

Integrating Arithmetic Learning Improves Mathematical Reasoning in Smaller Models

Neeraj Gangwar, Suma Bhat, Nickvash Kani

Electrical and Computer Engineering
University of Illinois Urbana-Champaign, IL, USA
{gangwar2, spbhat2, kani}@illinois.edu

Abstract

While large models pre-trained on high-quality data exhibit excellent performance on mathematical reasoning (e.g., GSM8k, MultiArith), it remains challenging to specialize smaller models for these tasks. Common approaches to address this challenge include knowledge distillation from large teacher models and data augmentation (e.g., rephrasing questions and generating synthetic solutions). Despite these efforts, smaller models struggle with arithmetic computations, leading to errors in mathematical reasoning. In this work, we leverage a synthetic arithmetic dataset generated programmatically to enhance the reasoning capabilities of smaller models. We investigate two key approaches to incorporate this dataset: (1) intermediate fine-tuning, in which a model is fine-tuned on the arithmetic dataset before training it on a reasoning dataset, and (2) integrating the arithmetic dataset into an instruction-tuning mixture, allowing the model to learn arithmetic skills alongside general instruction-following abilities. Our experiments on multiple reasoning benchmarks demonstrate that incorporating an arithmetic dataset, whether through targeted fine-tuning or within an instruction-tuning mixture, enhances models' arithmetic capabilities, thereby improving their mathematical reasoning performance.

Keywords: Mathematical Reasoning, Arithmetic Reasoning, Large Language Models

1. Introduction

Scaling the model and data sizes has had a tremendous effect on performance across various natural language processing (NLP) tasks (e.g., Grattafiori et al., 2024; Liu et al., 2024; Gemma Team et al., 2024; Qwen Team, 2024; Abdin et al., 2024; Lambert et al., 2024; Jaech et al., 2024; Chowdhery et al., 2023; Achiam et al., 2023; Touvron et al., 2023; Jiang et al., 2023). These pre-trained models can learn from a few demonstrations using in-context learning and do not require task-specific fine-tuning (Brown et al., 2020). They also benefit from generating a sequence of reasoning steps before arriving at the final answer. These strategies have been particularly effective for mathematical reasoning (Wei et al., 2022b; Nye et al., 2022; Fu et al., 2022; Zhou et al., 2022). While large models exhibit excellent performance on various reasoning benchmarks, smaller models remain essential due to their efficiency and adaptability. They require significantly fewer computational resources, making them ideal for deployment with limited infrastructure. They also enable faster inference and lower latency, which is crucial for real-time applications. Additionally, they can be fine-tuned more efficiently for specific tasks without the high costs associated with training massive models.

Traditionally, smaller pre-trained models are adapted for downstream tasks through supervised fine-tuning on a dataset or a combination of datasets formatted as instructions (also referred to as instruction tuning; Wei et al., 2022a). While

Question:

Dylan needed chicken sausages and fish sausages to make sausage buns at a party. He bought 38 chicken sausages and 6 more fish sausages than chicken sausages. How many sausages did Dylan buy in all?

No arithmetic dataset during training:

Dylan bought 38 chicken sausages.
He bought 6 more fish sausages than chicken sausages.
So he bought $6 + 38 = 42$ fish sausages.
In all, he bought $38 + 42 = 74$ sausages.
The answer is 74. ❌

Including arithmetic dataset during training:

Dylan bought 38 chicken sausages.
He bought 6 more fish sausages than chicken sausages.
So he bought $6 + 38 = 44$ fish sausages.
In total, he bought $38 + 44 = 82$ sausages.
The answer is 82. ✅

Figure 1: An example from the GSM8k test set and its solutions generated by (Top) FlanT5-Large directly fine-tuned on GSM8k, and (Bottom) FlanT5-Large fine-tuned on an arithmetic dataset before training it on GSM8k.

this approach is effective for simpler tasks, it falls short when applied to mathematical reasoning, as smaller models struggle to achieve strong performance. This challenge arises because math reasoning datasets, such as GSM8k (Cobbe et al., 2021), consist of a small number of reasoning problems, typically paired with a solution. The scarcity

of training examples makes it difficult for the model to capture the complexity of mathematical reasoning. To overcome this issue, a widely explored research direction is to distill knowledge from large pre-trained teacher models into smaller student models. Some methods use questions from existing training datasets and use prompting to generate multiple solutions per question for fine-tuning smaller models (Ho et al., 2023; Magister et al., 2023; Fu et al., 2023; Hsieh et al., 2023; Yue et al., 2024). Others use various techniques such as rephrasing questions to create more examples (Yu et al., 2024; Huang et al., 2024; Tong et al., 2024; Kaur et al., 2024) or generating multiple views of a solution (Liang et al., 2024), for better reasoning performance.

Although these methods boost mathematical reasoning performance in smaller models, the models may still struggle with arithmetic computations. In many cases, models generate the correct reasoning steps but arrive at an incorrect final answer due to numerical computation errors. See Figure 1 for an example. Some approaches address this issue by integrating external tools (Cobbe et al., 2021; Schick et al., 2023) or utilizing programs (Gao et al., 2023; Chen et al., 2023; Ye et al., 2023) to delegate these computations to external systems such as calculators. In this work, we explore whether the model performance can be improved by directly mitigating these errors without relying on external tools. Previous research has investigated ways to enhance arithmetic skills in Transformer-based models (McLeish et al., 2024; Liu and Low, 2023; Yang et al., 2023). But the effective transfer of these skills to downstream tasks, such as mathematical reasoning, remains largely unexplored. To bridge this gap, we examine how improving arithmetic skills strengthens a model’s mathematical reasoning.

In this work, we use a synthetic arithmetic dataset generated programmatically to enhance mathematical reasoning in small-frame models. We investigate two approaches for incorporating this dataset: (1) intermediate fine-tuning, in which a model is fine-tuned on the arithmetic dataset before training it on a reasoning dataset, and (2) integrating the arithmetic dataset into an instruction-tuning mixture. The first approach is inspired by transfer learning, as prior research shows that fine-tuning a model on a related dataset before training it on the target task can significantly improve its performance (Phang et al., 2018; Vu et al., 2020; Pruksachatkun et al., 2020). The second approach aligns with post-training techniques to refine pre-trained models by exposing them to diverse tasks, helping them acquire new skills and adapt better to various downstream tasks (Wei et al., 2021; Chung et al., 2024; Lambert et al., 2024).

Empirical observations using several mathematical datasets lead to the following key takeaways:

- Models fine-tuned on an arithmetic dataset before a reasoning dataset perform better than those fine-tuned directly on the reasoning dataset. Additionally, arithmetic datasets can be generated programmatically, eliminating the need for manual resources.
- Based on our observations with multiple datasets with varying mathematical reasoning tasks, we find that intermediate fine-tuning results in better out-of-domain generalization.
- Including an arithmetic dataset into the instruction-tuning mixture leads to better few-shot performance on multiple mathematical reasoning benchmarks.
- These models exhibit better robustness to numerical variations such as numerical substitution and digit expansion than models instruction-tuned on a mixture without the arithmetic dataset.

This work highlights the importance of explicit arithmetic training as a key factor in improving mathematical reasoning in smaller models. Our source code and datasets are publicly available.¹

2. Related Work

Model Specialization via Distillation. Adapting a pre-trained model for a downstream task has traditionally been done through task-specific fine-tuning. However, this approach does not work for mathematical reasoning tasks because datasets such as GSM8k do not contain enough examples to capture the complexity of mathematical reasoning. Several works have focused on distilling multi-step reasoning solutions from large language models (LLMs) to overcome this limitation. Fu et al. (2023) prompt Codex (Chen et al., 2021) to generate multiple multi-step solutions for examples in the GSM8k training set and fine-tune FlanT5 on the ones that lead to the correct final answer. Hsieh et al. (2023) use PaLM-540B (Chowdhery et al., 2023) for generating solutions and fine-tune T5 (Raffel et al., 2020) in a multi-task setting to generate the labels and rationale. Liu et al. (2023) use GPT-3.5-turbo to generate synthetic GSM8k-like examples. Yue et al. (2024) show that a hybrid of chain-of-thought and program-of-thought solutions performs better than using either format individually. In addition to using LLMs to generate more solutions, Yu et al. (2024) use LLM rephrasing and backward reasoning to augment questions and create a new dataset,

¹<https://github.com/neerajgangwar/reasoning-with-arith>

MetaMathQA. Several studies (Huang et al., 2024; Tong et al., 2024; Kaur et al., 2024) have built on these techniques by using larger teacher models to create more efficient reasoning datasets. In contrast, our approach focuses on a dataset that can be generated programmatically without using external models and addresses a specific weakness in the model.

Transfer Learning. Transfer learning has played a pivotal role in NLP. Vu et al. (2020) and Pruksachatkun et al. (2020) study the effect of intermediate fine-tuning on the model’s performance on a target task. Conneau and Lample (2019) explore cross-lingual model pre-training and show improvements in natural language inference and machine translation. Razdaibiedina et al. (2023) introduce progressive prompts, a continual learning approach with forward transfer without catastrophic forgetting. Training on large multi-task mixtures is also a common trend in NLP (Aribandi et al., 2022; Wei et al., 2022a; Chung et al., 2024; Lambert et al., 2024). Another research direction explores identifying relevant examples for a given downstream task from a huge collection of datasets, like P3 (Sanh et al., 2021). These methods create embeddings for all examples of interest using hidden states (Iverson et al., 2023) or gradients (Xia et al., 2024).

3. Methodology

In this study, we investigate two approaches for transferring arithmetic skills to the more complex domain of mathematical reasoning: (1) intermediate fine-tuning on an arithmetic dataset, and (2) incorporating the arithmetic dataset during instruction tuning.

Intermediate Fine-Tuning. Fine-tuning a model on an intermediate task before a downstream task can improve the model’s performance on the latter (Phang et al., 2018; Vu et al., 2020; Pruksachatkun et al., 2020). The downstream task is also referred to as the target task. This is called intermediate fine-tuning and can lead to successful knowledge transfer for similar intermediate and target tasks. Building on prior work on transfer learning, we use intermediate fine-tuning to transfer arithmetic abilities to mathematical reasoning.

Unlike natural language, mathematical computations are precise and do not contain redundancies, typically requiring more training examples for effective learning than NLP tasks. When a model is fine-tuned on a reasoning dataset, it must simultaneously learn to generate correct reasoning steps and arithmetic computations. Moreover, datasets such as GSM8k contain a limited number of examples, restricting both the quantity and diversity of arithmetic computations. This limitation often

leads to arithmetic errors during inference. To address these challenges, we adopt a two-step training approach. First, we fine-tune the model on an arithmetic dataset, allowing it to learn a broad range of numerical computations across diverse values. This intermediate fine-tuning ensures that the model develops a strong foundation in arithmetic. Following this, the model is fine-tuned on a reasoning dataset, where it focuses on applying its pre-learned arithmetic skills rather than learning them from limited examples. This two-step process yields fewer arithmetic errors during mathematical reasoning than models directly fine-tuned on a reasoning dataset, leading to more accurate and reliable performance.

Instruction Tuning. We also explore the impact of incorporating an arithmetic dataset during post-training, which involves additional training on vast corpora of text. In particular, we focus on instruction tuning, a post-training technique to enhance the instruction-following ability of pre-trained models (Wei et al., 2021; Chung et al., 2024). This process involves fine-tuning a pre-trained model on a diverse set of instructions and the corresponding responses. The fine-tuning mixture typically contains examples from a wide range of tasks, including the mathematical reasoning domain. To strengthen the arithmetic capabilities, we include a synthetic arithmetic dataset in this mixture. Similar to intermediate fine-tuning, this process improves the model’s arithmetic proficiency and enhances mathematical reasoning performance by enabling more accurate numerical computations within reasoning tasks.

4. Datasets

4.1. Arithmetic Dataset

We refer to Liu and Low (2023) for generating our dataset programmatically. The dataset from Liu and Low (2023) contains basic arithmetic operations: addition, subtraction, multiplication, and division. We extend it to include problems on fractions and percentages. The datasets used in this work do not require computations over large numbers; hence, using the GSM8k training set as a representative benchmark, we limit operand lengths to seven digits. Furthermore, we use log-uniform sampling to ensure that the dataset is not skewed towards numbers with more digits. This dataset contains nearly 1.3M examples. See Appendix A for samples from the generated dataset.

4.2. GSM8k Training Dataset

For the intermediate fine-tuning experiments, we use GSM8k (Cobbe et al., 2021) for model special-

ization. As the dataset does not have a validation set, we randomly sample 512 examples from the training set to create one. We use two versions of GSM8k in this work.

Original. In the first version, we use the remaining examples from the training set. This dataset contains 6961 examples. We refer to this dataset as GSM8k (Orig.).

Distilled. We generate a distilled dataset using the questions from GSM8k (Orig.) to evaluate if intermediate fine-tuning benefits tasks with large training datasets. This dataset is generated by prompting Mistral-7B (Jiang et al., 2023) using the prompt from Wei et al. (2022b). We generate 64 solutions per question and retain only those that lead to the correct final answer. Correctness is determined by matching the generated final answer with the ground truth. We do not verify reasoning chains, as manual verification is prohibitively resource-intensive. After removing duplicate solutions, the dataset contains approximately 175K examples. We refer to this dataset as GSM8k (Dist.).

4.3. Instruction Tuning Dataset

Following the recent work of Lambert et al. (2024), we use the TüLU 3 SFT mixture as the instruction tuning dataset for our work. It encompasses a broad spectrum of tasks, including general-purpose chat, complex mathematical reasoning, programming, instruction following, and persona-based synthetic data. This dataset contains nearly 1M examples, making it significantly smaller than the Flan collection (Longpre et al., 2023). However, it contains more mathematical reasoning examples than datasets such as Flan-mini (Ghosal et al., 2023).

5. Intermediate Fine-Tuning

5.1. Experimental Setup

Models and Training Details. We experiment with both encoder-decoder and decoder-only architectures. We use FlanT5 (base and large; Chung et al., 2024) and GPT2 (base, medium, and large; Radford et al., 2019). We first fine-tune these models on the arithmetic dataset for two epochs. To adapt these models for reasoning, we continue the training from these checkpoints on GSM8k. The models are fine-tuned for 20 and 100 epochs on GSM8k (Dist.) and GSM8k (Orig.), respectively. See Appendix B.1 for further details.

Baseline. For the baseline, we consider models that are directly fine-tuned on GSM8k without any intermediate fine-tuning.

Evaluation Tasks. We evaluate our approach on the GSM8k test set. We also test out-of-domain generalization using the following datasets: (1) MultiArith (Roy and Roth, 2015), which contains problems focused on basic arithmetic operations and is relatively simpler than GSM8k, (2) ASDiv (Miao et al., 2020) with diverse math problems focused on language usage patterns, and (3) SVAMP (Patel et al., 2021) with varying structures to ensure that a model cannot solve the problems by applying simple heuristics and ignoring question text.

Decoding. We use greedy decoding and self-consistency decoding during inference. For self-consistency decoding, nucleus sampling (Holtzman et al., 2019) is used with $T = 0.6$ and $p = 0.9$ to sample eight responses, and the most consistent final answer is chosen. As nucleus sampling is a stochastic decoding method, we repeat the evaluation three times and report the mean accuracy.

5.2. Results

In-Domain Performance. We first evaluate the models on the GSM8k test set. Table 1 shows the accuracy (%) achieved by various models. These results indicate that FlanT5 benefits from intermediate fine-tuning, significantly improving the GSM8k performance. Additionally, these results demonstrate that intermediate fine-tuning helps with both GSM8k (Orig.), which has a small training set, and GSM8k (Dist.), which already contains significantly more training examples.

For GPT2, we observe a slight decline in reasoning performance when the model is specialized in reasoning using GSM8k (Orig.) after fine-tuning it on the arithmetic dataset. However, this issue does not arise with GSM8k (Dist.), where we see performance gains comparable to those of FlanT5. We attribute this behavior to the fact that intermediate fine-tuning optimizes the model for arithmetic tasks, making it more challenging to adapt to other tasks compared to the original model. A larger fine-tuning dataset mitigates this issue, as demonstrated by the significant performance improvement when reasoning specialization is performed using GSM8k (Dist.).

Out-of-Domain Performance. Next, the models fine-tuned on GSM8k are evaluated on MultiArith, ASDiv, and SVAMP, and the results are shown in Table 1. These results indicate that intermediate fine-tuning does not hurt out-of-domain generalization. Conversely, the models fine-tuned on the arithmetic dataset first generalize better than those directly fine-tuned on GSM8k.

Arithmetic in Reasoning Context. While the models fine-tuned on the arithmetic dataset excel

Training Dataset	Model	IFT	Greedy Decoding				Self-Consistency Decoding			
			GSM8k	MultiArith	ASDiv	SVAMP	GSM8k	MultiArith	ASDiv	SVAMP
			Accuracy (Δ)				Accuracy (Δ)			
GSM8k (Orig.)	FlanT5-Base	×	7.7	17.2	8.5	6.6	9.1	17.4	8.6	7.7
		✓	10.5 (+2.8)	25.6 (+8.4)	12.5 (+4.0)	10.2 (+3.6)	12.0 (+2.9)	30.2 (+12.8)	14.1 (+5.5)	11.4 (+3.7)
	FlanT5-Large	×	12.9	28.9	15.7	12.1	14.7	29.1	16.8	12.6
		✓	17.1 (+4.2)	45.0 (+16.1)	20.8 (+5.1)	16.2 (+4.1)	18.6 (+3.9)	47.2 (+18.1)	21.4 (+4.6)	16.6 (+4.0)
	GPT2	×	5.5	12.8	4.6	3.9	6.9	18.3	5.2	4.7
		✓	6.6 (+1.1)	17.2 (+4.4)	6.9 (+2.3)	5.8 (+1.9)	7.7 (+0.8)	22.4 (+4.1)	8.4 (+3.2)	6.5 (+1.8)
	GPT2-Medium	×	7.9	20.0	8.8	5.9	9.0	23.7	9.3	6.5
		✓	7.4 (-0.5)	26.7 (+6.7)	10.1 (+1.3)	6.1 (+0.2)	7.7 (-1.3)	32.8 (+9.1)	11.2 (+1.9)	7.1 (+0.6)
	GPT2-Large	×	8.5	23.3	11.0	9.4	9.1	28.9	11.9	10.0
		✓	7.1 (-1.4)	26.7 (+3.4)	8.3 (-2.7)	7.7 (-1.7)	8.8 (-0.3)	25.2 (-3.7)	9.4 (-2.5)	8.9 (-1.1)
GSM8k (Dist.)	FlanT5-Base	×	17.5	31.1	23.6	20.4	19.9	33.9	24.6	20.2
		✓	21.4 (+3.9)	42.2 (+11.1)	29.6 (+6.0)	22.5 (+2.1)	23.6 (+3.7)	45.6 (+11.7)	31.2 (+6.6)	25.4 (+5.2)
	FlanT5-Large	×	22.4	45.0	29.1	23.2	24.9	43.7	30.2	25.0
		✓	27.7 (+5.3)	73.3 (+28.3)	40.0 (+10.9)	33.3 (+10.1)	30.1 (+5.2)	76.1 (+32.4)	41.4 (+11.2)	34.9 (+9.9)
	GPT2	×	14.8	44.4	24.2	17.9	18.8	46.9	25.5	18.7
		✓	19.6 (+4.8)	70.6 (+26.2)	32.4 (+8.2)	22.5 (+4.6)	22.2 (+3.4)	77.4 (+30.5)	35.3 (+9.8)	26.2 (+7.5)
	GPT2-Medium	×	21.4	55.0	27.7	18.8	24.8	56.1	29.8	23.0
		✓	25.9 (+4.5)	72.8 (+17.8)	37.7 (+10.0)	26.5 (+7.7)	29.6 (+4.8)	80.0 (+23.9)	40.6 (+10.8)	30.9 (+7.9)
	GPT2-Large	×	19.9	52.8	29.0	22.3	24.1	56.5	31.1	23.5
		✓	24.3 (+4.4)	75.0 (+22.2)	37.1 (+8.1)	30.2 (+7.9)	28.8 (+4.7)	81.7 (+25.2)	41.1 (+10.0)	33.7 (+10.2)

Table 1: Accuracy (%) of models fine-tuned on the GSM8k datasets, with and without prior intermediate fine-tuning (IFT) on the arithmetic dataset. We report results for both greedy decoding and self-consistency decoding. Performance on MultiArith, ASDiv, and SVAMP is included to demonstrate that IFT maintains out-of-domain generalization.

at basic arithmetic tasks compared to their original versions, do these skills transfer to reasoning tasks when they are fine-tuned on reasoning datasets? Moreover, are they the reason behind the better reasoning performance, as shown in Table 1? Accuracy on the test sets from various reasoning datasets, such as GSM8k and MultiArith, does not directly capture this, as an incorrect final answer can stem from multiple factors beyond arithmetic errors. To better understand the impact of arithmetic computations, we specifically look at them within the reasoning process, ensuring that all other reasoning steps remain correct.

We use the GSM8k test set and identify the arithmetic computations using the calculation annotations (enclosed within $\langle\langle\rangle\rangle$). Given a question and its solution up to an annotation, the models are prompted to generate the next tokens, which are then compared to the ground truth. We define the accuracy of these tokens as GSM8k arithmetic accuracy. See Appendix C for more details. Figure 2 shows how well the models from Table 1 handle arithmetic computations within reasoning contexts. These results suggest that intermediate fine-tuning reduces arithmetic errors compared to the models fine-tuned directly on GSM8k, leading to an average improvement of 11.7% in arithmetic computations. Even GPT2, when fine-tuned on GSM8k (Orig.), makes fewer arithmetic errors with intermediate fine-tuning. These results support the hypothesis that the slight decline in the reasoning

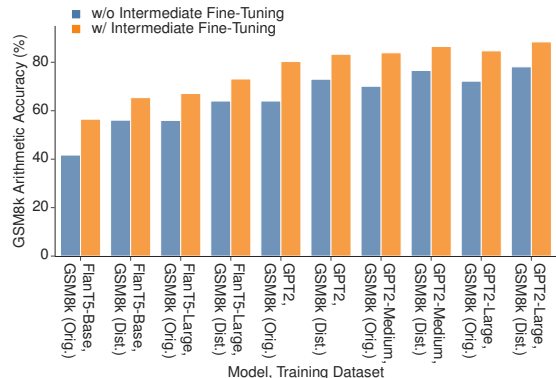


Figure 2: GSM8k arithmetic accuracy or the ability of the fine-tuned models to generate correct arithmetic computations in reasoning contexts. This evaluation is performed on the GSM8k test set.

performance of GPT2 fine-tuned on GSM8k (Orig.) after the intermediate fine-tuning results from the model being optimized for arithmetic tasks, potentially making it challenging to adapt the model for reasoning tasks with a smaller training set. A larger training dataset, GSM8k (Dist.) in this case, mitigates the issue.

Prolonged Intermediate Fine-Tuning. We next investigate the impact of extending the intermediate fine-tuning beyond two epochs. We find that fine-tuning models on the arithmetic dataset for

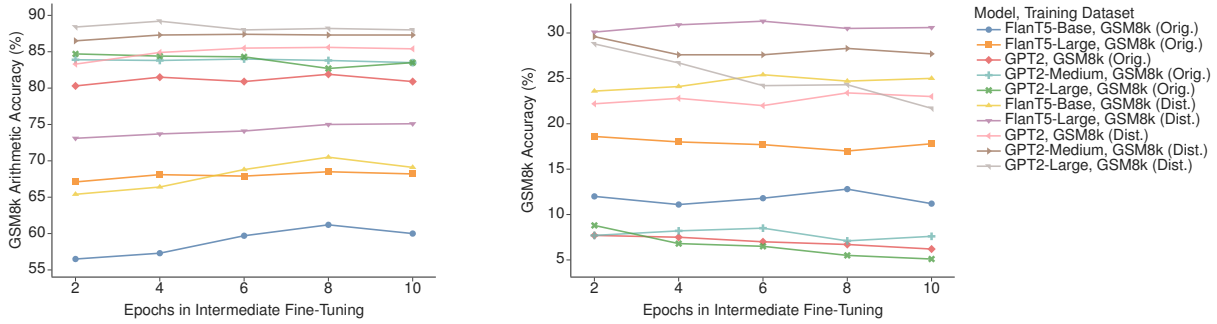


Figure 3: GSM8k arithmetic accuracy (*Left*) and GSM8k test accuracy (*Right*) of the models fine-tuned on GSM8k after the intermediate fine-tuning for different number of epochs. The GSM8k arithmetic performance saturates after two epochs of intermediate fine-tuning and yields no further improvement in the GSM8k performance.

additional epochs provides no benefit and even adversely affects the reasoning performance when models are later fine-tuned on a reasoning dataset. We attribute this outcome to two factors. First, prolonged fine-tuning on the arithmetic dataset further optimizes the model for arithmetic tasks, limiting its adaptability for other tasks. Second, the model learns the computations for solving GSM8k problems within the first few epochs, and further training does not improve this ability. Figure 3 illustrates GSM8k arithmetic accuracy and overall performance as a function of intermediate fine-tuning epochs. The results show that extended intermediate fine-tuning does not improve GSM8k arithmetic accuracy and specializes models for arithmetic tasks. This leads to no improvement or decline in GSM8k performance.

6. Instruction Tuning

6.1. Experimental Setup

Model and Training Details. As LLMs predominantly use the decoder-only architecture, we use GPT2 (large; Radford et al., 2019) for this experiment. The models are fine-tuned for five epochs. See Appendix B.2 for training details.

Baselines. We use the following two baselines: (1) the pre-trained model and (2) the model fine-tuned only on the TüLU 3 SFT mixture.

Evaluation Tasks. We use nine math reasoning datasets to evaluate the impact of including an arithmetic dataset in the instruction tuning mixture. Similar to the intermediate fine-tuning experiments, we evaluate the models on GSM8k, ASDiv, SVAMP, and MultiArith. Additionally, we consider the following datasets from MAWPS (Koncel-Kedziorski et al., 2016): (1) AddSub (Hosseini et al., 2014), which is a collection of addition and subtraction problems, (2) SingleEq (Roy et al., 2015), which contains sin-

gle equation problems, (3) SingleOp (Roy et al., 2015) with single operation arithmetic problems, and (4) SimulEq (Kushman et al., 2014) with multiple equation math word problems. AQuA (Ling et al., 2017), which contains algebraic problems in multiple-choice format, is also used for evaluation.

We also evaluate the robustness of the models against various perturbations. For this purpose, we use two datasets: GSM-Plus (Li et al., 2024) and GSM-Symbolic (Mirzadeh et al., 2024). GSM-Plus is an adversarial grade school math dataset that introduces five variations for each problem in the GSM8k test set: numerical variation, arithmetic variation, rephrasing, distractor insertion, and omissions of necessary statements. GSM-Symbolic contains 100 test problems from GSM8k for which variations can be systematically generated by altering numerical values or proper names. As this work focuses on arithmetic computations, we generate 50 variations by modifying the numerical values in the original problems for our experiments.

Decoding. We use few-shot prompting to evaluate the models. Four exemplars are used in the prompts due to the maximum sequence length limit in GPT2. We use exemplars from the prompts used in Wei et al. (2022b). See Appendix D for more information. We follow the same strategies for decoding as described in Section 5.

6.2. Results

Table 2 shows the results of instruction tuning GPT2-Large with and without including the synthetic arithmetic dataset in the fine-tuning mixture. The model fine-tuned on the TüLU 3 SFT mixture and synthetic arithmetic examples achieves better performance across math reasoning datasets, except AQuA, than the model fine-tuned only on the TüLU 3 SFT mixture, with both greedy decoding and self-consistency decoding. For self-consistency decoding, the former outperforms the latter in all

TüLU 3 SFT	Arith.	GSM8k	ASDiv	SVAMP	MAWPS	AQuA	GSM-Plus
<i>Greedy Decoding</i>							
✗	✗	1.5	2.3	3.3	2.7	16.5	2.0
✓	✗	15.0	24.5	19.9	30.3	18.1	7.0
✓	✓	16.3	36.9	28.1	43.3	22.8	9.2
<i>Self-Consistency Decoding</i>							
✗	✗	2.5	2.6	4.1	2.4	18.0	2.0
✓	✗	17.3	30.9	26.4	37.4	21.5	9.4
✓	✓	19.6	42.6	33.9	50.1	19.8	11.4

Table 2: Accuracy (%) achieved by GPT2-Large when instruction-tuned on the TüLU 3 SFT mixture with and without including the synthetic arithmetic dataset. The first rows under both greedy decoding and self-consistency decoding denote the pre-trained model.

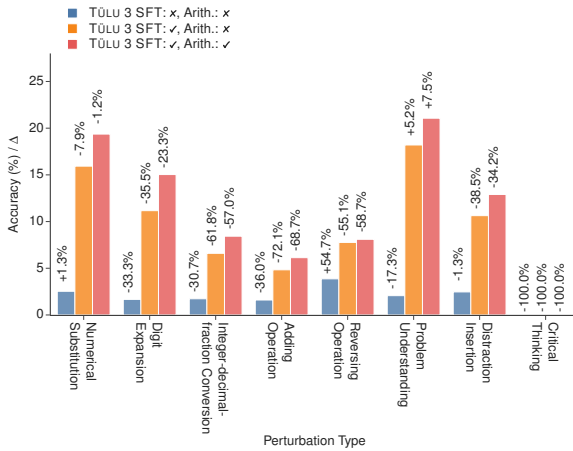


Figure 4: Performance of the pre-trained and instruction-tuned GPT2-Large models on GSM-Plus for different perturbation types using self-consistency decoding. The model fine-tuned on TüLU 3 SFT mixture and the arithmetic dataset performs better across different perturbation types. The percentages above the bars represent the performance drop relative to the original GSM8k dataset, as shown in Table 2.

three evaluation attempts across datasets. We also compute the GSM8k arithmetic accuracy for these models and find a 3% increase in accuracy by including the arithmetic dataset in the fine-tuning mixture. See Appendix E for the results on individual datasets in MAWPS.

Neither model performs well on AQuA, and the performance of all models is close to random choice. Randomly choosing an option leads to an average accuracy of $19.9\% \pm 2.7\%$ after 100 trials.

6.3. Robustness to Perturbations

GSM-Plus. The overall accuracy is shown in Table 2. The breakup of the overall performance by perturbation types is illustrated in Figure 4. The

TüLU 3 SFT	Arith.	GSM8k Acc. (%)	GSM-Symb. Acc. (%)	Δ (%)
✓	✗	15.3	8.5 ± 1.6	-44.4
✓	✓	22.0	16.3 ± 2.1	-25.9

Table 3: Performance of the instruction-tuned models using self-consistency decoding when evaluated on the original GSM8k problems vs the same problems in GSM-Symbolic. Δ indicates the performance drop.

Dataset	w/o Arith.	w/ Arith.
IFEval		
Strict (prompt level)	0.3549	0.3420
Strict (instruction level)	0.4964	0.4904
Loose (prompt level)	0.3826	0.3826
Loose (instruction level)	0.5252	0.5252
BigBench-Hard (Avg.)	19.04%	18.03%

Table 4: Performance of instruction-tuned GPT2-Large (from Section 6) on IFEval and BigBench-Hard. See Table 10 in the appendix for the detailed performance on BigBench-Hard.

model fine-tuned on the TüLU 3 SFT and arithmetic mixture performs better than the model fine-tuned only on the TüLU 3 SFT mixture across different perturbations. We further investigate the performance drop for the two models relative to the original GSM8k dataset. In particular, we are interested in numerical variations, which include numerical substitutions, digit expansions, and integer-decimal-fraction conversions. For these perturbations, we find that the model fine-tuned on the TüLU 3 SFT mixture and the arithmetic dataset sees a lower performance drop relative to the original GSM8k dataset than the model fine-tuned only on the TüLU 3 SFT mixture.

GSM-Symbolic. Table 3 shows the performance of the instruction-tuned models on GSM-Symbolic.

GSM8k	ASDiv	SVAMP	MAWPS	AQuA	GSM-Plus
12.6	36.3	29.7	47.1	22.1	7.4

Table 5: Accuracy (%) achieved by GPT2-Large (with greedy decoding) when instruction-tuned on a dataset composed of 90% of the TüLU 3 SFT examples combined with the arithmetic dataset.

Model	Greedy Decoding				Self-Consistency Decoding			
	GSM8k	MultiArith	ASDiv	SVAMP	GSM8k	MultiArith	ASDiv	SVAMP
FlanT5-Base	21.2	46.1	30.4	23.6	23.7	50.4	31.4	25.8
GPT2	19.4	67.8	31.2	20.9	23.1	75.4	34.3	25.6

Table 6: Accuracy (%) achieved by models fine-tuned on the GSM8k datasets after fine-tuning them on the decontaminated arithmetic dataset. Here, *decontamination* refers to the removal of arithmetic computations appearing in the GSM8K test set.

The performance of these models on the original problems is presented in the GSM8k column in the table. We observe that the model fine-tuned on the TüLU 3 SFT mixture with the arithmetic dataset sees a performance drop of 25.9% compared to a 44.4% drop for the model fine-tuned only on the TüLU 3 SFT mixture, indicating greater robustness to numerical changes in the original problems. We observe a similar pattern with greedy decoding. See Appendix F for details. These results indicate that including arithmetic examples in the fine-tuning mixture makes models more robust to numerical perturbations.

7. Discussion

7.1. Impact on General Abilities

In this section, we investigate the impact of incorporating the arithmetic examples in the training pipeline, given the models’ limited capacity. Because targeted fine-tuning has been shown to cause models to forget previously learned abilities, we exclude models from Section 5, which are directly fine-tuned on GSM8k, and focus solely on instruction-tuned models in this analysis. We evaluate the instruction-tuned models from Section 6 on IFEval (Zhou et al., 2023) and BigBench-Hard (Suzgun et al., 2022). The performance of these models is shown in Table 4. The IFEval results show that the impact on instruction-following abilities is minimal. On the other hand, the BigBench-Hard results indicate that model performance on other reasoning tasks is adversely impacted by including arithmetic examples in the SFT mixture.

7.2. Matching the Token Budget

The synthetic arithmetic dataset is one-tenth of the TüLU 3 SFT mixture in terms of the number of tokens. For this ablation study, we remove 10% of the

randomly selected examples from the TüLU 3 SFT mixture to match the token budget of the instruction-tuning run without the arithmetic dataset. We fine-tune GPT2-Large on a mixture of this dataset and the arithmetic dataset using the same pipeline as used in Section 6. Table 5 shows the performance of the resulting model. The performance on GSM8k and GSM-Plus degrades compared to the models instruction-tuned on all examples from the TüLU 3 SFT mixture (Table 2). These findings suggest that while the synthetic arithmetic dataset can help the model learn arithmetic computations, it cannot replace examples that help models develop broader reasoning abilities.

7.3. Data Contamination Analysis

Since the arithmetic dataset is generated programmatically, some of its arithmetic computations may overlap with those found in the downstream mathematical reasoning datasets. To ensure that the observed performance improvements are not simply due to such overlap, we investigate the extent to which potential data contamination may influence the results.

To eliminate dataset contamination as a possible reason for the improved performance in our experiments, we remove from the synthetic arithmetic dataset examples that contain arithmetic computations appearing in the GSM8k test set. The only exception to this is the computations involving single-digit numbers, as these are fundamental operations that the model is expected to memorize. We repeat the intermediate fine-tuning experiments described in Section 5, where we fine-tune FlanT5-Base and GPT2 on this dataset before subsequently fine-tuning them on GSM8k (Dist.). The results are shown in Table 6, which are largely similar to those reported in Table 1. These results indicate that the models do not merely memorize arithmetic computations, allowing us to rule out data contamination

as the cause of the performance improvements observed in our experiments.

8. Conclusion

In this work, we explored the impact of incorporating an arithmetic dataset through two approaches: intermediate fine-tuning and integration within the instruction-tuning mixture. Our experiments demonstrated that both approaches boost mathematical reasoning performance in smaller models. While intermediate fine-tuning can make subsequent fine-tuning on other tasks more challenging, the issue can be mitigated by using a larger dataset. Additionally, we found that models fine-tuned on the arithmetic-inclusive mixture demonstrate a clear performance advantage over those trained on a standard instruction set. These findings highlight the crucial role of explicit arithmetic training in strengthening mathematical reasoning in smaller models.

9. Limitations

While incorporating an arithmetic dataset improves a model's mathematical reasoning performance, smaller models still have considerable room for improvement. In this work, we include the arithmetic dataset in the training pipeline but do not investigate other factors, such as custom embedding schemes for arithmetic computation. A promising direction for future research is to incorporate insights from recent work, such as [McLeish et al. \(2024\)](#), into the model architecture. Another limitation pertains to the instruction-tuning mixture. While including the arithmetic dataset improves few-shot performance on mathematical reasoning benchmarks, data mixture ablations may be necessary to optimize the instruction-tuning mixture for better overall performance. Finally, although this study focuses on smaller models, its findings apply to larger models. The arithmetic capabilities of pre-trained models could be further enhanced by leveraging synthetic arithmetic datasets. We leave these explorations for future work.

10. Acknowledgments

This work was supported by the Illinois Campus Cluster Program and the Delta Advanced Computing and Data Resource. Delta is supported by the National Science Foundation (award OAC 2005572) and the State of Illinois and is a joint effort of the University of Illinois Urbana-Champaign and its National Center for Supercomputing Applications.

11. Bibliographical References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Hassan Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Singh Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Caio César Teodoro Mendes, Weizhu Chen, Vishrav Chaudhary, Parul Chopra, Allison Del Giorno, Gustavo de Rosa, Matthew Dixon, Ronen Eldan, Dan Iter, Abhishek Goswami, Suriya Gunasekar, Eman Haider, Junheng Hao, Russell J. Hewett, Jamie Huynh, Mojan Javaheripi, Xin Jin, Piero Kauffmann, Nikos Karampatziakis, Dongwoo Kim, Young Jin Kim, Mahoud Khademi, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuezhi Li, Chen Liang, Weishung Liu, Eric Lin, Zeqi Lin, Piyush Madan, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Liliang Ren, Corby Rosset, Sambudha Roy, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacrose, Shital Shah, Ning Shang, Hiteshi Sharma, Xianmin Song, Olatunji Ruwase, Praneetha Vadamanu, Xin Wang, Rachel Ward, Guanhua Wang, Philipp Andre Witte, Michael Wyatt, Can Xu, Jiahang Xu, Sonali Yadav, Fan Yang, Ziyi Yang, Donghan Yu, Cheng-Yuan Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yunan Zhang, Xiren Zhou, and Yifan Yang. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *ArXiv*, abs/2404.14219.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Vamsi Aribandi, Yi Tay, Tal Schuster, Jinfeng Rao, Huaixiu Steven Zheng, Sanket Vaibhav Mehta, Honglei Zhuang, Vinh Q Tran, Dara Bahri, Jianmo Ni, et al. 2022. Ext5: Towards extreme multi-task scaling for transfer learning. In *International Conference on Learning Representations*.
- Timothy J. Boerner, Stephen Deems, Thomas R. Furlani, Shelley L. Knuth, and John Towns. 2023. [Access: Advancing innovation: Nsf's advanced cyberinfrastructure coordination ecosystem: Services & support](#). In *Practice and Experience in Advanced Research Computing 2023: Computing for the Common Good*, PEARC '23, page

- 173–176, New York, NY, USA. Association for Computing Machinery.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. 2023. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *Transactions on Machine Learning Research*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. *Advances in neural information processing systems*, 32.
- Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and Tushar Khot. 2023. Specializing smaller language models towards multi-step reasoning. In *International Conference on Machine Learning*, pages 10421–10430. PMLR.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2022. Complexity-based prompting for multi-step reasoning. In *The Eleventh International Conference on Learning Representations*.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. PAL: Program-aided language models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10764–10799. PMLR.
- Gemma Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Deepanway Ghosal, Yew Ken Chia, Navonil Majumder, and Soujanya Poria. 2023. Flacuna: Unleashing the problem solving power of vicuna using flan fine-tuning.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Namgyu Ho, Laura Schmid, and Se-Young Yun. 2023. Large language models are reasoning teachers. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14852–14882.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. 2014. Learning to solve arithmetic word problems with verb categorization. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 523–533, Doha, Qatar. Association for Computational Linguistics.
- Cheng-Yu Hsieh, Chun-Liang Li, Chih-kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8003–8017.
- Yiming Huang, Xiao Liu, Yeyun Gong, Zhibin Gou, Yelong Shen, Nan Duan, and Weizhu Chen. 2024. Key-point-driven data synthesis with its enhancement on mathematical reasoning. *CoRR*.

- Hamish Ivison, Noah A Smith, Hannaneh Hajishirzi, and Pradeep Dasigi. 2023. Data-efficient fine-tuning using cross-task nearest neighbors. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9036–9061.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- 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. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Simran Kaur, Simon Park, Anirudh Goyal, and Sanjeev Arora. 2024. Instruct-skillmix: A powerful pipeline for llm instruction tuning. *arXiv preprint arXiv:2408.14774*.
- Rik Koncel-Kedziorski, Subhro Roy, Aida Amini, Nate Kushman, and Hannaneh Hajishirzi. 2016. Mawps: A math word problem repository. In *Proceedings of the 2016 conference of the north american chapter of the association for computational linguistics: human language technologies*, pages 1152–1157.
- Nate Kushman, Yoav Artzi, Luke Zettlemoyer, and Regina Barzilay. 2014. Learning to automatically solve algebra word problems. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 271–281.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, et al. 2024. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*.
- Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng Kong, and Wei Bi. 2024. Gsm-plus: A comprehensive benchmark for evaluating the robustness of llms as mathematical problem solvers. *Association for Computational Linguistics (ACL)*.
- Zhenwen Liang, Dian Yu, Xiaoman Pan, Wenlin Yao, Qingkai Zeng, Xiangliang Zhang, and Dong Yu. 2024. Mint: Boosting generalization in mathematical reasoning via multi-view fine-tuning. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11307–11318.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. Program induction by rationale generation: Learning to solve and explain algebraic word problems. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 158–167.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Bingbin Liu, Sebastien Bubeck, Ronen Eldan, Janardhan Kulkarni, Yuanzhi Li, Anh Nguyen, Rachel Ward, and Yi Zhang. 2023. Tinygsm: achieving 80% on gsm8k with one billion parameters. In *The 3rd Workshop on Mathematical Reasoning and AI at NeurIPS'23*.
- Tiedong Liu and Bryan Kian Hsiang Low. 2023. Goat: Fine-tuned llama outperforms GPT-4 on arithmetic tasks. *arXiv preprint arXiv:2305.14201*.
- 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. *arXiv preprint arXiv:2301.13688*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lucie Charlotte Magister, Jonathan Mallinson, Jakub Adamek, Eric Malmi, and Aliaksei Severyn. 2023. Teaching small language models to reason. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1773–1781.
- Sean McLeish, Arpit Bansal, Alex Stein, Neel Jain, John Kirchenbauer, Brian Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, Jonas Geiping, Avi Schwarzschild, and Tom Goldstein. 2024. [Transformers can do arithmetic with the right embeddings](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 108012–108041. Curran Associates, Inc.
- Shen-yun Miao, Chao-Chun Liang, and Keh-Yih Su. 2020. [A diverse corpus for evaluating and developing English math word problem solvers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 975–984, Online. Association for Computational Linguistics.

- Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad Farajtabar. 2024. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models. *arXiv preprint arXiv:2410.05229*.
- Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. 2022. Show your work: Scratchpads for intermediate computation with language models. In *Deep Learning for Code Workshop*.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. [Are NLP models really able to solve simple math word problems?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2080–2094. Online. Association for Computational Linguistics.
- Jason Phang, Thibault Févry, and Samuel R Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *arXiv preprint arXiv:1811.01088*.
- Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel Bowman. 2020. Intermediate-task transfer learning with pretrained language models: When and why does it work? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5231–5247.
- Qwen Qwen Team. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Anastasia Razdaibiedina, Yuning Mao, Rui Hou, Madian Khabsa, Mike Lewis, and Amjad Almahairi. 2023. Progressive prompts: Continual learning for language models. In *The Eleventh International Conference on Learning Representations*.
- Subhro Roy and Dan Roth. 2015. [Solving general arithmetic word problems](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1743–1752, Lisbon, Portugal. Association for Computational Linguistics.
- Subhro Roy, Tim Vieira, and Dan Roth. 2015. Reasoning about quantities in natural language. *Transactions of the Association for Computational Linguistics*, 3:1–13.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Stella Biderman, Leo Gao, Tali Bers, Thomas Wolf, and Alexander M. Rush. 2021. [Multitask prompted training enables zero-shot task generalization](#).
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36:68539–68551.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, , and Jason Wei. 2022. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*.
- Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. 2024. Dart-math: Difficulty-aware rejection tuning for mathematical problem-solving. *Advances in Neural Information Processing Systems*, 37:7821–7846.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Tu Vu, Tong Wang, Tsendsuren Munkhdalai, Alessandro Sordani, Adam Trischler, Andrew Mattarella-Micke, Subhransu Maji, and Mohit Iyer. 2020. [Exploring and predicting transferability across NLP tasks](#). In *Proceedings of the 2020*

- Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7882–7926, Online. Association for Computational Linguistics.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2022a. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. 2024. Less: Selecting influential data for targeted instruction tuning. *arXiv preprint arXiv:2402.04333*.
- Zhen Yang, Ming Ding, Qingsong Lv, Zhihuan Jiang, Zehai He, Yuyi Guo, Jinfeng Bai, and Jie Tang. 2023. Gpt can solve mathematical problems without a calculator. *arXiv preprint arXiv:2309.03241*.
- Xi Ye, Qiaochu Chen, Isil Dillig, and Greg Durrett. 2023. [SatLM: Satisfiability-aided language models using declarative prompting](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 45548–45580. Curran Associates, Inc.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng YU, Zhengying Liu, Yu Zhang, James Kwok, Zhengguo Li, Adrian Weller, and Weiyang Liu. 2024. [Metamath: Bootstrap your own mathematical questions for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhua Chen. 2024. Mammoth: Building math generalist models through hybrid instruction tuning. In *The Twelfth International Conference on Learning Representations*.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, et al. 2022. Least-to-most prompting enables complex reasoning in large language models. In *The Eleventh International Conference on Learning Representations*.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.

A. Synthetic Arithmetic Dataset

The dataset consists of problems of addition, subtraction, multiplication, division, fractions, and percentages, all with two operands. These problems do not require any reasoning and are simple arithmetic problems. Table 7 shows examples from each category.

B. Training Details

B.1. Intermediate Fine-Tuning

We use the AdamW optimizer (Loshchilov and Hutter, 2017) with a learning rate of 10^{-4} , a weight decay of 10^{-4} , and an effective batch size of 128. For FlanT5-Large and GPT2-Large, a learning rate warmup of 500 steps is used. The intermediate fine-tuning is performed for two epochs without validation. To adapt these models for reasoning, we continue the training from these checkpoints on GSM8k. The models are fine-tuned for 20 and 100 epochs on GSM8k (Dist.) and GSM8k (Orig.), respectively. The best checkpoint is selected based on the GSM8k validation performance.

B.2. Instruction Tuning

We use the AdamW optimizer (Loshchilov and Hutter, 2017) with 2×10^{-4} learning rate and a weight decay of 10^{-4} . A learning rate warmup of 500 steps is used. As examples in the TULU 3 SFT mixture have varied sequence lengths and differ significantly from the arithmetic examples in this regard, we use a variable batch size with approximately 0.5M tokens in each batch. The input to the model is truncated from the left to have at most 1024 tokens.

C. Arithmetic in Reasoning Context

We evaluate the frequency of numerical computation errors made by models within a reasoning context. For this evaluation, we use the GSM8k test set and identify arithmetic computations based on the provided calculation annotations (enclosed within $\langle\langle\rangle\rangle$). Figure 5 shows examples used for this evaluation. Given a question and its solution up to an annotation, models are prompted to generate five tokens. We extract the numerical value at the beginning of the generated text and compare it to the expected output to measure accuracy.

D. Prompt for Math Word Problems

We use four exemplars from the prompt from Wei et al. (2022b) to evaluate the models on math word problem datasets except AQuA. The prompt is

Type	Problem
Addition	Find the value of $3 + 1872$.
Subtraction	Solve $90352 - 19621$.
Multiplication	Compute the value of 552×4 .
Division	Determine $832 / 2$.
Fractions	$1 / 19$ of 573097 =
Percentages	Calculate 1% of 3200

Table 7: Samples from the synthetic arithmetic dataset.

Original Example:

Question: Dylan needed chicken sausages and fish sausages to make sausage buns at a party. He bought 38 chicken sausages and 6 more fish sausages than chicken sausages. How many sausages did Dylan buy in all?
Answer: He bought $38 + 6 = \langle\langle 38+6=44 \rangle\rangle 44$ fish sausages. Dylan bought $38 + 44 = \langle\langle 38+44=82 \rangle\rangle 82$ sausages in all. ##### 82

Examples for GSM8k Arithmetic Evaluation:

Model Input:

Question: Dylan needed chicken sausages and fish sausages to make sausage buns at a party. He bought 38 chicken sausages and 6 more fish sausages than chicken sausages. How many sausages did Dylan buy in all?
Let's think step by step
He bought $38 + 6 =$

Expected Output:

44

Model Input:

Question: Dylan needed chicken sausages and fish sausages to make sausage buns at a party. He bought 38 chicken sausages and 6 more fish sausages than chicken sausages. How many sausages did Dylan buy in all?
Let's think step by step
He bought $38 + 6 = 44$ fish sausages. Dylan bought $38 + 44 =$

Expected Output:

82

Figure 5: An example of how models are evaluated for arithmetic errors in reasoning contexts. This example has two arithmetic computations, each resulting in a test example for the GSM8k arithmetic evaluation.

shown in Listing 1. The prompt for AQuA is shown in Listing 2.

E. MAWPS Detailed Results

Table 8 shows the accuracy (%) achieved by the instruction-tuned models on different datasets within MAWPS.

Question: There are 15 trees in the grove. Grove workers will plant trees in the grove today. After they are done, there will be 21 trees. How many trees did the grove workers plant today?
 Let's think step by step
 There are 15 trees originally.
 Then there were 21 trees after some more were planted.
 So there must have been $21 - 15 = 6$.
 The answer is 6.

Question: If there are 3 cars in the parking lot and 2 more cars arrive, how many cars are in the parking lot?
 Let's think step by step
 There are originally 3 cars.
 2 more cars arrive.
 $3 + 2 = 5$.
 The answer is 5.

Question: Leah had 32 chocolates and her sister had 42. If they ate 35, how many pieces do they have left in total?
 Let's think step by step
 Originally, Leah had 32 chocolates.
 Her sister had 42.
 So in total they had $32 + 42 = 74$.
 After eating 35, they had $74 - 35 = 39$.
 The answer is 39.

Question: Jason had 20 lollipops. He gave Denny some lollipops. Now Jason has 12 lollipops. How many lollipops did Jason give to Denny?
 Let's think step by step
 Jason started with 20 lollipops.
 Then he had 12 after giving some to Denny.
 So he gave Denny $20 - 12 = 8$.
 The answer is 8.

Listing 1: Prompt for math word problem datasets except for AQuA.

Q: John found that the average of 15 numbers is 40. If 10 is added to each number then the mean of the numbers is?
 Answer Choices: (a) 50 (b) 45 (c) 65 (d) 78 (e) 64
 A: If 10 is added to each number, then the mean of the numbers also increases by 10. So the new mean would be 50. The answer is (a).

Q: If $a / b = 3/4$ and $8a + 5b = 22$, then find the value of a.
 Answer Choices: (a) $1/2$ (b) $3/2$ (c) $5/2$ (d) $4/2$ (e) $7/2$
 A: If $a / b = 3/4$, then $b = 4a / 3$. So $8a + 5(4a / 3) = 22$. This simplifies to $8a + 20a / 3 = 22$, which means $44a / 3 = 22$. So a is equal to $3/2$. The answer is (b).

Q: A person is traveling at 20 km/hr and reached his destiny in 2.5 hr then find the distance?
 Answer Choices: (a) 53 km (b) 55 km (c) 52 km (d) 60 km (e) 50 km
 A: The distance that the person traveled would have been $20 \text{ km/hr} * 2.5 \text{ hrs} = 50 \text{ km}$. The answer is (e).

Q: How many keystrokes are needed to type the numbers from 1 to 500?
 Answer Choices: (a) 1156 (b) 1392 (c) 1480 (d) 1562 (e) 1788
 A: There are 9 one-digit numbers from 1 to 9. There are 90 two-digit numbers from 10 to 99. There are 401 three-digit numbers from 100 to 500. $9 + 90(2) + 401(3) = 1392$. The answer is (b).

Listing 2: Prompt for AQuA.

TüLu 3 SFT	Arith.	MultiArith	AddSub	SingleOp	SingleEq	SimulEq
<i>Greedy Decoding</i>						
X	X	2.8	2.8	4.4	1.8	1.4
✓	X	39.4	11.9	45.3	43.1	6.8
✓	✓	50.0	36.7	64.2	59.6	4.8
<i>Self-Consistency Decoding</i>						
X	X	3.7	1.2	2.7	0.6	2.5
✓	X	56.1	18.7	51.4	47.7	5.3
✓	✓	65.9	41.3	68.3	65.7	5.5

Table 8: Accuracy (%) achieved by the instruction-tuned GPT2-Large models on datasets in MAWPS. The first rows under both greedy and self-consistency decoding denote the pre-trained model.

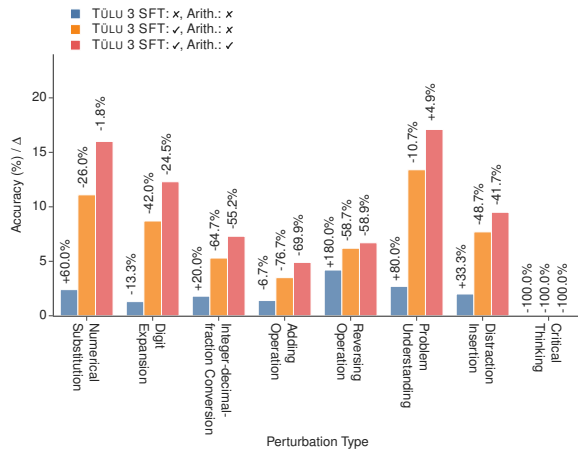


Figure 6: Performance of the pre-trained and instruction-tuned GPT2-Large models on GSM-Plus for different perturbation types using greedy decoding. The model fine-tuned on TüLU 3 SFT mixture and the arithmetic dataset performs better across different perturbation types. The percentages above the bars represent the performance drop relative to the original GSM8k dataset, as shown in Table 2.

TüLU 3 SFT	Arith.	GSM8k Acc. (%)	GSM-Symb. Acc. (%)	Δ (%)
✓	✗	10.0	6.8 ± 1.9	-32.0
✓	✓	17.0	13.3 ± 2.6	-21.8

Table 9: Performance of the post-trained models using greedy decoding when evaluated on the original GSM8k problems vs the same problems in GSM-Symbolic. The performance drop is indicated by Δ .

F. Robustness with Greedy Decoding

We also evaluate the robustness of post-trained models with greedy decoding. Our results show that the model fine-tuned on the TüLU 3 SFT mixture and the arithmetic dataset is more robust to numerical changes than the model fine-tuned only on the TüLU 3 SFT mixture. The results on GSM-Plus and GSM-Symbolic are illustrated in Figure 6 and Table 9, respectively.

G. Dataset Statistics

Table 11 shows the statistics of the datasets used in this work.

Task	w/o Arith.	w/ Arith.
boolean_expressions	55.20	53.20
causal_judgement	13.37	10.70
date_understanding	17.20	16.80
logical_deduction_five_objects	15.20	16.80
logical_deduction_seven_objects	5.60	0.40
logical_deduction_three_objects	27.20	28.80
movie_recommendation	8.00	16.40
navigate	58.40	42.40
object_counting	8.40	2.00
penguins_in_a_table	10.27	10.96
reasoning_about_colored_objects	6.40	12.00
snarks	13.48	20.22
sports_understanding	55.60	47.60
tracking_shuffled_objects_five_objects	11.60	10.80
tracking_shuffled_objects_seven_objects	0.40	0.40
tracking_shuffled_objects_three_objects	30.00	33.20
web_of_lies	4.80	1.60
word_sorting	1.60	0.40

Table 10: Detailed BigBench-Hard performance of GPT2-Large instruction-tuned on the TüLU 3 SFT mixture, with and without the inclusion of the arithmetic dataset.

Dataset	# Examples	Source
<i>Training and Validation</i>		
GSM8k (Orig.)	6,961	HF openai/gsm8k
GSM8k (Dist.)	175,668	-
Arithmetic	1,290,175	-
TüLU 3 SFT Mix.	896,090	HF allenai/tulu-3-sft-mixture
GSM8k (val.)	512	HF openai/gsm8k
<i>Test</i>		
GSM8k	1,319	HF openai/gsm8k
ASDiv	2,305	GH chaochun/nlu-asdiv-dataset
SVAMP	1,000	GH arkilpatel/SVAMP
MultiArith	180	HF ChilleD/MultiArith
AddSub	109	HF allenai/lila
SingleOp	159	HF allenai/lila
SingleEq	109	HF allenai/lila
SimulEq	146	HF allenai/lila
AQuA	254	GH google-deepmind/AQuA
GSM-Plus	10,552	HF qintongli/GSM-Plus
GSM-Symbolic	100 × 50	GH apple/ml-gsm-symbolic

Table 11: The number of examples in the datasets used in this work and their sources. HF and GH indicate HuggingFace and GitHub, respectively.