

Context Generation Improves Open Domain Question Answering

Dan Su^{‡,*}, Mostofa Patwary[§], Shrimai Prabhunoye[§], Peng Xu[§], Ryan Prenger[§],
Mohammad Shoeybi[§], Pascale Fung[‡], Anima Anandkumar[§], Bryan Catanzaro[§]

[‡]The Hong Kong University of Science and Technology, [§]NVIDIA

dsu@connect.ust.hk, mpatwary@nvidia.com

Abstract

Closed-book question answering (QA) requires a model to directly answer an open-domain question without access to any external knowledge. Prior work on *closed-book* QA either directly finetunes or prompts a pretrained language model (LM) to leverage the stored knowledge. However, they do not fully exploit the parameterized knowledge. To address this inefficiency, we propose a two-stage, *closed-book* QA framework which employs a *coarse-to-fine* approach to extract the relevant knowledge and answer a question. We first generate a related context for a given question by prompting a pretrained LM. We then prompt the same LM to generate an answer using the generated context and the question. Additionally, we marginalize over the generated contexts to improve the accuracies and reduce context uncertainty. Experimental results on three QA benchmarks show that our method significantly outperforms previous *closed-book* QA methods. For example on TriviaQA, our method improves exact match accuracy from 55.3% to 68.6%, and is on par with *open-book* QA methods (68.6% vs. 68.0%). Our results show that our new methodology is able to better exploit the stored knowledge in pretrained LMs without adding extra learnable parameters or needing finetuning, and paves the way for hybrid models that integrate pretrained LMs with external knowledge.

1 Introduction

Open-domain question answering (ODQA) produces an answer to a given question in the form of natural language, and the task has been extensively studied in recent years. Significant progress on ODQA has been made by developing the *open-book* QA methods (Chen et al., 2017; Lewis et al., 2020b; Guu et al., 2020; Izacard and Grave, 2021;

*This work was done when the first author was an intern at NVIDIA. Corresponding authors: Dan Su, Mostofa Patwary.

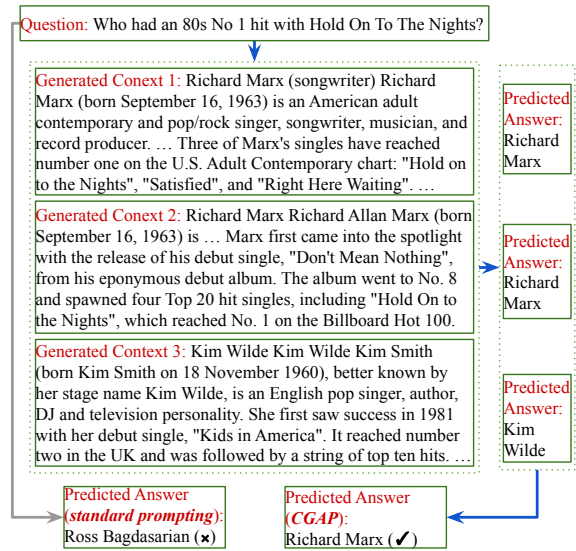


Figure 1: An example illustrating our two-stage, **CGAP** framework. **CGAP** generates more accurate answer (e.g. *Richard Marx*) compared to standard few-shot prompting (e.g. *Ross Bagdasarian*).

Lazaridou et al., 2022) that explicitly exploit external knowledge corpus via dense retrieval techniques like DPR (Karpukhin et al., 2020). However, learning a good retriever requires substantial resources, such as a large number of domain-specific pairs of question and contexts in the knowledge corpus (Karpukhin et al., 2020), or intensive compute resources (Lee et al., 2019). In addition, as the size of the knowledge corpus increases, it becomes harder to retrieve accurate contexts due to the high dimensionality of the search space (Reimers and Gurevych, 2021).

Another class of models, known as *closed-book* question answering (CBQA), were recently proposed (Roberts et al., 2020). CBQA tries to directly answer the open-domain questions without accessing any external knowledge sources, and instead leverages the parametric knowledge stored in the pretrained language models (LMs) (Raffel et al., 2020; Brown et al., 2020; Ye et al., 2020).

However, even with larger LMs, the *closed-book* methods are not competitive with the *open-book* methods in term of accuracy (Lewis et al., 2021).

While it has been shown that large pretrained LMs store an abundant amount of knowledge (Petroni et al., 2019; Roberts et al., 2020), we hypothesize the accuracy gaps are largely because the way of exploiting the parameterized knowledge are not sophisticated enough. Prior works on CBQA either finetune pretrained LM models on the entire QA datasets (Roberts et al., 2020; Ye et al., 2020), or they directly prompt those models using several few-shot QA pairs (Brown et al., 2020; Radford et al., 2019). On the contrary, *open-book* models use a two-stage pipeline. They first retrieve relevant contexts from external corpus, then they extract the answer based on the retrieved contexts.

Therefore, to better exploit the parameterized knowledge in pretrained LMs and bridge the large accuracy gaps between the *closed-book* and *open-book* methods, we propose a *coarse-to-fine*, two-stage method for CBQA task. The main idea is to leverage generated contexts as an intermediate bridge between the huge amount of parameterized knowledge stored in the LM and the answer that lies within this knowledge. To the best of our knowledge, no previous work has been conducted on generating context from large pretrained LMs for CBQA and leveraging them to predict answer.

Our proposed framework **CGAP** consists of two stages. It first performs **Context Generation** relevant to a given question by prompting a pretrained LM. It then prompts the same LM for **Answer Prediction** using the generated context and the question. In order to improve the accuracies and to reduce context uncertainties, we generate multiple contexts for each question and predict the final answer by majority voting. This step does not increase the inference cost as we generate the contexts in parallel by batching in a single inference call. Figure 1 illustrates how our two stage prompting and majority voting works. For the input question, CGAP generates 3 contexts and 3 predicted answers at the two stages respectively, and choose the most voted answer as the final answer. Note that we do not finetune the large pretrained LMs for context generation or answer prediction. This facilitates our approach to take advantage of large LMs such as GPT-3 (Brown et al., 2020), PALM (Chowdhery et al., 2022) or Megatron-Turing NLG 530B (Smith et al., 2022),

which are only available through APIs.

We conduct in-depth experimental studies on three open-domain QA benchmarks, Natural Questions (Kwiatkowski et al., 2019), WebQuestions (Berant et al., 2013), and TriviaQA (Joshi et al., 2017), and demonstrate significant improvements by our two stage prompting method. Our contributions are summarized as follows:

- We propose a simple yet effective few-shot prompting approach for ODQA that does not rely on any external knowledge sources or fine-tuning, but performs significantly better than existing *closed-book* approaches (e.g. exact matching 68.6% vs. 55.3%), and is on par with *open-book* methods (e.g. 68.6% vs. 68.0%).
- We show that the generated context can improve standard few-shot prompting based *closed-book* QA accuracy at various model scales (e.g. from 11.7% to 28.5%), and demonstrate that scaling up the context generation model further enlarges their accuracy gaps (e.g. 357M 28.5% vs. 530B 68.6%). To the best of our knowledge, we are the first to leverage generated context from large pretrained LMs for open-domain question answering.
- We show that generating multiple contexts without increasing the inference cost by batching can mitigate errors in answer prediction caused by variability in the unknown context (e.g. from 36.3% to 45.7%).

2 Methodology

Our proposed **Context Generation** and **Answer Prediction (CGAP)** framework is illustrated in Figure 2. CGAP consists of two stages. First, it generates relevant context to a given question by prompting a large pretrained LM. In the second stage, it predicts an answer using the generated context and the question by prompting the same LM. To accurately predict the answer, we generate multiple contexts. We run each of the two stages multiple times in parallel in batch for the same question, generating different contexts for each, and use majority voting to select the final answer.

Formally, for our task we have a question Q to be answered, and a support repository $\mathcal{D} = \{(c_1, q_1, a_1), \dots, (c_n, q_n, a_n)\}$ that consists of tuples of question q_i and answer a_i pairs with mapping to the context c_i . In our experiments, we use

the training sets of the corresponding datasets as \mathcal{D} .

2.1 Context Generation

As shown in Figure 2, in the first stage, given question Q , we select the m context generation prompts $S = \{(q_1, c_1), \dots, (q_m, c_m)\}$ from the support repository \mathcal{D} . We then use S with Q to prompt pretrained LM to generate k contexts, which are denoted by $C_{gen} = \{c_{gen}^1, c_{gen}^2, \dots, c_{gen}^k\}$.

Sample Selection Selecting appropriate samples for the prompts is the key to generate high-quality context relevant to a given question. Previous work has shown that leveraging relevant samples helps the LM to generate contextually relevant and factually correct context (Liu et al., 2021, 2022). We therefore use a similarity-based retriever to search relevant samples S from the corresponding supporting repository, \mathcal{D} . We use DPR (Karpukhin et al., 2020) in our framework. In our DPR setup, we represent the question and the samples in \mathcal{D} as 768-dimensional dense vector representations, computed via the BERT-based bi-encoder networks. We rank the documents according to their similarity score, calculated as:

$$Score(Q, (q_j, c_j)) = \text{BERT}(Q)^T \cdot \text{BERT}(q_j; c_j) \quad (1)$$

where $;$ denotes concatenation of the tokens of the question q_j and the context c_j . Finally, we get $S = \{(q_1, c_1), \dots, (q_m, c_m)\}$ which are the top- m retrieved samples for question Q .

We would like to emphasize that the selected samples from \mathcal{D} are used as examples in the few-shot prompting to the pretrained LM to generate context, not as the source of external knowledge containing the answer.

Prompts Construction Given the question Q and the set of question-context pair samples S selected, we use few-shot prompting to condition pretrained LMs on the samples. We use similar few-shot prompting technique for *closed-book* QA as in (Brown et al., 2020), that considers multiple <question, answer> pairs. The template we used to construct prompts is: $Q: \dots A: \dots$. Thus the constructed prompt $Prompt(Q)$ for a given question Q becomes:

$$Prompt(Q) = Q: q_m \setminus n A: c_m \setminus n \dots \\ Q: q_1 \setminus n A: c_1 \setminus n Q: Q \setminus n$$

We use ' $\setminus n$ ' to separate the question, context and the samples. We investigated the order of samples to optimize the prompt and find that using the retrieved samples in reversed order of similarity yields better accuracies across all datasets¹. We now pass $Prompt(Q)$ through a pretrained LM to generate the context as follows:

$$c_{gen} = \mathcal{LM}(Prompt(Q))$$

To generate a set of k contexts, $\{c_{gen}^1, \dots, c_{gen}^k\}$, we increase the inference batch size to k and generate all the k contexts in parallel in one inference call to the LM. Thus, the overall latency remains the same as using a single context.

2.2 Answer Prediction

In the second stage, we select m answer prediction prompts $S' = \{(q_1, a_1, c_1), \dots, (q_m, a_m, c_m)\}$ from \mathcal{D} and then we prompt the same LM using the generated context C_{gen} from the first stage, along with the question Q and S' . The LM predicts a set of k answers $A_p = \{a_p^1, a_p^2, \dots, a_p^k\}$ each corresponding to the k contexts in C_{gen} . The final answer A is selected by majority voting on A_p .

Sample Selection Constrained by the maximum sequence length of the LM, we can feed the LM only a few (c, q, a) samples. Thus, it could be difficult for the LM to learn how to predict the answer for the given question conditioned on the context, unless similar examples have been provided. For example, if we were asking the question 'who is the current director of the us mint?', the example that answering the question 'who is the fbi director of the united states?' from the provided context will be more helpful, than the example that is answering 'how many episodes are there in 'Dragon Ball Z'?' from the given context. We therefore use the same criteria for answer prediction as has been used for context generation. We use the same set of samples as selected in the first stage as described in Equation 1 and denote as $S' = \{(q_1, c_1, a_1), \dots, (q_m, c_m, a_m)\}$.

Prompt Construction We are prompting LMs with few-shot examples to predict answer for the question conditioned on the generated context. To equip the LM with this capability, we constructed intuitive prompts for the selected examples and feed them into the LM. Specifically, the template

¹We show an concrete example of $Prompt(Q)$ in Appendix Table 12

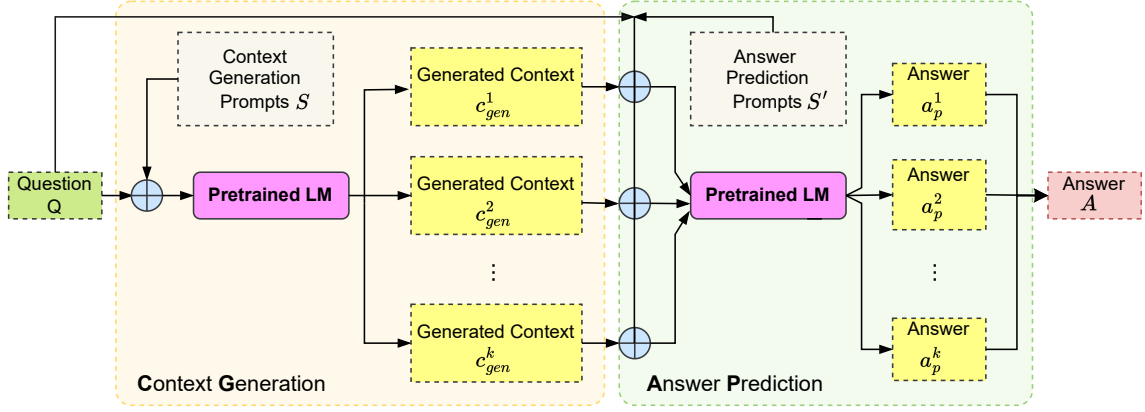


Figure 2: Overview architecture of our **CGAP** framework. It first does **Context Generation** by prompting large pretrained LMs, then it further prompts the LMs for **Answer Prediction** by feeding the generated context to the LM models alongside the question. k contexts are generated and the final answer A is chosen by majority voting. (If computation capability allows, it could prompt multiple (k) LMs in parallel at both two stages to speed up.)

we used to construct answer prediction prompts is: $C : \dots Q : \dots A : \dots$. Thus, the constructed prompt for a given question Q and the i -th generated context c_{gen}^i is:

$$\begin{aligned} Prompt(c_{gen}^i, Q) = & C : c_m \setminus n Q : q_m \setminus n A : a_m \setminus n \\ & \dots \\ & C : c_1 \setminus n Q : q_1 \setminus n A : a_1 \setminus n \\ & C : c_{gen}^i \setminus n Q : Q \setminus n \end{aligned} \quad (2)$$

We then feed $Prompt(c_{gen}^i, Q)$ into the pretrained LM to predict the answer:

$$a_p^i = \mathcal{LM}(Prompt(c_{gen}^i, Q)) \quad (3)$$

where we use a_p^i to denote the i -th answer predicted by the LM. The k generated contexts in c_{gen} will yield a set of answers $A_p = \{a_p^1, \dots, a_p^k\}$.

2.3 Context Marginalization

The large pretrained LM can generate impressively fluent and relevant context given input, it also has a tendency to generate factually incorrect statements, ranging from subtle inaccuracies to wild hallucinations (Shuster et al., 2021; Krishna et al., 2021; Su et al., 2022). Answers conditioned solely on hallucinated or erroneous statements are likely to be incorrect (Equation 3). Thus, we would like to remove the variability in the answer due to any particular generated context.

Ideally, we could marginalize over this unknown context by producing an answer for every possible context, weighting each answer by the probability of the context. Here we approximate this by

generating a set of contexts, and selecting the final answer based on majority voting. Suppose there are T unique answers $\{A_p^1, \dots, A_p^T\}$ from the k predicted answer from Equation 3 where $T \leq k$, then we select the J -th answer that receives the highest number of votes from the T different answers via:

$$J = \operatorname{argmax}_{j \in \{1, 2, \dots, T\}} \sum_{i=1}^k (\mathbb{1}(a_p^i = A_p^j)) \quad (4)$$

as the final answer A . As k gets larger, the final answer A will converge to the answer that would be produced marginalizing over all possible contexts. We refer to this majority vote over multiple generated contexts as context marginalization.

3 Experimental Setup

3.1 Datasets

We evaluated our experiments on three open-domain QA benchmark datasets: Natural Questions (NQ) (Kwiatkowski et al., 2019), TriviaQA (TQA) (Joshi et al., 2017), and WebQuestions (WQ) (Berant et al., 2013), using the same data splits for train, validation and test as in Lee et al. (2019); Izcard and Grave (2021).

NQ contains questions from Google search queries; TQA contains a collection of questions from trivia and quiz-league websites, and we use their unfiltered set; while questions of WQ were from Google Suggest API. For NQ and TQA, we use the processed data provided by Izcard and Grave (2021), in which each question-answer pair is accompanied by a 100-words Wikipedia passage

Model Type	Model	Method	NQ	TQA	WQ
Open-book	RAG (Lewis et al., 2020c)	<i>Finetuned</i>	44.5	68.0	45.5
	Fusion-in-Decoder (large) (Izcard and Grave, 2021)	<i>Finetuned</i>	51.4	67.6	-
	OB_{Google}^{PoE} (Lazaridou et al., 2022)	<i>Few-shot</i>	38.4	-	-
Closed-book	T5-11B (Roberts et al., 2020)	<i>Finetuned</i>	32.6	42.3	37.2
	T5-11B+SSM (Roberts et al., 2020)	<i>Finetuned</i>	34.8	51.0	40.8
	BART-large, pre-finetuned on PAQ (Lewis et al., 2021)	<i>Finetuned</i>	32.7	33.2	-
	LM-530B (API)	<i>Few-shot</i>	23.0	55.3	23.6
	CGAP (ours, 530B)	<i>Few-shot</i>	<u>42.0</u>	68.6	<u>41.8</u>

Table 1: Exact Match score for **CGAP** (highest accuracy configurations) in comparison to recent state-of-the-art *open-book* and *closed-book* based systems. Highest score indicated in **bold**, highest *closed-book* model underlined.

containing the answer. For WQ, we retrieved the corresponding context passage for each question from 2019/08/01 Wikipedia dump, using the DPR-based retriever that is trained jointly on the union of knowledge-intensive training data in KILT benchmark (Petroni et al., 2021).

3.2 Baselines

We compare our **CGAP** framework with the following baseline methods for *closed-book* QA.

Standard Few-shot Prompting We use the standard few-shot prompting technique similar to GPT-3 (Brown et al., 2020) in our evaluation on the *closed-book* QA datasets as described in Section 3.1. We consider this technique as the few-shot baseline in all our experiments. The baseline that is experimented using 530 billion (530B) parameterized LM is referred as **LM-530B**.

LM Fine-tuning Roberts et al. (2020) first proposed the *closed-book* QA task for open domain QA, and they directly fine-tuned T5 (Raffel et al., 2019) using the entire QA pairs in the training data, without access to any external knowledge corpus (referred as **T5-11B**). They also experimented with using ‘Salient Span-Masking’ (SSM) to continue pretraining the T5 checkpoints before fine-tuning for QA (referred as **T5-11B+SSM**). Lewis et al. (2021) pre-finetuned BART-large (Lewis et al., 2020a) on *Probably Asked Questions* (PAQ), a very large resource of 65M automatically generated QA-pairs, then further finetuned the model on corresponding training data (referred as **BART-large, pre-finetuned on PAQ**).

Open-book Few-shot Prompting Lazaridou et al. (2022) used few-shot prompting for open domain QA task, but they generate the answer via conditioning on retrieved documents from Google

Search API. (referred as OB_{Google}^{PoE})

3.3 State-of-the-art Open-book QA Models

We compare the state-of-the-art *open-book* QA models with **CGAP**. **Fusion-in-Decoder** (FiD) (Izcard and Grave, 2021) uses DPR (Karpukhin et al., 2020) to retrieve 100 passages from Wikipedia. Then they encode each passage independently and combine all outputs from the T5 encoder before passing them to the T5 decoder to generate a final answer. **RAG** (Lewis et al., 2020b) is an end-to-end retrieval-augmented generation model.

3.4 Implementation Details

To test how different model scales affect the performance of our approach, we train and experiment on a collection of decoder-only LMs using the Megatron-LM framework (Shoeybi et al., 2019), with 357 million (357m), 1.3 billion (1.3b), and 530 billion (530b) (Smith et al., 2022) parameters, at both context generator and answer prediction stage. We use top- p sampling with a value of 0.9 to generate diversified contexts. However, to handle the deterministic generation (e.g. short answer), we use greedy decoding at the answer prediction stage, similar to (Chowdhery et al., 2022; Wang et al., 2022).

For the prompt configuration at both stages, we choose 10 samples, constrained by the maximum sequence length of the LMs. We use DPR checkpoint from Huggingface² to select samples from the supporting repository.

3.5 Evaluation

For evaluating the open-domain QA task, we followed the recent works (Rajpurkar et al., 2016; Lee

²https://huggingface.co/facebook/dpr-ctx_encoder-multiset-base




Figure 3: Ablation on context generation LM size. The dash lines represent standard few-shot prompting baselines.

et al., 2019; Izcard and Grave, 2021) that use Exact Match (EM) as the evaluation metric. Each predicted answer is compared to the ground-truth after both are lowercased and stripped of articles, punctuation, and duplicate whitespace.

4 Results and Ablation Studies

We now show our main results as well as ablations to further analyze the effectiveness of our approach.

4.1 Main Results

Table 1 shows the EM score comparison between our CGAP-based method with existing *closed-book* baseline approaches³. We also compare with state-of-the-art *open-book* models at the upper section of the table.

As we can see, our CGAP based method outperforms other existing *closed-book* methods by large margin, especially on NQ and TQA datasets. The CGAP also outperforms the standard few-shot prompting baseline LM-530B on all three datasets (at least by 13.3 EM point).

Furthermore, CGAP obtains highest score on TriviaQA. The scores are also very close to the state-of-the-art *open-book* method RAG on NQ and WebQuestions, but only lose few points on NQ to FiD. While FiD uses 100 retrieved passages for answer prediction, CGAP only uses 8 generated contexts for approximate context marginalization.

4.2 Ablation Studies

We conducted a systematic ablation study to further investigate the contribution of the context generation model and the effect of context marginalization.

³GPT-3 API shows different results than reported in the paper (Brown et al., 2020). We therefore did not compare to it. Details are shown in Appendix A

4.2.1 Context Generation

While previous work (Roberts et al., 2020; Brown et al., 2020) demonstrated that the scale of the model sizes improves the answer accuracy of *closed-book* QA, there are also other findings showing that simply increasing the model size does not lead to substantive accuracy gains (Rae et al., 2021). Thus, we intend to investigate **how will the context generation LM affect the answer accuracy**.

We experimented by varying the LM sizes for context generation, and fix the answer generation LM. We used context generation LM sizes of 357m, 1.3B and 530B, and answer generation LM with 357m and 1.3B parameters. We also compare with standard few-shot prompting which has no context generation.

We plot the results in Figure 3. As we can see, there are huge accuracy gains from standard prompting, to CGAP method that has context generation. The accuracy increases by absolute 19.00% for NQ, 16.87% for TQA and 15.26% for WQ, when using 357M model for both standard prompting and CGAP approach. The answer accuracy continues to increase when we increase the LM size for context generation. Furthermore, we notice that the slopes of the accuracy gain curve using larger answer prediction model is steeper than using smaller one on all three datasets. This suggests the use of larger answer prediction LM to fully exploit the knowledge in generated context.

4.2.2 Context Marginalization

Since there will be some hallucinated content or erroneous statements in the generated context, we approximate context marginalization by sampling multiple contexts and selecting the final answer based on majority voting, as introduces in Section 2.3. Here, we investigate the performance gains brought in by context marginalization, and

AP LM Size	CG LM Size	Margin-alization	NQ	TQA	WQ
357M	357M	✗	22.9	28.5	26.2
		✓	25.7 (+2.8)	33.4 (+4.9)	29.6 (+3.4)
	1.3B	✗	23.1	29.7	28.3
		✓	26.1 (+3.0)	34.8 (+5.1)	31.3 (+3.0)
530B	530B	✗	26.3	36.3	31.2
		✓	28.9 (+2.7)	45.7 (+9.4)	34.0 (+2.8)
	530B	✗	29.5	56.3	28.3
		✓	42.0 (+12.5)	68.6 (+12.4)	41.8 (+13.5)

Table 2: Ablation on context marginalization. (AP and GP represent Answer Prediction and Context Generation, respectively.)




Figure 4: Ablation on k , the number of contexts for marginalization.

also the accuracy curves with varied number of sampled contexts used in the approximate marginalization k .

In Table 2, we show the accuracy comparisons w/ and w/o using marginalization ($k=8$), with different LM sizes. As we can see, **context marginalization improves the answer accuracy consistently on the three datasets**⁴, under all settings. Notably, there is much larger performance gains using marginalization when we scale up the model sizes to 530 billion parameters (i.e. increase EM score by 12.8% averaged on three datasets).

The larger the number of context samples k , the more accurately the majority vote reflects the true marginalization over all possible contexts. Therefore, we perform further ablation by changing the value k for 357M LM for both context generation and answer prediction. We plot the accuracy curves in Figure 4. We see that there are accuracy improvements when we use more context samples. As expected and curves plateau for larger values of k as the approximation approaches the true marginalization over all possible contexts.

⁴We show a concrete example in Appendix B Table 11

5 Analysis

Considering that it is the first time leveraging context generated by large pretrained LMs for ODQA, we also conducted further analysis.

We compare generated context with retrieved context in the two-stage, few-shot prompting based CBQA framework. It is a dominant paradigm to use retrieved context from external corpus together with the question for answer prediction for *open-book* QA (Chen et al., 2017; Lewis et al., 2020c; Izacard and Grave, 2021; Lazaridou et al., 2022).

5.1 Retrieved vs. Generated Context

In CBQA setting, we are not allowed to retrieve context from external knowledge sources. However, we can retrieve the contexts from the supporting repository based on their relevance to the given question. We use $c_r = \{c_r^1, c_r^2, \dots, c_r^m\}$ to represent the top- m relevant context for question Q . It can be obtained via Equation 1.

Let the top-1 retrieved context be $c_r^{\text{top-1}}$ for question Q . We use $c_r^{\text{top-1}}$ to compare with the generated context, c_{gen} . We use the same top- m prompts S' for answer prediction as introduced in Section 2.2. The answer a_p^r for the $c_r^{\text{top-1}}$ will be:

$$a_p^r = \mathcal{LM}(\text{Prompt}(c_r^{\text{top-1}}, Q)) \quad (5)$$

where $\text{Prompt}(c_r^{\text{top-1}}, Q)$ can be obtained via Equation 2.

The comparison between $c_r^{\text{top-1}}$ and c_{gen} is shown in Table 4. From the upper part of the table, we see that using $c_r^{\text{top-1}}$ gives slightly higher EM score than using c_{gen} generated by 357M and 1.3B LMs. However, c_{gen} gives higher EM scores than $c_r^{\text{top-1}}$ on all three datasets when we scale up the context generation LM size to 530B. This suggests the use of large pretrained LM for a better generated context.

Question: Which sitcom star appeared on the big screening 'The Object of My Affection'?	
Golden Answer: [Jennifer Anniston, Jen Aniston, ...]	
Predicted Answer (w/o c_{gen}):	Ross Hatley
Predicted Answer ($c_r^{\text{top-1}}$):	Laurie Metcalfe
Predicted Answer (c_{gen}):	Jennifer Anniston / Paul Rudd / Christine Baranski / Lisa Kudrow
Predicted Answer ($(c_{gen}^1, \dots, c_{gen}^k)$):	Jennifer Anniston

Table 3: Comparison of answers predicted w/o and w/ different context. Example from TriviaQA (Joshi et al., 2017) test set. Red and green colors denote in-correct and correct answer, respectively.

AP	Context	NQ	TQA	WQ
357M	$c_r^{\text{top-1}}$	25.1	32.2	28.3
	c_{gen} (357M LM)	22.9	28.5	26.2
	c_{gen} (1.3B LM)	23.1	29.7	28.3
	c_{gen} (530B LM)	26.3	36.3	31.2
530B	$c_r^{\text{top-1}}$	30.8	58.1	29.5
	c_{gen} (530B LM)	29.5	56.3	28.3

Table 4: Comparison of using retrieved top-1 context $c_r^{\text{top-1}}$, with few-shot generated context c_{gen} on *closed-book* QA task.

5.2 Multiple Retrievals vs. Context Marginalization

We notice that in Table 4, $c_r^{\text{top-1}}$ performs slightly better than c_{gen} when using 530B LM for answer prediction. We argue that this might be caused by the hallucination in c_{gen} . While we have shown in Section 4.2.2 that context marginalization could mitigate the problem and improve answer accuracy, we further facilitate c_{gen} (530B) with context marginalization and compare with retrieved context.

For fair comparison, we perform majority voting using the top- k retrieved context c_r , since Karpukhin et al. (2020) showed that the quality of the retrieved documents will also affect the final answer accuracy. Specifically, we replace $c_r^{\text{top-1}}$ with each retrieved context c_r^i in Equation 5 to predict answer $a_p^{r(i)}$ ($i = 1, \dots, k$), and use Equation 4 to select the most frequent answer as the final answer.

Furthermore, we replace $c_r^{\text{top-1}}$ with golden context c_{golden} in Equation 5. This will be the upper-bound of using retrieved/generated context in the two-stage, few-shot prompting CBQA task.

We show the results in Table 5. As we can see, using marginalization over c_{gen} consistently outperforms $c_r^{\text{top-1}}$, and also better than majority voting over multiple retrieved contexts c_r for answer prediction on all three datasets. Notably, marginal-

AP	Context	NQ	TQA	WQ
530B	c_{golden}	36.0	61.3	30.2
	$c_r^{\text{top-1}}$	30.8	58.1	29.5
	(c_r^1, \dots, c_r^k)	29.5	56.3	28.3
	c_{gen}	23.0	55.3	23.6
	$(c_{gen}^1, \dots, c_{gen}^k)$	42.0	68.6	41.8

Table 5: Comparison of using context marginalization ($c_{gen}^1, \dots, c_{gen}^k$), multiple retrievals (c_r^1, \dots, c_r^k), and golden context c_{golden} on *closed-book* QA task.

ization over c_{gen} yields higher EM score than using c_{golden} when using 530B LM for answer prediction. We observed similar trends when experimented on 357M and 1.3B parameter models. In Table 3, we show a concrete example that compares using different context for answer generation for better understanding⁵.

6 Related Works

Open-domain QA is the task of answering general-domain questions (Chen et al., 2017), in which the evidence is usually not given. Models that explicitly exploit an external corpus are referred as *open-book* models (Roberts et al., 2020). They typically index the corpus and then *retrieve-and-read* to extract the answer span from documents (Chen et al., 2017; Lee et al., 2019; Izacard and Grave, 2021; Lewis et al., 2020b; Lazaridou et al., 2022). Another recently proposed class of methods is *closed-book* QA models. Ye et al. (2020); Roberts et al. (2020) finetune pre-trained LMs such as T5 (Raffel et al., 2020) or BART (Lewis et al., 2020a) with QA pairs without access to any external knowledge or context.

Few-shot LM Prompting Radford et al. (2019); Brown et al. (2020) prompt GPT-2 (Radford et al., 2019) and GPT-3 (Brown et al., 2020) conditioned

⁵More concrete comparison examples are shown in Appendix B Table 9 and Table 10.

on several few-shot examples to predict the answer for ODQA. Most recent work by Lazaridou et al. (2022) further empower LM’s few-shot prompting abilities with information returned from the web using Google-Search API, and experimented on QA task. While Wei et al. (2022); Wang et al. (2022) use *chain of thought* few-shot prompting of LM to generate a coherent chain of short sentences that mimic the reasoning process of human might employ to solve reasoning tasks.

7 Conclusion

We propose a simple yet effective framework named **CGAP** for open-domain QA. CGAP performs **C**ontext **G**eneration followed by **A**nswer **P**rediction via two-stage prompting using large pre-trained LMs. It does not rely on external knowledge sources, and does not need finetuning or add extra learnable parameters. To the best of our knowledge, we are the first to leverage generated context from large pretrained LMs for open-domain QA. Experimental results on three QA benchmarks show that our method significantly outperforms previous *closed-book* QA methods and is par with *open-book* methods. We demonstrate our method up to 530B parameter models and showcase that larger models boost the accuracy by huge margins.

8 Limitations

As we show in the paper, **CGAP** has obtained satisfactory results on open-domain QA task. However, the method have limitations. The accuracy of CGAP will be affected by the size of LMs it uses, as we shown in Figure 3. In Section 4.1, our highest accuracy results reported in Table 1 used a large 530B pretrained LM, which is only accessible via API. Also, the generated context may contain hallucinated content.

References

Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1533–1544.

Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.

Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Realm: Retrieval-augmented language model pre-training. *arXiv preprint arXiv:2002.08909*.

Gautier Izacard and Édouard Grave. 2021. 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.

Mandar Joshi, Eunsol Choi, Daniel S 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.

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.

Kalpesh Krishna, Aurko Roy, and Mohit Iyyer. 2021. Hurdles to progress in long-form question answering. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4940–4957.

Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466.

Angeliki Lazaridou, Elena Gribovskaya, Wojciech Stokowiec, and Nikolai Grigorev. 2022. Internet-augmented language models through few-shot prompting for open-domain question answering. *arXiv preprint arXiv:2203.05115*.

Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. Latent retrieval for weakly supervised open domain question answering. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096.

- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020a. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- 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. 2020b. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020c. Retrieval-augmented generation for knowledge-intensive nlp tasks. *arXiv preprint arXiv:2005.11401*.
- Patrick Lewis, Yuxiang Wu, Linqing Liu, Pasquale Minervini, Heinrich Küttler, Aleksandra Piktus, Pontus Stenetorp, and Sebastian Riedel. 2021. Paq: 65 million probably-asked questions and what you can do with them. *Transactions of the Association for Computational Linguistics*, 9:1098–1115.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.
- Zihan Liu, Mostofa Patwary, Ryan Prenger, Shrimai Prabhumoye, Wei Ping, Mohammad Shoeybi, and Bryan Catanzaro. 2022. Multi-stage prompting for knowledgeable dialogue generation. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1317–1337.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, et al. 2021. Kilt: a benchmark for knowledge intensive language tasks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2523–2544.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog 1.8 (2019)*: 9.
- Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, Eliza Rutherford, Tom Hennigan, Jacob Menick, Albin Cassirer, Richard Powell, George van den Driessche, Lisa Anne Hendricks, Mari-beth Rauh, Po-Sen Huang, Amelia Glaese, Johannes Welbl, Sumanth Dathathri, Saffron Huang, Jonathan Uesato, John Mellor, Irina Higgins, Antonia Creswell, Nat McAleese, Amy Wu, Erich Elsen, Siddhant Jayakumar, Elena Buchatskaya, David Budden, Esme Sutherland, Karen Simonyan, Michela Paganini, Laurent Sifre, Lena Martens, Xiang Lorraine Li, Adhiguna Kuncoro, Aida Nematzadeh, Elena Gribovskaya, Domenic Donato, Angeliki Lazaridou, Arthur Mensch, Jean-Baptiste Lespiau, Maria Tsim-poukelli, Nikolai Grigorev, Doug Fritz, Thibault Sottiaux, Mantas Pajarskas, Toby Pohlen, Zhitao Gong, Daniel Toyama, Cyprien de Masson d’Autume, Yujia Li, Tayfun Terzi, Vladimir Mikulik, Igor Babuschkin, Aidan Clark, Diego de Las Casas, Aurelia Guy, Chris Jones, James Bradbury, Matthew Johnson, Blake Hechtman, Laura Weidinger, Iason Gabriel, William Isaac, Ed Lockhart, Simon Osindero, Laura Rimell, Chris Dyer, Oriol Vinyals, Kareem Ayoub, Jeff Stanway, Lorraine Bennett, Demis Hassabis, Koray Kavukcuoglu, and Geoffrey Irving. 2021. [Scaling language models: Methods, analysis & insights from training gopher](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- 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. *Journal of Machine Learning Research*, 21:1–67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Nils Reimers and Iryna Gurevych. 2021. [The curse of dense low-dimensional information retrieval for large index sizes](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 605–611, Online. Association for Computational Linguistics.
- 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.

Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*.

Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*.

Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. 2022. Using deep-speed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990*.

Dan Su, Xiaoguang Li, Jindi Zhang, Lifeng Shang, Xin Jiang, Qun Liu, and Pascale Fung. 2022. Read before generate! faithful long form question answering with machine reading. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 744–756.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.

Qinyuan Ye, Belinda Z Li, Sinong Wang, Benjamin Bolte, Hao Ma, Wen-tau Yih, Xiang Ren, and Madian Khabsa. 2020. Studying strategically: Learning to mask for closed-book qa. *arXiv preprint arXiv:2012.15856*.

A Standard Few-shot Prompting of GPT-3

Brown et al. (2020) adopted the standard few-shot prompting on GPT-3, and evaluated on the three open-domain QA datasets NQ (Kwiatkowski et al., 2019), WQ (Berant et al., 2013) and TQA (Joshi et al., 2017), for *closed-book* QA task. In order to compare with their reported results, we re-implement their method using the same few-shot configuration as described in the paper and query the OpenAI API.

Experimental Setups As OpenAI hasn’t officially release information about their API model sizes, we deduce the sizes of OpenAI API models based on their performances from EleutherAI’s blog⁶. Specifically, we query Ada and Babbage

models’ API, trying to reproduce the reported results for GPT-3 Medium (350M) and GPT-3 XL (1.3B) models, respectively.

We use two prompt formats to query the OpenAI API. The first prompt format is the one described in the paper (Brown et al., 2020) (referred as *GPT-3 format*): randomly draw 64 question-answer pairs from the corresponding supporting repository, and use ‘Q: ’ and ‘A: ’ respectively as prefix before each question and answer, to build the conditioning prompts. We also use the prompt format from EleutherAI’s language model evaluation harness github⁷ (referred as *EleutherAI*). Furthermore, we experiment using the same prompting format as we used in our standard prompting baseline (LM-530B) in Section 3.2 (referred as *Our format*), and prompting the LM of size 357M and 1.3B to compare.

Results We show the results of prompting GPT-3 under zero-shot, one-shot and few-shot settings in Table 6, Table 7 and Table 8 respectively. As we can see, no matter what prompting formats we use, the results reported in the GPT-3 paper (Brown et al., 2020) are almost always higher than our reproduced ones on all three datasets, over the two different LM sizes. The gaps become even larger at few-shot setting. Thus we conjecture that we are not able to reproduce the results reported by Brown et al. (2020) using GPT-3 (175B) on the three QA datasets. So we did not include their reported results to compare with our CGAP method in Table 1.

Furthermore, we notice that the results based on our baseline’s prompting configuration are always on par with the results from querying OpenAI API. Thus we believe that the **LM-530B** is a reliable and fair standard few-shot prompting baseline to compare with.

B Examples

We show three examples from NQ, TQA and WQ test set in Table 9, Table ?? and Table 10 respectively. In each table, we show the predicted answers from (1) standard prompting, (2) two-stage prompting using top-1 retrieved context $c_r^{\text{top-1}}$, (3) CGAP w/o marginalization, and (4) CGAP. All those predicted answers are based on LMs of size 530B.

We also show an example illustrate CGAP with 8 generated context and their corresponding pre-

⁶<https://blog.eleuther.ai/gpt3-model-sizes/>

⁷<https://github.com/EleutherAI/lm-evaluation-harness>

dicted answer in Table 11. As we can see, the contexts that contains lot of factually inaccurate or irrelevant content (e.g. generated context 1, 2, 4, 5, 8), thus the corresponding answer is wrong/inaccurate. However, the context generation LM also generates contexts that are more relevant and factual (e.g. generated context 3, 6, 7), and they help the answer prediction LM generate a correct answer. Therefore, CGAP can predict the final answer correctly based on marginalization over generated contexts.

Model Sizes	Model sources	Prompting format	zero-shot		
			NaturalQuestion	TriviaQA	WebQuestion
350M	GPT-3 Medium	GPT-3 paper (Brown et al., 2020)	1.75	7.61	3.20
	OpenAI API (Ada)	GPT-3 format	1.36	5.45	1.92
		EleutherAI	1.39	5.54	2.46
	LM-357M	Our format	1.41	5.04	2.12
1.3B	GPT-3 XL	GPT-3 paper (Brown et al., 2020)	4.40	19.70	4.63
	OpenAI API (Babbage)	GPT-3 format	2.27	9.84	2.12
		EleutherAI	2.47	12.77	5.22
	LM-1.3B	Our format	3.88	14.13	5.61

Table 6: Standard zero-shot prompting of GPT-3 for open-domain QA.

Model Sizes	Model sources	Prompting format	one-shot(k=1)		
			NaturalQuestion	TriviaQA	WebQuestion
350M	GPT-3 Medium	GPT-3 paper (Brown et al., 2020)	3.07	12.90	6.20
	OpenAI API (Ada)	GPT-3 format	1.83	10.26	5.07
		EleutherAI	1.77	10.02	5.61
	LM-357M	Our format	2.24	9.75	5.12
1.3B	GPT-3 XL	GPT-3 paper (Brown et al., 2020)	5.43	26.50	9.15
	OpenAI API (Babbage)	GPT-3 format	3.55	20.56	8.27
		EleutherAI	3.55	21.45	9.45
	LM-1.3B	Our format	4.71	21.21	8.76

Table 7: Standard one-shot prompting of GPT-3 for open-domain QA.

Model Sizes	Model sources	Prompting format	few-shot(k=64)		
			NaturalQuestion	TriviaQA	WebQuestion
350M	GPT-3 Medium	GPT-3 paper (Brown et al., 2020)	4.46	16.30	12.60
	OpenAI API (Ada)	GPT-3 format	3.43	12.46	10.73
		EleutherAI	3.71	12.46	10.29
	LM-357M	Our format	3.85	11.66	10.97
1.3B	GPT-3 XL	GPT-3 paper (Brown et al., 2020)	9.72	32.10	19.60
	OpenAI API (Babbage)	GPT-3 format	8.28	24.70	18.95
		EleutherAI	7.81	24.93	18.16
	LM-1.3B	Our format	7.87	24.88	17.52

Table 8: Standard few-shot (k=64) prompting of GPT-3 for open-domain QA.

Question: When is the next deadpool movie being released?	
Golden Answer:	[May 18, 2018]
Predicted Answer (standard prompting):	Prime availability TBD
Predicted Answer ($c_r^{\text{top-1}}$):	May 18, 2018
Predicted Answer (CGAP w/o marginalization):	May 18, 2018 / date21-May-2018 / May 29, 2019 /16th May 2018
Predicted Answer (CGAP):	May 18, 2018

Table 9: Example from NQ (Kwiatkowski et al., 2019) test set. Red and green colors denote in-correct and correct answer, respectively.

Question: Where George Lopez was born?	
Golden Answer:	[Mission Hills]
Predicted Answer (standard prompting):	San Fernando, California
Predicted Answer ($c_r^{\text{top-1}}$):	Washington Heights
Predicted Answer (CGAP w/o marginalization):	Mission Hills / Los Angeles/ San Fernando /Castle Hill /canada
Predicted Answer (CGAP):	Mission Hills

Table 10: Example from WebQuestions (Berant et al., 2013) test set.

Question: Where George Lopez was born? **Golden Answer:** [Mission Hills]

Generated Context 1: George Lopez They are American citizens, born in Los Angeles, to Gregorio Lopez (October 26, 1934 – January 25, 2011) and Maria Torres (October 12, 1938 – October 25, 2014), Mexican immigrants who were born in Reynosa, Tamaulipas. Lopez's paternal grandparents were Mexican, while his maternal grandparents were of Mexican and Spanish descent. Lopez's father was a worker for the Los Angeles Metropolitan Transit Authority. He married Maria in Los Angeles when he was 16 years old, and she was 15. Maria filed for divorce from Gregorio in the mid-1970s.

Predicted Answer: Los Angeles

Generated Context 2: George Lopez born April 23, 1961. The family eventually moved from Marin to Los Angeles in the mid-1960s. After Lopez graduated from San Fernando High School in 1979, he attended the California State University, Los Angeles. It was there where Lopez discovered his comedic talent after a video of him doing a stand-up routine made it into the hands of "Tonight Show" booker, Fred De Cordova, in 1985.

Predicted Answer: canada

Generated Context 3: George Lopez at 20th Century Fox Television until 2010, when he began appearing in his self-produced sitcom, George Lopez. Lopez was born on April 23, 1961, in Mission Hills, Los Angeles, California. He was raised by his maternal grandmother, Benita Gutierrez, a factory worker, and his maternal grandfather, Refugio Gutierrez, who worked as a janitor and metal fabricator at an ice skating rink.

Predicted Answer: Mission Hills

Generated Context 4: George Lopez in San Fernando, California. He is of Mestizo Mexican descent. His father, Saul Lopez, was from Sonora, and his mother, Elvira, was a native of Arizona. They were migrant workers who picked crops and moved often, sometimes taking their seven children with them. While growing up, Lopez did not speak English until he was seven. Lopez attended Our Lady of Lourdes, a Roman Catholic elementary school in Pico Rivera.

Predicted Answer: San Fernando

Generated Context 5: George Lopez A theater called "The Tower", which now houses a branch of Bank of America, stands on the former site of his childhood home. Lopez was born at Daniel Freeman Memorial Hospital in Inglewood, California and was raised by his maternal grandmother, Benita Gutierrez, a factory worker who took him to the movies every week. He grew up in the Echo Park neighborhood of Los Angeles, California and attended many different schools in the Los Angeles Unified School District.

Predicted Answer: Los Angeles

Generated Context 6: George Lopez López was born in Mission Hills, Los Angeles, California. He is the son of Frieda and Anatasio López, Mexican immigrants from Tepic, Nayarit, Mexico. Lopez's father was a farm worker, and his mother worked at a canning factory. Lopez has stated that he had a difficult childhood; his parents divorced when he was five, and his mother struggled to raise him and his two sisters.

Predicted Answer: Mission Hills

Generated Context 7: George Edward Lopez. Lopez was born in Mission Hills, Los Angeles. His parents, Frieda and Anatasio Lopez, were both born and raised in Texas. His mother was of Mexican descent and his father, also called "Charlie", was of Sicilian descent. Lopez grew up with two older sisters. He attended San Fernando High School in the Los Angeles neighborhood of San Fernando.

Predicted Answer: Mission Hills

Generated Context 8: George Lopez had in his mouth. He was raised by his mother, Frieda, a telephone company employee, and his grandmother, Benita Gutierrez, a factory worker. Lopez graduated from the Roman Catholic Saint Monica's High School in 1974, where he met his future wife, Ann Serrano. He moved to California and graduated from the California State University, Los Angeles, in 1986 with a degree in communications.

Predicted Answer: Castle Hill

Predicted Answer (CGAP): Mission Hills

Table 11: Example illustrates CGAP, in which we marginalize over 8 generated contexts.

Q: Who was the first U.S. president to be born in the United States of America?

A: with the issue of 1863 (displayed in First appearances) above). At that time, as aforesaid, he had already appeared on two different Confederate 200a2 stamps. Andrew Jackson appears on the following commemorative issues: Martin Van Buren (December 5, 1782 - 2013 July 24, 1862) was the eighth President of the United States from 1837 to 1841. Before his presidency, he served as the eighth Vice President (1833 2013 1837) and the 10th Secretary of State under Andrew Jackson. Van Buren was the first U.S. President to be born an American citizen. William Henry Harrison (February 9, 1773 2013 April 4, 1841) was the

Q: Which US President who was in office 1923-29, was the only one to be born on 4th July?

A: Calvin Coolidge John Calvin Coolidge Jr. (; July 4, 1872 2013 January 5, 1933) was an American politician and the 30th President of the United States from 1923 to 1929. A Republican lawyer from New England, born in Vermont, Coolidge worked his way up the ladder of Massachusetts state politics, eventually becoming governor. His response to the Boston Police Strike of 1919 thrust him into the national spotlight and gave him a reputation as a man of decisive action. Soon after, he was elected Vice President of the United States in 1920, and succeeded to the presidency upon the sudden

Q: What long running comic strip, introduced on Oct 2, 1950, was drawn by the same cartoonist, Charles M. Schultz, during its' entire run?

A: Peanuts Peanuts is a syndicated daily and Sunday American comic strip written and illustrated by Charles M. Schulz that ran from October 2, 1950, to February 13, 2000, continuing in reruns afterward. "Peanuts" is among the most popular and influential in the history of comic strips, with 17,897 strips published in all, making it "arguably the longest story ever told by one human being. At its peak in the mid-to late 1960s, "Peanuts" ran in over 2,600 newspapers, with a readership of around 355 million in 75 countries, and was translated into 21 languages. It helped to cement the"

Q: In the comic strip "Peanuts", whose father is a barber?

A: Franklin (Peanuts) Franklin Armstrong is a character in the long-running comic strip "Peanuts", created by Charles M. Schulz. Introduced on July 31, 1968, Franklin was the first African American character in the strip. He goes to school with Peppermint Patty and Marcie. In his first story arc, he met Charlie Brown when they were both at the beach. Franklin's father was a soldier fighting in Vietnam, to which Charlie Brown replied "My dad's a barber...he was in a war too, but I don't know which one." Franklin later paid Charlie Brown a visit and found some of Charlie Brown's other

Q: Which US President was given the nickname 'The Trust Buster'?

A: Howard Taft. A Progressive reformer, Roosevelt earned a reputation as a "trust buster" through his regulatory reforms and anti-trust prosecutions. His presidency saw the passage of the Pure Food and Drug Act, which established the Food and Drug Administration to regulate food safety, and the Hepburn Act, which increased the regulatory power of the Interstate Commerce Commission. Roosevelt took care, however, to show that he did not disagree with trusts and capitalism in principle, but was only against monopolistic practices. His "Square Deal" included regulation of railroad rates and pure foods and drugs; he saw it as a fair deal

Q: American cartoonist Thomas Nast is credited with creating the traditional image of which popular Christmas character?

A: Thomas Nast Thomas Nast (September 27, 1840 December 7, 1902) was a German-born American caricaturist and editorial cartoonist considered to be the "Father of the American Cartoon". He was the scourge of Democratic Representative "Boss" Tweed and the Tammany Hall Democratic party political machine. Among his notable works were the creation of the modern version of Santa Claus (based on the traditional German figures of Sankt Nikolaus and Weihnachtsmann) and the political symbol of the elephant for the Republican Party (GOP). Contrary to popular belief, Nast did not create Uncle Sam (the male personification of the United

Q: October 2, 1950 saw the debut of what comic strip which ran until creator Charles M. Schulz's death on Feb 12, 2000?

A: from a wide range of acclaimed contemporary artists and designers who have been inspired by the cartoon. Peanuts Peanuts is a syndicated daily and Sunday American comic strip written and illustrated by Charles M. Schulz that ran from October 2, 1950, to February 13, 2000, continuing in reruns afterward. "Peanuts" is among the most popular and influential in the history of comic strips, with 17,897 strips published in all, making it "arguably the longest story ever told by one human being". At its peak in the mid- to late 1960s, "Peanuts" ran in over 2,600 newspapers, with a readership of

Q: In which state was Charles Schulz born?

A: Charles M. Schulz Charles Monroe "Sparky" Schulz (November 26, 1922 February 12, 2000), nicknamed Sparky, was an American cartoonist. Schulz is known for the comic strip "Peanuts" (which featured the characters Charlie Brown and Snoopy, among others). He is widely regarded as one of the most influential cartoonists of all time, cited by cartoonists including Jim Davis, Bill Watterson, and Matt Groening. Born in Minneapolis, Minnesota, Schulz grew up in Saint Paul. He was the only child of Carl Schulz, who was born in Germany, and Dena Halverson, who had Norwegian heritage. His uncle called him "Sparky" after

Q: **Who was President when the first Peanuts cartoon was published?**

Table 12: *Prompt(Q)* Example. For the question "**Who was President when the first Peanuts cartoon was published?**" from TQA (Joshi et al., 2017), we selected 8 $\langle q_i, c_i \rangle$ samples from the supporting repository \mathcal{D} , and construct the *Prompt(Q)* as above. to prompt LMs for c_{gen} generation.