

MuggleMath: Assessing the Impact of Query and Response Augmentation on Math Reasoning

Anonymous ACL submission

Abstract

In math reasoning with large language models (LLMs), fine-tuning data augmentation by query evolution and diverse reasoning paths is empirically verified effective, profoundly narrowing the gap between open-sourced LLMs and cutting-edge proprietary LLMs. In this paper, we conduct an investigation for such data augmentation in math reasoning and are intended to answer: (1) What strategies of data augmentation are more effective; (2) What is the scaling relationship between the amount of augmented data and model performance; and (3) Can data augmentation incentivize generalization to out-of-domain mathematical reasoning tasks? To this end, we create a new dataset, AugGSM8K, by complicating and diversifying the queries from GSM8K and sampling multiple reasoning paths. We obtained a series of LLMs called MuggleMath by fine-tuning on subsets of AugGSM8K. MuggleMath substantially achieves new state-of-the-art on GSM8K (from 54% to 68.4% at the scale of 7B, and from 63.9% to 74.0% at the scale of 13B). A log-linear relationship is presented between MuggleMath’s performance and the amount of augmented data. We also find that MuggleMath is weak in out-of-domain math reasoning generalization to MATH, which suggests that augmentation on a single math subject could not help with overall math reasoning performance.

1 Introduction

The emergence of large language models (LLMs) (Ouyang et al., 2022; Anil et al., 2023; OpenAI, 2023) has profoundly revolutionized the field of natural language processing, exhibiting versatile performance in various tasks like code generation (Chen et al., 2021; Luo et al., 2023b), instruction following (Longpre et al., 2023), long context understanding (Tworkowski et al., 2023), and math reasoning (Wei et al., 2022; Taylor et al., 2022; Lewkowycz et al., 2022a). Math reasoning as

a representative reasoning task is widely studied to access the reasoning abilities in LLMs (Cobbe et al., 2021; Hendrycks et al., 2021). Proprietary LLMs, such as GPT-3.5, and GPT4 (OpenAI, 2023) have shown exceptional mathematical reasoning abilities, while there remains a substantial gap between open-source LLMs, such as GPT-J (Wang and Komatsuzaki, 2021) and LLaMA (Touvron et al., 2023a,b)) and the cutting-edge proprietary models.

To enable better mathematical reasoning abilities in open-sourced LLMs, they generally undergo a fine-tuning stage on supervised reasoning datasets. A series of efforts are committed to enhancing the mathematical reasoning capabilities of open-source LLMs, where a mainstream approach involves first augmenting new mathematical problems and answers, followed by supervised fine-tuning on the augmented dataset (Yuan et al., 2023a; Luo et al., 2023a; Yu et al., 2023). This type of approach has achieved good results, and in this paper, we would like to explore what are the key factors affecting the effectiveness of data augmentation for mathematical reasoning tasks and the scaling relationship between the amount of data augmentation and model performance. Specifically, with the help of proprietary models (GPT-3.5 and GPT-4), we applied five types of mathematical problem augmentation methods based on human experience in creating variations of mathematical problems similar to Luo et al. (2023b,a). We further generated multiple reasoning paths for each augmented problem since distinct reasoning paths can also enhance chain-of-thought reasoning (Huang et al., 2022; Zhu et al., 2023a; Yuan et al., 2023a). We obtained a new dataset called AugGSM8K after data augmentation on a widely used mathematical reasoning dataset GSM8K (Cobbe et al., 2021). By supervised fine-tuning on the open-source LLaMA (Touvron et al., 2023a) and LLaMA-2 (Touvron et al., 2023b) LLMs on different subsets of Aug-

GSM8K, we obtained a series of models dubbed MuggleMath. We find that with sufficient amounts of data, MuggleMath achieves a new state-of-the-art on GSM8K. In addition to this, we find a log-linear relationship between the performance of MuggleMath and the amount of data augmentation over a range of data volumes.

Although MuggleMath achieves strong performance on the GSM8K test set, the rationales for performance improvement by data augmentation remain unclear. We are therefore interested in the specific reason behind the performance improvement and whether it brings enhancement in LLMs’ mathematical reasoning capabilities generally.

To validate the generalization of MuggleMath, we conduct multi-task learning and analyze the transferability with AugGSM8K and MATH. We found that LLMs trained with supervised learning after data augmentation on GSM8K only bring marginal improvements to performance on MATH (Hendrycks et al., 2021). By visualizing the data distribution in the embedding space of LLaMA-2-7B, we observe that the embedding distribution of problems in AugGSM8K is very close to that of GSM8K, but significantly different from the problem distribution in the MATH dataset. This is the fundamental reason why performance improvements from data augmentation on GSM8K cannot be generalized to MATH.

The main contributions of our work can be summarized as follows:

- By augmenting GSM8K with various queries and multiple reasoning paths, we curated GSM8K to a new dataset named AugGSM8K.
- We utilize AugGSM8K for fine-tuning the LLaMA models to obtain MuggleMath, which greatly improves the in-domain performance of the open-sourced LLMs on GSM8K and achieves new state-of-the-art performances.
- We find a log-linear relationship between the accuracy of the model on the test set and the amount of data augmentation within a certain range while the coefficient is similar to augmenting new human-written samples.
- We demonstrate that the performance gains from data augmentation on GSM8K are difficult to generalize to out-of-domain dataset MATH due to distribution differences.

2 Related Works

Mathematical Reasoning for Large Language Models Mathematical reasoning is a crucial ability to examine large language models (Cobbe et al., 2021; Hendrycks et al., 2021; Wei et al., 2022; Yuan et al., 2023b). The mathematical reasoning ability of LLMs can be enhanced by math-related pre-training (Hendrycks et al., 2021; Lewkowycz et al., 2022a; Taylor et al., 2022; Lightman et al., 2023) and math-related supervised fine-tuning (Yuan et al., 2023a; Luo et al., 2023a; Yue et al., 2023; Yu et al., 2023). Query augmentation (Luo et al., 2023a; Yu et al., 2023) and response augmentation (Huang et al., 2022; Zelikman et al., 2022; Ni et al., 2023; Zhu et al., 2023b; Yuan et al., 2023a) are useful techniques to improve math in-domain performances during SFT. Query augmentation methods usually generate rephrased, easier, or harder problems and use proprietary LLMs to generate answers. Response augmentation methods generate new reasoning paths for problems in the training set. They could rely on answers in the training set to filter the generated reasoning paths. (Yuan et al., 2023a) invests the scaling relationship on supervised LLMs math performance with pre-train loss, supervised data amount, and augmented reasoning path amount. MetaMath (Yu et al., 2023) is a contemporary work that is similar to us in the augmentation method. The distinction lies in MetaMath’s focus on rewriting original questions to create new ones using the questions’ mathematical relationships. In contrast, our efforts are centered on generating new problems with equal or greater difficulty levels. Our method is capable of generating higher-quality data as evidenced by the fact that our fine-tuning performance on AugGSM8K(150K) surpasses that of MetaMath’s MetaMathQA-GSM8K(240K). Moreover, our work investigates the quantitative relationship between query and response augment amounts and in-domain and out-of-domain performances.

Data Augmentation for LLM Data augmentation is a common technique to improve downstream task performance in NLP (Feng et al., 2021). In the era of large language models, data augmentation is usually used for generating instruction following SFT datasets (Wang et al., 2023b; Taori et al., 2023). Queries (Ding et al., 2023; Xu et al., 2023) and responses (Mukherjee et al., 2023) of SFT datasets can both be augmented by prompting state-of-the-art proprietary LLMs. Compared

with their work, we are concentrated on augmenting math SFT dataset and we are more interested in scaling relationships on in-domain and out-of-domain generalizations.

Out-of-Distribution Generalization The challenge of out-of-distribution (OOD) generalization has garnered widespread attention across various domains (Karras et al., 2018; Wang et al., 2021; Zhou et al., 2023) in machine learning. This issue arises when the distribution of data encountered by a model during testing diverges from that of the training phase, leading to a decline in model performance. The OOD problem is multifaceted (Lipton et al., 2018; Schölkopf et al., 2012; Tran et al., 2022; Cai et al., 2023), with subcategories such as covariate shifts and concept shifts, among others. To mitigate the effects of OOD scenarios, a diverse array of strategies has been developed, including unsupervised domain generalization (Wang et al., 2021; Zhou et al., 2023), stable learning (Shen et al., 2020; Kuang et al., 2020), invariant representation learning (Creager et al., 2021), causal learning (Peters et al., 2015), and invariant risk minimization (Mao et al., 2023) and etc. Recent trends in the community have shown a growing preference for performance enhancement through data augmentation during the Self-supervised Fine-tuning (SFT) stage in large-scale models. However, the extent of OOD issues associated with this method and their severity remain underexplored. This study aims to fill this gap by conducting empirical experiments and providing a visual analysis to assess the impact of data augmentation on OOD generalization in the context of large models.

3 Experiments

We first introduce our experimental setup (3.1) and dataset augmentation method (3.2). Then we conduct analyses spanning several aspects of data augmentation to answer the three research questions in abstract.

3.1 Experimental Setup

Problem Definition We define the math reasoning SFT dataset as $\mathcal{D} = \{q_i, a_i\}_i$, where q_i is a question and a_i is a reasoning path with an answer. We augment the SFT dataset to a new dataset $\mathcal{D}' = \{q'_i, a'_i\}_i$. We apply SFT on the pre-trained language models and measure the augmented SFT dataset based on the accuracy of the in-domain

test set \mathcal{D}_{in} and out-of-domain test set \mathcal{D}_{out} . We calculate the accuracy based on greedy decoding.

In-domain Dataset GSM8K (Cobbe et al., 2021) is a dataset with elementary school math word problems with 7,473 training problems and 1,319 testing problems. This is viewed as our in-domain dataset \mathcal{D} and we will augment and train on it. The test set of GSM8K is viewed as \mathcal{D}_{in} .

Out-of-domain Dataset MATH (Hendrycks et al., 2021) is a dataset with challenging high-school math problems. Problems are classified into the following topics: Prealgebra, Algebra, Number Theory, Counting and Probability, Geometry, Intermediate Algebra, and Precalculus. Problems in MATH are harder and more diverse than in GSM8K. We use 500 test problems from (Lightman et al., 2023) as our out-of-domain test dataset \mathcal{D}_{out} . The reason we do not use MATH as the in-domain dataset is that it is hard to apply query and response augmentation for MATH-level problems since GPT-4 can only have an accuracy of 42.

Training We employ state-of-the-art open-source LLMs for fine-tuning, including LLaMA-1 7B (Touvron et al., 2023a), LLaMA-2 7B, LLaMA-2 13B, and LLaMA-2 70B (Touvron et al., 2023b), all of which undergo full fine-tuning. We adopt system prompt from (Taori et al., 2023) for fine-tuning and listed in A. We use AdamW for optimization. The training proceeds for three epochs with a learning rate of $2e-5$, a warmup ratio of 0.03, and a cosine learning rate scheduler. We do not apply early stops to choose checkpoints. In the training phase, to train the models that achieved the best results presented in 4 for MuggleMath, the computational resources utilized are as follows: the 7B, 7B-2, and 13B-2 models were trained for 5.5 hours, 5.5 hours, and 13.5 hours, respectively. For the 70B-2 model, we trained it for 20 hours

3.2 Dataset Augmentation

Query Augmentation To generate new queries, we use gpt-3.5-turbo-0613 and gpt-4-0613 as the generator. Inspired by Evol-Instruct (Luo et al., 2023b,a), we find that the diversity and complexity of queries in augmented datasets play a vital role in improving math reasoning benchmark performance. We employ human knowledge in modifying mathematical problems for query augmentation. Below are five query augmentation methods used in our experiments: **Change specific numbers; Introduce**

fractions or percentages; Combine multiple concepts; Include a conditional statement; Increase the complexity of the problem. The examples and detailed prompts we used for query augmentation are listed in B. The examples of augmented queries are shown in 6.

Response Augmentation Instead of using the trained SFT model proposed in (Yuan et al., 2023a), we use gpt-3.5-turbo-0613 and gpt-4-0613 to augment more reasoning paths. The main reason is that we can not filter out wrong reasoning paths without final answers. Thus we need to use a model that is as accurate as possible which is the state-of-the-art LLMs *ChatGPT*. We use a 1-shot prompt to ensure augmented response formats. The response prompt we used for query augmentation is listed in C. Augmented responses can result in some unconventional answers, such as excessively long reasoning paths and reasoning paths that do not contain an answer at their end. We devise manual rules to filter out these corresponding query-response pairs and manual rules are detailed in D. The examples of augmented responses are shown in 7.

Augmented Dataset The original GSM8K training set has 7,473 samples. We augment 5 more queries for each query in the training set and yield $7,473 \times 5 = 37,365$ augmented queries. We run this query augmentation three times with $\mathcal{D}_1, \mathcal{D}_3$ by GPT-3.5 and \mathcal{D}_2 by GPT-4, and $\|\mathcal{D}_i\| = 37,365$. Then we generate one response for each augmented query for \mathcal{D}_i and apply response filtering. We consider the query-response pairs after filtering as \mathcal{D}_i^j . We obtain approximately 30,000 query-response pairs for each \mathcal{D}_i^j . To explore the performance differences of different augmented settings, we generate five responses on the augmented queries \mathcal{D}_1 with GPT-4’s temperature set to 1.0 ($\mathcal{D}_1^1 \sim \mathcal{D}_1^5$), one response with GPT-4’s temperature set to 0.0 (\mathcal{D}_1^6), and one response with GPT-3.5’s temperature set to 1.0 (\mathcal{D}_1^7). We also try a zero-shot response generation named \mathcal{D}_1^8 . We use GPT-4 to augment responses as $\mathcal{D}_2^1, \mathcal{D}_3^1$. Since \mathcal{D}_2 is significantly larger than other subsets, we downsample it to $\hat{\mathcal{D}}_2^1$. We refer to the union of all augmented data and the original GSM8K training set as Aug-GSM8K, upon which we conduct experiments using various subsets. Detailed augmented dataset notations are listed in 1.

Subset	Query	Response	Temp.	Size (K)
\mathcal{D}	-	-	-	7.5
$\mathcal{D}_1^1 \sim \mathcal{D}_1^5$	GPT-3.5	GPT-4	1	30
\mathcal{D}_1^6	GPT-3.5	GPT-4	0	30
\mathcal{D}_1^7	GPT-3.5	GPT-3.5	1	25
\mathcal{D}_1^8	GPT-3.5	GPT-4	1	30
\mathcal{D}_2^1	GPT-4	GPT-4	0	35
$\hat{\mathcal{D}}_2^1$	GPT-4	GPT-4	0	30
\mathcal{D}_3^1	GPT-3.5	GPT-4	1	30

Table 1: The description of different subsets of the augmented in-domain dataset AugGSM8K.

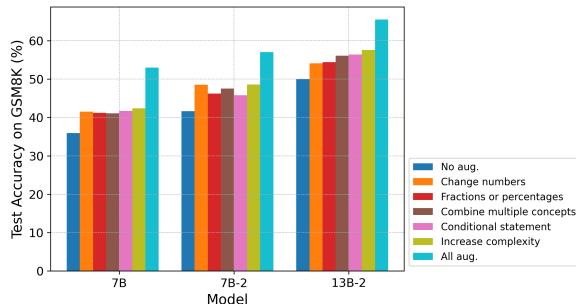


Figure 1: Comparison of test set accuracy on GSM8K for models of varying scales after fine-tuning on Aug-GSM8K subsets with different query augmentation strategies.

3.3 RQ1. What Strategies of Data Augmentation are More Effective

Query Augmentation Types We want to examine whether query augmentation works for math reasoning SFT since (Luo et al., 2023a) applies PPO which cannot make an apple-to-apple comparison. Each query in the original training dataset is augmented with 5 different types. We cluster these queries based on the query types. We apply SFT on the original training set (\mathcal{D}), each query type augmentation, and a combination of them ($\mathcal{D} + \mathcal{D}_1^1$). Results are shown in 1 and 10. Compared with no augment, each query augment method can improve the in-domain performance. **Increase complexity** augmentation method improves most among all of them. This suggests that enhancing the complexity of queries is one of the key factors influencing the sample efficiency of data augmentation. Combining five augment methods yields the best in-domain performance which proves the effectiveness of Evol-instruct under a fair setting.

Query and Response Sources Here we examine how the query and response quality influence the augmented model performance. We list results in 2, and draw the following conclusions: (a) Compar-

Model	7B	7B-2	13B-2
\mathcal{D}	35.9	41.6	50.0
$+\mathcal{D}_1^1 \times 0.8$	51.1	56.6	63.2
$+\mathcal{D}_1^1$	53.0	57.0	65.5
$+\mathcal{D}_1^6$	51.6	58.0	63.8
$+\mathcal{D}_1^7$	41.3	46.7	52.8
$+\mathcal{D}_1^8$	49.4	53.3	62.2
$+\hat{\mathcal{D}}_2^1$	52.3	57.8	63.3

Table 2: Performance of subsets of AugGSM8K with different query and response sources. $+\mathcal{D}_1^1$ is an omission of $\mathcal{D} + \mathcal{D}_1^1$, and the same notation is used in other tables in this paper.

ing $\hat{\mathcal{D}}_2^1$ and \mathcal{D}_1^6 , we find that the queries generated by GPT-4 and GPT-3.5 have no significant impact on SFT performance. (b) Comparing \mathcal{D}_1^1 and \mathcal{D}_1^6 , we can conclude that when using GPT-4 to generate responses, the temperature has no significant impact on SFT performance. (c) Comparing \mathcal{D}_1^1 and \mathcal{D}_1^8 , we can conclude that, compared to the zero-shots generation method, the response augmentation prompt we propose plays a substantial role in enhancing the quality of the generated data (+3.6 for LLaMA-7B, +3.7 for LLaMA-2-7B, +3.3 for LLaMA-2-13B). The main reason we consider this is our 1-shot setting stabilizes the response format. (d) Comparing \mathcal{D}_1^7 (25K) and $\mathcal{D}_1^1 \times 0.8$ (24K), we can conclude that, compared to GPT-3.5, the response augmented using GPT-4 yields significantly better results for SFT.

3.4 RQ2. What is the Scaling Relationship between the Amount of Augmented Data and Model Performance

Query Augmentation Amount We examine how query augmentation amount affects the in-domain performance. We examine seven data volume configurations including partitioning \mathcal{D}_1^1 into proportions of 0, 0.2, 0.4, 0.6, 0.8, 1.0, as well as $\mathcal{D}_1^1 + \hat{\mathcal{D}}_2^1$ and $\mathcal{D}_1^1 + \hat{\mathcal{D}}_2^1 + \mathcal{D}_3^1$ as the augmented datasets. Each augmented query only has one augmented response. They are mixed with GSM8K \mathcal{D} to apply SFT. From 2 and 11, we can find that within the data volume range of 13-97K, the in-domain performance exhibits a log-linear relationship with the query amount. We employ linear regression to approximate this relationship. As shown in 3, pre-training models with better initial math reasoning capabilities exhibit a smaller slope which is consistent with (Yuan et al., 2023a). This suggests it is harder to improve reasoning ability for a better pre-trained model. We also conduct validations on our fitted scaling law with an interpolate

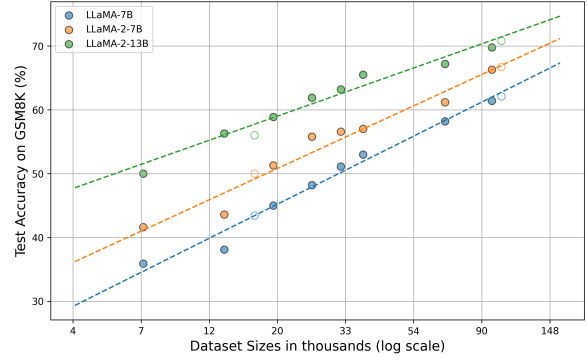


Figure 2: Comparison of test set accuracy on GSM8K for models of varying scales after fine-tuning on AugGSM8K subsets with different query augmentation amount.

point at a query amount of 17K ($\mathcal{D} + \mathcal{D}_1^1 \times 0.3$) and an extrapolate point at a query amount of 104K ($\mathcal{D} + \mathcal{D}_1^1 + \mathcal{D}_2^1 + \mathcal{D}_3^1$), discovering that the regression offers accurate predictions of model performances. We should notice this scaling law cannot be correct within all dataset size ranges since the test set accuracy is bounded.

Besides, the fitted regression shows when **query augmentation amount** doubles, LLaMA-7B models will improve $10.7 \times \log(2) = 7.4$, LLaMA2-7B will improve $9.8 \times \log(2) = 6.8$, and LLaMA2-13B will improve $7.6 \times \log(2) = 5.3$. As shown in Yuan et al. (2023a), it is estimated that when **human-written sample amount** doubles, LLaMA-7B models will improve 6.5 score, LLaMA2-7B will improve 6.6 score, and LLaMA2-13B models will improve 5.5 score. Query augmentation is similarly effective to human-written samples in term of in-domain performance. This demonstrates that query augmentation benefits from the performing proprietary LLMs on GSM8K, thus the sample quality generated by query augmentation is as high as those of human-written samples.

Response Augmentation Amount We further investigate under the data augmentation setting, if we keep the number of queries constant and increase the number of responses, how the in-domain performance changes. We use \mathcal{D}_1 as augmented queries and vary the response amount from 1 to 5 per augmented query. We also try majority voting (Wang et al., 2023a; Huang et al., 2022) to filter the augmented response since we cannot know the correct answer. In 3 and 12, we find for LLaMA-7B and LLaMA-2-7B models, once the response data volume reaches 97K (3 responses per query), further

Model	7B	7B-2	13B-2
Estimation	$y = 10.7 \log(x) + 13.2$	$y = 9.8 \log(x) + 21.3$	$y = 7.6 \log(x) + 36.3$
x = 17 prediction	43.4	49.2	57.7
x = 17 observation	43.4	50.0	56.0
x = 104 prediction	62.7	67.0	71.4
x = 104 observation	62.1	66.7	70.8

Table 3: The scaling law on amounts of augmented query in GSM8K.

increase the number of responses do not yield performance improvement. Before this point, model performance improves as the response amount increases. Thus, query augmentation with accurate responses seems more effective than only response augmentation with the augmented data size scales up. As for the LLaMA-2-13B model, within the data volume range of 37K to 157K, model performance consistently rises in a roughly linear fashion with increasing data volume, with a slower rate than the 7B model within the ascending interval. We then investigate the performance impact brought by majority voting filtering. If all responses have different answers, we discard the corresponding query-response. Surprisingly, we find that after applying majority voting, the model performance at the same data scale is generally lower than not applying filtering. A possible explanation is that even wrong responses generated by GPT-4 are useful for worse models (LLaMA) to improve their abilities. Another explanation is the reduction in the number of queries, as we discard the corresponding query when all response answers are different.

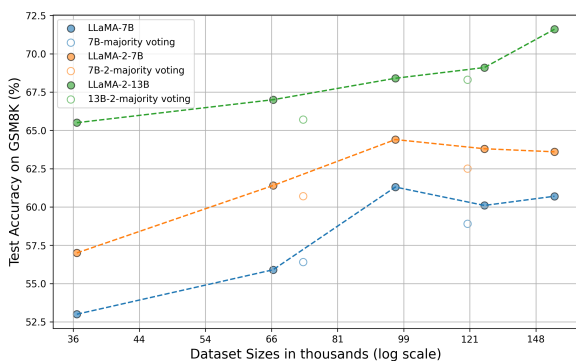


Figure 3: Comparison of test set accuracy on GSM8K for models of varying scales after fine-tuning on AugGSM8K subsets with different response augmentation amount.

Combination We investigate how the combination of query augmentation and response augmentation will affect the model’s performance. Re-

sults are listed in 4. We conduct SFT on $\mathcal{D} + \sum_{i=1}^3 \mathcal{D}_1^i + \hat{\mathcal{D}}_2^1 + \mathcal{D}_3^1$ and named MuggleMath effectively improving the in-domain accuracy of each model compared to using query or response augmentation only. These models outperform previous state-of-the-art open-sourced models with a very large margin for 7B and 13B models. **Model comparisons of MuggleMath and a broader range of state-of-the-art approaches are in 16.** Case studies of MuggleMath are listed in 8. It demonstrates that query augmentation and response augmentation can complement each other to a certain extent to improve in-domain performance.

3.5 RQ3. Can Data Augmentation Incentivize Generalization to Out-of-domain Mathematical Reasoning Tasks?

We have found that query and response augmentation significantly improves in-domain math reasoning performance. But we really interested in whether we can improve performances on out-of-domain distribution. We employ multi-task learning and transfer learning to see how models performed on the MATH dataset. We list results in 5. We find that (1) Multi-task learning and transfer learning outperform single-task supervised fine-tuning on LLaMA2-7/13B and do not improve on LLaMA-7B. (2) Although augmenting more query and response can improve GSM8K significantly, it has little to no help in improving MATH performance which indicates that in-domain augmentation is **not** helpful for out-of-domain generalization in this setting. Case studies of models performed on MATH are listed in 9.

To further investigate why AugGSM8K helps little on the MATH dataset, we use t-SNE in 4 to visualize the hidden representation of problems encoded by LLaMA2-7B, that is the 15-th layer of last token representation of the problem. We find GSM8K and MATH are separated in hidden space and only some of the problems in MATH are laid in the span of GSM8K. The augmented GSM8K problems are laid in the same span of

Model	7B	7B-2	13B-2	70B-2
\mathcal{D}	35.9	41.6	50.0	63.2
RFT (Yuan et al., 2023a)	49.1	51.2	55.3	64.8
WizardMath (Luo et al., 2023a)	-	54.9	63.9	81.6
$+\mathcal{D}_1^1 + \mathcal{D}_1^2 + \mathcal{D}_1^3$	61.3	64.4	68.4	-
$+\hat{\mathcal{D}}_1^1 + \hat{\mathcal{D}}_2^1 + \mathcal{D}_3^1$	61.4	66.3	69.8	-
$+\sum_{i=1}^3 \mathcal{D}_1^i + \hat{\mathcal{D}}_2^1 + \mathcal{D}_3^1$	65.2	67.4	72.6	80.1
$+\sum_{i=1}^3 \mathcal{D}_1^i + \hat{\mathcal{D}}_2^1 + \mathcal{D}_3^1$ (n=512)	65.4 (+16.3)	68.4 (+13.5)	74.0 (+10.1)	82.3 (+0.7)

Table 4: In-domain performance of combination on query and response augmentation. n=512 means that max decode token count is 512. For other experiments, we use max decode token count 256.

GSM8K which makes sense why it improves little for MATH.

To investigate if the proposed augmentation method improve performance on this subset of MATH, we find that while there is an overlap in the embedding space distribution between MATH and GSM8K, it is relatively small compared with that between GSM8K and AugGSM8K. In the transfer learning setting, training first on GSM8K and then on MATH with LLaMA-13B-2 does provide some benefits for certain subsets, such as Prealgebra, Algebra, and Geometry. However, if we train first on AugGSM8K and then on MATH, this benefit is not only marginal but may even lead to a decrease in performance on other subsets, like Geometry and Prealgebra, which could be related to the data proportions. Overall, the performance improvement on the MATH dataset from augmentation on GSM8K is minimal, even on subsets like Prealgebra, where there is some overlap. For 7B size and multi-task learning setting, we can draw the similar conclusion. Detailed results are listed in 18 to 23. This suggests if we want to improve math reasoning benchmark performances for LLMs, we can choose to apply augmentation on diverse math subjects.

4 Discussion

4.1 Training set vs. Test set accuracy

Query and response augmentation generate similar problems for the training set which leads to better training set accuracy. We have shown augmentations improve the accuracy of the in-domain test set. We want to investigate the relationship between the accuracy of the training set and the test set to find if the accuracy of the training set can be a performance indicator. We sample 500 samples from the original training set to calculate the accuracy. From 8, the training and test accuracy generally exhibit

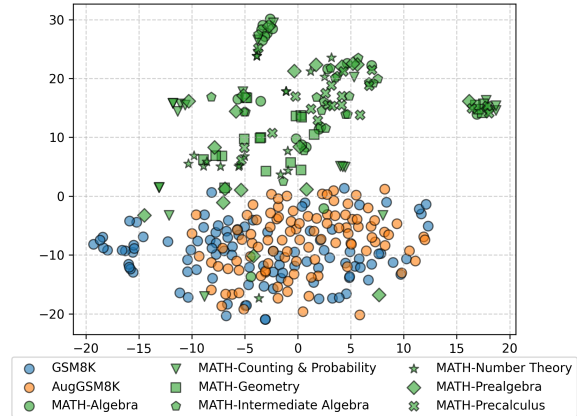


Figure 4: The embedding visualization of queries in GSM8K, MATH and AugGSM8K.

a positive correlation across different augmented data which shows the training accuracy could be an indicator of in-domain performance unless deliberately overfitting.

4.2 Make more augmentation on harder problems

During the query augmentation process, it is crucial to understand which kind of queries should be augmented. Augmenting too many easy problems may not be effective since the model may have mastered this level of problems. Here we examine if the model improves more when we augment more on harder or wrong problems. We define **hard** problems based on the number of equations, specifically, problems with fewer than three reasoning steps as easy, those with exactly three steps as medium, and those with more than three steps as hard (see more details in F). We define **wrong** problems if the SFT model solves them incorrectly. We apply SFT on subsets of AugGSM8K with augmented queries on easy, medium, hard, wrong, and random problems. From 5 and 14, it is evident that for LLaMA-2-7B and LLaMA-2-13B, the performance gain from

Training Setting	7B	7B-2	13B-2
<i>In Context Learning</i> on MATH	2.9	2.5	3.9
<i>Supervised Fine-tuning</i> on MATH	4.8	5.8	6.0
<i>Multi-task learning</i>			
+MATH	4.6	6.2	7.6
+ \mathcal{D}_1^1 +MATH	4.8	4.8	8.4
<i>Transfer learning</i>			
$\mathcal{D} \rightarrow$ MATH	4.4	6.0	9.4
+ $\mathcal{D}_1^1 \rightarrow$ MATH	6.2	5.6	7.8
+ $\sum_{i=1}^3 \mathcal{D}_1^i + \hat{\mathcal{D}}_2^1 + \mathcal{D}_3^1 \rightarrow$ MATH	5.6	8.4	9.4
+ $\sum_{i=1}^7 \mathcal{D}_1^i$ -majority voting+ $\mathcal{D}_2^1 + \mathcal{D}_3^1 \rightarrow$ MATH	5.6	6.0	9.0

Table 5: Comparison of test set accuracy on MATH. Multi-task learning means that we fine-tune the models on the mixed dataset of AugGSM8K subset and MATH. Transfer learning means that we first fine-tune the models on subsets of AugGSM8K and then fine-tune on MATH.

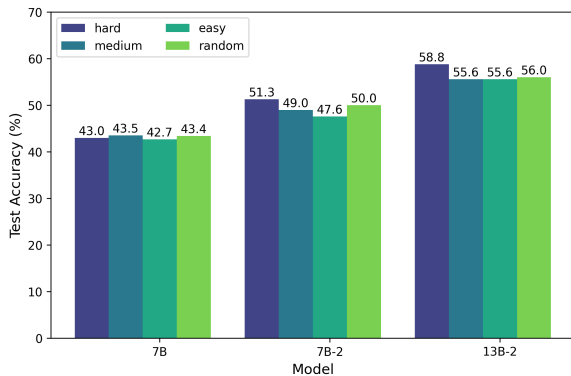


Figure 5: The performance of SFT models with different difficulty augmentation on GSM8K.

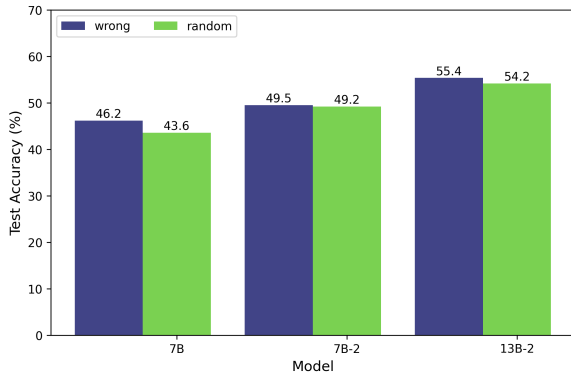


Figure 6: The performance of SFT models with wrong problem augmentation on GSM8K.

augmenting hard problems is significantly higher than that from augmenting other types of problems. From 6 and 15, we find that augmenting incorrect problems on three models consistently improves more than random query augmentation. In addition, we have conducted an analysis of our model’s performance on test set problems of varying difficulty.

Our analysis of the 7B-2 model’s performance on the test set, shows accuracy rates of 0.55, 0.42, and 0.21 for easy, medium, and hard problems, respectively, while MuggleMath achieves higher accuracy rates of 0.73, 0.70, and 0.64 for the same problem categories. This significant performance boost on difficult questions can be attributed to the fact that the augmented problems we generated are generally more complex than the original problems.

4.3 Result on the Perturbed Test Set

We have perturbed two new test sets based on the original GSM8K test set, Change-Test and Aug-Test in in 17.

Upon evaluating our model on these two perturbed test sets, we found that the performance of MuggleMath consistently and significantly exceeds that of the model fine-tuned on GSM8K alone. This observation suggests that our data augmentation techniques not only enhance the model’s ability to solve the original problems but also contribute to its improved performance on varied and perturbed inputs, thereby indicating a robust generalization capability in in-domain dataset.

5 Conclusion

In this paper, we investigate the scaling property of query and response augmentation with respect to math reasoning in-domain and out-of-domain performance. We find that query and response augmentation can improve in-domain performance very effectively which has a similar improvement to human-written query-response pairs augmentation. Although we can obtain state-of-the-art in-domain performance by such augmentation, we find it cannot generalize to out-of-domain math reasoning performances.

597 Limitations

598 In this study, we focused on the domain of math-
599 ematical reasoning—a key area of interest for the
600 large language model (LLM) research commu-
601 nity. We investigated the effectiveness of data aug-
602 mentation techniques on both in-domain (GSM8K
603 dataset) and out-of-domain (MATH dataset) per-
604 formance. While our work provides insights into
605 the performance scalability and generalizability
606 of Chain-of-Thought (COT) enhanced models in
607 mathematical reasoning, it’s important to acknowl-
608 edge certain limitations: (1)Scope of Mathematical
609 Datasets: Our research concentrated on specific
610 datasets, and the augmentation strategies were tai-
611 lored accordingly. The mathematical landscape is
612 vast, and our findings may not necessarily extrapo-
613 late to other datasets not included in our study. Fu-
614 ture work should consider a wider array of mathe-
615 matical datasets to validate and extend our findings.
616 (2)Depth of Generalizability Research: Although
617 our study is among the initial efforts to evaluate
618 the generalizability of COT-augmented models in
619 math reasoning, the study did not extensively ex-
620 plore solutions to enhance out-of-domain perfor-
621 mance across a broader spectrum of mathematical
622 reasoning.

623 References

624 Alibaba. 2023. Qwen technical report. *arXiv preprint*
625 *arXiv:2309.16609*.

626 Rohan Anil, Andrew M Dai, Orhan Firat, Melvin John-
627 son, Dmitry Lepikhin, Alexandre Passos, Siamak
628 Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng
629 Chen, et al. 2023. Palm 2 technical report. *arXiv*
630 *preprint arXiv:2305.10403*.

631 BaichuanInc. 2023. Baichuan 2. technical report. *arXiv*
632 *preprint arXiv:2309.16609*.

633 Tiffany Tianhui Cai, Hongseok Namkoong, and Steve
634 Yadlowsky. 2023. Diagnosing model performance
635 under distribution shift.

636 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming
637 Yuan, Henrique Ponde de Oliveira Pinto, Jared Kap-
638 lan, Harri Edwards, Yuri Burda, Nicholas Joseph,
639 Greg Brockman, Alex Ray, Raul Puri, Gretchen
640 Krueger, Michael Petrov, Heidy Khlaaf, Girish Sas-
641 try, Pamela Mishkin, Brooke Chan, Scott Gray,
642 Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz
643 Kaiser, Mohammad Bavarian, Clemens Winter,
644 Philippe Tillet, Felipe Petroski Such, Dave Cum-
645 mings, Matthias Plappert, Fotios Chantzis, Eliza-
646 beth Barnes, Ariel Herbert-Voss, William Hebgen
647 Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie

Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, 648
William Saunders, Christopher Hesse, Andrew N. 649
Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan 650
Morikawa, Alec Radford, Matthew Knight, Miles 651
Brundage, Mira Murati, Katie Mayer, Peter Welinder, 652
Bob McGrew, Dario Amodei, Sam McCandlish, Ilya 653
Sutskever, and Wojciech Zaremba. 2021. [Evaluating 654
large language models trained on code.](#) 655

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, 656
Maarten Bosma, Gaurav Mishra, Adam Roberts, 657
Paul Barham, Hyung Won Chung, Charles Sutton, 658
Sebastian Gehrmann, Parker Schuh, Kensen Shi, 659
Sasha Tsvyashchenko, Joshua Maynez, Abhishek 660
Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vin- 661
odkumar Prabhakaran, Emily Reif, Nan Du, Ben 662
Hutchinson, Reiner Pope, James Bradbury, Jacob 663
Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, 664
Toju Duke, Anselm Levskaya, Sanjay Ghemawat, 665
Sunipa Dev, Henryk Michalewski, Xavier Garcia, 666
Vedant Misra, Kevin Robinson, Liam Fedus, Denny 667
Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, 668
Barret Zoph, Alexander Spiridonov, Ryan Sepassi, 669
David Dohan, Shivani Agrawal, Mark Omernick, An- 670
drew M. Dai, Thanumalayan Sankaranarayanan Pil- 671
lai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, 672
Rewon Child, Oleksandr Polozov, Katherine Lee, 673
Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark 674
Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy 675
Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, 676
and Noah Fiedel. 2022. [Palm: Scaling language mod- 677
eling with pathways.](#) 678

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, 679
Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias 680
Plappert, Jerry Tworek, Jacob Hilton, Reiichiro 681
Nakano, et al. 2021. Training verifiers to solve math 682
word problems. *arXiv preprint arXiv:2110.14168*. 683

Elliot Creager, Jörn-Henrik Jacobsen, and Richard S. 684
Zemel. 2021. Environment inference for invariant 685
learning. pages 2189–2200. PMLR. 686

Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Zhi 687
Zheng, Shengding Hu, Zhiyuan Liu, Maosong Sun, 688
and Bowen Zhou. 2023. Enhancing chat language 689
models by scaling high-quality instructional conver- 690
sations. *arXiv preprint arXiv:2305.14233*. 691

Steven Y. Feng, Varun Gangal, Jason Wei, Sarath Chan- 692
dar, Soroush Vosoughi, Teruko Mitamura, and Ed- 693
uard Hovy. 2021. [A survey of data augmentation 694
approaches for NLP.](#) In *Findings of the Association 695
for Computational Linguistics: ACL-IJCNLP 2021*, 696
pages 968–988, Online. Association for Computa- 697
tional Linguistics. 698

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul 699
Arora, Steven Basart, Eric Tang, Dawn Song, and Ja- 700
cob Steinhardt. 2021. Measuring mathematical prob- 701
lem solving with the math dataset. *arXiv preprint 702
arXiv:2103.03874*. 703

Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, 704
Xuezhi Wang, Hongkun Yu, and Jiawei Han. 2022. 705
[Large language models can self-improve.](#) 706

707	Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive growing of gans for improved quality, stability, and variation. In <i>6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings</i> .	760
708		761
709		762
710		763
711		764
712		
713	Kun Kuang, Ruoxuan Xiong, Peng Cui, Susan Athey, and Bo Li. 2020. Stable prediction with model misspecification and agnostic distribution shift. pages 4485–4492. AAAI Press.	765
714		766
715		767
716		768
717	Ariel N. Lee, Cole J. Hunter, and Nataniel Ruiz. 2023. Platypus: Quick, cheap, and powerful refinement of llms .	769
718		770
719		771
720	Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. 2022a. Solving quantitative reasoning problems with language models .	772
721		773
722		774
723		775
724		776
725		777
726	Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay V. Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. 2022b. Solving quantitative reasoning problems with language models . In <i>NeurIPS</i> .	778
727		779
728		780
729		781
730		782
731		783
732		784
733	Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step . <i>arXiv preprint arXiv:2305.20050</i> .	785
734		786
735		787
736		788
737		789
738	Zachary C. Lipton, Yu-Xiang Wang, and Alexander J. Smola. 2018. Detecting and correcting for label shift with black box predictors. pages 3128–3136. PMLR.	790
739		791
740		792
741	Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. 2023. The flan collection: Designing data and methods for effective instruction tuning . <i>arXiv preprint arXiv:2301.13688</i> .	793
742		794
743		795
744		796
745		797
746	Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. 2023a. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct .	798
747		799
748		800
749		801
750		802
751	Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023b. Wizardcoder: Empowering code large language models with evol-instruct . <i>arXiv preprint arXiv:2306.08568</i> .	803
752		804
753		805
754		806
755		807
756	Yuzhou Mao, Liu Yu, Yi Yang, Fan Zhou, and Ting Zhong. 2023. Debiasing intrinsic bias and application bias jointly via invariant risk minimization (student abstract) . AAAI Press.	808
757		809
758		810
759		811
		812
		813
		814
		815
	Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. Orca: Progressive learning from complex explanation traces of gpt-4 . <i>arXiv preprint arXiv:2306.02707</i> .	
	Ansong Ni, Jeevana Priya Inala, Chenglong Wang, Alex Polozov, Christopher Meek, Dragomir Radev, and Jianfeng Gao. 2023. Learning math reasoning from self-sampled correct and partially-correct solutions . In <i>The Eleventh International Conference on Learning Representations</i> .	
	OpenAI. 2023. Gpt-4 technical report .	
	Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback .	
	Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon LLM: outperforming curated corpora with web data, and web data only .	
	Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. 2015. Causal inference using invariant prediction: identification and confidence intervals .	
	Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton-Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. Code llama: Open foundation models for code .	
	Bernhard Schölkopf, Dominik Janzing, Jonas Peters, Eleni Sgouritsa, Kun Zhang, and Joris M. Mooij. 2012. On causal and anticausal learning .	
	Zheyang Shen, Peng Cui, Tong Zhang, and Kun Kuang. 2020. Stable learning via sample reweighting . pages 5692–5699. AAAI Press.	
	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model . https://github.com/tatsu-lab/stanford_alpaca .	
	Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. 2022. Galactica: A large language model for science .	
	InternLM Team. 2023a. Internlm: A multilingual language model with progressively enhanced capabilities . https://github.com/InternLM/InternLM .	

816	MosaicML NLP Team. 2023b. Introducing mpt-7b: A new standard for open-source, commercially usable llms . Accessed: 2023-05-05.	
817		
818		
819	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier	
820	Martinet, Marie-Anne Lachaux, Timothée Lacroix,	
821	Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal	
822	Azhar, Aurelien Rodriguez, Armand Joulin, Edouard	
823	Grave, and Guillaume Lample. 2023a. Llama: Open and efficient foundation language models .	
824		
825	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-	
826	bert, Amjad Almahairi, Yasmine Babaei, Nikolay	
827	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti	
828	Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton	
829	Ferrer, Moya Chen, Guillem Cucurull, David Esiobu,	
830	Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller,	
831	Cynthia Gao, Vedanuj Goswami, Naman Goyal, An-	
832	thony Hartshorn, Saghar Hosseini, Rui Hou, Hakan	
833	Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa,	
834	Isabel Kloumann, Artem Korenev, Punit Singh Koura,	
835	Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Di-	
836	ana Liskovich, Yinghai Lu, Yuning Mao, Xavier Mar-	
837	tinnet, Todor Mihaylov, Pushkar Mishra, Igor Moly-	
838	bog, Yixin Nie, Andrew Poulton, Jeremy Reizen-	
839	stein, Rashi Rungta, Kalyan Saladi, Alan Schelten,	
840	Ruan Silva, Eric Michael Smith, Ranjan Subrama-	
841	nian, Xiaoqing Ellen Tan, Binh Tang, Ross Tay-	
842	lor, Adina Williams, Jian Xiang Kuan, Puxin Xu,	
843	Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan,	
844	Melanie Kambadur, Sharan Narang, Aurelien Ro-	
845	driguez, Robert Stojnic, Sergey Edunov, and Thomas	
846	Scialom. 2023b. Llama 2: Open foundation and fine-tuned chat models .	
847		
848	Dustin Tran, Jeremiah Z. Liu, Michael W. Dusen-	
849	berry, Du Phan, Mark Collier, Jie Ren, Kehang Han,	
850	Zi Wang, Zelda Mariet, Huiyi Hu, Neil Band, Tim	
851	G. J. Rudner, Karan Singhal, Zachary Nado, Joost	
852	van Amersfoort, Andreas Kirsch, Rodolphe Jenat-	
853	ton, Nithum Thain, Honglin Yuan, Kelly Buchanan,	
854	Kevin Murphy, D. Sculley, Yarin Gal, Zoubin Ghahra-	
855	mani, Jasper Snoek, and Balaji Lakshminarayanan.	
856	2022. Plex: Towards reliability using pretrained large	
857	model extensions.	
858	Szymon Tworowski, Konrad Staniszewski, Mikołaj	
859	Pacek, Yuhuai Wu, Henryk Michalewski, and Pi-	
860	otr Miłoś. 2023. Focused transformer: Con-	
861	trastive training for context scaling. <i>arXiv preprint</i>	
862	<i>arXiv:2307.03170</i> .	
863	Z. Lin Y. Sheng Z. Wu H. Zhang L. Zheng S. Zhuang	
864	Y. Zhuang J. Gonzalez I. Stoica W. Chiang, Z. Li	
865	and E. Xing. 2023. icuna: An open-source chatbot	
866	impressing gpt-4 with 90 <i>Technical Report</i> .	
867	Ben Wang and Aran Komatsuzaki. 2021. GPT-J-	
868	6B: A 6 Billion Parameter Autoregressive Lan-	
869	guage Model. https://github.com/kingoflolz/mesh-transformer-jax .	
870		
871	Jindong Wang, Cuiling Lan, Chang Liu, Yidong	
872	Ouyang, and Tao Qin. 2021. Generalizing to un-	
873	seen domains: A survey on domain generalization.	
874	pages 4627–4635.	
	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le,	875
	Ed H. Chi, Sharan Narang, Aakanksha Chowdhery,	876
	and Denny Zhou. 2023a. Self-consistency improves chain of thought reasoning in language models . In <i>The Eleventh International Conference on Learning Representations</i> .	877
		878
		879
		880
	Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa	881
	Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh	882
	Hajishirzi. 2023b. Self-instruct: Aligning language models with self-generated instructions . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.	883
		884
		885
		886
		887
		888
	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten	889
	Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou,	890
	et al. 2022. Chain-of-thought prompting elicits rea-	891
	soning in large language models. <i>Advances in Neural Information Processing Systems</i> , 35:24824–24837.	892
		893
	Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng,	894
	Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin	895
	Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions .	896
		897
	Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu,	898
	Zhengying Liu, Yu Zhang, James T. Kwok, Zhenguo	899
	Li, Adrian Weller, and Weiyang Liu. 2023. Meta-	900
	math: Bootstrap your own mathematical questions for large language models .	901
		902
	Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting	903
	Dong, Keming Lu, Chuanqi Tan, Chang Zhou, and	904
	Jingren Zhou. 2023a. Scaling relationship on learning mathematical reasoning with large language models .	905
		906
		907
	Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang,	908
	and Songfang Huang. 2023b. How well do large	909
	language models perform in arithmetic tasks? <i>arXiv preprint arXiv:2304.02015</i> .	910
		911
	Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wen-	912
	hao Huang, Huan Sun, Yu Su, and Wenhui Chen.	913
	Mammoth: Building math generalist models through	914
	hybrid instruction year = 2023, eprinttype = arXiv,	915
	eprint = 2309.05653,.	916
	Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wen-	917
	hao Huang, Huan Sun, Yu Su, and Wenhui Chen.	918
	2023. MAMmoTH: Building math generalist mod-	919
	els through hybrid instruction tuning. <i>arXiv preprint</i>	920
	<i>arXiv:2309.05653</i> .	921
	Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Good-	922
	man. 2022. STar: Bootstrapping reasoning with reasoning . In <i>Advances in Neural Information Processing Systems</i> .	923
		924
		925
	Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang,	926
	Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu,	927
	Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b:	928
	An open bilingual pre-trained model. <i>arXiv preprint</i>	929
	<i>arXiv:2210.02414</i> .	930

931 Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and
 932 Chen Change Loy. 2023. Domain generalization:
 933 A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*,
 934 pages 4396–4415.

935 Xinyu Zhu, Junjie Wang, Lin Zhang, Yuxiang Zhang,
 936 Yongfeng Huang, Ruyi Gan, Jiaying Zhang, and Yu-
 937 jiu Yang. 2023a. Solving math word problems via
 938 cooperative reasoning induced language models. In
 939 *Proceedings of the 61st Annual Meeting of the As-
 940 sociation for Computational Linguistics (Volume 1:
 941 Long Papers)*, pages 4471–4485, Toronto, Canada.
 942 Association for Computational Linguistics.

943 Xinyu Zhu, Junjie Wang, Lin Zhang, Yuxiang Zhang,
 944 Yongfeng Huang, Ruyi Gan, Jiaying Zhang, and Yu-
 945 jiu Yang. 2023b. Solving math word problems via
 946 cooperative reasoning induced language models. In
 947 *Proceedings of the 61st Annual Meeting of the As-
 948 sociation for Computational Linguistics (Volume 1:
 949 Long Papers)*, pages 4471–4485, Toronto, Canada.
 950 Association for Computational Linguistics.

A Instruction prompt for training and inference

951
952
953
954

Here is the instruction prompt used for the training and inference stage.

Fine-tuning system prompt

Below is an instruction that describes a task. Write a response that appropriately completes the request.### Instruction: ****Query.**** ### Response:

955

B Query augmentation prompt for GSM8K

956
957

Here is the query augmentation prompt we used. We require the models to generate five different augmented problems with our provided example. We use gpt-3.5-turbo-0613 and gpt-4-0613 APIs with a temperature of 1.0 to obtain augmented problems.

958
959
960
961
962
963

Query augmentation prompt

I want you to act as a math teacher. I will provide a grade school math question and you will help to create more challenging math questions by given ways. Given the question: “James writes a 3-page letter to 2 different friends twice a week. How many pages does he write a year?”, you will modify it by following ideas:

1. **Change specific numbers:** James writes a 2-page letter to 2 different friends 3 times a week. How many pages does he write in 4 years?
2. **Introduce fractions or percentages:** James writes a 3-page letter to 2 different friends twice a week. Each week, he adds 50% more pages to each letter. How many pages does he write in a month?
3. **Combine multiple concepts:** James writes a 3-page letter to 2 different friends twice a week. He uses both sides of the paper and each side can hold 250 words. If James writes 100 words per minute, how long does it take for him to write all the letters in a week?
4. **Include a conditional statement:** James writes a 3-page letter to 2 different friends twice a week. If it’s a holiday, he writes an additional 5-page letter to each friend.

964

Considering there are 10 holidays in a year, how many pages does he write in a year?

5. Increase the complexity of the problem: James writes a 3-page letter to 2 different friends twice a week. In addition, he writes a 5-page letter to 3 other friends once a week. How many pages does he write in a month, assuming there are 4 weeks in a month?

Now you are given the question:

****A new math problem here.****

C Response augmentation prompt for GSM8K

We use this prompt to generate responses to ensure the response format which can be viewed as 1-shot setting. We use gpt-3.5-turbo-0613 and gpt-4-0613 with temperature 0.0 or 1.0.

Response augmentation prompt

I want you to act as an excellent math solver. You will solve the given math question step by step. You need to reply with a python dictionary in the same format as the given examples. Retain decimals to three decimal places. The formulas in the process need to use the format: $48/2 = \ll 48/2=24 \gg 24$ clips. The end of response needs to be: ##### {answer}.

Examples: {"query": "Natalia sold clips to 48 of her friends in April, and then she sold half as many clips in May. How many clips did Natalia sell altogether in April and May?", "response": "Natalia sold $48/2 = \ll 48/2=24 \gg 24$ clips in May. Natalia sold $48+24 = \ll 48+24=72 \gg 72$ clips altogether in April and May.##### 72"}.

The given question:

****A new math problem here.****

D Response filter

We filter out generated responses by following rules.

- Delete the responses without an answer.
- Delete the responses that are excessively lengthy (> 1500).
- Remove superfluous characters beyond the reasoning path and the answer.

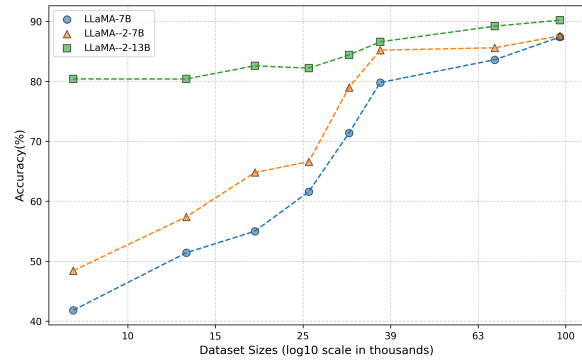


Figure 7: The relationship of accuracy on the training set and the amounts of augmentation data.

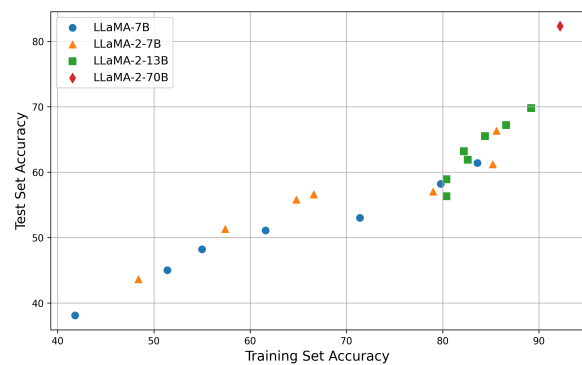


Figure 8: The relationship of accuracy on the GSM8K test set and the original GSM8K training set.

E Case Study

F Difficulty level definition on GSM8K

We conducted a statistical analysis of the reasoning paths required for 7,473 questions in the GSM8K training set, categorizing them as hard, medium, and easy. Specifically, we defined questions with more than three formulas, exactly three formulas, and less than three formulas as hard, medium, and easy, respectively. This categorization yielded a balanced distribution with 2,357 easy, 2,360 medium, and 2,756 hard problems. This approach ensures a relatively equal number of problems in each category.

G Detailed Experimental results

We list the detailed experimental results of different settings here.

Method	Query
Original	Weng earns \$12 an hour for babysitting. Yesterday, she just did 50 minutes of babysitting. How much did she earn?
Change specific numbers	Weng earns \$15 an hour for babysitting. Last week, she babysat for 3 hours and 45 minutes every day. How much did she earn in total last week?
Introduce fractions or percentages	Weng earns \$12 an hour for babysitting. However, for every hour she babysits beyond the first 3 hours, she earns an additional 20% per hour. If she babysat for 5 hours yesterday, how much did she earn?
Combine multiple concepts	Weng earns \$12 an hour for babysitting, but gets paid only 75% of the total at the time of service. The rest is paid at the end of the month. If she babysat for 3 hours yesterday, how much was she paid immediately and how much will she receive later?
Include a conditional statement	Weng earns \$12 an hour for babysitting. If it's a weekend, she gets a bonus of \$5. If she babysat for 2 hours last Sunday, how much did she earn?
Increase the complexity of the problem	Weng earns \$12 an hour for babysitting during weekdays and \$15 an hour on weekends. If she babysat for 2 hours on Tuesday and 3 hours on Saturday, how much did she earn in total?

Table 6: Examples of different methods for generating new queries.

Query	Weng earns \$12 an hour for regular babysitting but doubles her rate during weekends. Yesterday, she did 50 minutes of babysitting and it was a weekend. How much did she earn?
Response1	Weng earns \$12 per hour and doubles her rate during weekends, so she earned $\$12 * 2 = 24$ per hour yesterday. There are 60 minutes in an hour, so Weng did $50 \text{ minutes} / 60 \text{ minutes/hour} = 0.833$ hours of babysitting. Therefore, Weng earned $\$24 * 0.833 = 19.992$, which we can round to \$20.
Response2	Weng earns $\$12 * 2 = 24$ an hour during weekends. 50 minutes is $50/60 = 0.833$ of an hour. So, Weng earned $24 * 0.833 = 19.992$ dollars yesterday.
Response3	During weekends, Weng earns $12 * 2 = 24$ dollars per hour for babysitting. There are 60 minutes in an hour. So, 50 minutes is $50/60 = 0.833$ of an hour. Weng earned $24 * 0.833 = 20$ dollars for 50 minutes of babysitting during the weekend.
Response4	Her weekend rate is $12 * 2 = 24$ dollars per hour. There are 60 minutes in an hour, so she worked $50/60 = 0.833$ hours. So, Weng earned $24 * 0.833 = 19.992$ dollars.
Response5	Weng's rate during weekends is $12 * 2 = 24$ dollars per hour. Since an hour has 60 minutes, Weng earned $24 * (50/60) = 20$ dollars yesterday.
Response6	Weng earns $\$12 * 2 = 24$ an hour during weekends. She worked for $50/60 = 0.833$ hour yesterday. So, Weng earned $0.833 * \$24 = 19.992$ yesterday.
Response7	Weng earns \$12 an hour, but during the weekend, she doubles her rate which means she earns $12 * 2 = 24$ dollars an hour during the weekend. Babysitting for 50 minutes is equivalent to $50/60 = 0.833$ hours. Therefore, Weng earned $24 * 0.833 = 19.992$ dollars from babysitting yesterday.

Table 7: Different reasoning paths for the same query.

Query	Response
Janet's ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers' market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers' market?	Janet eats 3 eggs for breakfast and bakes 4 eggs into muffins, so she uses $3+4 = 7$ eggs every day. Janet's ducks lay 16 eggs per day, so after she uses 7 eggs, she has $16-7 = 9$ eggs left. She sells these 9 eggs at the farmers' market for \$2 per egg, so she makes $9*2 = 18$ dollars every day at the farmers' market.
A robe takes 2 bolts of blue fiber and half that much white fiber. How many bolts in total does it take?	The white fiber is half the amount of the blue fiber so it takes $2/2 = 1$ bolt of white fiber. All total it takes $2+1=3$ bolts of fiber.
Josh decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased the value of the house by 150%. How much profit did he make?	He bought the house for 80,000 and put 50,000 in repairs so he had $80,000+50,000 = \$130,000$ invested. The value of the house increased by 150% so it went up $130,000*1.5 = \$195,000$ in value. He had 130,000 invested and it increased in value by 195,000 so he made a profit of $195,000-130,000 = \$65,000$.
James decides to run 3 sprints 3 times a week. He runs 60 meters each sprint. How many total meters does he run a week?	He runs $3*60=180$ meters per day. So he runs $180*3=540$ meters.
Every day, Wendi feeds each of her chickens three cups of mixed chicken feed, containing seeds, mealworms and vegetables to help keep them healthy. She gives the chickens their feed in three separate meals. In the morning, she gives her flock of chickens 15 cups of feed. In the afternoon, she gives her chickens another 25 cups of feed. How many cups of feed does she need to give her chickens in the final meal of the day if the size of Wendi's flock is 20 chickens?	Every day, Wendi feeds her chickens a total of $15+25=40$ cups of feed. Thus, she must feed her chickens $40/3=13.333$ cups of feed in the final meal of the day.

Table 8: Some examples of MuggleMath-13B answering questions from the GSM8K test set. Wrong reasoning processes are labeled in red.

Instruction	Response
A regular hexagon can be divided into six equilateral triangles. If the perimeter of one of the triangles is 21 inches, what is the perimeter, in inches, of the regular hexagon?	Since the triangle is equilateral, each side length is 7 inches. The hexagon has six sides of equal length, so its perimeter is $6 \times 7 = \boxed{42}$ inches.
What is the smallest positive perfect cube that can be written as the sum of three consecutive integers?	The response provided is incomplete and does not provide a valid answer to the question.
The set of points (x, y, z) that satisfy $2x = 3y = -z$ is a line. The set of points (x, y, z) that satisfy $6x = -y = -4z$ is another line. Find the angle between these lines, in degrees.	The direction vector of the first line is $\langle 2, 3, -1 \rangle$, and the direction vector of the second line is $\langle 6, -1, -4 \rangle$. The cosine of the angle between these direction vectors is $\frac{\langle 2, 3, -1 \rangle \cdot \langle 6, -1, -4 \rangle}{\ \langle 2, 3, -1 \rangle \ \ \langle 6, -1, -4 \rangle \ } = \frac{12}{20} = \frac{3}{5}$. Therefore, the angle between the lines is $\boxed{75^\circ}$.
What is the distance, in units, between the points $(2, -6)$ and $(-4, 3)$? Express your answer in simplest radical form.	We use the distance formula: $\sqrt{(2 - (-4))^2 + (-6 - 3)^2} = \sqrt{16 + 9^2} = \sqrt{16 + 81} = \sqrt{97} = \boxed{3\sqrt{7}}$.
The expression $2 \cdot 3 \cdot 4 \cdot 5 + 1$ is equal to 121, since multiplication is carried out before addition. However, we can obtain values other than 121 for this expression if we are allowed to change it by inserting parentheses. For example, we can obtain 144 by writing $(2 \cdot (3 \cdot 4)) \cdot (5 + 1) = 144$. In total, how many values can be obtained from the expression $2 \cdot 3 \cdot 4 \cdot 5 + 1$ by inserting parentheses? (Note that rearranging terms is not allowed, only inserting parentheses).	We can obtain 121 by not inserting any parentheses, so we need to insert at least one pair of parentheses to obtain a different value. We can obtain 144 by inserting one pair of parentheses as shown above. If we insert two pairs of parentheses, we get $(2 \cdot (3 \cdot (4 \cdot 5))) + 1 = 121$. Thus, we can obtain 121 by inserting two pairs of parentheses, so we need to insert at least three pairs of parentheses to obtain a different value. We can obtain 168 by inserting three pairs of parentheses, since $(2 \cdot (3 \cdot (4 \cdot (5 + 1)))) = 168$. Thus, there are $\boxed{4}$ values that can be obtained from the expression $2 \cdot 3 \cdot 4 \cdot 5 + 1$ by inserting parentheses.

Table 9: Some examples of MuggleMath-13B answering questions from the MATH test set. Wrong reasoning processes are labeled in red.

Query aug. type	7B	7B-2	13B-2
No aug. (\mathcal{D})	35.9	41.6	50.0
Change numbers	41.5	48.5	54.1
Fractions or percentages	41.2	46.2	54.4
Combine multiple concepts	41.1	47.5	56.1
Conditional statement	41.7	45.8	56.4
Increase complexity	42.4	48.6	57.6
All aug. ($\mathcal{D} + \mathcal{D}_1^1$)	53.0	57.0	65.5

Table 10: Different query augmentation strategies on GSM8K performances.

Query aug.	7B	7B-2	13B-2
\mathcal{D}	35.9	41.6	50.0
$+\mathcal{D}_1^1 \times 0.2$	38.1	43.6	56.3
$+\mathcal{D}_1^1 \times 0.4$	45.0	51.3	58.9
$+\mathcal{D}_1^1 \times 0.6$	48.2	55.8	61.9
$+\mathcal{D}_1^1 \times 0.8$	51.1	56.6	63.2
$+\mathcal{D}_1^1$	53.0	57.0	65.5
$+\mathcal{D}_1^1 + \hat{\mathcal{D}}_2^1$	58.2	61.2	67.2
$+\mathcal{D}_1^1 + \hat{\mathcal{D}}_2^1 + \mathcal{D}_3^1$	61.4	66.3	69.8

Table 11: The performance of SFT with different amounts of augmented query on GSM8K.

Response aug.	7B	7B-2	13B-2
\mathcal{D}	35.9	41.6	50.0
$+\mathcal{D}_1^1$	53.0	57.0	65.5
$+\mathcal{D}_1^1 + \mathcal{D}_1^2$	55.9	61.4	67.0
$+\mathcal{D}_1^1 + \mathcal{D}_1^2 + \mathcal{D}_1^3$	61.3	64.4	68.4
$+\mathcal{D}_1^1 + \mathcal{D}_1^2 + \mathcal{D}_1^3 + \mathcal{D}_1^4$	60.1	63.8	69.1
$+\mathcal{D}_1^1 + \mathcal{D}_1^2 + \mathcal{D}_1^3 + \mathcal{D}_1^4 + \mathcal{D}_1^5$	60.7	63.6	71.6
$+\mathcal{D}_1^1 + \mathcal{D}_1^2 + \mathcal{D}_1^3$ - majority voting	56.4	60.7	65.7
$+\mathcal{D}_1^1 + \mathcal{D}_1^2 + \mathcal{D}_1^3 + \mathcal{D}_1^4 + \mathcal{D}_1^5$ - majority voting	58.9	62.5	68.3

Table 12: The performance of SFT with different amounts of augmented response on GSM8K.

Model	Data	\mathcal{D}	$+\mathcal{D}_1^1 \times 0.2$	$+\mathcal{D}_1^1 \times 0.4$	$+\mathcal{D}_1^1 \times 0.6$
7B	training set	56.4	41.8	51.4	55
	test set	35.9	38.1	45	48.2
7B-2	training set	65.2	48.4	57.4	64.8
	test set	41.6	43.6	51.3	55.8
13B-2	training set	75.4	80.4	80.4	82.6
	test set	50	56.3	58.9	61.9

(a) Part 1

Model	Data	$+\mathcal{D}_1^1 \times 0.8$	$+\mathcal{D}_1^1$	$+\mathcal{D}_1^1 + \hat{\mathcal{D}}_2^1$	$+\mathcal{D}_1^1 + \hat{\mathcal{D}}_2^1 + \mathcal{D}_3^1$
7B	training set	61.6	71.4	79.8	83.6
	test set	51.1	53	58.2	61.4
7B-2	training set	66.6	79	85.2	85.6
	test set	56.6	57	61.2	66.3
13B-2	training set	82.2	84.4	86.6	89.2
	test set	63.2	65.5	67.2	69.8

(b) Part 2

Table 13: The accuracy on the training dataset and test dataset for GSM8K.

Model	7B	7B-2	13B-2
\mathcal{D}	35.9	41.6	50.0
\mathcal{D}_1^1 on hard	43.0	51.3	58.8
\mathcal{D}_1^1 on medium	43.5	49.0	55.6
\mathcal{D}_1^1 on easy	42.7	47.6	55.6
\mathcal{D}_1^1 on random	43.4	50.0	56.0

Table 14: The performance of SFT with query augmentation with different difficulties on GSM8K.

Model	7B	7B-2	13B-2
\mathcal{D}_1^1 on wrong	46.2	49.5	55.4
\mathcal{D}_1^1 on random	43.6	49.2	54.2

Table 15: The performance of SFT with query augmentation with wrong problems or random problems on GSM8K.

closed-source models		
Model	#params	GSM8K
GPT-4(OpenAI, 2023)	-	92.0
GPT-3.5-Turbo(Ouyang et al., 2022)	-	80.8
Claude-2	-	85.2
PaLM (Chowdhery et al., 2022)	8B	4.1
PaLM	62B	33.0
PaLM	540B	56.5
PaLM-2(Anil et al., 2023)	540B	80.7
Flan-PaLM(Anil et al., 2023) 2	540B	84.7
Minerva(Lewkowycz et al., 2022b)	8B	16.2
Minerva	62B	52.4
Minerva	540B	58.8
open-source models (1-10B)		
LLaMA-1(Touvron et al., 2023a)	7B	11.0
LLaMA-2(Touvron et al., 2023b)	7B	14.6
MPT(Team, 2023b)	7B	6.8
Falcon	7B	6.8
Code-LLaMA(Rozière et al., 2023)	7B	25.2
InternLM(Team, 2023a)	7B	31.2
GPT-J(Wang and Komatsuzaki, 2021)	6B	34.9
ChatGLM-2(Zeng et al., 2022)	6B	32.4
Qwen(Alibaba, 2023)	7B	51.6
Baichuan-2(BaichuanInc, 2023)	7B	24.5
MAmooTH-CoT(Yue et al.)	7B	50.5
SFT	7B	41.6
RFT(Yuan et al., 2023a)	7B	50.3
WizardMath(Luo et al., 2023a)	7B	54.9
MetaMath(Yu et al., 2023)	7B	66.5
MuggleMath-7B	7B	68.4
open-source models (11-50B)		
LLaMA-1(Touvron et al., 2023a)	13B	17.8
LLaMA-1	33B	35.6
LLaMA-2(Touvron et al., 2023b)	13B	28.7
Platypus(Lee et al., 2023)	13B	25.7
LLaMA-2	34B	42.2
MPT(Team, 2023b)	30B	15.2
Falcon(Penedo et al., 2023)	40B	19.6
Vicuna(W. Chiang and Xing., 2023)	13B	27.6
Baichuan-2(BaichuanInc, 2023)	13B	52.8
MAmooTH-CoT(Yue et al.)	13B	56.3
Code-LLaMA(Rozière et al., 2023)	13B	36.1
SFT	13B	50.0
RFT(Yuan et al., 2023a)	13B	54.8
WizardMath(Luo et al., 2023a)	13B	63.9
MetaMath(Yu et al., 2023)	13B	72.3
MuggleMath-13B	13B	74.0
open-source models (51-70B)		
LLaMA-1(Touvron et al., 2023a)	65B	50.9
LLaMA-2(Touvron et al., 2023b)	70B	56.8
RFT(Yuan et al., 2023a)	70B	64.8
Platypus(Lee et al., 2023)	70B	70.6
MAmooTH-CoT(Yue et al.)	70B	71.4
WizardMath(Luo et al., 2023a)	70B	81.6
MetaMath(Yu et al., 2023)	70B	82.3
MuggleMath-70B	70B	82.3

	7B	7B-2	13B-2
Change-Test			
SFT	26.2	30.1	38.6
MuggleMath	60.1	62.8	67.1
Aug-Test			
SFT	14.2	17.2	22.4
MuggleMath	40.1	44.3	49.3

Table 17: We have perturbed two new test sets based on the original GSM8K test set. (A) Change-Test, is created by altering the numerical values in the GSM8K test set questions and correspondingly modifying the answers. There are 1211 query-response pairs in the Change-Test. (B) Aug-Test, is generated by augmenting the test set in the same manner as we did for the training set. There are 1378 query-response pairs in the Aug-Test.

Subject	math	GSM8k	GSM8K+ \mathcal{D}_1^i	GSM8K+ $\sum_{i=1}^3 \mathcal{D}_1^i + \mathcal{D}_2^1 + \mathcal{D}_3^1$
Counting & Probability	10.5	13.2	7.9	5.3
Algebra	7.3	12.1	12.9	16.9
Prealgebra	8.5	13.4	8.5	11.0
Geometry	2.4	9.8	4.9	2.4
Intermediate Algebra	6.2	5.2	3.1	5.2
Number Theory	3.2	6.5	6.5	8.1
Precalculus	3.6	5.4	7.1	7.1

Table 18: Transfer learning accuracy on subsets of MATH for LLaMA-13B-2

Subject	math	GSM8k	GSM8K+ \mathcal{D}_1^i	GSM8K+ $\sum_{i=1}^3 \mathcal{D}_1^i + \mathcal{D}_2^1 + \mathcal{D}_3^1$
Prealgebra	12.2	9.8	11.0	12.2
Number Theory	6.5	9.7	6.5	9.7
Algebra	7.3	5.6	5.6	15.3
Intermediate Algebra	2.1	4.1	4.1	1.0
Precalculus	3.6	1.8	3.6	1.8
Counting & Probability	5.3	7.9	5.3	13.2
Geometry	-	2.4	-	-

Table 19: Transfer learning accuracy on subsets of MATH for LLaMA-7B-2

Subject	math	GSM8k	GSM8K+ \mathcal{D}_1^i	GSM8K+ $\sum_{i=1}^3 \mathcal{D}_1^i + \mathcal{D}_2^1 + \mathcal{D}_3^1$
Prealgebra	7.3	9.8	7.3	14.6
Number Theory	6.5	3.2	1.6	3.2
Algebra	6.5	4.8	11.3	4.8
Intermediate Algebra	3.1	1.0	5.2	3.1
Precalculus	-	3.6	5.4	1.8
Counting & Probability	2.6	2.6	2.6	7.9
Geometry	4.9	4.9	2.4	2.4

Table 20: Transfer learning accuracy on subsets of MATH for LLaMA-7B

Subject	math	GSM8k	GSM8K+ \mathcal{D}_1^1
Prealgebra	8.5	12.2	14.6
Number Theory	3.2	12.9	3.2
Algebra	7.3	8.1	12.1
Intermediate Algebra	6.2	5.2	7.2
Precalculus	3.6	-	-
Counting & Probability	10.5	5.3	10.5
Geometry	2.4	7.3	4.9

Table 21: Multi-task learning accuracy on subsets of MATH for LLaMA-13B-2

Subject	math	GSM8k	GSM8K+ \mathcal{D}_1^1
Prealgebra	12.2	11.0	4.9
Number Theory	6.5	6.5	3.2
Algebra	7.3	5.6	11.3
Intermediate Algebra	2.1	6.2	-
Precalculus	3.6	-	-
Counting & Probability	5.3	5.3	2.6
Geometry	-	7.3	7.3

Table 22: Multi-task learning accuracy on subsets of MATH for LLaMA-7B-2

Subject	math	GSM8k	GSM8K+ \mathcal{D}_1^1
Prealgebra	7.3	6.1	7.3
Number Theory	6.5	3.2	3.2
Algebra	6.5	6.5	10.5
Intermediate Algebra	3.1	2.1	-
Precalculus	-	1.8	-
Counting & Probability	2.6	2.6	5.3
Geometry	4.9	9.8	2.4

Table 23: Multi-task learning accuracy on subsets of MATH for LLaMA-7B