

Towards Better Generalization in Open-Domain Question Answering by Mitigating Context Memorization

Zixuan Zhang¹, Revanth Gangi Reddy¹, Kevin Small², Tong Zhang¹, Heng Ji²

¹University of Illinois Urbana-Champaign, ²Amazon,
{zixuan11, revanth3, tozhang}@illinois.edu,
{smakevin, jihj}@amazon.com

Abstract

Open-domain Question Answering (OpenQA) aims at answering factual questions with an external large-scale knowledge corpus. However, real-world knowledge is not static; it updates and evolves continually. Such a dynamic characteristic of knowledge poses a vital challenge for these models, as the trained models need to constantly adapt to the latest information to make sure that the answers remain accurate. In addition, it is still unclear how well an OpenQA model can transfer to completely new knowledge domains. In this paper, we investigate the generalization performance of a retrieval-augmented QA model in two specific scenarios: 1) adapting to updated versions of the same knowledge corpus; 2) switching to completely different knowledge domains. We observe that the generalization challenges of OpenQA models stem from the reader’s over-reliance on memorizing the knowledge from the external corpus, which hinders the model from generalizing to a new knowledge corpus. We introduce *Corpus-Invariant Tuning* (CIT), a simple but effective training strategy, to mitigate the knowledge over-memorization by controlling the likelihood of retrieved contexts during training. Extensive experimental results on multiple OpenQA benchmarks show that CIT achieves significantly better generalizability without compromising the model’s performance in its original corpus and domain.¹

1 Introduction

Open-domain Question Answering (OpenQA) (Chen and Yih, 2020) aims at answering factual questions using a large-scale external knowledge corpus. This is in contrast to closed-book question answering (Roberts et al., 2020) wherein the model is expected to directly answer questions with no access to external knowledge. In general, closed-

book QA optimizes for memorization of knowledge in model parameters, while OpenQA focuses on retrieving relevant knowledge from an external corpus. OpenQA typically employs a retrieval-augmented approach (Karpukhin et al., 2020; Izacard and Grave, 2021a; Sachan et al., 2021), involving a two-stage process: a *retriever* to select relevant documents, followed by a *reader* to derive answers from these documents. It is more practical for real-world applications as it enables the use of extensive and varied knowledge sources for answering questions.

Retrieval-augmented OpenQA models rely on an external corpus to physically store the knowledge. However, real-world knowledge is not static; it updates and evolves continually. Therefore, it is essential to build models that are able to use fresh and real-time knowledge (Kasai et al., 2022; Vu et al., 2023), but the dynamic characteristic of knowledge poses a vital challenge as the trained models need to constantly adapt to the latest information to make sure that the answers remain relevant and accurate. In addition, closed-book QA systems have been proved limited in adapting to new information or domains due to their reliance on pre-existing knowledge, and updating their parametric knowledge requires extensive large-scale pre-training. Nevertheless, it is still unclear how well OpenQA systems can transfer to leveraging unseen corpora and domains during training.

In this paper, we first investigate how well state-of-the-art retrieval-augmented models, such as Atlas (Izacard et al., 2023), can adapt to new and diverse knowledge corpora. Specifically, we explore the model’s performance in two scenarios: 1) adapting to updated versions of the same corpus (in §2.2.1); 2) switching to completely different domains (in §2.2.2). Our investigation involves three settings: directly applying a pre-trained model, fine-tuning the model with the new corpora, and training the model afresh on the new corpora. Initial experi-

¹Programs and data are available for research purpose at <https://github.com/zhangzx-uiuc/CIT>

ments reveal that the model faces challenges in both scenarios. When directly transitioning to an updated corpus, there is a noticeable performance decline. Even additional tuning on the newer version doesn’t achieve the same effectiveness as training from scratch with the new data (56.9→59.5 vs 62.2 in Table 1). Similar outcomes are observed when shifting from a general domain, like Wikipedia, to a specialized one, such as biomedical (41.2→68.8 vs 69.7 in Table 2).

We hypothesize and validate that such generalization challenges stem from the reader’s over-reliance on memorizing the knowledge retrieved from the external corpus. This reliance primarily arises as the reader, with its primary training objective optimized for QA accuracy, often opts to hard-code a substantial amount of retrieved knowledge into its parametric memory. Such kind of over-memorization reduces the reader’s dependency on the retriever to choose more relevant contexts. This phenomenon hampers the model’s generalizability, particularly to updates in the knowledge corpus or changes in the knowledge domain. For instance, given a question *Who is the prime minister of the UK?*, if a model has already hard-coded an outdated answer *Boris Johnson* into its parameters (while being trained on an old corpus), it is harder to change its response even if the new information *Rishi Sunak* from an updated corpus is available.

To address this issue, we introduce *Corpus-Invariant Tuning* (CIT), a simple but effective training strategy to improve the corpus generalizability of retrieval-augmented text generation models. CIT aims to mitigate the reader’s tendency to memorize the documents retrieved from the corpus during training. This pushes the reader to rely more on retrieved documents to answer the input questions, rather than relying on memorizing the knowledge facts into its parameters. To achieve this, we propose a novel loss term to prevent memorization during training by controlling the likelihood of the retrieved documents. Through extensive experiments across various OpenQA benchmarks (Kwiatkowski et al., 2019; Joshi et al., 2017; Han et al., 2023), carried out in both zero-shot and continual fine-tuning scenarios, we demonstrate that a retrieval-augmented model trained using our proposed CIT loss exhibits considerably enhanced generalizability across different corpora. This is evident by the considerable improvements in exact match (EM) scores, reaching up to a 2.1% absolute gain.

Our contributions can be summarized as follows:

- We propose to mitigate knowledge over-memorization of the reader to improve the generalization ability of retrieval-augmented text generation models.
- We introduce *Corpus-Invariant Tuning* (CIT), a straightforward but effective training strategy that regularizes the reader’s likelihood of the retrieved documents to mitigate it from over-memorizing the corpus during training.
- Through extensive experiments on multiple benchmarks, we demonstrate that training models with CIT greatly improves the generalization of OpenQA models across both newer versions of the corpora and unseen domains.

2 Preliminaries

2.1 Problem Formulation

Open-domain QA aims to answer questions only using a large-scale unified corpus, where the background documents for each question is not specified in advance. Given a natural language question x , our objective is to build a model $f(\cdot)$ to predict an answer \hat{y} using a unified list of background documents Z , where $\hat{y} = f(x, Z)$. Such a setting is more practical for real-world applications because it mirrors the vast and unstructured nature of real-world knowledge.

Retrieval Augmentation Since the external corpus collectively stores all essential information for answering the questions, the typical strategy to tackle the OpenQA problem is to implement a retrieval-augmented approach with a two-stage framework: 1) a *retriever* to select a small subset of documents that are most relevant to the current question, and 2) a *reader* to seek for useful information from the retrieved documents and generate the answer. Specifically, the probability of a predicted answer \hat{y} is decomposed by

$$p(\hat{y} | x, Z) = \sum_{\mathcal{C} \subset Z} p(\mathcal{C} | x; \theta) \cdot p(\hat{y} | \mathcal{C}, x; \phi),$$

where \mathcal{C} denotes the set of retrieved documents, and θ and ϕ are the parameters of the retriever and the reader respectively. During training, the retriever (θ) and the reader (ϕ) are often jointly optimized to ensure their effective collaboration, where the optimization is typically conducted with iterative training (Izacard and Grave, 2021a) or Expectation-Maximization (EM) based approach to

train the model by treating the retrieved documents as hidden variables (Sachan et al., 2021).

2.2 Evaluation of Model Generalization

We aim to tackle the generalization challenge for retrieval-augmented models as discussed in Section 1. Specifically, we focus on the following two main research questions (RQs):

- *RQ1: How to improve the model’s generalization ability across different versions (temporal snapshots) of the same corpus?*
- *RQ2: How to improve the model’s generalization ability across the corpora in different domains?*

2.2.1 Evaluations of RQ1

We conduct proof-of-concept experiments to test whether current retrieval-augmented OpenQA models can remain effective when the external corpus is updated to a newer version. Specifically, we adopt the most recent retrieval-augmented model Atlas-XL (Izacard et al., 2023) and test it on the Natural Questions (NQ) benchmark² with two different versions of Wikipedia (Wiki-2017 and Wiki-2018)³ as the external corpus. We first fine-tune the Atlas-XL model on each version of Wikipedia, and then evaluate the model’s generalization ability by both *zero-shot testing* (train the model with Wiki-2017 and directly test it with Wiki-2018) and *continue fine-tuning* (train the model with Wiki-2017 and further fine-tune it with Wiki-2018). As shown in Table 1, we can first observe that the model performs better when initially fine-tuned with Wiki-2018, which shows that the updated KB can improve the performance.⁴ However, we can also observe a significant performance degradation when using the model trained with Wiki-2017 to directly test it on Wiki-2018. Despite subsequent fine-tuning efforts, the performance still falls short of the original results obtained from initially training and testing with Wiki-2018. These results indicate that the current retrieval-augmented models still struggle to effectively generalize when the background corpus undergoes evolution or updates.

²<https://ai.google.com/research/NaturalQuestions>

³The Wikipedia dumped in 2017 and 2018 respectively.

⁴The NQ benchmark is annotated in 2018, so Wiki-2018 is a more up-to-date background KB for the task.

Training Corpus	Testing Corpus	EM Score
Wiki-2017	Wiki-2018	56.9
Wiki-2017 → Wiki-2018	Wiki-2018	59.5
Wiki-2018	Wiki-2018	62.2

Table 1: Initial experiments with Atlas-XL on the NQ benchmark with different Wikipedia versions. Results are evaluated with the exact-match (EM) score (%).

2.2.2 Evaluations of RQ2

We conduct similar experiments with Atlas-XL to evaluate its generalization ability across different domains. We train the model on NQ with Wiki-2018 in the general domain, and test it on the *Biomedical* split in RobustQA⁵ with PubMed in the biomedical domain. The results presented in Table 2 reveal similar performance declines in both zero-shot and continual-fine-tuning settings, which indicates that the current OpenQA models also have inherent difficulty in generalizing across different domains.

Training Corpus	Testing Corpus	EM Score
Wiki2018(NQ)	PubMed(Bio)	41.2
Wiki2018(NQ)→PubMed(Bio)	PubMed(Bio)	68.8
PubMed(Bio)	PubMed(Bio)	69.7

Table 2: Initial results (%) with Atlas-XL to test its generalizability between NQ and BioASQ.

3 Corpus-Invariant Tuning

Motivated by the observed limitations in the generalization capabilities of retrieval-augmented models, we introduce *Corpus-Invariant Tuning* (CIT) to mitigate memorizing the lexical content of retrieved documents. Specifically, we posit that the generalization difficulties encountered by retrieval-augmented text generation models arises via excessive memorization the documents retrieved from the external corpus by the reader. In order to achieve higher question-answering accuracy during training, the reader tends to “hard-code” a large volume of retrieved documents rather than relying on an improved retriever for a better selection of relevant contexts, as is empirically validated in the document retrieval evaluations of Section 4.5. This limits the model’s ability to generalize because excessive memorization of documents by the reader

⁵<https://github.com/rujunhan/RobustQA-data>

dictates that when the external corpus is updated or transitions to a different domain, the model faces increased difficulty in adapting and correcting its knowledge compared to learning from scratch.⁶

Validation Here we provide an empirical validation on our hypothesis that the degradation of model generalization ability is caused by over-memorization of retrieved knowledge. We replace the retrieved contexts with ground-truth retrieval results on Wiki-2018, and conduct a stand-alone evaluation with the reader. We report both the EM score and the overlap rate, i.e., the percentage of incorrectly predicted answers that have overlaps with the ground-truth retrieval results in Wiki-2017. The results are shown in Table 3. We can observe that while the models transferred from Wiki-2017 perform slightly worse in terms of EM score, it has a lot more error cases that overlap with the retrieved documents on Wiki-2017. Such results directly show that the over-memorization of contexts is the primary cause of the degradation of model generalizability.

Dataset	Training Corpus	EM Score	Overlap Rate
NQ	Wiki-2017→2018	63.6	76.3
	Wiki-2018	65.2	30.2
TriviaQA	Wiki-2017→2018	78.1	80.8
	Wiki-2018	78.7	41.0

Table 3: Stand-alone evaluation results (%) with ground-truth retrieved documents on Wiki-2018. We report both the EM score, and the overlap rate between the incorrectly predicted answers with the retrieval results from the old corpus (Wiki-2017).

Corpus-Invariant Tuning (CIT) To solve this problem, we propose *Corpus-Invariant Tuning* (CIT), a straightforward but effective method to temper the reader’s tendency to over-memorize the contents of externally retrieved documents, thereby improving the model’s generalization abilities for downstream tasks like OpenQA. As depicted in Figure 1, the core idea of CIT is to control the reader’s memorization (likelihood) of the corpus to be “invariant” by introducing an additional loss term that ensures the reader’s likelihood of the retrieved documents does not increase during training. Specifically, for each training QA pair (x, y) and its retrieved document set \mathcal{C} , the loss term can be

⁶Such a phenomenon can be caused by the exposure bias problem as discussed in (Yu and Ji, 2023).

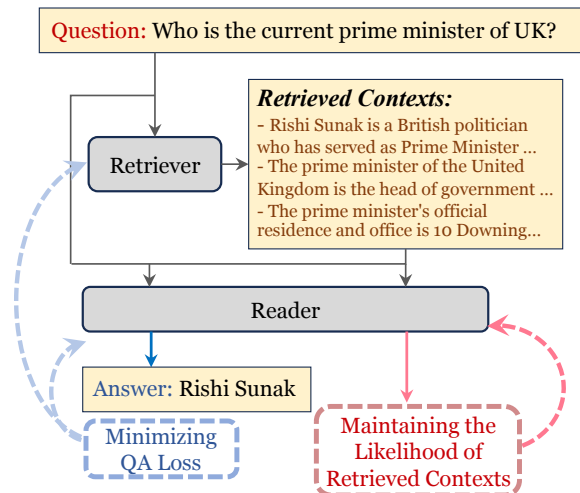


Figure 1: Our proposed Corpus-Invariant Tuning (CIT) Framework. In addition to the existing loss for question answering, we introduce an auxiliary CIT loss to make sure that the reader does not over memorize the retrieved contexts. Specifically, given each batch of QA pairs and the relevant documents retrieved from the corpus, the CIT loss makes sure that the reader’s likelihood of these documents does not increase.

written as

$$\mathcal{L}_{CIT} = \sum_{c \in \mathcal{C}} \|\log p_{\phi}(c) - \log p_{\phi_0}(c)\|^2, \quad (1)$$

where ϕ and ϕ_0 denote the current parameters and the original parameters⁷ of the reader respectively. We use $p_{\phi}(c)$ to represent the reader’s likelihood of a retrieved document c . In our experiments, we adopt the Masked Span Prediction (MSP) probability from the T5 model (Raffel et al., 2020) to maintain consistency with the Atlas architecture. Essentially, we randomly mask out a fixed number of spans of the input sentence and use the model’s probability of generating these spans in the correct order as the likelihood. The overall training objective is a combination of the original loss for question answering \mathcal{L}_{QA} and the CIT loss \mathcal{L}_{CIT} :

$$\mathcal{L} = \mathcal{L}_{QA} + \alpha \cdot \mathcal{L}_{CIT}, \quad (2)$$

where α is a configurable hyper-parameter.

Discussion Retrieval-augmented QA models typically maximize answer accuracy as the end-to-end training objective. However, given the distinct roles of retrievers and readers, there exist two distinct approaches through which this goal can be achieved: the model can either choose to enhance the retriever

⁷ ϕ_0 denotes the initial reader’s parameters before training.

to fetch more relevant documents or simply allow the reader to memorize pertinent knowledge. While both methods can contribute to performance improvements, the former approach increases generalization and CIT biases the model away from rigid memorization by the reader.

4 Experiments

4.1 Data

Our experiments are conducted on two general-domain OpenQA datasets, NaturalQuestion (NQ) (Kwiatkowski et al., 2019) and TriviaQA (Joshi et al., 2017), and a large cross-domain benchmark RobustQA (Han et al., 2023). Detailed statistics of these datasets are depicted in Table 4.

NQ and TriviaQA NQ and TriviaQA are the two most widely-used open-domain QA benchmarks, which contain factual question and answer pairs created and annotated based on Wikipedia. There are 79,168 and 78,785 QA pairs for training in NQ and TriviaQA respectively. In our experiments, we conduct both fully supervised training and few-shot training settings for model evaluation.

RobustQA RobustQA⁸ is a large-scale OpenQA evaluation benchmarks specifically designed for evaluating the cross-domain generalization capabilities of OpenQA models. RobustQA includes 8 distinct domains, each equipped with its own test set and a corresponding list of background documents. The QA pairs and the documents are adopted and annotated from FiQA,⁹ SearchQA (Dunn et al., 2017), BioASQ (Tsatsaronis et al., 2015), and LOTTE (Santhanam et al., 2022).

Benchmark	Domain	# Test Questions	Corpus Size
NQ	Wikipedia	3,610	-
TriviaQA	Wikipedia	11,313	-
RobustQA	Web Search	31,760	13,791,373
	Biomedical	1,956	15,559,026
	Finance	3,669	57,638
	Lifestyle	2,214	119,461
	Recreation	2,096	166,975
	Technology	2,115	638,509
	Science	1,426	1,694,164
	Writing	2,696	199,994

Table 4: Detailed statistics of OpenQA evaluation benchmarks used in our experiments.

⁸<https://github.com/aws-labs/robustqa-ac123>

⁹<https://sites.google.com/view/fiqa/home>

4.2 Baselines

We adopt the state-of-the-art retrieval-augmented language model Atlas-XL (Izacard et al., 2023) as our main baseline, which uses a Contriever (Izacard et al., 2022) as the retriever, and a Fusion-in-Decoder (FiD) model (Izacard et al., 2020) as the reader. The primary objective of our experiments is to evaluate whether the baseline model demonstrates improved performance when trained using our proposed CIT loss. Besides, we also introduce other most recent models Flan-T5 (Chung et al., 2022), RGF (Paranjape et al., 2022), ReAtt (Jiang et al., 2022), FiE+PAQ (Kedia et al., 2022), and FiD-KD (Izacard and Grave, 2021b) for comparison. Our model is labeled as *Atlas-XL+CIT* which applies an additional CIT loss to control the knowledge over-memorization of the reader.

4.3 RQ1: Different versions of the corpus.

Corresponding to the two research questions proposed in previous sections, we first focus on evaluating our model’s ability to generalize across different versions of the external corpus. We adopt the Wikipedia-domain benchmarks NQ and TriviaQA in our experiments, and test their cross-corpus generalization abilities on different Wikipedia versions. Similar to the setting in our preliminary experiments presented in Section 2.2, we use Wiki-2017 and Wiki-2018 as our background corpora, and we consider both zero-shot and fully-supervised settings to test the generalization ability of a trained model. Specifically, we use the following terms to denote different experiment settings:

- *Closed Book*: The model is trained and tested without a retriever. The reader is responsible to understand questions and provide answers.
- *Original*: Also known as the *Open Book* setting. The most typical experiment setting for retrieval-augmented models, where a retriever retrieves a set of documents from the external corpus, and the reader uses these documents to generate answer. We use the label *Original* to emphasize the absence of cross-corpus generalization in this setting, providing a baseline for comparison with the following settings to evaluate the model’s generalization ability.
- *Zero-shot Transfer*: The model is initially trained with the older version of a knowledge corpus, and directly tested with the updated

<i>Model</i>	<i>Setting</i>	<i>Training Corpus</i>	<i>Testing Corpus</i>	<i>NQ</i>	<i>TriviaQA</i>
Atlas-XL (Izcard et al., 2023)	<i>Closed Book</i>	-	-	30.2	41.6
FiD-KD (Izcard and Grave, 2021b)	<i>Original</i>	Wiki-2018	Wiki-2018	54.7	67.6
ReAtt (Jiang et al., 2022)	<i>Original</i>	Wiki-2018	Wiki-2018	54.7	-
FiE+PAQ (Kedia et al., 2022)	<i>Original</i>	Wiki-2018	Wiki-2018	58.4	72.6
Atlas-XL	<i>Original</i>	Wiki-2017	Wiki-2017	58.8	75.5
Atlas-XL + CIT		Wiki-2017	Wiki-2017	58.9	75.5
Atlas-XL	<i>Zero-shot</i>	Wiki-2017	Wiki-2018	56.9	75.1
Atlas-XL + CIT	<i>Transfer</i>	Wiki-2017	Wiki-2018	58.6	75.5
Atlas-XL	<i>Full-training Transfer</i>	Wiki-2017→Wiki-2018	Wiki-2018	59.5	76.8
Atlas-XL + CIT		Wiki-2017→Wiki-2018	Wiki-2018	61.6	77.4

Table 5: OpenQA results on the NQ and TriviaQA benchmarks, evaluated with the exact match (EM) score.

corpus version in a zero-shot manner without any additional fine-tuning.

- *Full-training Transfer*: As opposite to the zero-shot transfer setting, after being initially trained with the older version of a knowledge corpus, the model is further fine-tuned on the same training QA pairs with an updated version of the corpus, before being tested with the new corpus.

Main Results We conduct generalization experiments in both zero-shot and full-training settings on the NQ and the TriviaQA datasets, and the results are shown in Table 5. In the *Original* setting to train and test the model on the same corpus (Wiki-2017), we can observe that compared with the baseline model Atlas-XL, incorporating a CIT loss to reduce knowledge over-memorization will not diminish the task performance; in fact, it can even slightly enhance it in certain cases (NQ with Wiki-2017). This is probably because when the reader is discouraged from rigid knowledge memorization, the retriever still has enough room of improvement to retrieve better documents and enhance the performance. Besides, in both of the *Zero-shot Transfer* and *Full-training Transfer* settings, we can observe that the CIT loss can significantly improve the generalization performance of the model across different versions of a knowledge corpus. This is likely because when the reader is discouraged from hard-coding knowledge into its parameters, it becomes more receptive to assimilating and utilizing new information from the retriever. In summary, our proposed CIT loss significantly improves the model’s ability to generalize across different versions of external corpus, without compromising the absolute task performance for OpenQA.

4.4 RQ2: Different domains.

To address our second research question, in this section, we conduct experiments to evaluate whether our proposed CIT loss can help model better generalize across different knowledge domains. We conduct evaluations using the RobustQA benchmark, which encompasses eight diverse domains specifically tailored for OpenQA. We first assess the model’s ability to generalize from a general domain (Wikipedia) to these eight diverse domains. Subsequently, we evaluate the model’s effectiveness in generalizing interchangeably across these eight domains.

From Wikipedia to Specific Domains We first evaluate how well a model trained with Wikipedia can generalize on the eight specific domains in RobustQA. Specifically, all the models are first fine-tuned on the NQ dataset with Wiki-2018, and then tested with the eight domain-specific benchmarks. The results are presented in Table 6. In general, adding the CIT loss to an Atlas-XL model significantly improves the average F1 score across eight domains in RobustQA, creating a new state-of-the-art among all 3B(XL)-sized models. Our proposed CIT loss can also help on smaller sizes of models, like Atlas-base, with 2.2% absolute improvement of the F1 score. Within these domains, we can see that domains like *Life* show significant improvement. This is probably due to their larger overlaps with Wikipedia, which makes it more crucial to avoid over-memorization, so that the old Wikipedia knowledge will not affect generalization on the new domain. In contrast, domains like *Biomedical* exhibits less improvements. This is possibly because a *Biomedical* domain KB has a smaller overlap with Wikipedia, thereby reducing the neg-

<i>Method</i>	<i>Average</i>	<i>Biomedical</i>	<i>Search</i>	<i>Finance</i>	<i>Life</i>	<i>Recreation</i>	<i>Technology</i>	<i>Science</i>	<i>Writing</i>
RGF	23.5	33.8	49.0	13.2	20.2	19.1	17.1	15.5	20.3
Flan-T5-XL	32.1	43.1	70.9	14.6	25.5	25.4	21.3	23.9	32.1
Atlas-base	28.3	40.0	59.2	15.6	23.8	22.8	19.8	18.3	27.3
Atlas-base + CIT	30.5	40.3	60.9	15.5	26.7	24.9	19.9	23.6	32.3
Atlas-XL	33.2	41.2	61.0	19.9	32.0	27.9	22.2	24.8	36.7
Atlas-XL + CIT	35.4	43.7	71.5	20.1	33.8	28.1	22.2	26.9	37.0

Table 6: The evaluation results on RobustQA of how well the model can generalize across different domains. All the models are initially fine-tuned on the NQ dataset, and then directly tested on the 8 different domain-specific benchmarks in RobustQA. The performances are evaluated by F1 score to be consistent with the RobustQA paper.

ative impact of knowledge over-memorization on the model’s generalizability.

Cross-Domain Generalization In addition to testing the model trained from Wikipedia, we also assess the model’s capability to generalize between each pair of domains in RobustQA. We define the *cross-domain relative performance (CRP)* as the evaluation metric to intuitively characterize how well the model can generalize from a source domain s to a target domain t . Specifically, the $CRP(s, t)$ is defined as the ratio of cross-domain performance and intra-domain performance:

$$CRP(s, t) = \frac{score(s, t)}{score(t, t)} \quad (3)$$

where $score(s, t)$ is the performance (F1-score) of training the model in the source domain s and testing the model in the target domain t . In Figure 2, we set different CIT strength α , and visualize the model’s generalizability into a heatmap. Each heat value $h(s, t)$ stands for the absolute improvements of CRP over the baseline model Atlas-XL:

$$h(s, t) = CRP_{\alpha}(s, t) - CRP_{Atlas-XL}(s, t),$$

where darker green indicates larger improvements and darker red indicates larger declines. We can already observe improvements for most domain pairs with $\alpha = 0.1$, and while α reaches to 0.3, the improvements become much more significant and all domain pairs benefit from CIT in terms of cross-domain generalization.

4.5 Effect on Retrieval Performance

The proposed CIT training loss reduces the reader model’s memorization tendency, leading to greater reliance on the documents retrieved. This enhanced dependency during training on the retrieved documents appears to enhance the retriever’s performance, as seen from the improvements in retrieval

<i>Model</i>	<i>R@10</i>	<i>R@20</i>	<i>R@40</i>
Atlas-XL	79.7	84.3	88.4
Atlas-XL + CIT	85.2	88.9	91.5

Table 7: Retrieval performance (Recall@K) on the NQ benchmark.

performance observed in Table 7 upon integrating CIT loss. Additionally, we measured the coverage of the reader’s predicted answer in the retrieved documents for the NQ benchmark, noting an increase in coverage within the top 40 documents from 66.9% to 69.1%. This suggests that our proposed CIT training loss leads to an increased reliance by the reader on retrieved documents rather than corpus memorization.

4.6 Parameter Sensitivity

We then conduct a more in-depth study on the model sensitivity of the most important hyperparameter α , which controls the strength of the corpus-invariant tuning (as shown in Equation 2). By choosing different values of α , we compute the average cross-domain relative performance between both different corpus versions (RQ1) and different domains (RQ2), and the results are shown in Figure 3. We can observe that as α becomes larger, the generalization performance initially improves, and then starts to decline after reaching its peak. This trend is likely because if memorization is controlled too excessively, the reader may neglect memorizing some certain shared global knowledge that is actually beneficial for knowledge generalization. From Figure 3 we can observe that $\alpha = 0.2$ is best for generalization across different corpus versions (RQ1), and $\alpha = 0.3$ is the best for cross-domain setting (RQ2).




Figure 2: The result heatmaps for cross-domain generalization experiments. Each value in the heatmap represents the absolute improvement (compared with Atlas-XL) of cross-domain relative performance (CRP) defined in Equation 3. Darker green indicates larger improvements in cross-domain generalization.

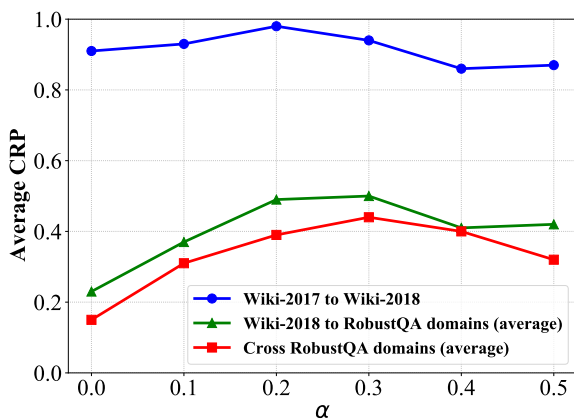


Figure 3: Parameter sensitivity on choices of α .

5 Related Work

Retrieval-Augmented Text Generation Enhancing text generation through the use of retrieved contexts has proven effective in a variety of knowledge-intensive downstream tasks. The most typical design for retrieval augmentation is to employ a retriever, which is jointly optimized with the reader in an end-to-end manner (Guu et al., 2020; Izacard et al., 2020; Karpukhin et al., 2020; Izacard and Grave, 2021b; Kedia et al., 2022; Jiang et al., 2022). The training methods include iterative training (Izacard and Grave, 2021a), and also EM-based algorithms to treat retrieved documents as hidden variables (Sachan et al., 2021). Retrieval augmentation can also act as plug-in modules (Yu et al., 2023). In this setting, the retrievers are not jointly trained with the reader, and the retrieved documents are used in an in-context manner (Ram et al., 2023; Huang et al., 2023). Recently, Atlas (Izacard et al., 2023) trains and releases a new set of retrieval-augmented models based on the T5 architectures (Raffel et al., 2020), which achieves

state-of-the-art performances on OpenQA tasks in few-shot settings.

Open-Domain Question Answering The task of OpenQA (Joshi et al., 2017; Kwiatkowski et al., 2019; Chen and Yih, 2020) aims at building models to answer questions without background documents. Given the high demand for external knowledge in this task, the standard approach involves integrating external corpora with a knowledge retriever to supply supporting evidence for answering questions (Izacard et al., 2020; Zhou et al., 2020). Recently, researchers also focus on new problem settings such as conversational QA (Fang et al., 2022), multi-hop QA (Xiong et al., 2021), and new knowledge-enhanced solutions like using knowledge graphs (Oduro-Afriyie and Jamil, 2023), and multi-hop reasoning (Hu et al., 2022).

Knowledge Generalization Because retrieval-augmented OpenQA models require an external KB to provide supporting documents, it is important to make sure that the model is robust and generalizable across different versions and domains of knowledge. There are a few previous studies that focus on the generalizability of OpenQA models. For example, Liu et al. (2022) focus on the question generalizability of OpenQA models, and conducts a detailed analysis on the generalization performances of current OpenQA models. Gangi Reddy et al. (2021); Reddy et al. (2022b,a) propose training with synthetic data to improve the robustness of retrieval models in OpenQA settings. Siriwardhana et al. (2023) focuses on the domain adaption problem of OpenQA, and proposes a reconstruction-based auxiliary loss to improve the model’s generalizability. Regarding dataset development, RobustQA (Han et al., 2023) creates a new benchmark

that involves multiple real-world domains. However, there are no previous studies that aim to tackle the model’s generalization ability on both different corpus and different knowledge domains. Also, we are the first to tackle this problem in a knowledge memorization perspective, enhancing the model’s generalization ability by reducing rigid memorization of the reader modules.

6 Conclusion

In this paper, we present *Corpus-Invariant Tuning* (CIT), a simple but effective training strategy to improve the generalization ability of a retrieval-augmented text generation model across different corpus versions and different knowledge domains. The main idea of CIT is to mitigate rigid knowledge memorization during training, so that the reader module can easily accept new knowledge in the retrieved documents and adapt to novel unseen domains. Specifically, we control the reader’s likelihood of the retrieved documents during training, to make sure that over-memorization of corpus knowledge is prevented. Extensive experiments are conducted on multiple OpenQA datasets in both zero-shot and fully-supervised training settings, and the results demonstrate that training the model with the proposed CIT loss significantly improves the model’s generalizability across different corpus versions and knowledge domains, without sacrificing the model’s inherent performance in its original domain.

7 Limitations

Although retrieval-augmented text generation models are effective for many knowledge-intensive tasks, they have an inherent limitation of the large requirement of computational memory. To ensure time efficiency in the retrieval process, an index of the external corpus, often vast in size, must be pre-constructed. Another notable limitation of CIT is that the extent of memorization mitigation depends on a hyper-parameter, which is experimented and chosen by humans. Ideally, the model should be able to automatically determine the best level of memorization mitigation to reach an optimal balance between parametric knowledge and retrieval augmentation. This is an exciting new research topic, and we will explore this as the future work.

Acknowledgement

We thank the anonymous reviewers for their constructive suggestions. This research is based upon work supported by U.S. DARPA KAIROS Program No. FA8750-19-2-1004. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Danqi Chen and Wen-tau Yih. 2020. [Open-domain question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 34–37, Online. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *CoRR*, abs/2210.11416.
- Matthew Dunn, Levent Sagun, Mike Higgins, V. Ugur Guney, Volkan Cirik, and Kyunghyun Cho. 2017. [Searchqa: A new q&a dataset augmented with context from a search engine](#).
- Hung-Chieh Fang, Kuo-Han Hung, Chao-Wei Huang, and Yun-Nung Chen. 2022. [Open-domain conversational question answering with historical answers](#). *CoRR*, abs/2211.09401.
- Revanth Gangi Reddy, Bhavani Iyer, Md Arafat Sultan, Rong Zhang, Avirup Sil, Vittorio Castelli, Radu Florian, and Salim Roukos. 2021. Synthetic target domain supervision for open retrieval qa. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1793–1797.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [REALM: retrieval-augmented language model pre-training](#). *CoRR*, abs/2002.08909.
- Rujun Han, Peng Qi, Yuhao Zhang, Lan Liu, Juliette Burger, William Yang Wang, Zhiheng Huang, Bing Xiang, and Dan Roth. 2023. [RobustQA: Benchmarking the robustness of domain adaptation for open-domain question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2023*,

- pages 4294–4311, Toronto, Canada. Association for Computational Linguistics.
- Ziniu Hu, Yichong Xu, Wenhao Yu, Shuohang Wang, Ziyi Yang, Chenguang Zhu, Kai-Wei Chang, and Yizhou Sun. 2022. [Empowering language models with knowledge graph reasoning for open-domain question answering](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 9562–9581. Association for Computational Linguistics.
- Jie Huang, Wei Ping, Peng Xu, Mohammad Shoeybi, Kevin Chen-Chuan Chang, and Bryan Catanzaro. 2023. [RAVEN: in-context learning with retrieval augmented encoder-decoder language models](#). *CoRR*, abs/2308.07922.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised dense information retrieval with contrastive learning](#). *Trans. Mach. Learn. Res.*, 2022.
- Gautier Izacard and Edouard Grave. 2021a. [Distilling knowledge from reader to retriever for question answering](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Gautier Izacard and Edouard Grave. 2021b. [Leveraging passage retrieval with generative models for open domain question answering](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880, Online. Association for Computational Linguistics.
- Gautier Izacard, Patrick S. H. Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. [Atlas: Few-shot learning with retrieval augmented language models](#). *J. Mach. Learn. Res.*, 24:251:1–251:43.
- Gautier Izacard, Fabio Petroni, Lucas Hosseini, Nicola De Cao, Sebastian Riedel, and Edouard Grave. 2020. [A memory efficient baseline for open domain question answering](#). *CoRR*, abs/2012.15156.
- Zhengbao Jiang, Luyu Gao, Zhiruo Wang, Jun Araki, Haibo Ding, Jamie Callan, and Graham Neubig. 2022. [Retrieval as attention: End-to-end learning of retrieval and reading within a single transformer](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2336–2349, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- 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 *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Jungo Kasai, Keisuke Sakaguchi, Yoichi Takahashi, Ronan Le Bras, Akari Asai, Xinyan Yu, Dragomir Radev, Noah A Smith, Yejin Choi, and Kentaro Inui. 2022. [Realtime qa: What’s the answer right now?](#) *arXiv preprint arXiv:2207.13332*.
- Akhil Kedia, Mohd Abbas Zaidi, and Haejun Lee. 2022. [FiE: Building a global probability space by leveraging early fusion in encoder for open-domain question answering](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4246–4260, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. 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](#). *Trans. Assoc. Comput. Linguistics*, 7:452–466.
- Linqing Liu, Patrick Lewis, Sebastian Riedel, and Pontus Stenetorp. 2022. [Challenges in generalization in open domain question answering](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2014–2029, Seattle, United States. Association for Computational Linguistics.
- Joel Oduro-Afriyie and Hasan Jamil. 2023. [Knowledge graph enabled open-domain conversational question answering](#). In *Flexible Query Answering Systems - 15th International Conference, FQAS 2023, Mallorca, Spain, September 5-7, 2023, Proceedings*, volume 14113 of *Lecture Notes in Computer Science*, pages 63–76. Springer.
- Bhargavi Paranjape, Matthew Lamm, and Ian Tenney. 2022. [Retrieval-guided counterfactual generation for QA](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1670–1686, Dublin, Ireland. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [In-context retrieval-augmented language models](#). *CoRR*, abs/2302.00083.

- Revanth Gangi Reddy, Md Arafat Sultan, Martin Franz, Avirup Sil, and Heng Ji. 2022a. Entity-conditioned question generation for robust attention distribution in neural information retrieval. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2462–2466.
- Revanth Gangi Reddy, Vikas Yadav, Md Arafat Sultan, Martin Franz, Vittorio Castelli, Heng Ji, and Avirup Sil. 2022b. Towards robust neural retrieval with source domain synthetic pre-finetuning. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1065–1070.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426.
- 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 *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 25968–25981.
- Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. [ColBERTv2: Effective and efficient retrieval via lightweight late interaction](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3715–3734, Seattle, United States. Association for Computational Linguistics.
- Shamane Siriwardhana, Rivindu Weerasekera, Tharindu Kaluarachchi, Elliott Wen, Rajib Rana, and Suranga Nanayakkara. 2023. [Improving the domain adaptation of retrieval augmented generation \(RAG\) models for open domain question answering](#). *Trans. Assoc. Comput. Linguistics*, 11:1–17.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- George Tsatsaronis, Georgios Balikas, Prodromos Malakasiotis, Ioannis Partalas, Matthias Zschunke, Michael R Alvers, Dirk Weissenborn, Anastasia Krithara, Sergios Petridis, Dimitris Polychronopoulos, et al. 2015. An overview of the bioasq large-scale biomedical semantic indexing and question answering competition. *BMC bioinformatics*, 16(1):1–28.
- Tu Vu, Mohit Iyyer, Xuezhi Wang, Noah Constant, Jerry Wei, Jason Wei, Chris Tar, Yun-Hsuan Sung, Denny Zhou, Quoc Le, et al. 2023. Freshllms: Refreshing large language models with search engine augmentation. *arXiv preprint arXiv:2310.03214*.
- Wenhan Xiong, Xiang Lorraine Li, Srini Iyer, Jingfei Du, Patrick S. H. Lewis, William Yang Wang, Yashar Mehdad, Scott Yih, Sebastian Riedel, Douwe Kiela, and Barlas Oguz. 2021. [Answering complex open-domain questions with multi-hop dense retrieval](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Pengfei Yu and Heng Ji. 2023. [Self information update for large language models through mitigating exposure bias](#).
- Zichun Yu, Chenyan Xiong, Shi Yu, and Zhiyuan Liu. 2023. [Augmentation-adapted retriever improves generalization of language models as generic plug-in](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2421–2436, Toronto, Canada. Association for Computational Linguistics.
- Mantong Zhou, Zhouxing Shi, Minlie Huang, and Xiaoyan Zhu. 2020. [Knowledge-aided open-domain question answering](#). *CoRR*, abs/2006.05244.

A Implementation Details

Experiment Details We train our model on 4 NVIDIA A100 GPUs with 80 GB memory, and the total training time is about 3.5 hours for the model to converge on an OpenQA dataset. During training, the index of the external corpus is pre-computed and equally sharded across all 4 GPUs. We adopt distributed data parallelism to make a copy of the model on each GPU and the data batch is splitted to the 4 devices.

Hyper-Parameter Settings Detailed hyper-parameter search range and choices are shown in Table 8 and Table 9 respectively. The choices are made by grid search.

Hyper-parameters	Searching Range
Maximum Length of FiD	[384, 512, 768]
# Retrieved Contexts	[20, 30, 40]
Generation Length	[10, 15, 20, 25, 30]
Masked Percentage for CIT	[0.1, 0.15, 0.2]
Strength of CIT α	[0.1, 0.2, 0.3, 0.4, 0.5, 0.6]
Learning Rate	[1e-5, 2e-5, 3e-5, 4e-5, 5e-5]
Batch Size	[4, 8, 12, 16]
Maximum Training Steps	[500, 1,000, 1,500]
Warm-up Steps	[50, 100, 150]
Weight Decay	[1e-2, 1e-3, 1e-4]
Retriever Dropout	[0.1, 0.2, 0.3]
Reader Dropout	[0.1, 0.2, 0.3]

Table 8: Detailed hyper-parameter search ranges for Corpus-Invariant Tuning.

Hyper-parameters	Values
Maximum Length of FiD	512
# Retrieved Contexts	40
Generation Length	20
Masked Percentage of CIT	0.15
Strength of CIT α	0.2 (RQ1); 0.3 (RQ2)
Learning Rate	4e-5
Batch Size	8
Maximum Training Steps	500
Warm-up Steps	50
Weight Decay	1e-2
Retriever Dropout	0.1
Reader Dropout	0.1

Table 9: Detailed hyper-parameter configurations for Corpus-Invariant Tuning.

B Case Study on Decoder-Only Models

The main idea of CIT is to control knowledge over-memorization to improve model’s generalization ability among different corpora and different domains. Such a training strategy is not only effective

for retrieval-augmented encoder-decoder models like Atlas, but can theoretically also apply to larger-scale auto-regressive foundation models. Therefore, we also conduct preliminary experiments with LLaMA-2-7b (Touvron et al., 2023) and Contriever as a case study to evaluate the robustness and ubiquity of our model among different language model architectures. Specifically, we freeze the retriever, and use the retrieved documents as the input prompt to generate the answer for each question. While maximizing the probabilities of the correct answers, we apply CIT in a similar way of maintaining the direct log-likelihood of these retrieved contexts C :

$$\mathcal{L}_{CIT} = \sum_{c \in C} \|\log p'_\phi(c) - \log p'_{\phi_0}(c)\|^2.$$

Different from the masked span prediction probability p_ϕ in Equation (1), the p'_ϕ here represents the language modeling probability of the entire passage c . Denoting the i -th word in c as c_i , then p'_ϕ can be formulated by

$$p'_\phi(c) = \prod_{i=1}^{N-1} p(c_{i+1} | c_{1:i}).$$

As shown in Table 10, we can observe that with auto-regressive language models like LLaMA, our proposed CIT can still achieves considerable improvements of the model’s generalization ability.

Dataset	Model	EM Score
NQ	LLaMA2-7b	54.2
	LLaMA2-7b + CIT	57.7
TriviaQA	LLaMA2-7b	69.9
	LLaMA2-7b + CIT	71.4

Table 10: Evaluation results of applying the idea of CIT on the LLaMA2-7b model with a fixed retriever (Contriever). The experiments are conducted in a zero-shot transfer setting between Wiki-2017 and Wiki-2018.