

# AUGMENTING MATH WORD PROBLEMS VIA ITERATIVE QUESTION COMPOSING

Haoxiong Liu<sup>1\*†</sup> Yifan Zhang<sup>1\*</sup> Yifan Luo<sup>2,4</sup> Andrew Chi-Chih Yao<sup>1,3,4</sup>

<sup>1</sup>IIS, Tsinghua University

<sup>2</sup>SCSS, Beijing University of Posts and Telecommunications

<sup>3</sup>Shanghai Artificial Laboratory

<sup>4</sup>Shanghai Qizhi Institute

{liuhx20, zhangyif21}@mails.tsinghua.edu.cn,

andrewcyao@tsinghua.edu.cn

## ABSTRACT

Despite the advancements in large language models (LLMs) for mathematical reasoning, solving competition-level math problems remains a significant challenge, especially for open-source LLMs without external tools. We introduce the MMIQC dataset, comprising a mixture of processed web data and synthetic question-response pairs, aimed at enhancing the mathematical reasoning capabilities of base language models. Models fine-tuned on MMIQC consistently surpass their counterparts in performance on the MATH benchmark across various model sizes. Notably, Qwen-72B-MMIQC achieves a 45.0% accuracy, exceeding the previous open-source state-of-the-art by 8.2% and outperforming the initial version GPT-4 released in 2023. Extensive evaluation results on Hungarian high school finals suggest that such improvement can generalize to unseen data. Our ablation study on MMIQC reveals that a large part of the improvement can be attributed to our novel augmentation method, Iterative Question Composing (IQC), which involves iteratively composing new questions from seed problems using an LLM and applying rejection sampling through another LLM.




Figure 1: Performance evaluation of various LLMs on MATH (Hendrycks et al., 2021a) and the 2023 Hungarian National High School Mathematics Finals (Paster, 2023b).

\*Equal contribution.

†Corresponding author.




Figure 2: The performance of base models and their fine-tuned versions on MATH benchmark. The models remarked with an \* are trained and evaluated by us. We can see that the models fine-tuned on MMIQC consistently outperform their counterparts by a clear margin.

## 1 INTRODUCTION

Although large language models have been demonstrated to be powerful in various applications (Chen et al., 2021; Brown et al., 2020; Ouyang et al., 2022; Park et al., 2023; Huang et al., 2022b), solving math problems that require complex reasoning skills remains a challenging task. On MATH (Hendrycks et al., 2021b), a competition-level math problem benchmark, open-source base LLMs such as the LLaMA family (Touvron et al., 2023a;b) fail to answer most of the problems correctly.

Previous work tries to enhance the mathematical reasoning abilities of base models by fine-tuning them on domain-specific data. Specifically, One line of work (Azerbayev et al., 2023; Lewkowycz et al., 2022) collects math corpora from the web and fine-tunes the models on them, which is also known as the procedure of continual pre-training (Cossu et al., 2022). Another line of work focuses on constructing synthetic data through rejection sampling (Yuan et al., 2023), distilling from GPT-4/GPT-3.5 (Yue et al., 2023) or question bootstrapping (Yu et al., 2023), and then use the generated question-response pairs to perform supervised fine-tuning in the way described in Taori et al. (2023); Ouyang et al. (2022). However, there still exists a large performance gap between these fine-tuned models and the most advanced close-source models such as GPT-4 (OpenAI, 2023) and Gemini-Ultra (Team et al., 2023). Given that simply adding more data does not always lead to better performance as shown in (Yu et al., 2023), how to bridge the gap remains an open challenge.

This work tackles the challenge by combining the two lines of work. On one hand, we reuse the high-quality corpora used in the pre-training stage during fine-tuning. Specifically, MMIQC contains around 1200k question-response pairs we filtered and pre-processed from the web pages at math.stackexchange.com, which are included in the RedPajama dataset (Computer, 2023). On the other hand, for the synthetic data part of MMIQC, we increase the diversity by using multiple kinds of augmentation methods listed below: 1) Prompting GPT-4 with an integrated version of the question bootstrapping prompts used in Yu et al. (2023), and do rejection sampling with GPT-3.5-Turbo on both seed and augmented problems. 2) Using a modified prompt presented in Liu et al. (2023) to ask GPT-4 to generate similar problems with answers given seed problems of the training set of MATH. Although the generated answers can be wrong, we perform rejection sampling on these problems as well. 3) Performing IQC (Iterative Question Composing) with 4 iterations in total. We iteratively ask GPT-4 to compose new questions from the given seed problems and do rejection sampling to filter those problems with answers aligned with GPT-3.5-turbo’s answers. 4) Filtering a 204k subset of MetaMathQA (Yu et al., 2023) and adding it to the MMIQC dataset (More details on MMIQC will be introduced in Section 4).

We fine-tune several base models on MMIQC, resulting in models consistently achieving a large margin compared to their counterparts when evaluated on MATH, as shown in Figure 2. Specifically, the models Mistral-7B-MMIQC, Llemma-34B-MMIQC, DeepSeek-67B-MMIQC and Qwen-72B-MMIQC, which are obtained by fine-tuning Mistral-7B (Jiang et al., 2023), Llemma-34B (Azerbayev et al., 2023) and DeepSeek-67B (Bi et al., 2024) on MMIQC, achieve 36.0%, 38.6%, 41.0% and 45.0% accuracy on MATH, 5.8%, 3.8%, 4.2% and 3.3% higher than the counterpart models that

are fine-tuned on MetaMathQA, respectively. We also evaluate the models on the 2023 Hungarian national high school finals in mathematics Paster (2023a). The results in Figure 1 suggest that the mathematical reasoning abilities the models acquire through being fine-tuned on MMIQC can generalize to unseen held-out problems.

We highlight our contributions as follows:

- We propose IQC (Iterative Question Composing), a data augmentation method that can iteratively generate diverse data starting from a seed dataset of math word problems.
- We release MMIQC, a mixture of processed web data and synthetic question-response pairs. In different model sizes, the models fine-tuned on MMIQC consistently outperform their counterparts by a clear margin on the MATH test set. Notably, Qwen-72B-MMIQC achieves a 45.0% accuracy, exceeding the previous open-source state-of-the-art<sup>1</sup> by 8.2% and outperforming the initial version GPT-4 released in 2023. Such improvement can generalize to unseen held-out data, e.g., Hungarian high school finals.
- Our results show that reusing the high-quality data in the pre-training corpora during the fine-tuning stage can improve the model performance, successfully combining the two lines of work of continual pre-training and supervised fine-tuning.
- Our results also show that using multiple augmentation methods to construct datasets for fine-tuning is an efficient way to boost the performance of LLMs.

## 2 RELATED WORK

**Base Large Language Models.** Base large language models (LLMs) trained on massive corpora (e.g. 1.4T tokens of text for Llama (Touvron et al. (2023a))) from various sources with a simple auto-regressive next token prediction loss have achieved great success in various natural language processing tasks Radford et al. (2019); Brown et al. (2020); Touvron et al. (2023a;b); Jiang et al. (2023). Although these pre-trained models are not intended to serve for solving complex mathematical problems, Wei et al. (2023) show that few-shot prompting can help the models answer a certain fraction of problems correctly. Nevertheless, to achieve better performance, fine-tuning the base LLMs on domain-specific data is required.

**Fine-tuning Base LLMs on Mathematical Datasets.** Current practice of fine-tuning base LLMs on mathematical datasets can be classified into two kinds: 1) continual pretraining (Lewkowycz et al., 2022; Azerbayev et al., 2023). This line of work typically collects billion-tokens level mathematical text data from the web, such as mathematical sub-sites of Stack Exchange and ArXiv, and fine-tune the model in the same way as that in the pre-training stage. 2) SFT (Supervised Fine-Tuning) Yuan et al. (2023); Yu et al. (2023); Yue et al. (2023); Gou et al. (2023). Works in this line collect question-response pairs via various methods and train the models on their dataset in an Alpaca style. Due to the scarcity of publicly available high-quality question-response pairs datasets and the costly nature of manually composing math word problems, how to augment new data from the existing datasets becomes the focus of these works.

Our work is located in the middle between these two: MMIQC is a mixture of filtered pre-training corpus and question-response pairs generated using various augmentation methods.

**Reasoning Frameworks for Solving Mathematical Problems.** Much effort has been devoted to achieving a higher accuracy on math word problem benchmarks by designing different procedures of using the given LLMs to obtain the answers, which we refer to as *reasoning frameworks*. Among them, *Prompting-based* methods Radford et al. (2019); Wei et al. (2023); Fu et al. (2022) play a significant role in activating the potential reasoning abilities for base LLMs through carefully designing the prompts shown to the models. Self-consistency (Wang et al., 2023b) samples multiple rationale paths for a model and then decides the answer by majority voting. In contrast of self-consistency, Cobbe et al. (2021); Uesato et al. (2022); Lightman et al. (2023) use Outcome Reward Models (ORM) and Process Reward Models (PRM) trained on human annotations as verifiers to help select the answer with the highest score from the sampled reasoning paths of LLMs. Getting rid of the need of manual

<sup>1</sup>As of the time of writing in January 2024, to the best of our knowledge, the open-source SOTA on MATH is the DeepSeek-67B-MetaMathQA model reported in Wang et al. (2023a), which achieves 36.8% accuracy without external tool usage.

---

**Algorithm 1:** Iterative Question Composing

---

**Require:** Question composing model  $\pi_q$ , rejection sampling model  $\pi_r$ , answer extractor defining  $\simeq$ , text templater  $x(\cdot, \cdot)$  with inverse  $x^{-1}(\cdot)$ , initial seed dataset  $S_0 = \{(q_i, a_i)\}_{i=1}^n$ , total iterations  $K$ , question composing prompts  $p_1, p_2, \dots, p_K$ , rejection sampling prompt  $p_r$ , maximum rejection samples per problem  $m$

- 1: **for**  $k = 1$  **to**  $K$  **do**
- 2:   Initialize  $S_k \leftarrow \{\}$ ,  $R_k \leftarrow \{\}$
- 3:   **for all**  $(q, a) \in S_{k-1}$  **do**
- 4:     Sample  $x' \sim \pi_q(\cdot | p_k \oplus x(q, a))$
- 5:     Decompose  $(q', a') \leftarrow x^{-1}(x')$
- 6:     Append  $S_k \leftarrow S_k \cup \{(q', a')\}$
- 7:     **for**  $j = 1$  **to**  $m$  **do**
- 8:       Sample  $a^{(j)} \sim \pi_r(\cdot | p_r \oplus q')$
- 9:       **if**  $a^{(j)} \simeq a'$  **then**
- 10:          Append  $R_k \leftarrow R_k \cup \{(q', a^{(j)})\}$
- 11:       **end if**
- 12:     **end for**
- 13:   **end for**
- 14:   Combine  $D_k \leftarrow S_k \cup R_k$
- 15: **end for**

**output** Collections  $D_1, D_2, \dots, D_K$

---

annotation, Wang et al. (2023a) score a given reasoning step by estimating the potential of that step to lead to a correct answer automatically.

Some frameworks also include the use of plug-in tools and external APIs. Program-aided prompting (Gao et al., 2022; Yue et al., 2023) provides in-context samples containing Python codes for LLMs and uses code interpreters to execute the output to facilitate reasoning. Further, Gou et al. (2023) interleave natural language rationales with Sympy code and fine-tune the model on trajectories sampled from GPT-4 to follow their framework in two steps, namely imitation learning and output space shaping.

We note that our results in Figure 2 do not include multiple times of sampling, use of verifiers or code interpreters, thus cannot be directly compared with the results reported in these works.

### 3 ITERATIVE QUESTION COMPOSING

Traditional data augmentation methods primarily concentrate on modifying either the questions or answers while retaining their original meanings, or generating similar problems, as discussed in (Yu et al., 2023) and Liu et al. (2023). These methods, however, are limited in their diversity as they aim to create nearly identical problems. Our approach, termed **IQC (Iterative Question Composing)**, deviates from this by iteratively constructing more complex problems. It augments the initial problems, adding additional reasoning steps without altering their intrinsic logical structure. This ensures that the newly formed problems are organically linked to the original problem and elaborately tries to not include extraneous elements induced by a large transition of the reasoning process.

**Notations.** In our description, we refer to the combination of an LLM, its tokenizer, encoding/decoding methods, and a fixed generation configuration (inclusive of generation strategy, sampling temperature, and stopping criteria) simply as ‘an LLM’. For an LLM  $\pi$ , we denote the output distribution given prompt  $p \in \mathcal{A}^*$  as  $\pi(\cdot | p)$ . The concatenation of two text paragraphs  $p_1$  and  $p_2$  is represented as  $p_1 \oplus p_2$ .

The IQC process begins with specifying an LLM  $\pi_q$  for question composing and another model  $\pi_r$  for rejection sampling. An answer extractor is needed to derive answers from responses. Two responses  $r_1$  and  $r_2$  are considered equivalent, denoted  $r_1 \simeq r_2$ , if the same answer can be extracted from both. The process initiates with a seed dataset  $S_0 = \{(q_i, a_i)\}_{i=1}^n$ .

In iteration #1, we prompt  $\pi_q$  with  $p_1 \oplus x(q, a)$  for each  $(q, a) \in S_0$ , where  $x(\cdot, \cdot)$  is a text template transforming a question-response pair into text, and  $p_1$  solicits a new question-answer composition. This yields a new dataset

$$S_1 = \{(q'_i, a'_i)\}_{i=1}^n,$$

where  $(q'_i, a'_i) = x^{-1}(x'_i)$  and  $x'_i \sim \pi_q(\cdot | p_1 \oplus x_i)$  is the output for the  $i$ th sample. We further enhance  $S_1$  by rejection sampling from  $\pi_r$ , resulting in

$$R_1 := \{(q'_i, a_i^{(j)}) | a_i^{(j)} \simeq a'_i, i \in [n], j \in [m]\},$$

where  $a_i^{(j)}$  are the sampled responses from  $\pi_r(\cdot | p_r \oplus q'_i)$ . The dataset  $D_1$  is then formed by uniting  $S_1$  and  $R_1$ :

$$D_1 := S_1 \cup R_1.$$

For each subsequent iteration # $k$ , the aforementioned procedure is repeated using  $S_{k-1}$  as the seed dataset, with varying question composing prompts  $p_k$ . The complete IQC process is delineated in Algorithm 1.

**Seed Question:**  
Evaluate

$$(5a^2 - 13a + 4)(2a - 3)$$

for  $a = 1\frac{1}{2}$ .

---

**Iter # 1 Question:**  
If  $b = 2a - 3$  and  $a = 1\frac{1}{2}$ , what is the value of  $(5a^2 - 13a + 4)b$ ?

---

**Iter # 2 Question:**  
Given  $b = 2a - 3$ ,  $a = 1\frac{1}{2}$ , and  $c = 3b + 5$ , find the value of  $c(5a^2 - 13a + 4)$ .

---

**Iter # 3 Question:**  
Given  $b = 2a - 3$ ,  $a = 1\frac{1}{2}$ ,  $c = 3b + 5$ , and  $d = c^2 - 4c$ , find the value of  $d + c(5a^2 - 13a + 4)$ .

---

**Iter # 4 Question:**  
Given  $b = 2a - 3$ ,  $a = 1\frac{1}{2}$ ,  $c = 3b + 5$ ,  $d = c^2 - 4c$ , and  $e = d^3 + 2cd - 7$ , find the value of  $e + c(5a^2 - 13a + 4) + d$ .

Figure 3: An example of the questions composed via IQC by GPT-4 given 1 seed problem in MATH training set.

You will be provided with 1 math problem and its solution and answer (which are not guaranteed to be right). Please generate 1 new problem that (implicitly) contains the original problem as a subproblem or substep.

Your response should only contain one line text with 3 fields "problem", "solution" and "answer" in the same format as the given problem. The solution to the generated problem should be as brief as possible and \*\*should not quote the conclusion of the original problem\*\*. Ensure there is only one latex box in the solution and the answer is completely the same with the content in the box.

\*\*Please use two backslashes to represent one in the strings in order that it can be properly read in python.\*\* For example, you should write "`\cdot`" as "`\\cdot`".

Figure 4: The prompt we use to perform question composing in IQC. The brown part is not used in iteration #1.

## 4 THE MMIQC DATASET

In this section, we introduce how each part of MMIQC is constructed in detail.

**Subset of MetaMathQA.** The original MetaMathQA dataset is constructed by sampling GPT-3.5 for  $k = 20$  times under a  $T = 0.7$  temperature for each problem in the training set of MATH (Hendrycks et al., 2021a) and GSM8K (Cobbe et al., 2021) dataset, or its bootstrapped versions. We restrict the number of samples for each completely same question to be 3 and 1 for MATH and GSM8K, respectively, to obtain a subset of MetaMathQA. This subset contains 112.2K GSM8K question-response pairs and 91.5K MATH pairs.

You will be provided with 1 math problem in newline-delimited json format. Please augment 5 diverse problems from the given problem.

The way you augment a problem can be:

- Rephrase the problem.
- Change the scenario without modifying specific quantities.
- Set 1 number in the problem to an unknown variable, put the answer in the problem and ask what is the value of the variable. Ensure the generated problem is reasonable. Otherwise, skip this method.
- Other approaches that can ensure the correctness of the answer you provide to the augmented problem.

Your response should only contain text in newline-delimited json format, keeping the same with the given problem. Please use two backslashes to represent one in the strings.

Figure 5: The prompt we use to perform question bootstrapping for asking GPT-4.

**Answer Augmentation and Question Bootstrapping.** We integrate the question bootstrapping methods used in Yu et al. (2023) into a single prompt shown in Figure 5. Our motivation is that given GPT-4 is highly capable of natural language understanding, a few-shot prompting style used in Yu et al. (2023) might suppress the diversity of the augmented questions. The seed dataset is constructed by the samples in the training set of MATH that do not contain Asymptote language in their question statements. We perform rejection sampling from GPT-3.5 on both the seed dataset and generated questions using the prompt shown in Figure 6, obtaining 66.5K question-response pairs. We use a temperature  $T = 1.0$  for both question bootstrapping and rejection sampling.

**Augmented Similar Problems.** Liu et al. (2023) asks GPT-3.5 to generate 10 similar questions given 1 seed problem in the training set of GSM8K. In our practice, we find that GPT tends to generate several almost the same problems regardless of the given seed problem when asked to generate up to 10 new problems. Thus, we only ask GPT to generate 3 problems (with a solution, for rejection sampling) each time, using the prompt in Figure 7. Considering rejection sampling needs the answer to the problem better to be correct, we use the stronger GPT-4 instead of GPT-3.5. To control the cost, our prompt emphasizes that the solution should be as brief as possible. The total number of the augmented similar problems and the question-response pairs rejection sampled from them is 38.2K. The rejection sampling prompt is the same one in Figure 6 as well. We use a temperature  $T = 1.0$  for both procedures.

**Iterative Question Composing.** We perform Iterative Question Composing for 4 iterations as described in Section 3. Specifically, we use GPT-4 for question composing model  $\pi_q$  with a  $T = 0.7$  temperature and GPT-3.5 for rejection sampling model  $\pi_r$  with a  $T = 1.0$  temperature. The question composing prompts and rejection sampling prompt are shown in Figure 4 and Figure 6, respectively. The text templater  $x(\cdot, \cdot)$  we use is a program that transforms each question-response pair into JSON text format, with fields ‘problem’ and ‘solution’. The resulting dataset has 55.1K samples in total.<sup>2</sup> We provide an example of the generated questions in different iterations corresponding to the same seed problem in Figure 3. We note that although some of the questions are not rigorously a sub-problem or sub-step of the corresponding problem in the previous iteration as required in our prompt, they are still valid questions that can increase the diversity of the dataset.

<sup>2</sup>A part of the samples are generated by performing IQC for 2 iterations using a legacy version of prompts.

You will be presented a mathematical problem. You should solve the problem step-by-step carefully. Present the final answer in latex boxed format, e.g.,  $\boxed{63\pi}$ .

Figure 6: The prompt we use to do rejection sampling from GPTs.

You will be provided with 1 math problem in newline-delimited json format. Please generate 3 diverse new problems similar to the given problem.

Your response should only contain text in newline-delimited json format, keeping the same with the given problem. The solutions to the generated problems should be as brief as possible. Ensure there is only one box in the solution and the answer is completely the same with the content in the box. Please use two backslashes to represent one in the strings.

Figure 7: The prompt we use to generate questions similar to the seed problems for asking GPT-4.

**Mathematics Stack Exchange.** We observe that in the OpenWebMath (Paster et al., 2023) dataset, the data from Mathematics Stack Exchange shows high quality and is most related to competition-level math. Motivated by this, we extract the data collected from Mathematics Stack Exchange in RedPajama (Computer, 2023) and pre-process it into question-response pairs. For each Mathematics Stack Exchange page, we only retain the answer ranked first by RedPajama. Then we filter out the answer that does not contain a formula environment symbol ‘\$’. This results in a dataset with 1203.6K question-response pairs.

Table 1 shows the make-up of MMIQC. When fine-tuning the models MMIQC contains 3 repetitions of the subsets mentioned above, except for the Mathematics Stack Exchange part. We shuffle the order of samples after combining the subsets.

## 5 EXPERIMENTS

### 5.1 FINE-TUNING SETUP

Our fine-tuning strategy mainly follows the practice of Taori et al. (2023), except that we use a different prompt template to transform the question-response pairs. For a sample from Mathematics Stack Exchange, the corresponding prompt fed into the model during training is a simple concatenation of the question and response with two new-line symbols. For a sample from other subsets, we additionally add a prefix ‘Please solve the following problem and put your answer at the end with ‘The answer is: ’.’ to the question-response concatenation.

We use the HuggingFace transformers library Wolf et al. (2019) for our fine-tuning experiments. We fine-tune all models on MMIQC for 1 epoch, using a 3% warm-up ratio linear learning rate schedule. For the choice of maximum learning rate, we do a simple hyper-parameter selection experiment shown in Table 2 and determine it to be  $1e-5$ . We use the BFloat16 numerical format during training. Employing the DeepSpeed Zero-3 Stage Rajbhandari et al. (2020), we fine-tune 7B models on one node of 8xA800 GPUs with micro batch-size at 8, and gradient accumulation at 4, 34B models on 2 nodes with micro batch-size at 4 and gradient accumulation at 4 and  $\sim 70B$  models on 4 nodes with micro batch-size at 4 and gradient accumulation at 2, maintaining an effective batch size of 256. It takes around 14 hours, 61 hours and 90 hours to fine-tune 7B, 34B and  $\sim 70B$  models under the setups stated above, respectively.

### 5.2 MODEL EVALUATION

For a fair comparison, we first evaluate the fine-tuned models on MATH (Hendrycks et al., 2021a), a competition-level math word problems benchmark with 5000 test problems in a **zero-shot** setting. We prompt all our fine-tuned models with the test question with the prefix ‘Please solve the following problem and put your answer at the end with ‘The answer is: ’.’, and extract the answer from the output using a modified version of the answer extractor provided in Lewkowycz et al. (2022). We

Table 1: The composition of MMIQC.

DATA	# SAMPLES	#REPETITIONS	RATIO
METAMATHQA	203.7K	3	26.6%
ANSAUG & QB	66.5K	3	8.7%
AUGSIMILAR	38.2K	3	5.0%
<b>IQC</b>	<b>55.1K</b>	3	<b>7.2%</b>
MATHSTEX	1203.6K	1	52.5%

Table 2: Ablation study on the optimal learning rate. We fine-tune Mistral-7B on MMIQC with different maximal learning rate values and evaluate the fine-tuned models on MATH to decide the best candidate.

LR	1E-6	5E-6	<b>1E-5</b>	2E-5	5E-5	1E-4
MATH(%)	32.3	35.1	<b>36.0</b>	35.4	31.5	27.1

use a series of rules to infer whether the extracted answer is the same as the ground-truth answer, including a comparison using SymPy Meurer et al. (2017). The complete results of our evaluation on MATH and a comparison with existing models are shown in Table 3.

For the evaluation on 2023 Hungarian national high school finals in mathematics, we use the few-shot prompt used in Paster (2023a). We manually assess the grades for every model according to the examiner instructions. The results shown in Figure 1 are the grades under a full mark of 117.

### 5.3 ABLATION STUDY ON SUBSETS OF MMIQC

In order to understand the ratio of contribution to the improvement revealed in Table 3 of different subsets of MMIQC, we fine-tune Mistral-7B with a series of training sets constructed by gradually adding the subsets. When MathStackExchange is not added, we fine-tune for 3 epochs. When MathStackExchange is added to the training dataset, we mix 3 repetitions of other data with 1 repetition of the MathStackExchange, and fine-tune for only 1 epoch. It can be seen from Table 4 that

- Although our filtered subset of MetaMathQA is only half the size of the original dataset (which has 395K samples, more than the total number of samples of our synthetic data), the performance drop is only 1.8%. This shows that the  $k = 20$  strategy in Yu et al. (2023) results in some redundancy.
- Our Answer Augmentation & Question Boosting data help the fine-tuned model beat Mistral-7B-MetaMathQA, verifying our hypothesis that directly asking GPT to perform question bootstrapping is more efficient than providing few-shot examples to them.
- Our IQC method leads to a significant 3.1% improvement from a high accuracy of 31.5% with only 55.1K samples, showing its efficiency. Moreover, the later iterations of IQC also account for a certain ratio of improvement, proving that IQC is a method that can continuously generate new data that can help increase the diversity when added to the data generated in previous iterations.

### 5.4 CONTAMINATION TEST

We check the  $n$ -gram matches for MMIQC to ensure that the improvement is not a result of direct memorization. We use the script provided by Azerbayev et al. (2023) to check the  $n$ -gram matches between the synthetic part of the MMIQC and MATH test set. It turns out that for a 30-gram match check, there are 44 hits of match between the ‘solution’ field of MATH test set and the ‘output’ field of MMIQC, far fewer than the 168 hits between that of MATH test set and MATH training set. Moreover, we manually check these 44 hits and find that 43 among them belong to the case where intermediate steps of the solutions to similar but different questions collide, with the only exception being the question ‘A regular polygon has interior angles of 144 degrees. How many sides does the polygon have?’. This almost rules out the possibility that fine-tuned models get memorization of solutions to the problems in the test set, indicating a very low risk of data contamination for MMIQC.



Table 3: A comparative analysis of the accuracies achieved by various models on the MATH benchmark. The models marked with an asterisk(\*) are fine-tuned and evaluated by us. Other results, unless otherwise cited, are derived from Wang et al. (2023a). This comparison highlights the significant improvements our fine-tuned models demonstrate over existing solutions in mathematical problem-solving accuracy.

MODEL	FT-DATASET	TOOL USAGE?	EVAL METHOD	MATH(%)
<b>PROPRIETARY MODELS</b>				
MINERVA-540B (UESATO ET AL., 2022)	ARXIV+WEB	No	MAJ1@64	50.3
GPT-4 (2023-0314) (BUBECK ET AL., 2023)	-	No	PASS@1	42.5
GEMINI-ULTRA (TEAM ET AL., 2023)	-	No	PASS@1	53.2
<b>~7B MODELS</b>				
LLAMA-2-7B (TOUVRON ET AL., 2023B)	-	No	PASS@1	2.5
QWEN-7B (BAI ET AL., 2023)	-	No	PASS@1	11.6
LLEMMA-7B (AZERBAYEV ET AL., 2023)	PROOF-PILE-2	No	PASS@1	18.0
METAMATH-7B (YU ET AL., 2023)	METAMATHQA	No	PASS@1	19.8
MISTRAL-7B-METAMATHQA (YU ET AL., 2023)	METAMATHQA	No	PASS@1	28.2
<b>MISTRAL-7B-MMIQC*</b>	<b>MMIQC</b>	No	PASS@1	<b>36.0</b>
MAMMOTH-CODER-7B (YUE ET AL., 2023)	MATHINSTRUCT	CODE	PASS@1	35.2
TORA-CODE-7B (GOU ET AL., 2023)	TORA-CORPUS	CODE	PASS@1	44.6
<b>~34B MODELS</b>				
CODELLAMMA-34B	-	CODE	PASS@1	25.0
LLEMMA-34B-METAMATHQA	METAMATHQA	No	PASS@1	34.8
<b>LLEMMA-34B-MMIQC*</b>	<b>MMIQC</b>	No	PASS@1	<b>38.6</b>
LLEMMA-34B-METAMATHQA	METAMATHQA	MATH-SHEPHERD	MAJ+VERIFY1@256	47.3
TORA-CODE-34B (GOU ET AL., 2023)	TORA-CORPUS	CODE	PASS@1	50.8
<b>~70B MODELS</b>				
LLAMA-2-70B (TOUVRON ET AL., 2023B)	-	No	PASS@1	13.5
DEEPSEEK-67B (BI ET AL., 2024)	-	No	PASS@1	18.7
DEEPSEEK-67B-METAMATHQA	METAMATHQA	No	PASS@1	36.8
<b>DEEPSEEK-67B-MMIQC*</b>	<b>MMIQC</b>	No	PASS@1	<b>41.0</b>
DEEPSEEK-67B-METAMATHQA	METAMATHQA	No	MAJ1@256	45.4
DEEPSEEK-67B-METAMATHQA	METAMATHQA	MATH-SHEPHERD	MAJ+VERIFY1@256	48.1
QWEN-72B (BAI ET AL., 2023)	-	No	PASS@1	35.2
QWEN-72B-METAMATHQA*	METAMATHQA	No	PASS@1	41.7
<b>QWEN-72B-MMIQC*</b>	<b>MMIQC</b>	No	PASS@1	<b>45.0</b>

Table 4: How different subsets of MMIQC affect the accuracy of the finetuned model on MATH. The base model is Mistral-7B, and the accuracy baseline is the reported 28.2% of Mistral-7B-MetaMathQA.

DATA	# SAMPLES	MATH(%)
METAMATHQA	395K	28.2
METAMATHQA (SUBSET)	203.7K	26.4 (-1.8)
+ ANSAUG & QB	+66.5K	30.1 (+1.9)
+ AUGSIMILAR	+38.2K	31.5 (+3.3)
+ IQC ITER #1	+21.8K	33.0 (+4.8)
+ IQC ITER #2	+16.0K	33.7 (+5.5)
+ IQC ITER #3 & #4	+17.3K	34.4 (+6.2)
+ MATHSTACKEXCHANGE	+1203.6K	36.0 (+7.8)

## 6 CONCLUSION

In this work, we introduce a novel data augmentation method for math word problem datasets called IQC (Iterative Question Composing) and use it in the construction of our MMIQC dataset. Our evaluation results show that the models fine-tuned on MMIQC achieve new SOTAs on the MATH benchmark. The improvements of our models benefit from the diverse data sources of MMIQC and the effectiveness of IQC.

For future directions, we are interested in how to equip open-source models with the ability to compose questions, in order to perform IQC in a self-evolution style, similar to that in Huang et al. (2022a). Besides, how to integrate the verification systems (Wang et al., 2023a; Liu et al., 2023) that are originally used to improve the accuracy during inference time into the procedure of IQC, is also an attractive topic.

## REFERENCES

- Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen McAleer, Albert Q Jiang, Jia Deng, Stella Biderman, and Sean Welleck. Llemma: An open language model for mathematics. *arXiv preprint arXiv:2310.10631*, 2023.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, et al. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*, 2024.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *ArXiv*, abs/2005.14165, 2020.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrike, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Together Computer. Redpajama: An open source recipe to reproduce llama training dataset, April 2023. URL <https://github.com/togethercomputer/RedPajama-Data>.
- Andrea Cossu, Tinne Tuytelaars, Antonio Carta, Lucia Passaro, Vincenzo Lomonaco, and Davide Bacciu. Continual pre-training mitigates forgetting in language and vision. *arXiv preprint arXiv:2205.09357*, 2022.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting for multi-step reasoning. *arXiv preprint arXiv:2210.00720*, 2022.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. Pal: Program-aided language models. *arXiv preprint arXiv:2211.10435*, 2022.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yujiu Yang, Minlie Huang, Nan Duan, Weizhu Chen, et al. Tora: A tool-integrated reasoning agent for mathematical problem solving. *arXiv preprint arXiv:2309.17452*, 2023.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2021a.

- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *NeurIPS*, 2021b.
- Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. Large language models can self-improve. *arXiv preprint arXiv:2210.11610*, 2022a.
- Wenlong Huang, P. Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. *ArXiv*, abs/2201.07207, 2022b. URL <https://api.semanticscholar.org/CorpusID:246035276>.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.
- Aitor Lewkowycz, Anders Johan Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Venkatesh Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. *arXiv preprint arXiv:2305.20050*, 2023.
- Bingbin Liu, Sebastien Bubeck, Ronen Eldan, Janardhan Kulkarni, Yuanzhi Li, Anh Nguyen, Rachel Ward, and Yi Zhang. Tinygsm: achieving 80% on gsm8k with small language models. *arXiv preprint arXiv:2312.09241*, 2023.
- Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, AMiT Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz. Sympy: symbolic computing in python. *PeerJ Computer Science*, 3:e103, January 2017. ISSN 2376-5992. doi: 10.7717/peerj-cs.103. URL <https://doi.org/10.7717/peerj-cs.103>.
- OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023. URL <https://api.semanticscholar.org/CorpusID:257532815>.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*, 2022.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–22, 2023.
- Keiran Paster. Testing language models on a held-out high school national finals exam. [https://huggingface.co/datasets/keirp/hungarian\\_national\\_hs\\_finals\\_exam](https://huggingface.co/datasets/keirp/hungarian_national_hs_finals_exam), 2023a.
- Keiran Paster. Testing language models on a held-out high school national finals exam. [https://huggingface.co/datasets/keirp/hungarian\\_national\\_hs\\_finals\\_exam](https://huggingface.co/datasets/keirp/hungarian_national_hs_finals_exam), 2023b.
- Keiran Paster, Openwebmath: An open dataset of high-quality mathematical web text. *arXiv preprint, forthcoming*, 2023.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 2019.

- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–16. IEEE, 2020.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca), 2023.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023a.
- Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony S. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *ArXiv*, abs/2307.09288, 2023b. URL <https://api.semanticscholar.org/CorpusID:259950998>.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process- and outcome-based feedback, 2022.
- Peiyi Wang, Lei Li, Zhihong Shao, RX Xu, Damai Dai, Yifei Li, Deli Chen, Y Wu, and Zhifang Sui. Math-shepherd: A label-free step-by-step verifier for llms in mathematical reasoning. *arXiv preprint arXiv:2312.08935*, 2023a.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=1PL1NIMMrw>.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.

Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T. Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*, 2023.

Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. Scaling relationship on learning mathematical reasoning with large language models. *arXiv preprint arXiv:2308.01825*, 2023.

Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. Mammoth: Building math generalist models through hybrid instruction tuning. *arXiv preprint arXiv:2309.05653*, 2023.