long context distillation (§ 3.4) to obtain a question-specific LoRA module.

ing the training phase, the teacher model, equipped with rich contextual information, generates short contexts through context compression. The student model, in turn, learns to automatically generate these short contexts by leveraging its own LLM capabilities.

As shown in the left part of Figure 2, to help LLMs fully utilize the knowledge and generate compressed text, we first pretrain the context generator on our collected question-compressed document pairs. By leveraging symbolic distillation, we employ the long context compression method LongLLMLingua (Jiang et al., 2023) to condense a large corpus of retrieved documents. These compressed texts c' then serve as fine-tuning data alongside specific prompts p_q (B.2) and question-answer pairs for the **context generator**³ G_θ (θ represents the model’s parameter), which guides the model to think about its knowledge and generate a short dummy document:

$$d = G_\theta(p_q(q; c')) \quad (1)$$

where d is the dummy document generated from the **context generator**. This process enables LLMs to conceive compressed knowledge that robustly parallels the question’s knowledge requirements.

3.3 Implicit Imagination with Hypernetwork

Generally speaking, richer context can help LLM better answer questions. That is, the representation

of questions and the internal state of LLM when utilizing rich context are the better states. Therefore, in the absence of context, we should focus on building models to awaken LLM to achieve this better state and as a better QA model.

We utilize the hypernetwork⁴ to learn to generate a better QA model, which does not directly optimize the LoRA (Hu et al., 2021) module but generates specific LoRA adapter weights using the inputs for QA (bottom part of Figure 2). This is akin to repeating the question in the prompt (Xu et al., 2023b) and incorporating certain topical cues to stimulate the model’s recall of relevant questions (Wang et al., 2023c). However, the distinction lies in the fact that they serve as wake-up features, whereas we are generating model parameters.

The hypernetwork architecture for generating LoRA weights is exhaustively outlined in Figure 3. Specifically, D_k^q and U_k^q represent the low-rank projections of layer k correlated with the \mathcal{Q} , while D_k^v and U_k^v correspond to those associated with the \mathcal{V} . Hypernetwork represented as g_D and g_U , takes $\text{concat}(f, i_k^{\{q,v\}})$ as input, where f is the feature vectors that use the model’s encoder to obtain and using whitening algorithm (Su, 2021) for dimensionality reduction, $i_k^{\{q,v\}} \in \{0, \dots, 2 \times \#\text{blocks}\}$ signifies the positional embedding differentiating between layers and between $\mathcal{Q}\mathcal{V}$. Each hypernetwork is defined by weights W_d and W_u which represent the down and up projections respectively.

⁴We conduct a detailed analysis of the reasons behind the hypernetwork in the A.3.

³We discuss the role of context generator in the A.2.

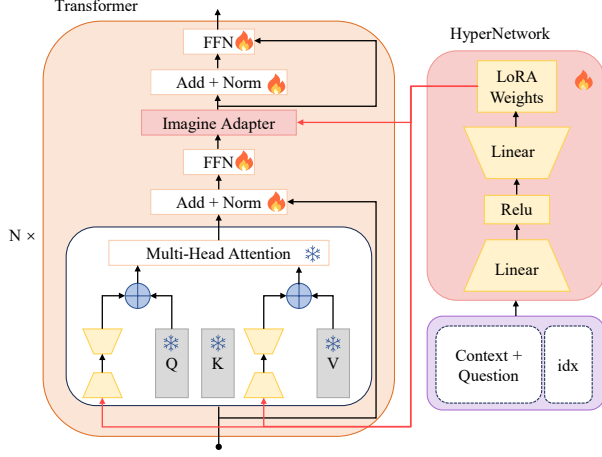


Figure 3: The Architecture of hypernetwork. Hypernetwork generates LoRA adapter weights for each question. During training, only Hypernetwork, FFN, and Norm weights are updated.

Finally, the hypernetwork equations for $D^{\{q,v\}}$ and $U^{\{q,v\}}$ can be expressed as:

$$f = \text{whitening}(\text{Encoder}(q; d)) \quad (2)$$

$$x = \text{concat}(f, i_k^{\{q,v\}} \mid i_k^q = 2k, i_k^v = 2k + 1) \quad (3)$$

$$D^{\{q,v\}}, U^{\{q,v\}} = g_D(x), g_U(x) \quad (4)$$

where the **Encoder** represents the encoder of the model, **whitening** is a dimensionality reduction algorithm, and **concat** means to splice the content. g_D and g_U denote the descending and ascending dimensions of hypernetwork. More formally,

$$g(x) = \text{MM}(\text{ReLU}(\text{MM}(W_d, W_u))) \quad (5)$$

where MM stands for matrix multiplication, ReLU is a activation function.

3.4 Training with Long Context Distillation

Within the framework of knowledge distillation, components such as hidden representations (Jiao et al., 2020), attention dependencies (Wang et al., 2020), and relations among representations (Park et al., 2021) are regarded as valuable knowledge for transfer. In this paper, we consider long context distillation (LCD) as the contextualized knowledge that mainly guides the student.

Specifically, the teacher model FiD (Izacard and Grave, 2021), which utilizes longer contextual inputs and theoretically contains more information (richer context). It will activate more specific internal knowledge and serve as a supervisory model. The teacher model assists the student model T5

(Roberts et al., 2020a), which has the same size as the teacher and leverages short contextual inputs. This aids in activating richer feature representations and knowledge. The optimization objective for the student model at each mini-batch $z_r = (x_r, y_r)$ is:

$$\mathcal{L}_s(\theta_s, \theta_t, z_r) = \alpha \mathcal{L}_{\text{ce}}(y_r, S(x_r; \theta_s)) + (1 - \alpha) \mathcal{L}_{\text{ce}}(T(x_r; \theta_t), S(x_r; \theta_s)) \quad (6)$$

where we have a teacher model denoted as $T(\cdot; \theta_t)$ and a student model denoted as $S(\cdot; \theta_s)$. The corresponding model parameters are θ_t and θ_s .

As shown in the right of Figure 2, we perform an additional representation alignment for better knowledge transfer. In our distillation, both teacher model and student model have the L layers, we feed the text into them and can obtain the corresponding output hidden states $\{H_l^t\}_{l=0}^L$, $\{H_l^s\}_{l=0}^L$, and attention matrices $\{A_l^t\}_{l=1}^L$, $\{A_l^s\}_{l=1}^L$. We suppose the student's l -th layer is aligned with the teacher's l -th layer, then the outputs of the student (i.e., H_l^s and A_l^s) should be close to the teacher's (i.e., H_l^t and A_l^t). For aligning hidden states, following (Park et al., 2021), we use cosine distance COS to calculate the proximity between the hidden states of the teacher and the student:

$$\mathcal{L}_{\text{hid}} = -\text{COS}(H_l^s, H_l^t) \quad (7)$$

While for aligning attention dependencies, we follow (Jiao et al., 2020) to optimize the mean square error (MSE) between the attention matrices of the teacher and the student:

$$\mathcal{L}_{\text{attn}} = -\text{MSE}(A_l^s, A_l^t) \quad (8)$$

The overall objective for knowledge transfer is:

$$\mathcal{L}_{\text{align}}(H_l^s, H_l^t, A_l^s, A_l^t) = \mathcal{L}_{\text{attn}} + \mathcal{L}_{\text{hid}} \quad (9)$$

The overall objective for training IMcQA is the weighted sum of the two objectives:

$$\mathcal{L} = \mathcal{L}_s + \lambda \mathcal{L}_{\text{align}} \quad (10)$$

4 Experiment

In this section, we conduct experiments to demonstrate the effectiveness and efficiency of IAG and IMcQA on QA. The experiment mainly answers four research questions (RQs):

RQ1: Can IAG and IMcQA achieve knowledge augmentation for QA over LLMs? (§ 4.4)

| Models | # Docs | NQ | | TriviaQA | | WebQ | |
|---|-----------------|--------------|-------------------------|--------------|-------------------------|--------------|-------------------------|
| | | Large (800M) | XL (3B) | Large (800M) | XL (3B) | Large (800M) | XL (3B) |
| <i># Closed-book Setting</i> | | | | | | | |
| T5 (Roberts et al., 2020a) | 0 | 28.5* | 28.30 | 28.7* | 33.92 | 30.6* | 34.43 |
| LoRA (Hu et al., 2021) | 0 | 17.70 | 23.15 | 23.87 | 32.16 | 29.13 | 35.24 |
| IMcQA (Ours) | 0 | 29.32 | 29.59 | 30.11 | 35.71 | 32.68 | 37.40 |
| <i># Retrieval Augmented Setting (compared with RAG)</i> | | | | | | | |
| DPR* (Karpukhin et al., 2020) (110M) | 100 | 41.5 | - | 56.8 | - | 41.1 | - |
| RAG* (Lewis et al., 2020) | 10 | 44.5 | - | 56.1 | - | 45.2 | - |
| FiD* (Izacard and Grave, 2021) | 10 | 46.7 | 50.1 | 61.9 | 66.3 | 48.1 | 50.8 |
| FiD (Izacard and Grave, 2021) | 100 | <u>51.4*</u> | 55.2[‡] | 67.6* | 72.9[‡] | 50.5 | 52.9[‡] |
| EAR (Chuang et al., 2023) | 10 | 39.6 | 42.3* | 60.0 | 64.6* | - | - |
| RFiD (Wang et al., 2023a) | 10 | 48.3 | 50.5 | 63.4 | 67.8 | - | - |
| FILCO* (Wang et al., 2023d) | 1 | - | 44.7 | - | 59.0 | - | - |
| IMcQA (Ours) | 10 | <u>49.9</u> | <u>50.9[‡]</u> | <u>69.7</u> | <u>70.3[‡]</u> | <u>51.5</u> | <u>52.8[‡]</u> |
| IMcQA (Ours) | 30 | 53.1 | - | 70.5 | - | 52.0 | - |
| <i># Generation Augmented Setting (compared with GAG)</i> | | | | | | | |
| GENREAD (sampling)* (Yu et al., 2023) | 10 [†] | 40.3 | 42.6 | 67.8 | 69.6 | 51.5 | 52.6 |
| GENREAD (clustering)* (Yu et al., 2023) | 10 [†] | 43.5 | 45.6 | 70.2 | <u>71.6</u> | 53.5 | 54.4 |
| IMcQA (Ours) | 10 [†] | 48.8 | 49.2[‡] | 70.9 | 72.2[‡] | 54.5 | 55.6[‡] |
| <i># Imagination Augmented Setting</i> | | | | | | | |
| LoRA (Hu et al., 2021) | 1 [†] | 40.1 | 44.2 | 62.8 | 66.9 | 43.7 | 48.2 |
| IMcQA (Ours) | 1 [†] | 42.3 | 46.5 | 65.5 | 68.4 | 45.3 | 50.5 |

Table 1: QA performances of different methods with different settings. The first part (closed-book setting) indicates that only utilize questions; The latter three parts utilize explicit documents. The best results are in bold, while the second-best are underlined. * means that those results are from existing papers, [†] denotes that the documents were generated ([‡] indicates that the number of documents is reduced due to insufficient memory for distillation).

RQ2: Does IMcQA have a good out-of-distribution generalization ability? (§ 4.5)

RQ3: Does IAG have advantages in effectiveness and efficiency compared to RAG and GAG? (§ 4.6)

RQ4: What is the role of explicit and implicit imagination modules in IMcQA? (§ 4.7)

4.1 Datasets

We evaluate the proposed approach on three public question answering datasets: NaturalQuestions (NQ) (Kwiatkowski et al., 2019), WebQuestions (WQ) (Berant et al., 2013) and TriviaQA (TQA) (Joshi et al., 2017). To evaluate the model performance, we use the exact match (EM) score for evaluating predicted answers (Rajpurkar et al., 2016). We provide dataset details in the B.4.

4.2 Baselines

Both the moderately sized language model (< 1B) and the large language model (LLM) ($\geq 3B$) are under consideration. T5 (Roberts et al., 2020a) is selected as the backbone for our moderately sized language models. We evaluate our proposed IMcQA against several knowledge-enhanced approaches, which include RAG models such as DPR (Karpukhin et al., 2020), RAG (Lewis et al., 2020), EAR (Chuang et al., 2023), RFiD (Wang et al.,

2023a), FILCO (Wang et al., 2023d) and FiD (Izacard and Grave, 2021), as well as the GAG model GENREAD (Yu et al., 2023), and parameters efficient fine-tuning method LoRA (Hu et al., 2021).

To demonstrate the plug-and-play capability of IMcQA on the zero-shot settings of LLMs ($\geq 3B$), we use Llama2-7B and -13B (Touvron et al., 2023) as the basic model. We evaluate with 6 diverse settings: without retrieval, with retrieval, with LoRA, RECITE (Sun et al., 2023), HICL (Wang et al., 2024) and using the proposed IMcQA.

4.3 Implementations

In the pretraining stage, the **context generator** initialized with T5-large utilizes the generated question-compressed pairs. During the second stage, the teacher model employs a FiD reader with different sizes (FiD-l and FiD-xl) that are fine-tuned on the training split of target datasets. The student model freezes the backbone and updates solely the hypernetwork⁵.

4.4 Main Results

4.4.1 Supervised Setting

Table 1 shows the performance results, full results including T5-Base are in the C.1. When juxtaposed

⁵B.3 contains more implementation and baseline details.




Figure 4: Zero-Shot results (EM, %) of Llama2-7B on three open-domain QA datasets. The number in parentheses indicates the number of documents used.

with **closed-book** models, **RAG**, and **GAG** methods, our proposed **IAG** framework **IMcQA** method exhibits state-of-the-art (SOTA) performance with the equivalent magnitude of document numbers.

In the closed-book setting (in the upper part of the table), our method outperforms the baseline by an average of +2% EM score, indicating its excellence in utilizing internal knowledge with imagination. It’s especially noteworthy that as the model size expands, the performance advantages of the imagination become ever more evident.

The following three parts show the experimental results under the open domain setting⁶. **Although our method only deals with one short dummy document, it can still achieve results similar to or better than the RAG and GAG methods, which handle 10 documents.** The findings reveal that IMcQA exploits imagined condensed text to strike a balance between efficiency and overhead. Beyond this, when IMcQA implements 10 retrieved documents under RAG, it overtakes RFiD’s performance, posting 1.6% higher in NQ, 4.4% in TQA, and 2.7% in WQ. IMcQA employing 10 generated documents under GAG outstrips GENREAD (clustering)’s performance, documenting a 4.5% increment in NQ, 0.7% in TQA, and 1.1% in WQ.

4.4.2 Zero-shot Setting

Figure 4 illustrates the zero-shot results for LLMs implementing IMcQA with a frozen Llama2-7B and -13B. This research seeks to explore the possibility of enhancing LLMs via IAG. Due to the high computational demands of training, we only fine-tuned the hypernetwork on a mixed dataset without LCD in this experiment and evaluated per-

⁶Due to memory constraints from online distillation, we only compare the IMcQA-1 under the RAG setting using 30 documents with FiD using 100 documents.

formance in a zero-shot setting⁷. More zero-shot setting results can be seen in C.3.

We discerned that Llama2’s performance can be enhanced by imagining knowledge autonomously. While leveraging explicit imagined context could amplify the average EM +1%, this is not as significant as the improvement achieved by retrieving 10 documents, indicating the limitations of relying solely on prompt cues for triggering corresponding knowledge. IMcQA can enhance knowledge via two main imagination processes, escalating EM by +15.33% for NQ, +11.97% for TQA, and +16.38% for WQ. Compared to two other advanced RAG methods, IMcQA using a single document, performs only 1 EM lower than the HICL method (Wang et al., 2024) on the TQA but achieves +10% EM on the NQ and +5% EM on the WQ. With IAG, Llama2-7B demonstrated an average improvement of +14% across the three datasets. This trend is also observed in Llama2-13B’s results (Figure 6). This implies that even in zero-shot settings, our method can still offer substantial benefits to LLMs.

| Models | # Docs | Base (220M) | | | Large (800M) | | |
|--------|----------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | | NQ | TQA | WQ | NQ | TQA | WQ |
| T5 | 0 | 22.16 | 3.18 | 4.12 | 28.5* | 3.18 | 4.12 |
| IMcQA | 0 | 23.89 | 6.21 | 10.94 | 29.32 | 10.17 | 14.06 |
| FiD | 10 | 46.81 | 53.93 | 24.02 | 46.7* | 57.93 | 25.12 |
| IMcQA | 10 | <u>47.01</u> | <u>55.74</u> | <u>24.13</u> | <u>49.92</u> | <u>60.03</u> | <u>25.79</u> |
| LoRA | 1 [†] | 5.43 | 3.15 | 4.02 | 4.42 | 6.50 | 7.87 |
| IMcQA | 1 [†] | 40.14 | 46.61 | 18.92 | 42.32 | 54.80 | 22.05 |

Table 2: **IID and OOD results.** The performance on three open-domain datasets for the model trained on NQ is reported, with the underlined values indicating the IID performance. Full OOD results and details of three datasets are shown in the C.2.

4.5 Out-Of-Distribution (OOD) Performance

To further demonstrate the generalizations of the IMcQA method and the importance of hypernetwork, we also evaluate its performance in OOD generalizations. Table 2 shows the IID and OOD performance of FiD, and IMcQA methods with different document settings when training on NQ (From NQ generalization to the other two datasets).

It is patently clear that an increment in document provision leads to better OOD performance, likely due to the presence of answer-oriented content within these documents. Remarkably, IMcQA can come within a relatively narrow 5% gap of FiD, even when utilizing a single imagined document as

⁷Detailed prompts information can be found in the B.2.

opposed to 10 retrieved documents.

Simultaneously, IMcQA generally showcases superior performance in OOD when provided with 10 retrieved documents. This superiority can be traced back to the pivotal role played by hypernetwork in generating LoRA adapters’ weights based on questions. This equips models with the capability to invoke and access internal knowledge based on context-specific discourse rather than confining to resolving distinct questions.

| Models | Training Params | # Documents | # Avg Tokens | Inference Time | GPU Memory | Training Time |
|---------|-----------------|-------------|--------------|----------------|------------|---------------|
| T5 | 220M | 0 | 19.8 | 79.8s | 2828M | 0.9h |
| IMcQA | 139.3M | 0 | 19.8 | 82.3s | 2710M | 1.2h |
| IMcQA | 139.3M | 1 | 522.1 | 214.6s | 2882M | 1.7h |
| FiD | 220M | 10 | 1748.3 | 683.3s | 4358M | 2.3h |
| GENREAD | 220M | 10 | 1912.5 | 704.8s | 4412M | - |
| FiD | 220M | 100 | 16625.7 | 1293.2s | 19048M | 5.8h |

Table 3: Training and inference cost on the NQ. The backbone model is T5-Base.

4.6 Training Cost and Inference Speed-up

We proceeded to measure the inference speed documented in GPU time and training time for 5000 steps on the NQ dataset using T5-Base. The experiments were conducted on a single RTX 3090 GPU, maintaining a standard batch size of 8 during training and 1 during inference.

As evident from Table 3, the proposed method’s advantage lies in its diminished requirement for parameter updates, which can be attributed to the shared hypernetwork’s utilization that generates LoRA adapters, thereby negating the necessity of individual LoRA adapters’ setup. Despite the lack of a training advantage due to distillation constraints, IMcQA achieves efficient reasoning through an extremely lightweight design, saving more than half the training time compared to methods using a large number of documents ($\times 0.3$). Compared to the other two methods, the processed tokens are significantly decreased, while either outperforming them or showing negligible differences in performance. This represents an optimal trade-off between efficiency and computational demand. Moreover, unlike GAG, our approach incurs no financial costs associated with API calls, and the reduced model size facilitates faster generation.

4.7 Ablation Experiment

This study introduces two key imagination processes to stimulate LLMs’ internal knowledge: Explicit Imagination (EI) and Implicit Imagination (II). We particularly examined the influence of different imagination types on performance.

Figure 5 demonstrates that both EI and II are important for IMcQA. Omitting either one results in a considerable reduction in performance, with a drop exceeding 10% observed when EI is neglected. This is harmonious with the initial observation that performance improvement becomes more noticeable when relevant documents are available, thus underscoring EI’s superiority.

The outcomes of Long Context Distillation (LCD) and the application of EI in the hypernetwork also make marginal contributions to the overall results. This validates the previous assertion that a more extensive context tends to optimize performance, although with limited gains. The experiments and analysis above demonstrate the importance of each component and the effectiveness of our IMcQA method.




Figure 5: Ablation experiment results (%) of T5-Base on three open-domain QA datasets.

5 Conclusion and Future Work

This study proposes a novel knowledge-augmented strategy for LLMs, namely Imagine Augmented Generation (IAG), and a specific method IMcQA for open domain question answering. The proposed method effectively activates and utilizes intrinsic knowledge within LLMs through two imaginations: explicit imagination, and implicit imagination. Experimental results demonstrate a significant improvement in QA performance while remaining relatively lightweight. Although the main focus of this method is on one specific task, we believe these findings can offer a novel perspective on how to better harness the potential of LLMs. In the future, we plan to apply IAG to more NLP tasks and explore multimodal knowledge-augmented generation.

586 Limitations

587 While this study has demonstrated significant
588 achievements in QA tasks, there are notable limita-
589 tions:

590 **Tasks.** The proposed methods in the study are spe-
591 cialized specifically for QA. It remains unknown
592 how effective they would be in other types of
593 knowledge-intensive tasks, such as fact-checking
594 or dialogue systems. Further validation is needed
595 to assess the generalizations and applicability of
596 this approach.

597 **Multimodal.** We have only considered imagined
598 text and hidden representations. In future work, it
599 is imperative to explore multimodal information
600 including the impact of imagining images on per-
601 formance.

602 **Method.** Our method relies on the knowledge
603 learned by LLMs in the pre-training phase, which
604 may limit the model’s ability to quickly adapt to
605 new information. The dependency on internal
606 knowledge activation in IAG may lead to a less
607 transparent decision-making process in the model,
608 making it challenging to explain the logic behind
609 the generated answers. In the future, there is a need
610 to continue exploring adaptive knowledge enhance-
611 ment methods to optimize results further.

612 Ethical Considerations

613 In this paper, we proposed a novel knowledge en-
614 hancement method aimed at leveraging the knowl-
615 edge of LLMs. However, LLMs may generate
616 inappropriate or discriminatory knowledge. Our
617 approach does not introduce ethical concerns. The
618 datasets we used are public, and there are no pri-
619 vacy issues.

620 References

- 621 Abdelrahman Abdallah and Adam Jatowt. 2023.
622 Generator-retriever-generator: A novel approach to
623 open-domain question answering. *arXiv preprint*
624 *arXiv:2307.11278*.
- 625 Syeda Nahida Akter, Aman Madaan, Sangwu Lee, Yim-
626 ing Yang, and Eric Nyberg. 2024. *Self-imagine: Ef-*
627 *fective unimodal reasoning with multimodal models*
628 *using self-imagination*.
- 629 J. Berant, A. Chou, R. Frostig, and P. Liang. 2013. Se-
630 mantic parsing on Freebase from question-answer
631 pairs. In *Empirical Methods in Natural Language*
632 *Processing (EMNLP)*.

- Chandra Bhagavatula, Ronan Le Bras, Chaitanya
633 Malaviya, Keisuke Sakaguchi, Ari Holtzman, Han-
634 nah Rashkin, Doug Downey, Wen tau Yih, and Yejin
635 Choi. 2020. *Abductive commonsense reasoning*. In
636 *International Conference on Learning Representa-*
637 *tions*. 638
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim
639 Sturge, and Jamie Taylor. 2008. *Freebase: a col-*
640 *laboratively created graph database for structuring*
641 *human knowledge*. In *Proceedings of the 2008 ACM*
642 *SIGMOD International Conference on Management*
643 *of Data, SIGMOD ’08*, page 1247–1250, New York,
644 NY, USA. Association for Computing Machinery. 645
- Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai,
646 Zhijian Liu, Song Han, and Jiaya Jia. 2023. *Longlora:*
647 *Efficient fine-tuning of long-context large language*
648 *models*. 649
- Yung-Sung Chuang, Wei Fang, Shang-Wen Li, Wen tau
650 Yih, and James Glass. 2023. *Expand, rerank, and*
651 *retrieve: Query reranking for open-domain question*
652 *answering*. 653
- Giacomo Frisoni, Alessio Cocchieri, Alex Presepi, Gian-
654 luca Moro, and Zaiqiao Meng. 2024. *To generate or*
655 *to retrieve? on the effectiveness of artificial contexts*
656 *for medical open-domain question answering*. 657
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia,
658 Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu Guo,
659 Meng Wang, and Haofen Wang. 2024. *Retrieval-*
660 *augmented generation for large language models: A*
661 *survey*. 662
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat,
663 and Mingwei Chang. 2020. *Retrieval augmented*
664 *language model pre-training*. In *Proceedings of the*
665 *37th International Conference on Machine Learning*,
666 volume 119 of *Proceedings of Machine Learning*
667 *Research*, pages 3929–3938. PMLR. 668
- David Ha, Andrew Dai, and Quoc V. Le. 2016. *Hyper-*
669 *networks*. 670
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015.
671 *Distilling the knowledge in a neural network*. *arXiv*
672 *preprint arXiv:1503.02531*. 673
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan
674 Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and
675 Weizhu Chen. 2021. *Lora: Low-rank adaptation of*
676 *large language models*. 677
- Hamish Ivison, Akshita Bhagia, Yizhong Wang, Han-
678 naneh Hajishirzi, and Matthew Peters. 2023. *Hint:*
679 *Hypernetwork instruction tuning for efficient zero-*
680 *shot generalisation*. *ACL*. 681
- Gautier Izacard and Edouard Grave. 2021. *Leveraging*
682 *passage retrieval with generative models for open do-*
683 *main question answering*. In *Proceedings of the 16th*
684 *Conference of the European Chapter of the Associ-*
685 *ation for Computational Linguistics: Main Volume*,
686 pages 874–880, Online. Association for Computa-
687 tional Linguistics. 688

| | | |
|-----|---|-----|
| 689 | Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023. Longllmlingua: Accelerating and enhancing llms in long context scenarios via prompt compression. | 744 |
| 690 | | 745 |
| 691 | | 746 |
| 692 | | 747 |
| | | 748 |
| 693 | Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. TinyBERT: Distilling BERT for natural language understanding. In <i>Findings of the Association for Computational Linguistics: EMNLP 2020</i> , pages 4163–4174, Online. Association for Computational Linguistics. | 749 |
| 694 | | 750 |
| 695 | | 751 |
| 696 | | 752 |
| 697 | | |
| 698 | | 753 |
| 699 | | 754 |
| | | 755 |
| 700 | Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In <i>Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics. | 756 |
| 701 | | 757 |
| 702 | | 758 |
| 703 | | 759 |
| 704 | | |
| 705 | | 760 |
| 706 | | 761 |
| | | 762 |
| 707 | 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 <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 6769–6781, Online. Association for Computational Linguistics. | 763 |
| 708 | | 764 |
| 709 | | |
| 710 | | 765 |
| 711 | | 766 |
| 712 | | |
| 713 | | 767 |
| | | 768 |
| 714 | Amirhossein Kazemnejad, Mehdi Rezagholizadeh, Prasanna Parthasarathi, and Sarath Chandar. 2023. Measuring the knowledge acquisition-utilization gap in pretrained language models. In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 4305–4319, Singapore. Association for Computational Linguistics. | 769 |
| 715 | | 770 |
| 716 | | 771 |
| 717 | | 772 |
| 718 | | 773 |
| 719 | | 774 |
| 720 | | |
| 721 | Zixuan Ke, Weize Kong, Cheng Li, Mingyang Zhang, Qiaozhu Mei, and Michael Bendersky. 2024. Bridging the preference gap between retrievers and llms. | 775 |
| 722 | | 776 |
| 723 | | 777 |
| | | 778 |
| 724 | Jaehyung Kim, Jaehyun Nam, Sangwoo Mo, Jongjin Park, Sang-Woo Lee, Minjoon Seo, Jung-Woo Ha, and Jinwoo Shin. 2024. Sure: Summarizing retrievals using answer candidates for open-domain QA of LLMs. In <i>The Twelfth International Conference on Learning Representations</i> . | 779 |
| 725 | | 780 |
| 726 | | |
| 727 | | 781 |
| 728 | | 782 |
| 729 | | 783 |
| | | 784 |
| | | 785 |
| 730 | Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: A benchmark for question answering research. <i>Transactions of the Association for Computational Linguistics</i> , 7:452–466. | 786 |
| 731 | | 787 |
| 732 | | 788 |
| 733 | | |
| 734 | | 789 |
| 735 | | 790 |
| 736 | | 791 |
| 737 | | 792 |
| 738 | | 793 |
| | | 794 |
| 739 | Yejoon Lee, Philhoon Oh, and James Thorne. 2023. Knowledge corpus error in question answering. In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 9183–9197, Singapore. Association for Computational Linguistics. | 795 |
| 740 | | 796 |
| 741 | | 797 |
| 742 | | 798 |
| 743 | | 799 |
| | | |
| | Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In <i>Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20</i> , Red Hook, NY, USA. Curran Associates Inc. | 744 |
| | | 745 |
| | | 746 |
| | | 747 |
| | | 748 |
| | | 749 |
| | | 750 |
| | | 751 |
| | | 752 |
| | Yucheng Li, Bo Dong, Chenghua Lin, and Frank Guerin. 2023. Compressing context to enhance inference efficiency of large language models. | 753 |
| | | 754 |
| | | 755 |
| | Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023a. Lost in the middle: How language models use long contexts. | 756 |
| | | 757 |
| | | 758 |
| | | 759 |
| | Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023b. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. <i>ACM Comput. Surv.</i> , 55(9). | 760 |
| | | 761 |
| | | 762 |
| | | 763 |
| | | 764 |
| | Jesse Mu, Xiang Lisa Li, and Noah Goodman. 2023. Learning to compress prompts with gist tokens. | 765 |
| | | 766 |
| | Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. | 767 |
| | | 768 |
| | | 769 |
| | | 770 |
| | | 771 |
| | | 772 |
| | | 773 |
| | | 774 |
| | Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Menglin Xia, Xufang Luo, Jue Zhang, Qingwei Lin, Victor Rühle, Yuqing Yang, Chin-Yew Lin, H. Vicky Zhao, Lili Qiu, and Dongmei Zhang. 2024. Llmlingua-2: Data distillation for efficient and faithful task-agnostic prompt compression. | 775 |
| | | 776 |
| | | 777 |
| | | 778 |
| | | 779 |
| | | 780 |
| | Geondo Park, Gyeongman Kim, and Eunho Yang. 2021. Distilling linguistic context for language model compression. In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 364–378. | 781 |
| | | 782 |
| | | 783 |
| | | 784 |
| | | 785 |
| | Jason Phang, Yi Mao, Pengcheng He, and Weizhu Chen. 2022. Hypertuning: Toward adapting large language models without back-propagation. | 786 |
| | | 787 |
| | | 788 |
| | Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In <i>Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing</i> , pages 2383–2392, Austin, Texas. Association for Computational Linguistics. | 789 |
| | | 790 |
| | | 791 |
| | | 792 |
| | | 793 |
| | | 794 |
| | Leonardo Ranaldi and Andre Freitas. 2024. Aligning large and small language models via chain-of-thought reasoning. In <i>Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages | 795 |
| | | 796 |
| | | 797 |
| | | 798 |
| | | 799 |

| | | |
|-----|--|-----|
| 800 | 1812–1827, St. Julian’s, Malta. Association for Computational Linguistics. | |
| 801 | | |
| 802 | Adam Roberts, Colin Raffel, and Noam Shazeer. 2020a. How much knowledge can you pack into the parameters of a language model? In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 5418–5426, Online. Association for Computational Linguistics. | |
| 803 | | |
| 804 | | |
| 805 | | |
| 806 | | |
| 807 | | |
| 808 | Adam Roberts, Colin Raffel, and Noam Shazeer. 2020b. How much knowledge can you pack into the parameters of a language model? In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 5418–5426, Online. Association for Computational Linguistics. | |
| 809 | | |
| 810 | | |
| 811 | | |
| 812 | | |
| 813 | | |
| 814 | Devendra Singh Sachan, Siva Reddy, William L. Hamilton, Chris Dyer, and Dani Yogatama. 2021. End-to-end training of multi-document reader and retriever for open-domain question answering . In <i>Advances in Neural Information Processing Systems</i> . | |
| 815 | | |
| 816 | | |
| 817 | | |
| 818 | | |
| 819 | Sagi Shaiyer, Lawrence E Hunter, and Katharina von der Wense. 2024. Desiderata for the context use of question answering systems . | |
| 820 | | |
| 821 | | |
| 822 | Jianlin Su. 2021. You probably don’t need bert-flow: A linear transformation comparable to bert-flow . | |
| 823 | | |
| 824 | Haitian Sun, Tania Bedrax-Weiss, and William Cohen. 2019. PullNet: Open domain question answering with iterative retrieval on knowledge bases and text . In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)</i> , pages 2380–2390, Hong Kong, China. Association for Computational Linguistics. | |
| 825 | | |
| 826 | | |
| 827 | | |
| 828 | | |
| 829 | | |
| 830 | | |
| 831 | | |
| 832 | | |
| 833 | Zhiqing Sun, Xuezhi Wang, Yi Tay, Yiming Yang, and Denny Zhou. 2023. Recitation-augmented language models . | |
| 834 | | |
| 835 | | |
| 836 | Tianyi Tang, Yushuo Chen, Yifan Du, Junyi Li, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Learning to imagine: Visually-augmented natural language generation . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 9468–9481, Toronto, Canada. Association for Computational Linguistics. | |
| 837 | | |
| 838 | | |
| 839 | | |
| 840 | | |
| 841 | | |
| 842 | | |
| 843 | Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, | |
| 844 | | |
| 845 | | |
| 846 | | |
| 847 | | |
| 848 | | |
| 849 | | |
| 850 | | |
| 851 | | |
| 852 | | |
| 853 | | |
| 854 | | |
| 855 | | |
| 856 | | |
| 857 | | |
| | Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models . | 858 |
| | | 859 |
| | | 860 |
| | | 861 |
| | | 862 |
| | | 863 |
| | | 864 |
| | | 865 |
| | Cunxiang Wang, Haofei Yu, and Yue Zhang. 2023a. RFiD: Towards rational fusion-in-decoder for open-domain question answering . In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 2473–2481, Toronto, Canada. Association for Computational Linguistics. | 866 |
| | | 867 |
| | | 868 |
| | | 869 |
| | | 870 |
| | | 871 |
| | Peifeng Wang, Zhengyang Wang, Zheng Li, Yifan Gao, Bing Yin, and Xiang Ren. 2023b. Scott: Self-consistent chain-of-thought distillation . | 872 |
| | | 873 |
| | | 874 |
| | Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers . <i>Advances in Neural Information Processing Systems</i> , 33:5776–5788. | 875 |
| | | 876 |
| | | 877 |
| | | 878 |
| | | 879 |
| | Xinyi Wang, Wanrong Zhu, Michael Saxon, Mark Steyvers, and William Yang Wang. 2023c. Large language models are latent variable models: Explaining and finding good demonstrations for in-context learning . | 880 |
| | | 881 |
| | | 882 |
| | | 883 |
| | | 884 |
| | Yifan Wang, Qingyan Guo, Xinzhe Ni, Chufan Shi, Lema Liu, Haiyun Jiang, and Yujiu Yang. 2024. Hint-enhanced in-context learning wakes large language models up for knowledge-intensive tasks . | 885 |
| | | 886 |
| | | 887 |
| | | 888 |
| | Zhiruo Wang, Jun Araki, Zhengbao Jiang, Md Rizwan Parvez, and Graham Neubig. 2023d. Learning to filter context for retrieval-augmented generation . | 889 |
| | | 890 |
| | | 891 |
| | Peter West, Chandra Bhagavatula, Jack Hessel, Jena Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. Symbolic knowledge distillation: from general language models to commonsense models . In <i>Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 4602–4625, Seattle, United States. Association for Computational Linguistics. | 892 |
| | | 893 |
| | | 894 |
| | | 895 |
| | | 896 |
| | | 897 |
| | | 898 |
| | | 899 |
| | | 900 |
| | | 901 |
| | Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023a. Re-comp: Improving retrieval-augmented lms with compression and selective augmentation . | 902 |
| | | 903 |
| | | 904 |
| | Xiaohan Xu, Chongyang Tao, Tao Shen, Can Xu, Hongbo Xu, Guodong Long, and Jian guang Lou. 2023b. Re-reading improves reasoning in language models . | 905 |
| | | 906 |
| | | 907 |
| | | 908 |
| | Yunzhi Yao, Shaohan Huang, Li Dong, Furu Wei, Huajun Chen, and Ningyu Zhang. 2022. Kformer: Knowledge injection in transformer feed-forward layers . <i>ArXiv</i> , abs/2201.05742. | 909 |
| | | 910 |
| | | 911 |
| | | 912 |

| | | | |
|-----|---|--|------|
| 913 | Yunzhi Yao, Peng Wang, Shengyu Mao, Chuanqi Tan, | may encounter knowledge it has not yet learned. | 963 |
| 914 | Fei Huang, Huajun Chen, and Ningyu Zhang. 2023. | Thus, IAG does not rely on external large models or | 964 |
| 915 | Knowledge ruminations for pre-trained language mod- | retrievers for external reasoning and can complete | 965 |
| 916 | els . In <i>Proceedings of the 2023 Conference on Em-</i> | reasoning independently. | 966 |
| 917 | <i>pirical Methods in Natural Language Processing</i> , | | |
| 918 | pages 3387–3404, Singapore. Association for Com- | | |
| 919 | putational Linguistics. | | |
| 920 | Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, | A.3 Hypernetwork | 967 |
| 921 | Mingxuan Ju, Soumya Sanyal, Chenguang Zhu, | Hypernetworks have gained significant attention | 968 |
| 922 | Michael Zeng, and Meng Jiang. 2023. Generate | in recent years due to their potential to enhance | 969 |
| 923 | rather than retrieve: Large language models are | various aspects of neural network performance. In | 970 |
| 924 | strong context generators. In <i>International Confer-</i> | this section, we analyze the reasons for employing | 971 |
| 925 | <i>ence for Learning Representation (ICLR)</i> . | hypernetworks in detail: | 972 |
| 926 | Yukun Zhao, Lingyong Yan, Weiwei Sun, Guoliang | Hypernetworks, which are neural networks de- | 973 |
| 927 | Xing, Chong Meng, Shuaiqiang Wang, Zhicong | signed to generate the weights of other networks, | 974 |
| 928 | Cheng, Zhaochun Ren, and Dawei Yin. 2023. Know- | allow for dynamic adjustment of model param- | 975 |
| 929 | ing what llms do not know: A simple yet effective | eters. This adaptability enables the model to better | 976 |
| 930 | self-detection method . | suit different tasks and datasets, thereby improv- | 977 |
| 931 | | ing overall performance. By utilizing hypernet- | 978 |
| | A Method | works, the number of models that need to be trained | 979 |
| 932 | A.1 Comparison of Three Paradigms | individually can be significantly reduced. Tradi- | 980 |
| 933 | Compared to RAG and GAG, our method has cer- | tional methods require separate models for each | 981 |
| 934 | tain limitations, such as requiring a more complex | task, whereas hypernetworks can generate weights | 982 |
| 935 | training process and the necessity of training a | for multiple tasks. This capability enhances train- | 983 |
| 936 | model. Similar to the GAG method, which uses a | ing efficiency. In our task, we use hypernetworks to | 984 |
| 937 | master’s degree in law as a knowledge base, our | generate adapters for the question and input, which | 985 |
| 938 | method also struggles to generate content when | are then inserted into the model. This helps the | 986 |
| 939 | encountering new and unknown world knowledge, | model incorporate the knowledge targeted by the | 987 |
| 940 | which presents a challenge that needs to be ad- | question, corresponding to implicit imagination. | 988 |
| 941 | ressed. Additionally, the knowledge base might | Compared to traditional efficient fine-tuning, this | 989 |
| 942 | be affected by knowledge gaps in low-resource | process is more aligned with the goal of imagina- | 990 |
| 943 | settings where there is a lack of a comprehensive | tion. | 991 |
| 944 | knowledge base. | Hypernetworks can capture the commonalities | 992 |
| 945 | Next, we compare IAG, RAG, and GAG across | and differences between various tasks by learn- | 993 |
| 946 | four criteria for a more intuitive understanding. | ing to generate weights. This ability to general- | 994 |
| 947 | From the table 4, it can be observed that the docu- | ize across tasks improves the model’s performance | 995 |
| 948 | ment relevance obtained by IAG and GAG is higher, | on unseen data, making it more robust in diverse | 996 |
| 949 | while RAG heavily relies on the retriever and exter- | scenarios. In multi-task learning or meta-learning | 997 |
| 950 | nal knowledge base. In terms of document length | scenarios, hypernetworks can considerably reduce | 998 |
| 951 | usage, IAG only needs to use a virtual document, | the need for storing multiple independent models. | 999 |
| 952 | greatly reducing the number of tokens. Therefore, | A hypernetwork only needs to store a single gen- | 1000 |
| 953 | IAG is superior to the other two methods in terms | erating network and some shared parameters, thus | 1001 |
| 954 | of reasoning time. | significantly decreasing the storage space required. | 1002 |
| 955 | A.2 Context Generator | Hypernetworks can quickly generate new weights | 1003 |
| 956 | There are two main goals in the pre-training of the | to adapt to new tasks as they arise. This rapid | 1004 |
| 957 | model in the first stage of IAG (context generator): | adaptation capability is particularly useful in ap- | 1005 |
| 958 | first, to improve its ability as a document generator | plications that require frequent updates or expan- | 1006 |
| 959 | by learning to generate rich and concise documents; | sions. In our experiments 4.5, we also found that | 1007 |
| 960 | second, to introduce some external knowledge that | using a hypernetwork can significantly enhance | 1008 |
| 961 | the model initially does not possess. It’s worth | the generalization ability for tasks. This is be- | 1009 |
| 962 | noting that the second goal is crucial, as the model | cause it not only retains knowledge within the | 1010 |
| | | domain-specific modules but also learns to gen- | 1011 |
| | | erate question-targeted knowledge to be inserted | 1012 |

| | Document Relevance | Context Length | Inference Time | Inference Dependence |
|-----|--------------------|----------------|----------------|----------------------------|
| RAG | Medium | Too Long | Very High | Retriever |
| GAG | High | Long | High | Larger Model (InstructGPT) |
| IAG | High | Short | Low | None |

Table 4: Comparison of Different Paradigms

into the model.

B Experimental Settings

B.1 Background

Our task formulation follows retrieval augmented models for QA (Guu et al., 2020; Sachan et al., 2021). Let \mathcal{V}^* denote the infinite set, encompassing all potential strings over the tokens in vocabulary \mathcal{V} , and this includes the empty string. An instance within a QA dataset is defined as a triplet (q, a, c) comprising question q , answer a , and context c , where $q, a, c \in \mathcal{V}^*$. Conventionally, the context c is drawn from the knowledge corpus \mathcal{Z} , like Wikipedia, whereby $\mathcal{Z} \subset \mathcal{V}^*$.

The goal of QA is to learn a distribution function, represented as $p(a|q)$, wherein the models decode a string a that serves as an abstractive answer to a given query q . In a closed-book setting, LLMs directly encode the given question and predict the answer (Roberts et al., 2020b). Specifically, considering the context c as the empty string, the reliance is solely on the model parameters, i.e., $\hat{a} = \arg \max_{a \in \mathcal{V}^*} p(a|q, \theta)$, where θ represents the LLMs’ parameters. However, employing a direct approach of requesting models to output answers frequently results in subpar performance, primarily attributable to omitting a substantial amount of world knowledge during the process. Therefore, a popular approach is open domain setting, which marginalizes $p(a|q, c)$ over contexts c in the knowledge corpus (Lewis et al., 2020; Sachan et al., 2021) or generated from models (Yu et al., 2023). Given the computational infeasibility of calculating probabilities for all contexts, $p(a|q, c)$ is approximated to the sum of probabilities for top k contexts, i.e., $p(a|q, c) = \sum_{c_i \in c} p(a|q, c_i)p(c_i|q)$, where $c_i \in c$ and $c \in \text{Top}K(q)$. $\text{Top}k(q)$ denotes the set of resulting top k passages after the retrieval or generated with a query q .

B.2 Prompts for Explicit Imagine with LLMs

The prompt for explicit imagination of the context generator to imagine a short dummy useful docu-

ment is:

Imagine contexts based on the question: \n input
\n Contexts: \n

Table 11 shows the full prompts for zero-shot results on LLM that we use for open domain QA: NQ, TQA, WQ.

| Models | Docu-ments | Steps | Lr | Batch Size |
|-----------|----------------|-------|------|------------|
| T5 | 0 | 40000 | 1e-4 | 8 |
| LoRA-Base | 0 | 40000 | 5e-4 | 8 |
| IMcQA | 0 | 50000 | 1e-3 | 8 |
| LoRA-l | 0 | 40000 | 1e-4 | 4 |
| IMcQA-l | 0 | 50000 | 5e-4 | 4 |
| FiD-3b | 0 | 40000 | 1e-4 | 2 |
| LoRA-3b | 0 | 40000 | 1e-4 | 4 |
| IMcQA | 0 | 50000 | 1e-4 | 1 |
| LoRA-Base | 0 [†] | 40000 | 5e-4 | 8 |
| IMcQA | 0 [†] | 50000 | 1e-3 | 8 |
| LoRA-l | 0 [†] | 40000 | 1e-4 | 4 |
| IMcQA-l | 0 [†] | 50000 | 5e-4 | 4 |
| LoRA-3b | 0 [†] | 40000 | 1e-4 | 2 |
| IMcQA-3b | 0 [†] | 50000 | 1e-4 | 1 |
| IMcQA | 10 | 50000 | 5e-4 | 1 |
| IMcQA-l | 10 | 50000 | 5e-4 | 1 |
| FiD-3b | 10 | 40000 | 1e-4 | 1 |
| IMcQA-3b | 10 | 50000 | 1e-4 | 1 |

Table 5: Hyperparameter Settings.

B.3 Implementations

In this section, we describe the implementation of our experiments in detail, including the baseline methods, backbone models, and hyperparameters. Our model is built based on the T5 (Roberts et al., 2020a). Differing from fine-tuning all model parameters θ of the updated Pre-trained Language Model (LLM), LoRA (Hu et al., 2021) freezes all pre-trained Transformer parameters and optimizes only the parameters of each LoRA adapter. We employ LoRA to train a parameter-efficient fine-tuning baseline. Drawing from this, our ap-

| Models | # Documents | NQ | | | TQA | | | WQ | | |
|-----------|----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | | NQ | TQA | WQ | NQ | TQA | WQ | NQ | TQA | WQ |
| T5 | 0 | 22.16 | 3.18 | 4.12 | 2.65 | 21.8 | 3.15 | 0.88 | 2.95 | 28.3 |
| LoRA-Base | 0 | 16.17 | 4.71 | 6.89 | 3.15 | 21.16 | 0.00 | 1.33 | 3.04 | 26.38 |
| IMcQA | 0 | 23.89 | 6.21 | 10.94 | 5.31 | 22.69 | 6.30 | 3.23 | 5.10 | 30.31 |
| LoRA-Base | 1 [†] | 37.17 | 45.20 | 15.62 | 19.57 | 55.37 | 12.50 | 14.15 | 30.89 | 28.88 |
| IMcQA | 1 [†] | 40.14 | 46.61 | 18.92 | 24.78 | 60.75 | 12.82 | 17.70 | 35.24 | 41.06 |
| FiD | 10 | 46.81 | 53.93 | 24.02 | 28.57 | 63.32 | 17.83 | 18.81 | 41.88 | 41.78 |
| IMcQA | 10 | 47.01 | 55.74 | 24.13 | 31.77 | 64.95 | 19.52 | 24.43 | 48.10 | 46.36 |
| T5-l | 0 | 28.5* | 3.18 | 4.12 | 2.65 | 28.7* | 3.15 | 0.88 | 2.95 | 30.6* |
| LoRA-l | 0 | 17.70 | 7.49 | 8.66 | 3.54 | 23.87 | 4.72 | 0.00 | 5.65 | 29.13 |
| IMcQA-l | 0 | 29.32 | 10.17 | 14.06 | 7.02 | 30.11 | 7.81 | 2.65 | 7.06 | 32.68 |
| LoRA-l | 1 [†] | 37.61 | 48.50 | 20.71 | 20.54 | 62.71 | 14.81 | 15.36 | 33.83 | 39.37 |
| IMcQA-l | 1 [†] | 42.32 | 54.80 | 22.05 | 26.11 | 65.48 | 18.11 | 18.58 | 47.46 | 45.28 |
| FiD-l | 10 | 46.7* | 57.93 | 25.12 | 34.29 | 61.9* | 19.64 | 27.65 | 53.87 | 48.1* |
| IMcQA-l | 10 | 49.92 | 60.03 | 25.79 | 34.35 | 69.67 | 20.28 | 30.19 | 54.94 | 51.52 |

Table 6: **OOD results.** The primary row in the table header delineates the dataset trained, while the underscored secondary row demonstrates the in-distribution performance. IMcQA attains optimal performance both in-distribution and OOD under diverse document configurations.

proach updates only the parameters of the Hypernetwork to generate the weights for each LoRA adapter. This method is adopted based on LongLoRA’s (Chen et al., 2023) recommendations and experimental findings, demonstrating improved performance when the normalization and FFN layers components are updated. This is because: 1) dynamically generating LoRA weights enhances generalization and parameter sharing, and 2) LoRA performs comparably to fine-tuning but mitigates the risk of catastrophic forgetting.

For the baseline, most of the hyperparameters are the default parameters of FiD (Izacard and Grave, 2021). For LoRA (Hu et al., 2021), add the LoRA module only to the QV of the attention layers and also release the normalization and FFN layers.

We consider conducting experiments using three different sizes of T5, namely T5-base, T5-large, T5-3b, and Llama2-7B, Llama2-13B (Touvron et al., 2023). Due to memory constraints and online distillation limitations, A100 supports processing 20 documents for T5-3b, while Llama2 does not support distillation. All experiments with T5-3b are conducted on 2 A100 GPUs, T5-large on 2 A6000 GPUs, and T5-Base on 2 RTX 3090 GPUs. However, experiments with Llama2-7b and 13b, except for IMcQA on 2 A100 GPUs, are tested on 8 RTX 3090 GPUs.

B.3.1 Hyperparameters

The detailed hyperparameter setting is as shown in Table 5. For the LoRA modules, we set the α 32 and the *lora rank* 32.

B.3.2 Baselines

DPR (Karpukhin et al., 2020) generates by searching for the most relevant documents through dense vector space representation.

FiD (Izacard and Grave, 2021) retrieve relevant documents and send them separately to the Encoder, then fuse the information in the Decoder.

RFiD (Wang et al., 2023a) uses the encoder of FiD to distinguish between causal and incidental features, and guides the decoder to generate answers based on this distinction.

EAR (Chuang et al., 2023) significantly enhances the traditional sparse retrieval method BM25 by connecting query expansion models and retrievers.

FILCO (Wang et al., 2023d) identifies useful context based on lexical and information-theoretic methods.

GENREAD (Yu et al., 2023) prompt LLMs like InstructGPT (Ouyang et al., 2022) to generate a large number of relevant documents and let the reader process them.

LoRA We use LoRA (Hu et al., 2021) to obtain an efficiently fine-tuned baseline and compare it with our method.

B.3.3 Evaluation

For QA datasets, we choose the exact match (EM) score (Rajpurkar et al., 2016) as the evaluation metric. An answer is deemed correct if it aligns with any of the responses in the list of acceptable answers after normalization. Normalization involves transforming the text into lowercase, omitting articles, punctuation, and eliminating redundant spaces.

| Dataset | Train | Dev | Test |
|---------|--------|-------|--------|
| WebQ | 3,417 | 361 | 2,032 |
| NQ | 79,168 | 8,757 | 3,610 |
| TQA | 78,785 | 8,837 | 11,313 |

Table 7: Open-Domain QA dataset statistics

B.4 Downstream Evaluation Datasets

We use the following three Open-Domain QA for the experiments (§ 4.1).

- NaturalQuestions (Kwiatkowski et al., 2019) contains questions corresponding to Google search queries. The open-domain version of this dataset is obtained by discarding answers with more than 5 tokens, each accompanied by a Wikipedia article containing the answer.
- TriviaQA (Joshi et al., 2017) contains questions gathered from trivia and quiz-league websites. The unfiltered version of TriviaQA is used for open-domain question answering, each question is accompanied by pages from web and Wikipedia searches that may contain the answer.
- WebQuestions (Berant et al., 2013) contains questions from web queries matched to corresponding entries in FreeBase (Bollacker et al., 2008).

Table 7 presents detailed statistics of the dataset sizes, including the training, development, and test sets. We note that all our models are trained exclusively on the training data, and we did not include the development data in our training process. Therefore, the performance numbers reported in the paper for the dev and test data are independent of the training data.




Figure 6: Zero-Shot results (EM, %) of Llama2-13B on three open-domain QA datasets. The number in parentheses indicates the number of documents used.

C Full Experimental Results

C.1 Supervised Performance

As shown in Table 12, our initial observations indicate that regardless of the method implemented, supplying a certain quantity of related documents can expedite improvement and enhance performance in QA. FiD (Izacard and Grave, 2021) model outclasses all baseline models in performance. Notably, utilizing FiD-xl with a mere 10 documents yields performance on par with that attained through the use of FiD-l with 100 documents. Larger models not only encapsulate more knowledge but also demonstrate a superior ability to activate and apply this knowledge efficiently.

Additionally, in comparison with LoRA (Hu et al., 2021) methods, IMcQA enhances EM scores by an average of +2.2%. In the closed-book setting, the LoRA method manifests a substantial decrease in performance, likely attributable to the inadequacy of learning sufficient knowledge via questions for storage in the LoRA module. On the other hand, IMcQA harnesses both explicit and implicit imaginations to exploit knowledge for improved outcomes. These results indicate that the knowledge stored in the LLMs' parameters can still be further exploited.

C.2 OOD Results

Table 6 shows the full OOD results in QA. It can be observed that our method has the best OOD generalization ability on all three benchmarks. Although LoRA performs well on the in-distribution part, its performance is generally poor on OOD, with some even showing negative performance. This highlights the importance of the domain adaptability of the implicit imagination Hypernetwork in our method, which generates LoRA adapter weights




Figure 7: Zero-Shot results (Best_Subspan EM, %) of Llama2-7B on three open-domain QA datasets.




Figure 8: Zero-Shot results (Best_Subspan EM, %) of Llama2-13B on three open-domain QA datasets.

based on input.

C.3 Zero-Shot Results

LLMs have limited capacity to utilize extensive context effectively and are prone to generating illusions and redundant content. Best_subspan EM assesses whether the answer is included in the output. Previous studies have corroborated that LLMs encapsulate a considerable volume of knowledge and exhibit robust performance in QA.

Here, we report the Best_Subspan_EM values of Llama2-7B and Llama2-13B on three QA datasets. From Figure 7 and Figure 8, it can be observed that Best_Subspan_EM significantly improves, but the EM values are relatively small. This indicates that LLMs may not effectively utilize retrieval documents and are prone to outputting a lot of irrelevant information. Therefore, there is an urgent need to explore efficient techniques that leverage external information and internal knowledge.

However, the model did exhibit a weak adherence to instructions, often failing to output the exact answer. Remarkably, Llama2-13B displayed a decline in EM with increased document length on the WQ dataset, whereas the Best_Subspan_EM value augmented. Contrarily, our method excelled in extracting key information by using text imagination during the compression phase.

C.4 OOD and Ablation Experiment Results

Here, we supplement the experimental results of LoRA and IMcQA under supervised fine-tuning in closed-book settings and the ablation results of feedforward neural network (FFN) and Long Con-

| Model | NQ | TriviaQA | WebQ |
|---------------------------------|------|----------|------|
| LLaMA-2-7B | | | |
| Zero-shot | 8.6 | 14.5 | 2.6 |
| DPR + ICL | 18.3 | 32.5 | 15.6 |
| DPR + RECITE (Sun et al., 2023) | 16.8 | 43.9 | 24.8 |
| DPR + HICL (Wang et al., 2024) | 25.1 | 47.5 | 28.1 |
| DPR + ImcQA (Ours) | 33.7 | 44.5 | 31.9 |

Table 8: Zero-shot results of Llama2-7B

text Distillation (LCD). It can be observed that our method like LoRA, belongs to parameter-efficient fine-tuning, and because we share the Hypernetwork to generate LoRA adapter weights, we fine-tune fewer parameters.

From Table 9, it can be seen that releasing FFN can bring more performance improvement, possibly because adding LoRA in Attention cannot fully utilize enough knowledge (Yao et al., 2022). With the support of LCD, performance is further improved, with an average increase in EM of +5%. This also proves the effectiveness of our proposed LCD. In comparison with IMcQA and LoRA, it becomes more evident that LoRA tends to transfer knowledge to the LoRA module, resulting in low generalization. Our method enhances knowledge activation through dynamic generation, showing significant effects not only in-distribution but also in OOD.

C.5 Error Analysis

Using LLM as a knowledge base inevitably leads to hallucinations, which is a significant area of research in LLM development. In our quality analysis, we sampled 100 generated documents. As shown in Table 10, we found that hallucinations occurred with a probability of 4%, while the occur-

1258 rence of meaningless text, such as repeated values,
1259 was 6%. Consequently, the impact of hallucina-
1260 tions in our method is relatively minor.

| Models | # Documents | Trainable Params | NQ | | | TQA | | | WQ | | |
|-------------|-------------|------------------|-------|-------|-------|------|-------|------|------|-------|-------|
| | | | NQ | TQA | WQ | NQ | TQA | WQ | NQ | TQA | WQ |
| T5 | 0 | 220M | 22.16 | 3.18 | 4.12 | 2.65 | 21.8 | 3.15 | 0.88 | 2.95 | 28.3 |
| LoRA-Base | 0 | 28.3M | 5.43 | 3.15 | 4.02 | 0.00 | 9.60 | 0.00 | 0.22 | 1.77 | 20.47 |
| w FFN | 0 | 141.5M | 16.17 | 4.71 | 6.89 | 3.15 | 21.16 | 0.00 | 1.33 | 3.04 | 26.38 |
| w FFN & LCD | 0 | 141.5M | 21.37 | 2.82 | 6.89 | 1.99 | 17.94 | 3.74 | 0.00 | 2.82 | 32.50 |
| IMcQA | 0 | 26.1M | 5.31 | 3.82 | 5.71 | 0.22 | 10.34 | 2.12 | 0.55 | 2.30 | 16.58 |
| w FFN | 0 | 139.3M | 21.05 | 4.52 | 6.50 | 3.51 | 19.08 | 3.15 | 2.11 | 3.84 | 28.17 |
| w FFN & LCD | 0 | 141.5M | 23.89 | 6.21 | 10.94 | 5.31 | 22.69 | 6.30 | 3.23 | 5.10 | 30.31 |
| T5-l | 0 | 770M | 28.5* | 3.18 | 4.12 | 2.65 | 28.7* | 3.15 | 0.88 | 2.95 | 30.6* |
| LoRA-l | 0 | 42.5M | 4.42 | 6.50 | 7.87 | 3.98 | 10.03 | 3.94 | 1.99 | 6.71 | 18.11 |
| w FFN | 0 | 445.1M | 17.70 | 7.49 | 8.66 | 3.54 | 23.87 | 4.72 | 0.00 | 5.65 | 29.13 |
| w FFN & LCD | 0 | 445.1M | 28.32 | 4.52 | 10.94 | 5.31 | 25.71 | 6.12 | 1.75 | 4.52 | 29.92 |
| IMcQA-l | 0 | 34.8M | 7.08 | 8.90 | 9.45 | 4.42 | 13.14 | 8.66 | 2.43 | 10.17 | 17.72 |
| w FFN | 0 | 437.5M | 23.01 | 8.33 | 11.02 | 3.51 | 20.08 | 3.15 | 3.51 | 5.65 | 31.50 |
| w FFN & LCD | 0 | 437.5M | 29.32 | 10.17 | 14.06 | 7.02 | 30.11 | 7.81 | 2.65 | 7.06 | 32.68 |

Table 9: OOD and ablation experiment results. * denotes the results are from the existing papers and LCD denotes Long Context Distillation.

| Hallucinations | Meaningless |
|--|---|
| 4% | 6% |
| Question: When is the next Deadpool movie being released? Document: " Deadpool (film) Deadpool is a 2016 American superhero film based on the Marvel Comics character of the same name, produced by Marvel Studios and distributed by Walt Disney Studios Motion Pictures. | Question: Who got the first Nobel Prize in Physics? Document: The Nobel Prize is not a prize in itself. |
| Correct answer: May 18, 2018 | Correct answer: Wilhelm Conrad Röntgen |

Table 10: Hallucinations and Meaningless Analysis.

| Methods | Prompt |
|-----------|--|
| CBQA | Please write a high-quality answer for the given question using your knowledge. Only give me the answer and do not output any other words. Question: {question} Answer: |
| | Please write a high-quality answer for the given question using only the provided search results (some of which might be irrelevant). Only give me the answer and do not output any other words. Context: {context} Answer the question based on the given passages. Question: {question} Answer: |
| Retrieval | Please write a high-quality answer for the given question using your knowledge and the provided imagined compressed results (some of which might be irrelevant). Only give me the answer and do not output any other words. Imagined Context: {context} Answer the question based on your knowledge and the given imagined context. Question: {question} Answer: |

Table 11: Prompts for different methods on Zero-Shot setting. **CBQA** denotes closed-book QA that just prompts the model with the question.

| Models | Reader Params | # Documents | NQ | TriviaQA | WebQ |
|--|---------------|-----------------|--------------|--------------|--------------|
| <i># Closed-book Setting</i> | | | | | |
| T5* (Roberts et al., 2020a) | 220M | 0 | 25.9 | 23.8 | 27.9 |
| T5-l* (Roberts et al., 2020a) | 770M | 0 | 28.5 | 28.7 | 30.6 |
| T5-xl (Roberts et al., 2020a) | 3b | 0 | 28.30 | 33.92 | 34.43 |
| LoRA-Base | 220M | 0 | 16.17 | 21.16 | 26.38 |
| LoRA-l | 770M | 0 | 17.70 | 23.87 | 29.13 |
| LoRA-xl | 3b | 0 | 23.15 | 32.16 | 35.24 |
| IMcQA (Ours) | 220M | 0 | 23.89 | 22.69 | 30.31 |
| IMcQA-l (Ours) | 770M | 0 | 29.32 | 30.11 | 32.68 |
| IMcQA-xl (Ours) | 3b | 0 | 29.59 | 35.71 | 37.40 |
| <i># Retrieval Augmented Generation</i> | | | | | |
| DPR* (Karpukhin et al., 2020) | 110M | 100 | 41.5 | 56.8 | 41.1 |
| RAG* (Lewis et al., 2020) | 400M | 10 | 44.5 | 56.1 | 45.2 |
| FiD* (Izacard and Grave, 2021) | 220M | 100 | 48.2 | 65.0 | 46.71 |
| FiD-l* (Izacard and Grave, 2021) | 770M | 100 | 51.4 | 67.6 | 50.52 |
| FiD-xl (Izacard and Grave, 2021) | 3b | 20 | 55.18 | 72.92 | 52.85 |
| FiD-l* (Izacard and Grave, 2021) | 770M | 10 | 46.7 | 61.9 | 48.1 |
| FiD-xl* (Izacard and Grave, 2021) | 3b | 10 | 50.1 | 66.3 | 50.8 |
| EAR-l (Chuang et al., 2023) | 770M | 10 | 39.6 | 60.0 | - |
| EAR-xl* (Chuang et al., 2023) | 3b | 10 | 42.3 | 64.6 | - |
| RFiD-l (Wang et al., 2023a) | 770M | 10 | 48.3 | 63.4 | - |
| RFiD-xl (Wang et al., 2023a) | 3b | 10 | 50.5 | 67.8 | - |
| FILCO-xl* (Wang et al., 2023d) | 3b | 1 | 44.7 | 59.0 | - |
| IMcQA (Ours) | 220M | 10 | 47.01 | 64.95 | 46.36 |
| IMcQA-l (Ours) | 770M | 10 | 49.92 | 69.67 | 51.52 |
| IMcQA-xl (Ours) | 3b | 5 [‡] | 50.87 | 70.34 | <u>52.78</u> |
| IMcQA-l (Ours) | 770M | 30 | <u>53.1</u> | <u>70.5</u> | 52.0 |
| <i># Generation Augmented Generation</i> | | | | | |
| GENREAD-l (sampling)* (Yu et al., 2023) | 770M | 10 [†] | 40.3 | 67.8 | 51.5 |
| GENREAD-l (clustering)* (Yu et al., 2023) | 770M | 10 [†] | 43.5 | 70.2 | 53.5 |
| GENREAD-xl (sampling)* (Yu et al., 2023) | 3b | 10 [†] | 42.6 | 69.6 | 52.6 |
| GENREAD-xl (clustering)* (Yu et al., 2023) | 3b | 10 [†] | 45.6 | <u>71.6</u> | 54.4 |
| IMcQA (Ours) | 220M | 10 [†] | 46.22 | 66.70 | 51.43 |
| IMcQA-l (Ours) | 770M | 10 [†] | <u>48.83</u> | 70.85 | <u>54.52</u> |
| IMcQA-xl (Ours) | 3b | 5 ^{†‡} | 49.23 | 72.18 | 55.39 |
| <i># Imagination Augmented Generation (Ours)</i> | | | | | |
| LoRA-Base | 220M | 1 [†] | 34.51 | 54.05 | 32.28 |
| LoRA-l | 770M | 1 [†] | 40.05 | 62.81 | 43.70 |
| LoRA-xl | 3b | 1 [†] | 44.15 | 66.92 | 48.23 |
| IMcQA | 220M | 1 [†] | 40.14 | 60.75 | 41.73 |
| IMcQA-l | 770M | 1 [†] | 42.32 | 65.48 | 45.28 |
| IMcQA-xl | 3b | 1 [†] | 46.51 | 68.38 | 50.45 |

Table 12: Full QA performances (%) of different methods on three datasets. The first part (closed-book setting) indicates that explicit documentation was not utilized; The latter three parts utilize explicit augmented documents. The best results are in bold. * means that those results are from existing papers, † denotes that the number of documents is generated (‡ indicates that the number of documents is reduced due to insufficient memory for distillation).