

Solving Quantitative Reasoning Problems with Language Models

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*

Google Research

Abstract

Language models have achieved remarkable performance on a wide range of tasks that require natural language understanding. Nevertheless, state-of-the-art models have generally struggled with tasks that require quantitative reasoning, such as solving mathematics, science, and engineering problems at the college level. To help close this gap, we introduce Minerva, a large language model pretrained on general natural language data and further trained on technical content. The model achieves state-of-the-art performance on technical benchmarks without the use of external tools. We also evaluate our model on over two hundred undergraduate-level problems in physics, biology, chemistry, economics, and other sciences that require quantitative reasoning, and find that the model can correctly answer nearly a third of them.

1 Introduction

Artificial neural networks have seen remarkable success in a variety of domains including computer vision, speech recognition, audio and image generation, translation, game playing, and robotics. In particular, large language models have achieved excellent performance across a variety of natural language tasks including common-sense reasoning, question answering, and summarization (Raffel et al., 2019; Brown et al., 2020; Rae et al., 2021; Smith et al., 2022; Chowdhery et al., 2022). However, these models have struggled with tasks that require quantitative reasoning, such as solving mathematics, science, and engineering problems (Hendrycks et al., 2021; Cobbe et al., 2021).

Quantitative reasoning problems are an interesting domain of application for language models because they test the capability of models on several fronts. They require the solver to correctly parse a natural language input, potentially recall world knowledge that pertains to the problem, and apply an algorithm or series of computations to the information provided in order to arrive at a correct solution. They also require that the solver is able to correctly parse and generate precise sequences of mathematical tokens, as well as apply a computational procedure to tokens via symbolic and numerical manipulation. Finally, such problems are a proving ground for research toward robust quantitative reasoning solvers that are useful in supporting the work of humans in scientific and technical fields.

Previous research has shown that large language models achieve impressive performance on math and programming questions after training on domain specific datasets (Chen et al., 2021; Austin et al., 2021;

*Equal leadership and advising contribution

†Equal contribution

Question: A line parallel to $y = 4x + 6$ passes through $(5, 10)$. What is the y -coordinate of the point where this line crosses the y -axis?

Target: Since the line is parallel to $y = 4x + 6$, we know that its slope is 4. Thus the equation of the line is $y = 4x + b$ where b is the y -intercept. Plugging $(5, 10)$ into this gives $10 = 4 \cdot 5 + b \Rightarrow b = \boxed{-10}$ which is what we wanted.

Model output: The line is parallel to $y = 4x + 6$, which means that it has the same slope as $y = 4x + 6$. Since the line passes through $(5, 10)$, we can use the point-slope form of a line to find the equation of the line:

$$y - 10 = 4(x - 5)$$

$$y - 10 = 4x - 20$$

$$y = 4x - 10$$

Now that we have the equation of the line, we can find the y -coordinate of the point where the line crosses the y -axis by substituting $x = 0$ into the equation:

$$y = 4 \cdot 0 - 10 = \boxed{-10}.$$

Figure 1: Question from the MATH dataset (Algebra 4) with a solution generated by Minerva 62B. The solution was selected by majority vote over 256 samples. This and other MATH questions were prompted with the same 4-shot prompt. The model derives the solution by writing down a line equation, simplifying it, substituting x , and solving for y . The model outputs valid \LaTeX along with natural language explanations. The approach taken by the model is different from the one used in the ground truth target.

Drori et al., 2021). In this work, we apply this approach to quantitative reasoning problems, in which the model is required to provide a complete and self-contained solution, without relying on external tools. Such tasks include mathematics word problems, competition mathematics evaluations, and many problems in science and engineering.

1.1 Our Contribution

We present Minerva, a language model that achieves strong performance on many quantitative reasoning tasks. Our model is able to process scientific and mathematical questions formulated in natural language, and generate step-by-step solutions using correct \LaTeX notation. Figures 1, 2 and 3 show example responses from Minerva to mathematics and physics questions.

Minerva is based on the PaLM general language models Chowdhery et al. (2022) that are further trained on a high-quality dataset containing scientific and mathematical data. In particular, we start with 8B, 62B, and 540B parameter pretrained models, and continue training them on our technical content dataset. We achieve state-of-the-art performance on MATH Hendrycks et al. (2021), GSM8k Cobbe et al. (2021), and a STEM subset of the MMLU Hendrycks et al. (2020) dataset, all of which include math and science questions formulated in natural language. Notably, our models achieve robust performance in a few-shot setting without being explicitly trained on these evaluation datasets.

The main novelty of this paper is a large training dataset that juxtaposes natural language with the correct use of formal mathematical language, such as equations and diagrams. The data is collected from the arXiv preprint server and from web pages that we carefully process to minimise the loss of mathematical content. In addition to establishing new state-of-the-art results, this work establishes a new baseline for the performance achievable on quantitative reasoning benchmarks by increasing data quality and model size.

To further increase the coverage of quantitative reasoning benchmarks, we build a dataset of over 200

Question: Assume that the variance of the first n natural numbers is 10, and the variance of the first m even natural numbers is 16. Compute $m + n$.

Model output: Let σ^2 be the variance of the first m even natural numbers, and τ^2 be the variance of the first n natural numbers. Then, we have $\tau^2 = \frac{n^2-1}{12}$ and $\sigma^2 = \frac{m^2-1}{3}$. We solve the system of equations:

$$\frac{n^2 - 1}{12} = 10$$

$$\frac{m^2 - 1}{3} = 16.$$

This translates into the system

$$n^2 = 121$$

$$m^2 = 48 + 1 = 49.$$

Therefore, $n = \boxed{11}$ and $m = \boxed{7}$, so $n + m = \boxed{18}$.

Question: For every $a, b, b \neq a$ prove that

$$\frac{a^2 + b^2}{2} > \left(\frac{a+b}{2}\right)^2.$$

Model output:

$$\frac{a^2+b^2}{2} > \left(\frac{a+b}{2}\right)^2$$

$$\iff \frac{a^2+b^2}{2} > \frac{a^2+b^2+2ab}{4}$$

$$\iff a^2 + b^2 > \frac{a^2+b^2+2ab}{2}$$

$$\iff 2a^2 + 2b^2 > a^2 + b^2 + 2ab$$

$$\iff a^2 + b^2 > 2ab$$

$$\iff a^2 + b^2 - 2ab > 0$$

$$\iff (a - b)^2 > 0$$

which is true, because the square of a real number is positive.

Figure 2: Example questions with solutions from Minerva 540B. **Left:** Question from the Joint Entrance Examination Main Math 2020 exam taken each year by almost 2M Indian high-school students intending to study engineering and similar fields. **Right:** Question from the National Math Exam in Poland (May 2022). The exam is taken by approximately 270K high-school students every year.

undergraduate-level questions in science and mathematics from MIT’s OpenCourseWare (OCW). This provides a measure of our model’s quantitative reasoning abilities in a chain-of-thought context beyond a pure mathematical setting.

1.2 Related Works

Solving quantitative reasoning problems expressed in natural language has been an active area of study (Koncel-Kedziorski et al., 2015; Hosseini et al., 2014). Prompting language models using scratchpad Nye et al. (2021) or chain-of-thought Wei et al. (2022) solutions can lead them to output step-by-step solutions to unseen problems. The GSM8k work Cobbe et al. (2021) showed that training verifiers to rerank model outputs can lead to improved performance. The original version of GSM8k included special syntax for algebraic calculations, which were processed by a calculator. In this work we focus on self-contained models without access to external tools.

The standard method for evaluating language models on generative tasks is to greedily sample one solution per problem. Recent works (Chen et al., 2021; Li et al., 2022; Thoppilan et al., 2022; Wang et al., 2022) have shown that it is advantageous to sample multiple solutions per problem, and then filter those down to a final answer. We find that majority voting (Wang et al., 2022) significantly improves performance over greedy decoding.

The work Drori et al. (2021) includes an evaluation of davinci-002, OpenAI’s latest publicly available language

Question: A uniform solid disk starts from rest and rolls down an inclined plane without slipping. After some time, what fraction of the disk's total kinetic energy is rotational kinetic energy? What of the following is the right choice? Explain your answer.

(A) 1/4 (B) 1/3 (C) 1/2 (D) 2/3

Model output: The disk has both translational and rotational kinetic energy. Let v be the speed of the center of mass of the disk, I be the moment of inertia of the disk, M be the mass of the disk, and R be the radius of the disk. The translational kinetic energy of the disk is given by

$$K_t = \frac{1}{2}Mv^2.$$

The rotational kinetic energy of the disk is given by

$$K_r = \frac{1}{2}I\omega^2 = \frac{1}{2}\left(\frac{1}{2}MR^2\right)\left(\frac{v}{R}\right)^2 = \frac{1}{4}Mv^2.$$

Therefore, the ratio of the rotational kinetic energy to the total kinetic energy is

$$\frac{K_r}{K_t + K_r} = \frac{\frac{1}{4}Mv^2}{\frac{1}{2}Mv^2 + \frac{1}{4}Mv^2} = \frac{\frac{1}{4}Mv^2}{\frac{3}{4}Mv^2} = \frac{1}{3}.$$

Final Answer: The final answer is (B).

Question: A particle moves so that it is at $(3 \sin \frac{t}{4}, 3 \cos \frac{t}{4})$ at time t . Find the speed of the particle, measured in unit of distance per unit of time.

Target: The parametric curve $(x, y) = (3 \sin \frac{t}{4}, 3 \cos \frac{t}{4})$ describes a circle with radius 3. Furthermore, it makes a full revolution at time $t = 8\pi$. (figure) Therefore, the speed of the particle is $\frac{6\pi}{8\pi} = \frac{3}{4}$.

Model output: The speed is the magnitude of the velocity vector. Therefore, we must first find the velocity vector, which is the derivative of the position vector. We have

$$\mathbf{v} = \frac{d}{dt} \left(3 \sin \frac{t}{4}, 3 \cos \frac{t}{4} \right) = 3 \left(\cos \frac{t}{4}, -\sin \frac{t}{4} \right) \frac{1}{4} = \left(\frac{3}{4} \cos \frac{t}{4}, -\frac{3}{4} \sin \frac{t}{4} \right).$$

Therefore, the magnitude of the velocity vector is

$$\begin{aligned} |\mathbf{v}| &= \sqrt{\left(\frac{3}{4} \cos \frac{t}{4}\right)^2 + \left(-\frac{3}{4} \sin \frac{t}{4}\right)^2} = \sqrt{\frac{9}{16} \cos^2 \frac{t}{4} + \frac{9}{16} \sin^2 \frac{t}{4}} \\ &= \frac{3}{4} \sqrt{\cos^2 \frac{t}{4} + \sin^2 \frac{t}{4}} = \frac{3}{4}. \end{aligned}$$

Therefore, the speed of the particle is $\frac{3}{4}$.

Figure 3: Physics questions with their Minerva solutions. **Top:** Example question from MMLU (College Physics) and the solution generated by Minerva 540B. The model introduces notation to describe the moving object's physical properties, writes down the relevant physical relationships, recalls the moment of inertia of a solid disk, and then correctly computes the desired energy ratio. **Bottom:** Question from the MATH dataset (Precalculus 4). The model correctly takes the derivative of a position vector to compute the velocity. It then computes the norm of the resulting vector, and uses a trigonometric identity to obtain a final numerical answer. Minerva takes a completely different approach from the ground truth solution.

model, on a subset of 90 problems from the MATH dataset. Due to the focus on a subset of questions, as well as changes made to the way questions are formatted, it is difficult to directly compare our results with those of [Drori et al. \(2021\)](#). In Section 3, we compare OpenAI davinci-002 with our models under the same experimental conditions.

Code generation. Applying code generating models to mathematical problems has been an active area of exploration. PaLM [Chowdhery et al. \(2022\)](#) showed that a large language model with code in its training dataset can achieve good performance on a code version of GSM8k. Furthermore, the Codex model ([Chen et al., 2021](#)) can generate code solutions to MATH problems [Drori et al. \(2021\)](#). These solutions often rely on external libraries to perform mathematical operations such as solving equations or taking limits. This is a complementary approach to ours, in which we directly probe the model’s ability to arrive at an answer by relying only on its own reasoning capability.

Formal mathematics. Mathematics developed as a discipline based in natural language, but its axiomatic fundamentals make it possible to simulate mathematical thinking. This can be achieved using specialized programming languages that facilitate the simulation of logical and mathematical thinking using a computer, such as Coq ([development team, 2022](#)), Isabelle ([Wenzel et al., 2008](#)), HOL4 ([Harrison, 1996](#)), Lean ([de Moura et al., 2015](#)), Metamath ([Megill and Wheeler, 2019](#)) and Mizar ([Grabowski et al., 2010](#)). Work on automation of proof assistants and automated theorem provers such as E ([Schulz, 2013](#)), leanCoP ([Otten, 2008](#)), and Vampire ([Kovács and Voronkov, 2013](#)) has substantially benefited from integration with machine learning methods ([Alemi et al., 2016](#); [Goertzel et al., 2021](#); [Li et al., 2021](#); [Polu and Sutskever, 2020](#); [Kaliszyk et al., 2018](#)).

Language models applied to formal and synthetic mathematical problems. Previous work trained language models to predict mathematical expressions [Rabe et al. \(2021\)](#); [Li et al. \(2021\)](#); [Polu and Sutskever \(2020\)](#); [Wu et al. \(2021\)](#); [Han et al. \(2022\)](#); [Polu et al. \(2022\)](#); [Jiang et al. \(2022\)](#); [Wu et al. \(2022\)](#). In turn, such a predictive model can be used to guide a proof search, as done by [Polu and Sutskever \(2020\)](#). Large language models excel in modelling natural language, though in the case of formal languages, models that facilitate retaining information about the graph structure of a given mathematical formula, such as GNNs, are still very competitive.

Modelling mathematics as a discipline of natural language. New benchmark datasets ([Hendrycks et al., 2021](#); [Welleck et al., 2021](#)) cover more advanced mathematical topics. In this domain language models are facing limited competition from other classes of models.

2 Training and Evaluation

2.1 Mathematical Training Dataset

Our models were trained on a dataset of 38.5B tokens from webpages filtered for mathematical content and from papers submitted to the arXiv preprint server. In addition, the dataset includes general natural language data, which is the same dataset that was used for pretraining PaLM. Our mathematical webpage dataset was constructed by collecting pages that contain mathematical expressions in MathJax format. The pages underwent a cleaning process that removes most HTML tags but preserves mathematical notation, including L^AT_EX symbols and formatting. The result is that mathematical formulae like $e^{\pi i} + 1 = 0$ or $E = mc^2$ are presented in full to the model during training. This procedure makes it possible for the model to perform well on tasks that require calculation and symbolic manipulation. Table 1 provides a breakdown of the training dataset. See Appendix B for more details.

Table 1: Proportion of data, and number of tokens, from each source in the technical training dataset. The General Natural Language dataset is a subset of the dataset used to pretrain the model.

Data source	Proportion of data	Tokens	Present during pretraining
Math Web Pages	47.5%	17.5B	No
arXiv	47.5%	21.0B	No
General Natural Language Data	5%	>100B	Yes

2.2 Models and Training Procedure

Our approach is to start with the PaLM pretrained decoder-only transformer language models [Chowdhery et al. \(2022\)](#), and further train (finetune) them on our mathematical dataset using an autoregressive objective. Table 2 contains the main model and training hyperparameters. The largest model, with 540B parameters, was finetuned on 26B tokens. While this model is highly undertrained compared to the 8B and 62B models, it still achieves superior performance. Additional details can be found in Appendix C.

Table 2: Model architecture and continued training hyperparameters. Model training was resumed from the pretrained PaLM models, and the number of steps quoted refers only to continued training on our technical dataset.

Model	Layers	Heads	d_{model}	Parameters	Steps	Tokens
Minerva 8B	32	16	4096	8.63B	624k	164B
Minerva 62B	64	32	8192	62.50B	416k	109B
Minerva 540B	118	48	18 432	540.35B	399k	26B

2.3 Evaluation Datasets

We mainly focus on few shot evaluation, though see Appendix E.3 for a discussion of finetuned evaluation. For evaluation, we truncate the inputs from the left to 1024 tokens and we use the model to generate up to 512 tokens. When sampling once per problem, we sample greedily. When sampling multiple times per problem we use nucleus sampling ([Holtzman et al., 2019](#)) with temperature $T = 0.6$, $p = 0.95$. For generative tasks, the model produces a chain-of-thought answer and demarcates a final answer. We evaluate a solution as correct if the final answer matches the ground truth solution, independent of the quality of the chain-of-thought preceding it. To evaluate correctness, we parse the final answers and compare them using the SymPy library ([Meurer et al., 2017](#)). This is done in order to correctly identify answers that are mathematically equivalent such as $1/\sqrt{3}$ and $\sqrt{3}/3$. See Appendix D.1 for further details.

The existing datasets on which we focus are:

- MATH: a dataset of 12K middle school and high school mathematics problems [Hendrycks et al. \(2021\)](#). Problem statements are written in L^AT_EX. We prompt the model with a fixed 4-shot prompt (listed in Appendix D.2). This prompt includes four random examples from the training dataset whose ground truth targets are not too long.
- GSM8k: middle school math word problems [Cobbe et al. \(2021\)](#). Models are evaluated using the chain-of-thought prompt from Wei et al. [Wei et al. \(2022\)](#). Previous models evaluated on GSM8k made use of an external calculator. In this work, our model does not have access to any external tools.

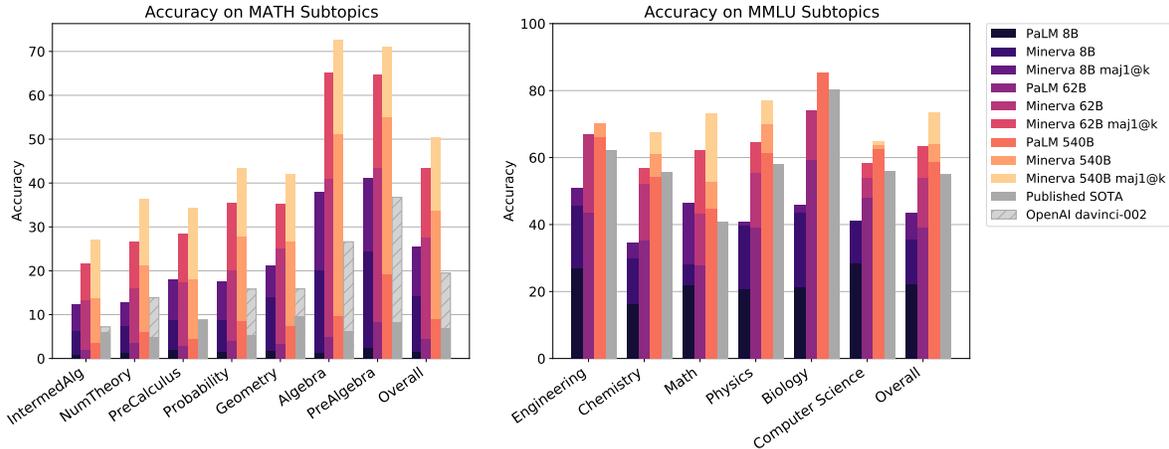


Figure 4: Performance on MATH and MMLU-STEM by subtopic. Minerva achieves state-of-the-art results on both datasets. `maj1@k` denotes evaluations where k samples were generated for each problem and only the most common answer was selected (Wang et al., 2022). For MATH, $k = 256$ for Minerva 8B and 62B, and $k = 64$ for 540B. For MMLU-STEM, $k = 16$. `davinci-002` is the latest publicly available language model from OpenAI.

- MMLU-STEM: subset of the MMLU dataset (Hendrycks et al., 2020) focused on science, technology, engineering, and mathematics (STEM). For the original version, we use the 5-shot prompt from the development set for each task. We also consider chain-of-thought prompting for this task, where we prompt the model with examples that include step-by-step solutions. We use a multiple-choice version of the MATH prompt for topics that involve mathematical reasoning, and add step-by-step solutions to the standard 5-shot prompts for the rest of the topics. See Appendix G for more details.

2.4 Undergraduate-Level STEM Problems

To evaluate the scientific reasoning capabilities of Minerva, we harvested a set of STEM problems at the undergraduate level, most of which involve multi-step reasoning, which we refer to in this paper as OCWCourses. Using publicly-available course materials offered by MIT (OpenCourseWare), we collected problems with automatically-verifiable solutions (either numeric or symbolically verifiable via `SYMPY`) from courses including “solid-state chemistry”, “information and entropy”, “differential equations”, and “special relativity.” These problems were processed by contractors to be self-contained and to have a clearly-delineated final answer. Problems asking for a proof or open-ended short answer were not included. In total we curated 272 problems, 191 of which have numeric solutions and 81 have symbolic solutions. In Appendix F, we detail the contributions from each course, and the process of converting these course materials into a format suitable for processing by language models. We also provide the text of all problems. We plan to release these as part of an open-source dataset which will be detailed in an upcoming manuscript.

2.5 Inference-Time Techniques

We find that we can considerably outperform greedy decoding by sampling $k > 1$ solutions (with a non-zero temperature) and selecting one using majority voting Wang et al. (2022). This consists of grouping predictions with respect to their final answer and selecting the most common answer. We denote this as `maj1@k`, following Li et al. (2022). A variation of this algorithm, denoted `majn@k`, involves selecting the n most

Table 3: **Model performance on several quantitative reasoning datasets.** For majority voting we use $k = 256$ (64 for 540B) samples for MATH, $k = 64$ for OCWCourses, $k = 100$ (40 for 540B) for GSM8k and $k = 16$ for MMLU-STEM. The PaLM GSM8k results do not use a calculator and were reported in (Chowdhery et al., 2022). We evaluated datasets that did not have published results on recent models on OpenAI davinci-002. Despite MMLU-STEM being a multiple choice task, we can apply majority vote by prompting the model to generate a rationale prior to the final answer, sampling multiple times, and then using majority vote on the final answers. Superscripts denote results that are quoted from previous work: ^a GPT-2 Hendrycks et al. (2021), ^b PaLM 540B maj1@40 Wang et al. (2022), and ^c Chinchilla Hoffmann et al. (2022).

	MATH	OCWCourses	GSM8k	MMLU-STEM
PaLM 8B	1.5%	1.5%	4.1%	22.0%
Minerva 8B	14.1%	7.7%	16.2%	35.6%
Minerva 8B, maj1@k	25.4%	12.5%	28.4%	43.4%
PaLM 62B	4.4%	5.9%	33.0%	39.1%
Minerva 62B	27.6%	12.9%	52.4%	53.9%
Minerva 62B, maj1@k	43.4%	23.5%	68.5%	63.5%
PaLM 540B	8.8%	7.1%	56.5%	58.7%
Minerva 540B	33.6%	17.6%	58.8%	63.9%
Minerva 540B, maj1@k	50.3%	30.8%	78.5%	75.0%
OpenAI davinci-002	19.1%	14.8%	-	-
Published SOTA	6.9% ^a	-	74.4% ^b	54.9% ^c

common answers. Intuitively, the reason majority voting improves performance is that while there are many ways to answer a question incorrectly, there are typically very few ways to answer correctly.

Contrast majority voting with `pass@k`, where a task is considered solved if any single sample solves it out of k samples. See Section 4.2 for more details on `pass@k` performance. In Appendix E.1, we report on how performance depends on k for different metrics. We find that while `pass@k` continues to improve as k is increased, majority voting performance saturates faster: 97% of the large k accuracy is achieved at $k = 64$ for MATH and $k = 16$ for GSM8k. This is likely because majority voting selects the most common answer in the modeled distribution, and the error of this estimate decreases with increasing k . This is in contrast to `pass@k` where the performance improvement comes from the tail of the distribution, which can keep improving as k is increased.

Log-likelihood is another metric that can be used to rerank samples. We found that majority voting performs significantly better than log-likelihood reranking (see Appendix E.2).

3 Results

Table 3 summarizes the results for Minerva models and other models, on the evaluation datasets described in Section 2.3. Figure 4 presents a breakdown of the MATH dataset results by subtopic. For MMLU evaluations, unless otherwise noted, performance is measured by using the standard 5-shot prompt per topic and picking the answer with the highest score. When evaluating MMLU with majority voting, we sample $k = 16$ model answers using a chain-of-thought prompt.

We present model output samples in Figures 1, 2 and 3, and additional output samples are listed in the Appendix. In addition, we evaluated Minerva 62B on the National Math Exam in Poland and found that it achieves a score of 57%, which happened to be the national average in 2021 (CKE, 2021, p. 23). The 540B

model achieves 65%.

We include results on the latest publicly available language model from OpenAI, davinci-002, evaluated using the OpenAI API with temperature set to the official recommendation ($T = 0.2$). The combination of training data, scale and inference techniques yields state of the art results on all the technical tasks that we considered. For all tasks (with the exception of GSM8k), the improvement with respect to previous results is considerable.

While our main focus is on few shot evaluation, we also tried to finetune Minerva on MATH. While we did not observe any improvement, we found that finetuning PaLM on MATH did give a significant improvement, which suggests that the marginal utility of standard finetuning decreases as the quality and diversity of the unsupervised training dataset improves. Further details can be found in Appendix E.3.

3.1 Basic arithmetic

In Appendix H, we study the performance of Minerva 540B on simple arithmetic tasks. The model achieves over 80% accuracy on 10-digit addition and over 20% accuracy on 18-digit addition.

4 Performance Analysis

4.1 Model Mistakes

To better understand the types of mistakes our models make, we compare the performance of Minerva 8B and Minerva 62B on 216 problems with high confidence majority decisions of both models. Specifically, we selected examples where the top answer received at least 15% of votes, and that either Minerva 8B was correct and Minerva 62B was incorrect (15 samples), or vice versa (201 samples). The categories and examples for each category are described in Appendix I.2.

As shown in Table 4, the prevailing errors of the 8B model were related to incorrect reasoning or calculations. Many of the calculation errors were relatively benign arithmetic mistakes. Solutions that were too short were relatively rare (in these cases, the model immediately produces an incorrect answer without any intermediate reasoning steps). Finally, in a few cases, the model hallucinates an equation or mathematical fact that is not real.

In the samples where the 62B model was incorrect, the dominating failure modes were again incorrect reasoning and incorrect calculations. In summary, we find that the 62B Minerva model retains most of the skills of the 8B model and improves upon both reasoning and calculation robustness.

Table 4: Failure modes of the 8B Minerva model, out of 201 samples which the 62B model solved correctly and the 8B model did not.

Type of mistakes	Occurrences
Incorrect reasoning	82
Incorrect calculation	70
Misunderstands question	22
Uses incorrect fact	16
Solution too short	4
Hallucinated math objects	4
Other mistakes	3

4.2 False Positives

In our approach to solving quantitative reasoning problems, we are able to automatically verify whether the final answer to a problem is correct, but we do not have an automatic way to verify the model’s chain of reasoning. This leaves open the possibility of false positives: samples which have the correct final answer, but for which the reasoning is incomplete or incorrect.

We selected 100 random questions from MATH (20 per difficulty level), along with answers sampled at zero temperature from the 62B model. We then manually inspected the answers to determine the false positive rate, which is the ratio between number of false positive examples and number of examples for which the final answer is correct; see Table 5. We found that the overall false positive rate is low, though it does increase with difficulty level.

Our focus on `pass@1` and majority voting as the primary evaluation metrics is due in part to the fact that they are less susceptible to false positives than `pass@k` (Li et al., 2022). While the `pass@256` accuracy is 84.5% for the 62B model, false positives account for part of it. We inspected the samples that failed in majority voting but passed on `pass@k` due to a single correct answer, and estimate the false positive rate for `pass@256` to be 30% among samples selected in this way. After removing false positives, we estimate that the `pass@256` accuracy to be bigger than 68%; see Appendix I.3 for details.

Table 5: Estimated false positive rates of the 62B model on the MATH dataset, by difficulty level. The average is the estimated false positive rate on the MATH dataset, given by the average of per-level false positive rates weighted by positive rates.

	Difficulty level					Average
	1	2	3	4	5	
False positive rate	< 5%	10%	< 5%	15%	30%	8%

5 Memorization

A central question in interpreting Minerva’s solutions is whether performance reflects genuine analytic capability or instead rote memorization. This is especially relevant as there has been much prior work indicating that language models often memorize some fraction of their training data (Trinh and Le, 2018; Radford et al., 2019; Carlini et al., 2022). When examining model solutions, we find that memorization of intermediate facts, such as numerical values of square roots or trigonometric identities, are crucial elements of model solutions. Truly strong performance would combine recall of intermediate facts with genuine solution synthesis. We would like to investigate a strong form of memorization, where model performance is a result of memorizing the explicit problems and solutions in our evaluation set, but also a weaker form, where the model has memorized alternate answers to the same questions.

In order to evaluate the degree to which our models solve problems by recalling information memorized from training data, we conduct three analyses on the MATH dataset. First we directly search for problems and solutions in our training corpus. Next, we generate modified versions of problems and evaluate our models’ robustness to these changes. Finally, we measure the degree of overlap between the ground truth solutions and solutions generated by our model and measure the effect of this similarity on model performance. Overall, we find little evidence that the model’s performance can be attributed to memorization.

5.1 Training and Evaluation Dataset Overlap

We selected the problems for which our 62B parameter model produced a correct answer, and filtered them to the 100 problems with the highest majority vote score, expecting that problems with a high majority vote score are more likely to have been memorized. For each of these question-answer pairs, we compute the BLEU score across chunks of 500 characters in our Math Web Pages dataset (a histogram of the BLEU scores is shown in Appendix Figure 10). We then manually inspect the 250 documents with the highest BLEU scores. While many of the top matches were from homework help sites with math questions and solutions, none of the questions matched the questions in the subset of MATH under consideration. We have included these 250 segments in Appendix J.1. We note that some problems from MATH can be found on the web. Nevertheless, this analysis concludes that these problems did not make it through our data collection process.

5.2 Performance on Modified MATH Problems

To further investigate memorization, we randomly selected twenty problems which the 62B model answered correctly under majority voting. We manually modified each problem either by introducing minor changes to problem wording (framing) or by changing the numbers which appeared in the problem and modifying the solution accordingly. We then compared the accuracy over sampled solutions before and after the modification. Results are shown in Figure 5. In both cases the accuracy before and after modifications are correlated, with no clear bias in favor of the original formulation. This is suggestive of minimal memorization. The modified problems are listed in Appendix J.2.

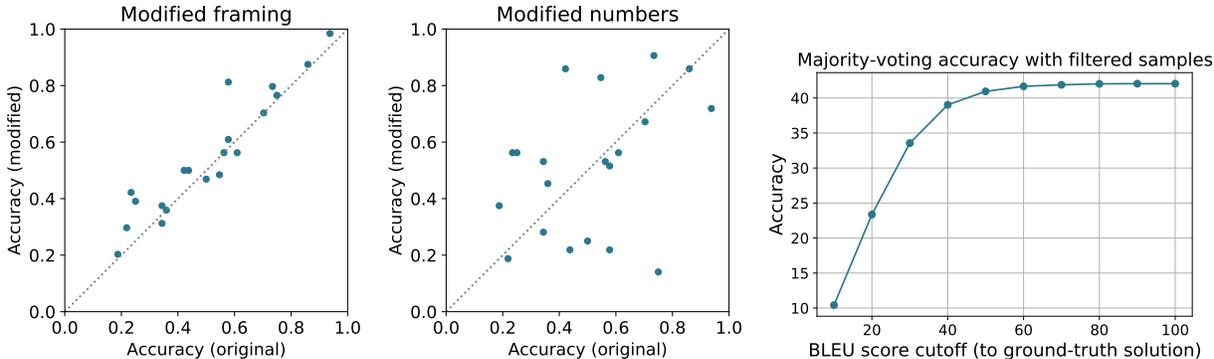


Figure 5: **Results indicating lack of memorization on MATH.** **Left, Center:** Accuracy of original questions from the MATH dataset and their modified versions. Each point represents a question. The x axis is accuracy on the original question, and the y axis is accuracy on the modified one. **Right:** Majority vote accuracy, computed only on samples with BLEU score to the ground truth solution less than or equal to the x -axis value.

5.3 BLEU Score Between Ground Truth and Generated Solutions

We seek to detect memorization of solutions by computing BLEU score between ground truth answers and model generated answers. We use the 62B model and analyze 256 samples per problem in the MATH dataset. First, we compute overlap statistics for all correct samples. We find that 160 out of 5,000 test questions have a sample with a BLEU score greater than or equal to 80 (see Appendix J.3). We note that they tend to be short solutions. To understand the effect of answer similarity on performance, we remove model samples above a certain BLEU score threshold, and recompute the majority vote accuracy. We find that majority vote performance is robust even down to relatively low similarities (see Figure 5), indicating that performance cannot be attributed to model outputs that are very similar to ground truth answers.

6 Conclusions and Discussion

In this work, we take an approach to quantitative reasoning that relies on solving problems using mathematical reasoning expressed in natural language. We show that by training a large language model on a high quality mathematical dataset, we are able to achieve strong performance on tasks that require logical reasoning, numerical calculation, and symbolic manipulation. Our model does not make use of external tools, and at inference time relies exclusively on autoregressive sampling to achieve this performance. Complementary approaches to quantitative reasoning include code-generating models and formal methods. These are all different routes toward a common goal: an agent that can reason about and solve quantitative problems. We believe that such an agent should combine useful elements from all of these approaches.

6.1 Limitations of Our Approach

Our approach to quantitative reasoning has several limitations. First, we have no automatic way of verifying the correctness of the model’s answers. This is in contrast to formal approaches, for which automatic verification is intrinsic. Second, our model has no access to external tools such as a calculator or a Python interpreter. It is therefore limited in its ability to perform quantitative reasoning tasks that require complicated numerical calculations. Third, because our model was trained on a large amount of data, we have little direct control over the specific capabilities that the model acquired.

6.2 Societal Impact

Artificial neural networks capable of solving quantitative reasoning problems in a general setting have the potential of substantial societal impact. Minerva, while a step in this direction, is still far from achieving this goal, and its potential societal impact is therefore limited. The model’s performance is still well below human performance, and furthermore, we do not have an automatic way of verifying the correctness of its outputs. If these issues could be solved, we expect the impacts of this model to be broadly positive. A direct application could be an accessible and affordable math tutor which could help improve educational inequalities.

7 Acknowledgments

We thank David Andre, Jacob Austin, Maarten Bosma, Aakanksha Chowdhery, Sergey Ioffe, Colin Raffel, Charles Sutton, and Christian Szegedy for helpful discussions.

References

- Alemi, A. A., Chollet, F., Een, N., Irving, G., Szegedy, C., and Urban, J. (2016). Deepmath - Deep Sequence Models for Premise Selection. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS’16, pages 2243–2251, USA. Curran Associates Inc.
- Austin, J., Odena, A., Nye, M., Bosma, M., Michalewski, H., Dohan, D., Jiang, E., Cai, C., Terry, M., Le, Q., et al. (2021). Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Carlini, N., Ippolito, D., Jagielski, M., Lee, K., Tramer, F., and Zhang, C. (2022). Quantifying memorization across neural language models.

- Chen, M., Tworek, J., Jun, H., Yuan, Q., Pinto, H. P. d. O., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., Ryder, N., Pavlov, M., Power, A., Kaiser, L., Bavarian, M., Winter, C., Tillet, P., Such, F. P., Cummings, D., Plappert, M., Chantzis, F., Barnes, E., Herbert-Voss, A., Guss, W. H., Nichol, A., Paino, A., Tezak, N., Tang, J., Babuschkin, I., Balaji, S., Jain, S., Saunders, W., Hesse, C., Carr, A. N., Leike, J., Achiam, J., Misra, V., Morikawa, E., Radford, A., Knight, M., Brundage, M., Murati, M., Mayer, K., Welinder, P., McGrew, B., Amodei, D., McCandlish, S., Sutskever, I., and Zaremba, W. (2021). Evaluating large language models trained on code.
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H. W., Sutton, C., Gehrmann, S., Schuh, P., Shi, K., Tsvyashchenko, S., Maynez, J., Rao, A., Barnes, P., Tay, Y., Shazeer, N., Prabhakaran, V., Reif, E., Du, N., Hutchinson, B., Pope, R., Bradbury, J., Austin, J., Isard, M., Gur-Ari, G., Yin, P., Duke, T., Levskaya, A., Ghemawat, S., Dev, S., Michalewski, H., Garcia, X., Misra, V., Robinson, K., Fedus, L., Zhou, D., Ippolito, D., Luan, D., Lim, H., Zoph, B., Spiridonov, A., Sepassi, R., Dohan, D., Agrawal, S., Omernick, M., Dai, A. M., Pillai, T. S., Pellat, M., Lewkowycz, A., Moreira, E., Child, R., Polozov, O., Lee, K., Zhou, Z., Wang, X., Saeta, B., Diaz, M., Firat, O., Catasta, M., Wei, J., Meier-Hellstern, K., Eck, D., Dean, J., Petrov, S., and Fiedel, N. (2022). Palm: Scaling language modeling with pathways.
- CKE (2021). *Skale centylowe wyników - matura 2021*.
- Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., Hesse, C., and Schulman, J. (2021). Training verifiers to solve math word problems.
- de Moura, L. M., Kong, S., Avigad, J., van Doorn, F., and von Raumer, J. (2015). The lean theorem prover (system description). In Felty, A. P. and Middeldorp, A., editors, *CADE*, volume 9195 of *Lecture Notes in Computer Science*, pages 378–388. Springer.
- development team, T. C. (2022). *The Coq reference manual*. Version 8.15.
- Drori, I., Zhang, S., Shuttlesworth, R., Tang, L., Lu, A., Ke, E., Liu, K., Chen, L., Tran, S., Cheng, N., Wang, R., Singh, N., Patti, T. L., Lynch, J., Shporer, A., Verma, N., Wu, E., and Strang, G. (2021). A neural network solves, explains, and generates university math problems by program synthesis and few-shot learning at human level.
- Goertzel, Z. A., Chvalovský, K., Jakubuv, J., Olsák, M., and Urban, J. (2021). Fast and slow enigmas and parental guidance. *CoRR*, abs/2107.06750.
- Grabowski, A., Kornilowicz, A., and Naumowicz, A. (2010). Mizar in a nutshell. *J. Formalized Reasoning*, 3(2):153–245.
- Han, J. M., Rute, J., Wu, Y., Ayers, E., and Polu, S. (2022). Proof artifact co-training for theorem proving with language models. In *International Conference on Learning Representations*.
- Harrison, J. (1996). Hol light: A tutorial introduction. In Srivas, M. and Camilleri, A., editors, *Formal Methods in Computer-Aided Design*, pages 265–269, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D., and Steinhardt, J. (2020). Measuring massive multitask language understanding. *CoRR*, abs/2009.03300.
- Hendrycks, D., Burns, C., Kadavath, S., Arora, A., Basart, S., Tang, E., Song, D., and Steinhardt, J. (2021). Measuring mathematical problem solving with the math dataset.
- Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., Casas, D. d. L., Hendricks, L. A., Welbl, J., Clark, A., Hennigan, T., Noland, E., Millican, K., Driessche, G. v. d., Damoc, B., Guy, A., Osindero, S., Simonyan, K., Elsen, E., Rae, J. W., Vinyals, O., and Sifre, L. (2022). Training compute-optimal large language models.

- Holtzman, A., Buys, J., Du, L., Forbes, M., and Choi, Y. (2019). The curious case of neural text degeneration.
- Hosseini, M. J., Hajishirzi, H., Etzioni, O., and Kushman, N. (2014). Learning to solve arithmetic word problems with verb categorization. *EMNLP*, 523533.
- Jiang, A. Q., Li, W., Tworkowski, S., Czechowski, K., Odrzygózd, T., Milos, P., Wu, Y., and Jamnik, M. (2022). Thor: Wielding hammers to integrate language models and automated theorem provers. *CoRR*, abs/2205.10893.
- Kaliszyk, C., Urban, J., Michalewski, H., and Olšák, M. (2018). Reinforcement learning of theorem proving. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Koncel-Kedziorski, R., Hajishirzi, H., Sabharwal, A., Etzioni, O., and Ang, S. D. (2015). Parsing algebraic word problems into equations. *Transactions of the Association for Computational Linguistics*, 3.
- Kovács, L. and Voronkov, A. (2013). First-order theorem proving and Vampire. In *CAV*.
- Li, W., Yu, L., Wu, Y., and Paulson, L. C. (2021). Isarstep: a benchmark for high-level mathematical reasoning. In *ICLR*.
- Li, Y., Choi, D., Chung, J., Kushman, N., Schrittwieser, J., Leblond, R., Eccles, T., Keeling, J., Gimeno, F., Lago, A. D., Hubert, T., Choy, P., d’Autume, C. d. M., Babuschkin, I., Chen, X., Huang, P.-S., Welbl, J., Goyal, S., Cherepanov, A., Molloy, J., Mankowitz, D. J., Robson, E. S., Kohli, P., de Freitas, N., Kavukcuoglu, K., and Vinyals, O. (2022). Competition-level code generation with alphacode.
- Megill, N. D. and Wheeler, D. A. (2019). *Metamath: A Computer Language for Pure Mathematics*. Lulu Press, Morrisville, North Carolina. <http://us.metamath.org/downloads/metamath.pdf>.
- Meurer, A., Smith, C. P., Paprocki, M., Čertík, O., Kirpichev, S. B., Rocklin, M., Kumar, A., Ivanov, S., Moore, J. K., Singh, S., Rathnayake, T., Vig, S., Granger, B. E., Muller, R. P., Bonazzi, F., Gupta, H., Vats, S., Johansson, F., Pedregosa, F., Curry, M. J., Terrel, A. R., Roučka, v., Saboo, A., Fernando, I., Kulal, S., Cimrman, R., and Scopatz, A. (2017). Sympy: symbolic computing in python. *PeerJ Computer Science*, 3:e103.
- Nye, M., Andreassen, A. J., Gur-Ari, G., Michalewski, H., Austin, J., Bieber, D., Dohan, D., Lewkowycz, A., Bosma, M., Luan, D., Sutton, C., and Odena, A. (2021). Show Your Work: Scratchpads for Intermediate Computation with Language Models. *arXiv e-prints*, page arXiv:2112.00114.
- Otten, J. (2008). leancop 2.0 and ileancop 1.2: High performance lean theorem proving in classical and intuitionistic logic (system descriptions). In Armando, A., Baumgartner, P., and Dowek, G., editors, *Automated Reasoning*, pages 283–291, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Polu, S., Han, J. M., Zheng, K., Baksys, M., Babuschkin, I., and Sutskever, I. (2022). Formal mathematics statement curriculum learning. *CoRR*, abs/2202.01344.
- Polu, S. and Sutskever, I. (2020). Generative language modeling for automated theorem proving. *arXiv preprint arXiv:2009.03393*.
- Rabe, M. N., Lee, D., Bansal, K., and Szegedy, C. (2021). Mathematical reasoning via self-supervised skip-tree training.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Rae, J. W., Borgeaud, S., Cai, T., Millican, K., Hoffmann, J., Song, F., Aslanides, J., Henderson, S., Ring, R., Young, S., et al. (2021). Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*.

- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2019). Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Roberts, A., Chung, H. W., Levskaya, A., Mishra, G., Bradbury, J., Andor, D., Narang, S., Lester, B., Gaffney, C., Mohiuddin, A., Hawthorne, C., Lewkowycz, A., Salcianu, A., van Zee, M., Austin, J., Goodman, S., Soares, L. B., Hu, H., Tsvyashchenko, S., Chowdhery, A., Bastings, J., Bulian, J., Garcia, X., Ni, J., Chen, A., Kenealy, K., Clark, J. H., Lee, S., Garrette, D., Lee-Thorp, J., Raffel, C., Shazeer, N., Ritter, M., Bosma, M., Passos, A., Maitin-Shepard, J., Fiedel, N., Omernick, M., Saeta, B., Sepassi, R., Spiridonov, A., Newlan, J., and Gesmundo, A. (2022). Scaling up models and data with t5x and `seqio`.
- Schulz, S. (2013). System Description: E 1.8. In McMillan, K., Middeldorp, A., and Voronkov, A., editors, *Proc. of the 19th LPAR*, volume 8312 of *LNCS*. Springer.
- Smith, S., Patwary, M., Norick, B., LeGresley, P., Rajbhandari, S., Casper, J., Liu, Z., Prabhume, S., Zerveas, G., Korthikanti, V., et al. (2022). Using deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990*.
- Thoppilan, R., De Freitas, D., Hall, J., Shazeer, N., Kulshreshtha, A., Cheng, H.-T., Jin, A., Bos, T., Baker, L., Du, Y., Li, Y., Lee, H., Zheng, H. S., Ghafouri, A., Menegali, M., Huang, Y., Krikun, M., Lepikhin, D., Qin, J., Chen, D., Xu, Y., Chen, Z., Roberts, A., Bosma, M., Zhao, V., Zhou, Y., Chang, C.-C., Krivokon, I., Rusch, W., Pickett, M., Srinivasan, P., Man, L., Meier-Hellstern, K., Morris, M. R., Doshi, T., Santos, R. D., Duke, T., Soraker, J., Zevenbergen, B., Prabhakaran, V., Diaz, M., Hutchinson, B., Olson, K., Molina, A., Hoffman-John, E., Lee, J., Aroyo, L., Rajakumar, R., Butryna, A., Lamm, M., Kuzmina, V., Fenton, J., Cohen, A., Bernstein, R., Kurzweil, R., Aguera-Arcas, B., Cui, C., Croak, M., Chi, E., and Le, Q. (2022). *Lamda: Language models for dialog applications*.
- Trinh, T. H. and Le, Q. V. (2018). A simple method for commonsense reasoning. *arXiv preprint arXiv:1806.02847*.
- Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., Narang, S., Chowdhery, A., and Zhou, D. (2022). Self-consistency improves chain of thought reasoning in language models.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E., Le, Q., and Zhou, D. (2022). Chain of thought prompting elicits reasoning in large language models.
- Welleck, S., Liu, J., Bras, R. L., Hajishirzi, H., Choi, Y., and Cho, K. (2021). Naturalproofs: Mathematical theorem proving in natural language. *CoRR*, abs/2104.01112.
- Wenzel, M., Paulson, L. C., and Nipkow, T. (2008). The isabelle framework. In Mohamed, O. A., Muñoz, C., and Tahar, S., editors, *Theorem Proving in Higher Order Logics*, pages 33–38, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Wu, Y., Jiang, A. Q., Li, W., Rabe, M. N., Staats, C., Jamnik, M., and Szegedy, C. (2022). Autoformalization with large language models. *CoRR*, abs/2205.12615.
- Wu, Y., Rabe, M. N., Li, W., Ba, J., Grosse, R. B., and Szegedy, C. (2021). LIME: learning inductive bias for primitives of mathematical reasoning. In Meila, M. and Zhang, T., editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 11251–11262. PMLR.

Appendix

A Detailed Contributions

Aitor prepared the Mathematical web pages dataset and **Aitor** and **David** prepared the arXiv dataset used to train Minerva.

Aitor trained the Minerva models presented in the paper, and he, along with **David** and **Vedant**, conducted ablation studies.

Aitor, **Ambrose**, and **David** built the experimental infrastructure for training and evaluating Minerva. They, along with **Anders**, **Ethan**, **Henryk**, **Vinay**, and **Vedant** collected the evaluation datasets and conducted model evaluations.

Aitor, **Anders**, **Behnam**, **Ethan**, **Guy**, and **Vedant** conducted experiments and ablation studies on inference-time techniques.

Vedant and **Vinay** collected the OCWCourses dataset and supervised the contractors' work.

Aitor, **Ambrose**, **Anders**, **David**, **Ethan**, **Guy**, **Henryk**, **Theo**, **Vedant**, **Vinay**, and **Yuhuai** analyzed the models' results, including sample explorations to categorize model mistakes and identify false positives.

Aitor, **Anders**, and **Cem** conducted fine-tuning evaluation experiments.

Ethan, **Vedant**, and **Vinay** designed and conducted the memorization experiments.

Aitor, **Anders**, **Ethan**, **Guy**, **Henryk**, **Imanol**, **Vedant**, and **Yuhuai** wrote the paper.

Aitor, **Behnam**, **Guy**, and **Vedant** advised and led the project throughout its life cycle.

B Training Dataset Details

The two main data sources for our training dataset are arXiv papers and web pages that contain mathematics. Here we present additional details on how the data from each source was collected and processed.

B.1 arXiv

The arXiv dataset contains 2M arXiv papers up to February 2021, in \LaTeX format. If multiple \LaTeX files were present, they were concatenated. Comments were removed, and anything before the first section header or after an appendix/bibliography header was removed. The title and abstract of each paper were added to the document from the arXiv metadata. In order to retain high quality documents and maximize the information per token, papers were filtered out if they were longer than 75k tokens, had on average more than 0.6 tokens per character, had no `\section` headers, or ended up being empty after processing. The final arXiv dataset after processing includes 1.2M papers totalling 58GB of data.

B.2 Mathematical web pages

We started with a collection of web pages that included the string "`<math`" or "`MathJax-Element-`" in the raw HTML, which we used as our filter for pages that include mathematical content. We considered pages as of January 2022. We then used several heuristics to process the pages. We found empirically that these are sufficient to extract most of the available mathematical content in either \LaTeX format or ASCII-math format. The majority of the documents (about 80% of documents) have one of these two formats:

1. A majority of these HTML documents contain math in TeX or AsciiMath format inside tags of the form `<script type="math/latex">` or `<script type="math/asciimath">`.
2. Another common appearance of L^AT_EX happens with `<annotation encoding="application/x-tex">` tags inside `<math>` MathML blocks. We extract the content of these `<annotation>` blocks but do not include other content from inside the `<math>` blocks.

The remaining documents (about 20%) generally have math in MathML format, which we discarded. After extracting the content in any of the previous two forms, we removed all other content that was inside `<math>` or `` blocks, because these blocks often encode the MathML version of TeX or AsciiMath content. After filtering, processing, and selecting only English documents, the final dataset size is 60GB.

C Model and Training Procedure Details

We start with pretrained PaLM models, and perform unsupervised finetuning on our technical dataset to obtain Minerva. The models have context length 2048. They are trained with batch size 128 (except for the 540B model which was trained with batch size 32) and without dropout.

The learning rate schedule was reciprocal square-root decay, which continued the schedule of the pretrained models. The 8B model was pretrained for 1M steps and further trained for 600k additional unsupervised finetuning steps. The 62B model was pretrained for 520k steps and further trained for 400k additional unsupervised finetuning steps. The 540B model was pretrained for 257k steps and was further trained for 383k additional steps during unsupervised finetuning.

Finally, the learning rate was dropped 10x and all models were then trained for 4% additional steps. We note that these models had a significantly larger batch size during pretraining.

We used the `tf5x` framework (Roberts et al., 2022) and trained our models with v4 TPU on Google Cloud. The 8B model was trained for 14 days on a v4-128, the 62B model was trained for 17 days on a v4-512, and the 540B model was trained for 29 days on a v4-1024.

D MATH Evaluation Details

D.1 MATH Answer Normalization

Extracting and evaluating the correctness of answers to math questions is non-trivial because answers can often be presented in many different ways, both in terms of formatting (e.g. answers can be underlined, or surrounded by a box) and in terms of mathematical content (a large number can be equivalently represented as 1,000 or 1000, answers about currency potentially have the currency symbol attached to them, etc.). Here we describe how final answers are extracted and normalized. After normalization, answers are compared using SymPy (see below). Failing to normalize answers properly will typically lead to falsely identifying correct answers as incorrect (“false negatives”), and therefore to underestimate the model’s accuracy.

We first extract the final answer from the full model response, which potentially includes chain-of-thought reasoning. In the few-shot prompt, we used the format "Final Answer: The final answer is ANSWER. I hope it is correct." for every final answer. We look for this pattern in the model output and extract ANSWER.

We then apply a normalization function to this answer, shown in Listing 1. In order to develop it we manually inspected ground truth targets, samples from Minerva, and samples from OpenAI davinci-002. We were especially careful to avoid changes in the format of the ground truth target that might produce false positives.

```

1 SUBSTITUTIONS = [
2     ('an ', ''), ('a ', ''), ('.$', '$'), ('\\$', ''), (r'\ ', ''),
3     (' ', ''), ('mbox', 'text'), ('\text{and}', '\text{and}'),
4     ('\text{and}', '\text{and}'), ('\text{m}', '\text{m}')
5 ]
6 REMOVED_EXPRESSIONS = [
7     'square', 'ways', 'integers', 'dollars', 'mph', 'inches', 'ft',
8     'hours', 'km', 'units', '\\dots', 'sue', 'points', 'feet',
9     'minutes', 'digits', 'cents', 'degrees', 'cm', 'gm', 'pounds',
10    'meters', 'meals', 'edges', 'students', 'childrentickets', 'multiples',
11    '\\text{s}', '\\text{.}', '\\text{ns}', '\\text{}^2',
12    '\\text{}^3', '\\text{\\n}', '\\text{\\}', r'\mathrm{th}',
13    r'^\circ', r'\{\circ}', r'\;', r'\!', '{,}', '"', '\\dots'
14 ]
15
16 def normalize_final_answer(final_answer: str) -> str:
17     """Normalize a final answer to a quantitative reasoning question."""
18     final_answer = final_answer.split('=')[-1]
19
20     for before, after in SUBSTITUTIONS:
21         final_answer = final_answer.replace(before, after)
22     for expr in REMOVED_EXPRESSIONS:
23         final_answer = final_answer.replace(expr, '')
24
25     # Extract answer that is in LaTeX math, is bold,
26     # is surrounded by a box, etc.
27     final_answer = re.sub(r'(.*) (\$) (.*) (\$) (.*)', '$\\3$', final_answer)
28     final_answer = re.sub(r'(\text\{.\} (.*) (\})), '\2', final_answer)
29     final_answer = re.sub(r'(\textbf\{.\} (.*) (\})), '\2', final_answer)
30     final_answer = re.sub(r'(\overline\{.\} (.*) (\})), '\2', final_answer)
31     final_answer = re.sub(r'(\boxed\{.\} (.*) (\})), '\2', final_answer)
32
33     # Normalize shorthand TeX:
34     # \frac{a}{b} -> \frac{a}{b}
35     # \frac{abc}{bef} -> \frac{abc}{bef}
36     # \frac{abc}{b} -> \frac{a}{b}c
37     # \sqrt{a} -> \sqrt{a}
38     # \sqrt{a}b -> \sqrt{a}b
39     final_answer = re.sub(
40         r'(frac) ([^{}]) (.*)', 'frac{\\2}{\\3}', final_answer)
41     final_answer = re.sub(
42         r'(sqrt) ([^{}])', 'sqrt{\\2}', final_answer)
43     final_answer = final_answer.replace('$', '')
44
45     # Normalize 100,000 -> 100000
46     if final_answer.replace(',', '').isdigit():
47         final_answer = final_answer.replace(',', '')
48
49     return final_answer

```

Listing 1: Python code used to normalize final answers.

After applying this normalization function, we checked whether the formatted target and prediction strings are SymPy-equivalent. SymPy equivalence is determined by parsing the answers via `sympy.parsing.latex.parse_latex` and then checking whether subtracting the two resulting SymPy objects and applying `sympy.simplify` gives zero. We set a timeout of 5s when calling `sympy.simplify`, and labeled strings as nonequivalent if this timeout was exceeded.

For MATH problems, SymPy equivalence improved overall accuracy by around 1%. See Table 6 for the accuracies in MATH with only exact string match vs. SymPy equivalence.

Table 6: Comparing MATH accuracy when evaluating results with and without SymPy processing.

	MATH Accuracy	
	without SymPy	with SymPy
Minerva 8B	13.3	14.1
Minerva 8B Majority	24.6	25.4
Minerva 62B	26.5	27.6
Minerva 62B Majority	42.2	43.4
OpenAI davinci-002	18.7	19.1

D.2 MATH Prompt

Listing 2 shows the 4-shot prompt used when sampling answers to MATH questions. We picked it by choosing 8 random examples from MATH and selecting examples which did not include Asymptote plotting commands. We chose four examples so that most problems fit within a context length of 1024, to enable comparisons with a wide range of models.

```

1 Problem:
2 Find the domain of the expression  $\frac{\sqrt{x-2}}{\sqrt{5-x}}$ .
3
4 Solution:
5 The expressions inside each square root must be non-negative. Therefore,
6  $x-2 \ge 0$ , so  $x \ge 2$ , and  $5-x \ge 0$ , so  $x \le 5$ . Also, the denominator
7 cannot be equal to zero, so  $5-x > 0$ , which gives  $x < 5$ . Therefore, the domain of
8 the expression is  $\boxed{[2,5)}$ .
9 Final Answer: The final answer is  $[2,5)$ . I hope it is correct.
10
11 Problem:
12 If  $\det \mathbf{A} = 2$  and  $\det \mathbf{B} = 12$ , then find
13  $\det (\mathbf{A} \mathbf{B})$ .
14
15 Solution:
16 We have that  $\det (\mathbf{A} \mathbf{B}) = (\det \mathbf{A})(\det \mathbf{B})$ 
17  $= (2)(12) = \boxed{24}$ .
18 Final Answer: The final answer is  $24$ . I hope it is correct.
19
20 Problem:
21 Terrell usually lifts two 20-pound weights 12 times. If he uses two 15-pound
22 weights instead, how many times must Terrell lift them in order to lift the
23 same total weight?
24
25 Solution:
26 If Terrell lifts two 20-pound weights 12 times, he lifts a total of
27  $2 \cdot 12 \cdot 20 = 480$  pounds of weight. If he lifts two 15-pound
28 weights instead for  $n$  times, he will lift a total of  $2 \cdot 15 \cdot n = 30n$ 
29 pounds of weight. Equating this to 480 pounds, we can solve for  $n$ :
30 
$$30n = 480$$

31 
$$\rightarrow n = 480/30 = \boxed{16}$$

32
33 Final Answer: The final answer is  $16$ . I hope it is correct.
34
35 Problem:
36 If the system of equations
37 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

38
39 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

40
41 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

42
43 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

44
45 Solution:
46 If we multiply the first equation by  $-\frac{3}{2}$ , we obtain
47

```

```

48 $$6y-9x=-\frac{3}{2}a.$$Since we also know that $6y-9x=b$, we have
49
50 $$-\frac{3}{2}a=b\Rightarrow\frac{a}{b}=\boxed{-\frac{2}{3}}.$$
51 Final Answer: The final answer is $-\frac{2}{3}$. I hope it is correct.

```

Listing 2: 4-shot prompt used for MATH problems.

E Additional Evaluation Experiments

E.1 Dependence of performance on number of generated samples

We study the dependence of performance on the number of generated samples per question on MATH and GSM8k. Table 7 shows results for $\text{maj1}@k$ and $\text{maj5}@k$, and Figure 6 shows the dependence on k for $\text{pass}@k$ and majority voting. We observe that while $\text{pass}@k$ continues to improve, majority voting saturates quickly.

Table 7: Performance on MATH ($k = 256$) and GSM8k ($k = 100$) when generating k samples per task.

	MATH	GSM8k
Minerva 8B, $\text{maj1}@k$	25.4%	28.4%
Minerva 8B, $\text{maj5}@k$	47.6%	56.8%
Minerva 62B, $\text{maj1}@k$	43.4%	67.5%
Minerva 62B, $\text{maj5}@k$	64.9%	89.0%
Published SOTA	6.9%	74.5%

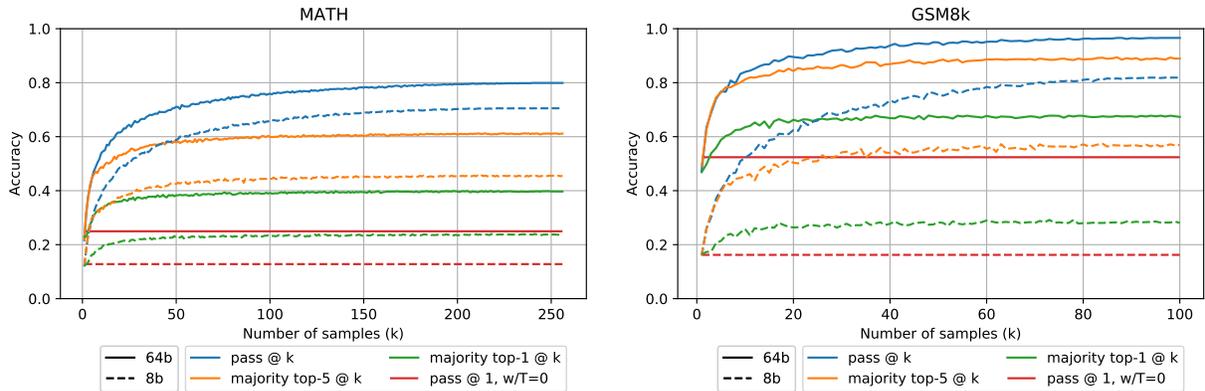


Figure 6: Accuracy as a function of k , the number of samples per task. Majority voting performance saturates quickly while $\text{pass}@k$ seems to continue improving slowly. Accuracies were computed using exact string match (without SymPy processing).

E.2 Log-Likelihood Reranking

Table 8 compares majority voting with reranking based on the log-likelihood that the model assigns to each response. We observe that majority voting is significantly better.

Table 8: A comparison of the majority voting results presented in the main text with log-likelihood reranking. We do not use SymPy processing here.

	MATH
Minerva 62B, pass1 $T = 0.0$	26.5%
Minerva 62B, Majority Voting 1@k	42.0%
Minerva 62B, pass1 $T = 0.6$	21.8%
Minerva 62B, Log-likelihood 1@k	23.8%

E.3 Finetuning on MATH

Most of our results involve few-shot prompting Minerva on MATH and other datasets on which the model was not explicitly trained. In this section we discuss finetuning our models on the training split of the MATH dataset, and then evaluating on the test split as before. We finetune both the PaLM and Minerva 8B for 3000 steps with 2048 tokens per batch with batch size 128 and dropout of 0.1. Similar to Li et al. (2022), we found that the accuracy for PaLM kept improving despite the test loss increasing. We picked the model with the best test accuracy after 50 training steps.

We finetuned using a few different prompts: A 0-shot prompt, our custom 4-shot prompt, and a prompt containing 4 random examples. Each model was evaluated using the same prompt as was used during finetuning, except for the random prompt model, which was evaluated using the fixed 4-shot prompt that we used for the non-finetuned models.

The results can be found in Table 9. Standard finetuning does not seem to improve the performance of Minerva. On the other hand, it does lead to measurable improvements in PaLM, though this performance still lagged behind Minerva. These results suggest that the marginal utility of supervised finetuning decreases as one improves the quality and diversity of the unsupervised pretraining or unsupervised finetuning dataset.

Table 9: We finetune PaLM and Minerva using different finetuning methods. We find that while finetuning helps considerably for the PaLM, it does not help for Minerva.

	MATH Accuracy	
	PaLM 8B	Minerva 8B
Few Shot	1.5%	14.1%
Custom prompt finetuning	5.6 %	13.4%
Random prompt finetuning	4.4 %	12.9%
No prompt finetuning	5.6 %	13.0%

E.4 Majority Voting Thresholds

From Figure 6, we see how majority voting saturates rather quickly at some k , while $\text{pass}@k$ keeps improving. Here we analyze the asymptotic behavior of majority voting at large k .

Let c_i denote the sorted number of counts for answer i when we sample N times and let there be a total for A_N answers. In other words, $\sum_{i=1}^{A_N} c_i = N$, $c_i > c_{i+1}$. We expect that when sampling $k \ll N$ samples, we can model the sampling distribution as a multinomial distribution with probabilities $p_i = \frac{c_i}{N}$. This approximation will have the error of attributing $p_i = 0$ to any answer which doesn't appear in N draws, so we can't really

resolve probabilities smaller than $1/N$. This issue will not matter for our purposes as long as the maximum probability p_1 is significantly higher than $1/N$.

If we draw k samples from this multinomial distribution, we expect to *not* be able to identify the majority answer with 95% confidence as long as

$$p_1 - 2\sqrt{p_1(1-p_1)}/\sqrt{k} < p_2 + 2\sqrt{p_2(1-p_2)}/\sqrt{k} \quad (1)$$

$$k < \frac{4(\sqrt{p_1(1-p_1)} + \sqrt{p_2(1-p_2)})}{(p_1 - p_2)^2} <_{p_1=p_2=0.5} \frac{4}{(p_1 - p_2)^2} \quad (2)$$

For $k = 64$, this bound implies that the resolution for $p_1 - p_2$ is 0.25, but this is a very rough estimate. However, this exercise quantifies why and how majority voting saturates even if `pass@k` doesn't.

Another point is that in order to obtain the majority solution with 95% confidence, we need

$$p_1 - 2\sqrt{p_1(1-p_1)}/\sqrt{k} > 0 \rightarrow k > 4(1/p_1 - 1), p_1 > \frac{1}{k/4 + 1} \quad (3)$$

for $k = 64$, we can probe up to $p_1 > 0.06$.

F OCWCourses Evaluation Dataset Details

F.1 Breakdown of courses

Table 10 shows the breakdown of problems in our dataset by course. See Table 11 for a breakdown of problems by solution type.

Table 10: Problems of the OCWCourses dataset broken down by course.

Course	No. problems
Solid State Chemistry	97
Introduction to Astronomy	53
Differential equations	48
Dynamics and Control	26
Principles of Microeconomics	18
Special Relativity	11
Physical Chemistry	11
Ecology	5
Information and Entropy	3

F.2 Contractor instructions

Figure 7 shows the instructions provided to our contractor workforce.

Table 11: Answer types in OCWCourses

Answer Type	No. Problems
Numeric	191
Symbolic	81
Total	272

We would like to build a dataset of clean self-contained STEM problems and solutions written in clean and correct LaTeX code.

This dataset should have the following properties:

- **Self-contained problems with no external references:** A human should be able to solve each problem and understand the given solution without having to reference any other sources. For example, some problems reference lecture notes or a textbook. These problems should be rewritten to include the referenced information. If it takes you more than roughly five minutes to find the referenced material, please delete the problem; do not include it in the final submission.
- **No extraneous material:** The raw dataset contains extraneous data, such as headers, footers, problem numbers, and point values for problems. All of this data should be removed, so that each problem/solution pair contains only the content of the problem.
- **Clearly marked final answers:** For some problems, the solution ends in a specific value that constitutes the final answer. For example, a problem might ask the student to compute the value of an integral. In this case, the steps for computing the integral are part of the solution, but the expression that represents the antiderivative is the final answer (or in the case of a definite integral, the numerical value). When a problem has such a final answer, we ask that you annotate it using a special annotation. If such a final answer is not available, we ask that you try to define one yourself that represents the solution to the problem (though in some cases this will not be possible).
- **Including images and annotating non-essential images** If there are images in the problem, please include them with a single-line `includegraphics` command in the same way that they appear in the raw input files. To make the image render nicely, you can add a `[scale=...]` modifier; just make sure the command is on one line.

Figure 7: Instructions provided to contractors who worked on OCWCourses.

F.3 OCWCourses Prompt

```
1 Problem:
2 Subproblem 0: What is the net charge of arginine in a solution of  $\mathrm{pH}$  1.0?
3 Please format your answer as +n or -n.
4 Solution:
5 The answer is +2.
6 Final answer: The final answer is +2. I hope it is correct.
7
8 Problem:
9 Subproblem 0: Let  $z = 1 + \sqrt{3}i$ . Find  $a, b$  that satisfy the equation
10  $z^4 = a + bi$ . Express your answer as the ordered pair  $(a,b)$ .
11 Solution:
12  $z^4$  has argument  $4\pi/3$  and radius 16, so it's equal to  $-8-8\sqrt{3}i$ .
13 Thus  $a = -8, b = -8\sqrt{3}$ , and our answer is  $\boxed{(-8, -8\sqrt{3})}$ .
14 Final answer: The final answer is  $(-8, -8\sqrt{3})$ . I hope it is correct.
15
16 Problem:
17 Preamble: For each Laplace Transform  $\mathcal{Y}(s)$ , find the function  $\mathcal{Y}(t)$ :
18 Subproblem 0:
19  $\mathcal{Y}(s) = \frac{1}{(s+a)(s+b)}$ 
20 Solution:
21 We can simplify with partial fractions:
22  $\frac{1}{(s+a)(s+b)} = \frac{C}{s+a} + \frac{D}{s+b}$  find the constants
23  $C$  and  $D$  by setting  $(s=-a)$  and  $(s=-b)$ 
24 [
25 \begin{aligned}
26 \frac{1}{(s+a)(s+b)} &= \frac{C}{s+a} + \frac{D}{s+b} \\
27 1 &= C(s+b) + D(s+a) \\
28 C &= \frac{1}{b-a} \\
29 D &= \frac{1}{a-b}
30 \end{aligned}
31 ]
32 therefore
33  $\mathcal{Y}(s) = \frac{1}{b-a} \frac{1}{s+a} - \frac{1}{b-a} \frac{1}{s+b}$ 
34 [
35 By looking up the inverse Laplace Transform of  $\frac{1}{s+b}$ , we find the total
36 solution  $\mathcal{Y}(t)$ 
37 [
38  $y(t) = \frac{1}{b-a} \left( e^{-at} - e^{-bt} \right)$ 
39 ]
40 Final answer: The final answer is  $\frac{1}{b-a} \left( e^{-at} - e^{-bt} \right)$ .
41 I hope it is correct.
42
43 Problem:
44 Preamble: The following subproblems refer to the differential equation
45  $\ddot{x} + b\dot{x} + x = 0$ .
46 Subproblem 0: What is the characteristic polynomial  $p(s)$  of
47  $\ddot{x} + b\dot{x} + x = 0$ ?
48 Solution:
49 The characteristic polynomial is  $p(s) = \boxed{s^2 + bs + 1}$ .
50 Final answer: The final answer is  $s^2 + bs + 1$ . I hope it is correct.
```

Listing 3: Prompt used for OCWCourses.

F.4 Problems in OCWCourses

We provide the problems in OCWCourses as a separate file.

F.5 OCWCourses evaluation

As with the MATH dataset, special care must be taken in order to correctly extract answers and evaluate them for correctness. Here we describe how final answers are extracted and normalized. See Listing 4 for the code. During dataset creation, contractors annotated all automatically-verifiable solutions as belonging to one of several types: `symbolicexpression`, `symbolicequation`, or `numeric`. For `symbolicexpression` and `symbolicequation` answers, our approach is to convert the answer strings into SymPy quantities, and check equality programmatically. For `numeric` quantities, we first remove any units from the answer string,

then convert the answer string to a float. If either numeric quantity is close to zero, our equality condition is that the absolute value of their difference is less than a threshold (0.01) of their mean; otherwise, we use the `numpy.isclose()` comparison.

As with MATH, we first extract the final answer from the full model response, which potentially includes chain-of-thought reasoning. In the few-shot prompt, we used the format `Final Answer: The final answer is ANSWER. I hope it is correct. for every final answer.` We look for this pattern in the model output and extract ANSWER.

```

1 import numpy as np
2
3 def get_answer(s: str) -> str:
4     end_str = "I hope it is correct"
5     start_str = "Final answer: "
6     replacement_str = "The final answer is "
7
8     scrub_periods = lambda x: x.strip().rstrip('.').strip()
9
10    try:
11        ans = s.split(end_str)[0].split(start_str)[1].strip().replace(replacement_str, "")
12        ans = scrub_periods(ans)
13        return ans
14    except:
15        print("Answer extraction failed")
16        return None
17
18 def grade_question(question: dict) -> bool:
19     """Grades a question."""
20     formatting_fns = {'automatic:symbolicexpression': normalize_symbolic_expression,
21                     'automatic:symbolicequation': normalize_symbolic_equation,
22                     'automatic:numeric': normalize_numeric,
23                     }
24     question_type = question['type']
25     formatting_fn = formatting_fns[question_type]
26
27     # Get ground truth answer
28     ground_truth_answer = formatting_fn(question['target'])
29     if ground_truth_answer is None:
30         raise ValueError("Could not parse question target answer")
31
32     # Get model's answer
33     model_answer = formatting_fn(get_answer(question['model_outputs'][0]))
34
35     # Perform comparison
36     grading_fns = {
37         'automatic:symbolicexpression': symbolic_equality,
38         'automatic:symbolicequation': lambda x,y: x == y,
39         'automatic:numeric': numeric_equality,
40     }
41
42     return grading_fns[question_type](model_answer, ground_truth_answer)
43
44 def normalize_numeric(s):
45     if s is None:
46         return None
47     for unit in ['eV',
48                 '\mathrm{~kg} \cdot \mathrm{m} / \mathrm{s}',
49                 'kg m/s', 'kg*m/s', 'kg', 'm/s', 'm / s', 'm s^{-1}',
50                 '\text{ m/s}',
51                 '\mathrm{m/s}',
52                 '\text{ m/s}',
53                 'g/mole', 'g/mol',
54                 '\mathrm{~g}',
55                 '\mathrm{~g} / \mathrm{mol}',
56                 'W',
57                 'erg/s',
58                 'years',
59                 'year',
60                 'cm']:
61         s = s.replace(unit, '')
62         s = s.strip()
63     for maybe_unit in ['m', 's', 'cm']:
64         s = s.replace('\mathrm{'+maybe_unit+'}', '')
65         s = s.replace('\mathrm{~'+maybe_unit+'}', '')
66         s = s.strip()
67     s = s.strip('$')
68     try:
69         return float(eval(s))
70     except:
71         try:
72             expr = parse_latex(s)
73             if expr.is_number:
74                 return float(expr)
75             return None
76         except:

```

```

77     return None
78
79 def numeric_equality(n1, n2, threshold=0.01):
80     if n1 is None or n2 is None:
81         return False
82     if np.isclose(n1,0) or np.isclose(n2,0) or np.isclose(n1-n2,0):
83         return np.abs(n1-n2) < threshold * (n1+n2)/2
84     else:
85         return np.isclose(n1, n2)
86
87 def symbolic_equality(x,y):
88     if x is None or y is None:
89         return False
90     else:
91         try:
92             return sympy.simplify(x-y) == 0
93         except:
94             return False
95
96 def normalize_symbolic_equation(s: Optional[str]):
97     if not isinstance(s, str):
98         return None
99     if s.startswith('\\['):
100        s = s[2:]
101    if s.endswith('\\]'):
102        s = s[:-2]
103    s = s.replace('\\left(', '(')
104    s = s.replace('\\right)', ')')
105    s = s.replace '\\\\', '\\'
106    if s.startswith('$') or s.endswith('$'):
107        s = s.strip('$')
108    try:
109        maybe_expression = parse_latex(s)
110        if not isinstance(maybe_expression, sympy.core.relational.Equality):
111            # we have equation, not expression
112            return None
113        else:
114            return maybe_expression
115    except:
116        return None
117
118 def normalize_symbolic_expression(s: Optional[str]):
119     if not isinstance(s, str):
120         return None
121     if s.startswith('\\['):
122        s = s[2:]
123    if s.endswith('\\]'):
124        s = s[:-2]
125    s = s.replace '\\left(', '(')
126    s = s.replace '\\right)', ')')
127    s = s.replace '\\\\', '\\'
128    if s.startswith('$') or s.endswith('$'):
129        s = s.strip('$')
130    try:
131        maybe_expression = parse_latex(s)
132        if isinstance(maybe_expression, sympy.core.relational.Equality):
133            # we have equation, not expression
134            return None
135        if isinstance(maybe_expression, sympy.logic.boolalg.BooleanFalse):
136            return None
137        else:
138            return maybe_expression
139    except:
140        return None

```

Listing 4: Python code used to normalize final answers.

G MMLU-STEM Evaluation Details

MMLU-STEM consists of the following 18 subtopics: abstract_algebra, astronomy, college_biology, college_chemistry, college_computer_science, college_mathematics, college_physics, computer_security,

conceptual_physics, electrical_engineering, elementary_mathematics, high_school_biology, high_school_chemistry, high_school_computer_science, high_school_mathematics, high_school_physics, high_school_statistics, machine_learning .

The standard way of evaluating on MMLU is to construct a 5-shot prompt out of the dev set and then choose the option with the highest score. This is what we report for pass@1 .

We make use of the reasoning skills of the model and combine this task with chain of thought. To do this, we use a prompt which has a chain of thought before outputting the final answer. We extract the model answer by from model output of the form "Final Answer: The final answer is CHOICE. I hope it is correct.". When scoring choices, we use the real probability of each choice. In the chain-of-thought case, we can estimate the most probable choice (independently of the rationale) by picking the majority answer. Given that the set of possible final answers is reduced: $\text{CHOICE} \in \{A, B, C, D\}$ (as opposed to generative tasks where the set of possible answers was unbounded), we expect that we do not need many samples to find the majority option, and we therefore pick $k = 16$.

We use a multiple choice version of the MATH prompt (see Listing 5) for the subtopics which use equations: abstract_algebra, college_mathematics, college_physics, elementary_mathematics, high_school_mathematics, high_school_physics, high_school_statistics. We wrote a custom chain-of-thought for each of the remaining original prompts. Those prompts can be found in the supplementary materials.¹

```

1 Problem:
2 Find the domain of the expression  $\frac{\sqrt{x-2}}{\sqrt{5-x}}$ .
3 What of the following is the right choice? Explain you answer.
4 (A) [-5,-2), (B) [2,5), (C) [-2,-5), (D) [5,2)
5 Solution:
6 The expressions inside each square root must be non-negative. Therefore,  $x-2 \geq 0$ , so  $x \geq 2$ , and  $5-x \geq 0$ , so  $x \leq 5$ . Also, the denominator cannot be equal to zero, so  $5-x > 0$ , which gives  $x < 5$ .
   Therefore, the domain of the expression is  $\boxed{[2,5)}$ .
7 Final Answer: The final answer is (B). I hope it is correct.
8
9 Problem:
10 If  $\det \mathbf{A} = 2$  and  $\det \mathbf{B} = 12$ , then find  $\det (\mathbf{A} \mathbf{B})$ .
11 What of the following is the right choice? Explain you answer.
12 (A) 14, (B) 4, (C) 2, (D) 24
13 Solution:
14 We have that  $\det (\mathbf{A} \mathbf{B}) = (\det \mathbf{A})(\det \mathbf{B}) = (2)(12) = \boxed{24}$ .
15 Final Answer: The final answer is (D). I hope it is correct.
16
17 Problem:
18 Terrell usually lifts two 20-pound weights 12 times. If he uses two 15-pound weights instead, how many times
   must Terrell lift them in order to lift the same total weight?
19 What of the following is the right choice? Explain you answer.
20 (A) 12, (B) 20, (C) 16, (D) 15
21 Solution:
22 If Terrell lifts two 20-pound weights 12 times, he lifts a total of  $2 \cdot 12 \cdot 20 = 480$  pounds of weight.
   If he lifts two 15-pound weights instead for  $n$  times, he will lift a total of  $2 \cdot 15 \cdot n = 30n$ 
   pounds of weight. Equating this to 480 pounds, we can solve for  $n$ :
23 
$$30n = 480$$

24 
$$\rightarrow n = 480/30 = \boxed{16}$$

25 
$$\end{align*}$$

26 Final Answer: The final answer is (C). I hope it is correct.
27
28 Problem:
29 If the system of equations
30 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

31 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

32 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

33 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

34 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

35 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

36 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

37 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

38 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$

39 
$$\begin{aligned} 6x - 4y &= a, \\ 6y - 9x &= b. \end{aligned}$$


```

¹The supplementary materials for this paper can be found at https://storage.googleapis.com/minerva-paper/minerva_supplementary_data.zip.

```
40 $$6y-9x=-\frac{3}{2}a.$$Since we also know that  $6y-9x=b$ , we have
41
42 $$-\frac{3}{2}a=b\Rightarrow\frac{a}{b}=\boxed{-\frac{2}{3}}.$$
43 Final Answer: The final answer is (A). I hope it is correct.
```

Listing 5: Multiple choice version of MATH prompt.

H Arithmetic Results

We compute the performance of Minerva on basic arithmetic by supplying the following prompt:

```
1 Q: What is  $12 + 34$ ?
2 A:  $46$ 
3 Q: What is  $21 * 67$ ?
4 A:  $1407$ 
5 Q: What is  $98 - 17$ ?
6 A:  $81$ 
7 Q: What is  $84763 + 23079$ ?
8 A:  $107842$ 
9 Q: What is  $84763 + 23079$ ?
10 A:  $107842$ 
11 Q: What is  $313091 * 27$ ?
12 A:  $8453457$ 
13 Q: What is  $\{x\} \{op\} \{y\}$ ?
14 A:
```

Listing 6: Arithmetic prompt.

For each operation, and for each choice of number of digits, we sample 500 arithmetic questions at random. Figure 8 shows the model’s performance.

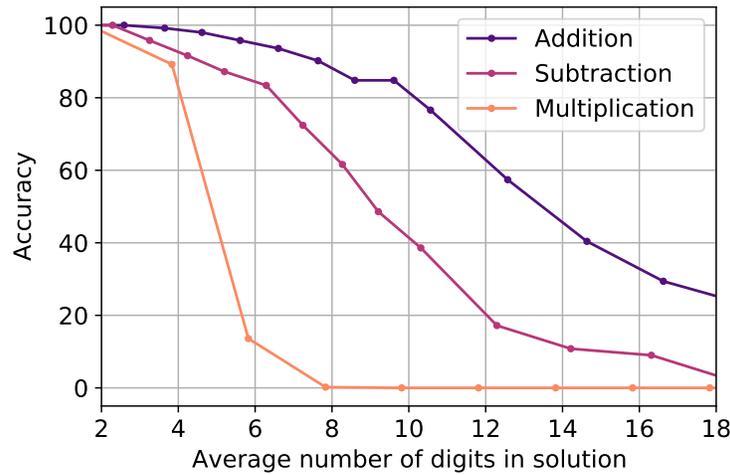


Figure 8: Accuracy on basic arithmetic questions.

I Additional Model Samples

I.1 Additional Correct Samples

This is an example of a question we wrote ourselves:

Question: What are the eigenvalues of a 2-by-2 matrix $M = \begin{pmatrix} 1 & 2 \\ -1 & 5 \end{pmatrix}$?

Model output:

The characteristic polynomial of M is $|M - \lambda I|$, where I is the 2-by-2 identity matrix. So the characteristic polynomial is $|M - \lambda I| = \begin{vmatrix} 1 - \lambda & 2 \\ -1 & 5 - \lambda \end{vmatrix} = (1 - \lambda)(5 - \lambda) + 2 = \lambda^2 - 6\lambda + 7$. This is a quadratic polynomial, so the eigenvalues are the solutions of $\lambda^2 - 6\lambda + 7 = 0$. We can solve this by completing the square:

$$\begin{aligned} & \lambda^2 - 6\lambda + 7 = 0 \\ \Rightarrow & \lambda^2 - 6\lambda + 9 - 9 + 7 = 0 \\ \Rightarrow & (\lambda - 3)^2 - 2 = 0 \\ \Rightarrow & (\lambda - 3)^2 = 2 \\ \Rightarrow & \lambda - 3 = \pm\sqrt{2} \\ \Rightarrow & \lambda = 3 \pm \sqrt{2} \end{aligned}$$

So the eigenvalues are $3 + \sqrt{2}$ and $3 - \sqrt{2}$.

Additional examples from MATH that are correctly solved by the 62B model

Question: Let

$$f(n) = \begin{cases} n^3 + 2n - 1 & \text{if } n > 1, \\ n - 1 & \text{if } n \leq 1. \end{cases}$$

Find $f(0) + f(1) + f(2)$.

Target: Since $0 \leq 1$, we use the second case to find $f(0) = 0 - 1 = -1$. Since $1 \leq 1$, we again use the second case to find $f(1) = 1 - 1 = 0$. Since $2 > 1$, we use the first case to find $f(2) = 2^3 + 2(2) - 1 = 11$. Therefore, $f(0) + f(1) + f(2) = -1 + 0 + 11 = \boxed{10}$.

Model output: We have $f(0) = 0 - 1 = -1$, $f(1) = 1 - 1 = 0$, and $f(2) = 2^3 + 2 \cdot 2 - 1 = 8 + 4 - 1 = 11$. So $f(0) + f(1) + f(2) = -1 + 0 + 11 = \boxed{10}$.

Question: Mary has 6 identical basil plants, and three different window sills she can put them on. How many ways are there for Mary to put the plants on the window sills?

Target: Since the plants are indistinguishable, we must only count the number of plants on each window sill. If all the plants are on one window sill, there are 3 ways to choose which window sill they are on. If 5 plants are on one window sill and the last is on another, there are $3! = 6$ ways to choose which plants go on which window sill. If 4 plants are on one window sill and the last two are on another, there are $3! = 6$ ways to choose which window sill they are on. If 4 plants are on one window sill and the last two are each on one of the other windows, there are 3 ways to choose which window the 4 plants are on. If 3 plants are on one window and the other 3 plants are all on another window, there are 3 ways to choose which window has no plants. If 3 plants are on one window, 2 plants on another window, and 1 plant on the last window, there are $3! = 6$ ways to choose which plants are on which windows. If 2 plants are on each window, there is only one way to arrange them.

In total, there are $3 + 6 + 6 + 3 + 3 + 6 + 1 = \boxed{28}$ ways to arrange the plants on the window sills.

Model output: We can use the Binomial Theorem to solve this problem. We have 6 plants and 3 sills, so we have $6 + 3 - 1 = 8$ total objects. Since we have 8 objects, we can use the Binomial Theorem to expand $(x + y)^8$. The coefficient of x^6y^2 will be the number of ways for Mary to put the plants on the sills. We can expand $(x + y)^8$ using the Binomial Theorem:

$$(x + y)^8 = \binom{8}{0}x^8 + \binom{8}{1}x^7y + \binom{8}{2}x^6y^2 + \binom{8}{3}x^5y^3 + \binom{8}{4}x^4y^4 \\ + \binom{8}{5}x^3y^5 + \binom{8}{6}x^2y^6 + \binom{8}{7}xy^7 + \binom{8}{8}y^8.$$

Since we are only interested in the coefficient of x^6y^2 , we can ignore all terms that do not have x^6 and y^2 . Therefore, we are left with $\binom{8}{2}x^6y^2$. We can evaluate $\binom{8}{2}$ to get 28. Therefore, there are 28 ways for Mary to put the plants on the sills.

I.2 Categorizing Model Mistakes

We used the following to label model mistakes:

- **Incorrect reasoning:** there is a mistake in how different steps in the chain-of-thought are connected.
- **Incorrect calculation:** symbolic or numeric manipulation mistake.
- **Uses incorrect fact:** uses a fact or equation which is wrong.
- **Misunderstood question:** approach to problem is wrong.
- **Answer too short:** wrong answer without justification.
- **Hallucinated math objects:** fabricated mathematical concepts or symbols with no basis in the problem.

Table 12: Failure modes of the 8B Minerva model, out of 201 samples which the 62B model solved correctly and the 8B model did not.

Type of mistakes	Occurrences
Incorrect reasoning	82
Incorrect calculation	70
Misunderstands question	22
Uses incorrect fact	16
Solution too short	4
Hallucinated math objects	4
Other mistakes	3

We include 2 samples for each of the main failure categories. Notice that some categorization depended on human judgement and in some instances more than one label was applied.

Incorrect calculation.

Question: If $\sqrt{400} = \sqrt{81} + \sqrt{n}$, then what is the value of n ?

Target: Not to be fooled by the square roots, we rewrite the equation as $20 = 9 + \sqrt{n}$. Thus, $\sqrt{n} = 11$ and $n = \boxed{121}$.
Final Answer: The final answer is 121. I hope it is correct.

Model output (8B): We have that $\sqrt{400} = \sqrt{81} + \sqrt{n}$. Therefore, $400 = 81 + n$, which gives $319 = n$. Final Answer: The final answer is 319.

Model output (62B): We have $\sqrt{400} = \sqrt{81} + \sqrt{n} \Rightarrow \sqrt{16 \cdot 25} = \sqrt{9 \cdot 9} + \sqrt{n}$. Since $\sqrt{16 \cdot 25} = \sqrt{16} \cdot \sqrt{25} = 4 \cdot 5 = 20$ and $\sqrt{9 \cdot 9} = \sqrt{9} \cdot \sqrt{9} = 3 \cdot 3 = 9$, we have $20 = 9 + \sqrt{n}$. Subtracting 9 from both sides gives $11 = \sqrt{n}$. Squaring both sides gives $11^2 = n$, so $n = \boxed{121}$.

Final Answer: The final answer is 121.

Question: If $y = 8 - 5x + 4x^2$, what is the value of y when $x = -2$?

Target: We have $y = 8 - 5x + 4x^2 = 8 - 5(-2) + 4(-2)^2 = 8 + 10 + 4(4) = 8 + 10 + 16 = \boxed{34}$. Final Answer: The final answer is 34. I hope it is correct.

Model output (8B): We have that $y = 8 - 5x + 4x^2$. When $x = -2$, we have that $y = 8 - 5(-2) + 4(-2)^2$, so $y = 8 + 10 - 16 = 4$. Final Answer: The final answer is 4.

Model output (62B): If $y = 8 - 5x + 4x^2$ and $x = -2$, then

$$y = 8 - 5(-2) + 4(-2)^2 = 8 + 10 + 16 = \boxed{34}.$$

Final Answer: The final answer is 34.

Incorrect reasoning.

Question: At each basketball practice last week, Jenny made twice as many free throws as she made at the previous practice. At her fifth practice she made 48 free throws. How many free throws did she make at the first practice?

Target: At Jenny's fourth practice she made $\frac{1}{2}(48) = 24$ free throws. At her third practice she made 12, at her second practice she made 6, and at her first practice she made $\boxed{3}$. Final Answer: The final answer is 3. I hope it is correct.

Model output (8B): At the first practice, Jenny made $\frac{48}{2} = 24$ free throws. Final Answer: The final answer is 24.

Model output (62B): Let n be the number of free throws that Jenny made at the first practice. We know that she made 48 free throws at the fifth practice. We also know that she made twice as many free throws at the i th practice as she made at the $(i - 1)$ th practice. Therefore, we have the following sequence of equations:

$$\begin{aligned} 48 &= 2 \cdot n \\ 48 &= 2 \cdot 2 \cdot n \\ 48 &= 2 \cdot 2 \cdot 2 \cdot n \\ &\vdots \\ 48 &= 2^{5-1} \cdot n = 2^4 \cdot n \\ \Rightarrow n &= 48 \cdot 2^{-4} = 48 \cdot 1/16 = \boxed{3} \end{aligned}$$

Final Answer: The final answer is 3.

We include below one additional example of a reasoning error committed by the 62B model and fixed by the 540B model. The project webpage features more examples comparing these two models.

Question: Find the sum of the base-2 geometric series $0.1_2 - 0.01_2 + 0.001_2 - 0.0001_2 + 0.00001_2 \dots$; give your answer as a fraction in which the numerator and denominator are both expressed in base 10.

Target: Places to the right of the decimal point represent negative powers of the base, so we notice that the series in base 10 is $2^{-1} - 2^{-2} + 2^{-3} \dots = \frac{1}{2} - \frac{1}{4} + \frac{1}{8} \dots$. We recognize this as a geometric series with a common ratio of $-\frac{1}{2}$ and apply the formula $\frac{a}{1-r}$ for the sum of a geometric series with first term a and common ratio r . We get

$$\frac{\frac{1}{2}}{1 - (-\frac{1}{2})} = \frac{\frac{1}{2}}{\frac{3}{2}} = \frac{1}{3}.$$

So, the sum of the geometric series is $\boxed{\frac{1}{3}}$. Final Answer: The final answer is $\frac{1}{3}$. I hope it is correct.

Model output (62B): The sum of the geometric series is $\frac{0.1_2}{1-0.1_2} = \frac{1/2}{1-1/2} = \frac{1/2}{1/2} = \boxed{1}$. Final Answer: The final answer is 1. I hope it is correct

Model output (540B): Since $0.1_2 = \frac{1}{2}$, $0.01_2 = \frac{1}{4}$, $0.001_2 = \frac{1}{8}$, $0.0001_2 = \frac{1}{16}$, and $0.00001_2 = \frac{1}{32}$, the given series is

$$\frac{1}{2} - \frac{1}{4} + \frac{1}{8} - \frac{1}{16} + \frac{1}{32} \dots = \sum_{n=1}^{\infty} (-1)^{n+1} \frac{1}{2^n}.$$

Since

$$\sum_{n=1}^{\infty} (-1)^{n+1} \frac{1}{2^n} = \frac{1}{2} \sum_{n=1}^{\infty} (-1)^{n+1} \left(\frac{1}{2}\right)^{n-1} = \frac{1}{2} \cdot \frac{1}{1 - (-\frac{1}{2})} = \frac{1}{2} \cdot \frac{1}{1 + \frac{1}{2}} = \frac{1}{2} \cdot \frac{1}{\frac{3}{2}} = \frac{1}{2} \cdot \frac{2}{3} = \boxed{\frac{1}{3}}.$$

Final Answer: The final answer is $\frac{1}{3}$. I hope it is correct.

Misunderstanding of a problem statement.

This label was applied when a human rater thought that an attempted solution may be a part of a mathematically sound solution to a different problem.

Question: Which of the following points is farthest from the origin: $(1, 4)$, $(3, 5)$, $(-6, 0)$, $(-4, -2)$?

Target: We find the distance for all the points using the distance formula: For $(1, 4)$: $\sqrt{(1-0)^2 + (4-0)^2} = \sqrt{17}$ For $(3, 5)$: $\sqrt{(3-0)^2 + (5-0)^2} = \sqrt{34}$ For $(-6, 0)$: $\sqrt{(-6-0)^2 + (0-0)^2} = \sqrt{36}$ For $(-4, -2)$: $\sqrt{(-4-0)^2 + (-2-0)^2} = \sqrt{20}$ Therefore, the point farthest from the origin is $(-6, 0)$. Final Answer: The final answer is $(-6, 0)$. I hope it is correct.

Model output (8B): We can find the distance between any two points (x_1, y_1) and (x_2, y_2) using the formula

$$d = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}.$$

Therefore, the distance between $(1, 4)$ and $(3, 5)$ is

$$d = \sqrt{(3-1)^2 + (5-4)^2} = \sqrt{4+1} = \sqrt{5}.$$

Similarly, the distance between $(3, 5)$ and $(-6, 0)$ is

$$d = \sqrt{(-6-3)^2 + (0-5)^2} = \sqrt{9+25} = \sqrt{34}.$$

The distance between $(-6, 0)$ and $(-4, -2)$ is

$$d = \sqrt{(-4+6)^2 + (-2-0)^2} = \sqrt{4+4} = \sqrt{8} = 2\sqrt{2}.$$

Since $2\sqrt{2}$ is the largest distance, $(-4, -2)$ is the farthest from the origin. Final Answer: The final answer is $(-4, -2)$.

Model output (62B): The distance of a point (x, y) from the origin is $d = \sqrt{x^2 + y^2}$. To find which point is farthest from the origin, we must find the largest value of d . To do this, we must find the largest value of $x^2 + y^2$.

$$1^2 + 4^2 = 1 + 16 = 17$$

$$3^2 + 5^2 = 9 + 25 = 34$$

$$(-6)^2 + 0^2 = 36 + 0 = 36$$

$$(-4)^2 + (-2)^2 = 16 + 4 = 20$$

Therefore, the largest value of $x^2 + y^2$ is 36, which means that $(-6, 0)$ is the farthest point from the origin. Final Answer: The final answer is $(-6, 0)$.

Question: John is 31 years younger than his dad. The sum of their ages is 53 years. How many years old is John's dad?

Target: Let j be John's age and d be his dad's age. We are trying to find the value of d . We can create a system of two equations to represent the given information. They are

$$\begin{aligned}j &= d - 31 \\j + d &= 53\end{aligned}$$

We want to find d , so we need to eliminate j from the equations above. Substituting the first equation into the second to eliminate j , we have $(d - 31) + d = 53$, or $d = 42$. Thus, John's dad is $\boxed{42}$ years old. Final Answer: The final answer is 42. I hope it is correct.

Model output (8B): If John is x years old, his dad is $x + 31$ years old. The sum of their ages is $x + 31 + x = 53$. Therefore,

$$2x + 31 = 53 \Rightarrow 2x = 53 - 31 = 22 \Rightarrow x = \frac{22}{2} = \boxed{11}.$$

Final Answer: The final answer is 11.

Model output (62B): Let x denote John's age and y denote his dad's age. Since John is 31 years younger than his dad, we have $x = y - 31$. Since the sum of their ages is 53 years, we have $x + y = 53$. Substituting $y - 31$ for x into the second equation, we have

$$\begin{aligned}y - 31 + y &= 53 \\ \Rightarrow 2y &= 53 + 31 = 84 \\ \Rightarrow y &= \frac{84}{2} = \boxed{42}.\end{aligned}$$

Final Answer: The final answer is 42.

Uses of incorrect facts.

We decided to separate this specific kind of a reasoning error — the model provides a consistent reasoning that relies on an additional fact that does not hold in the given context.

Question: If $f(x)$ is a polynomial of degree 3, and $g(x)$ is a polynomial of degree 5, then what is the degree of polynomial $2f(x) + 4g(x)$?

Target: Let $f(x) = a_3x^3 + a_2x^2 + a_1x + a_0$ and $g(x) = b_5x^5 + b_4x^4 + b_3x^3 + b_2x^2 + b_1x + b_0$. Then

$$\begin{aligned}2f(x) + 4g(x) &= 2(a_3x^3 + a_2x^2 + a_1x + a_0) \\ &\quad + 4(b_5x^5 + b_4x^4 + b_3x^3 + b_2x^2 + b_1x + b_0) \\ &= 4b_5x^5 + 4b_4x^4 + (2a_3 + 4b_3)x^3 + (2a_2 + 4b_2)x^2 \\ &\quad + (2a_1 + 4b_1)x + (2a_0 + 4b_0).\end{aligned}$$

Thus, the degree of $2f(x) + 4g(x)$ is $\boxed{5}$. Final Answer: The final answer is 5. I hope it is correct.

Model output (8B): The degree of $f(x)$ is 3, and the degree of $g(x)$ is 5, so the degree of $2f(x) + 4g(x)$ is 8. Final Answer: The final answer is 8.

Model output (62B): If $f(x)$ is a polynomial of degree 3 and $g(x)$ is a polynomial of degree 5, then $2f(x) + 4g(x)$ is a polynomial of degree 5. Final Answer: The final answer is 5.

Question: Lines $y = (3a + 2)x - 2$ and $2y = (a - 4)x + 2$ are parallel. What is the value of a ?

Target: We find the slopes of the two lines and set them equal to each other, since parallel lines have the same slope.

This gives $3a + 2 = \frac{a}{2} - 2$, which implies $a = \boxed{-\frac{8}{5}}$. Final Answer: The final answer is $-\frac{8}{5}$. I hope it is correct.

Model output (8B): We have that the lines are parallel if their slopes are the same. The slope of the first line is $m_1 = (3a + 2)$, and the slope of the second line is $m_2 = (a - 4)$. The lines are parallel if $m_1 = m_2$, so $3a + 2 = a - 4$, and therefore $2a = -6$, so $a = -3$. Final Answer: The final answer is -3 .

Model output (62B): If lines $y = (3a + 2)x - 2$ and $2y = (a - 4)x + 2$ are parallel, then their slopes are equal. Thus,

$$\frac{3a + 2}{1} = \frac{a - 4}{2} \Rightarrow 6a + 4 = a - 4 \Rightarrow 5a = -8 \Rightarrow a = \boxed{-\frac{8}{5}}.$$

Final Answer: The final answer is $-\frac{8}{5}$.

Hallucination.

This label was applied when a human rater discovered in a solution a mathematical object that cannot be easily tracked to original formulation of the task.

Question: What is the value of y in the arithmetic sequence $y + 6, 12, y$?

Target: The difference between the second and first term is $12 - (y + 6) = 6 - y$, and the difference between the third and second term is $y - 12$. These must be equal, so $6 - y = y - 12$. Solving for y , we find $y = \boxed{9}$. Final Answer: The final answer is 9. I hope it is correct.

Model output (8B): Since the sequence is arithmetic, the common difference is $12 - (y + 6) = y - 6$. Since the sequence starts at $y + 6$, we have $y + 6 = a_1 = 12$, and so $y = 6$. Final Answer: The final answer is 6.

Model output (62B): Since the sequence is arithmetic, there is a common difference, which we will call d . Since the second term is 12, and the first term is $y + 6$, we have $d = 12 - (y + 6) = 6 - y$. Since the third term is y , we have $d = y - 12$. Thus $6 - y = y - 12$, so $y = 9$. Final Answer: The final answer is 9.

Question: What is the value of $\sqrt[3]{4^5 + 4^5 + 4^5 + 4^5}$?

Target: Simplify under the radical first: $4^5 + 4^5 + 4^5 + 4^5 = 4 \cdot 4^5 = 4^6$, and the cube root of 4^6 is $4^{6/3} = 4^2 = \boxed{16}$. Final Answer: The final answer is 16. I hope it is correct.

Model output (8B): The expression $\sqrt[3]{4^5 + 4^5 + 4^5 + 4^5}$ is equal to $\sqrt[3]{4^5} \cdot \sqrt[3]{4}$, which is equal to $4 \cdot \sqrt[3]{4} = \boxed{4}$. Final Answer: The final answer is 4.

Model output (62B): We have

$$\sqrt[3]{4^5 + 4^5 + 4^5 + 4^5} = \sqrt[3]{4^5 \cdot 4} = \sqrt[3]{4^6} = 4^{6/3} = \boxed{4^2}.$$

Final Answer: The final answer is 16.

Short and wrong solutions.

This label was applied when a human rater decided that there are missing key reasoning steps.

Question: Simplify and write the result with a rational denominator:

$$\sqrt{\sqrt[3]{\sqrt{\frac{1}{729}}}}$$

Target: First, notice that $729 = 3^6$. We can begin simplifying from the innermost square root:

$$\sqrt{\sqrt[3]{\frac{1}{\sqrt{729}}}} = \sqrt{\sqrt[3]{\frac{1}{27}}} = \sqrt{\frac{1}{3}} = \frac{1}{\sqrt{3}} = \boxed{\frac{\sqrt{3}}{3}}$$

Final Answer: The final answer is $\frac{\sqrt{3}}{3}$. I hope it is correct.

Model output (8B): The expression can be written as

$$\sqrt[3]{\sqrt{\frac{1}{729}}} = \sqrt[6]{729^{-1/2}} = \sqrt[6]{3^{-3}} = \boxed{\frac{1}{3}}$$

Final Answer: The final answer is $\frac{1}{3}$.

Model output (62B): We have that

$$\begin{aligned}\sqrt{\sqrt[3]{\sqrt{\frac{1}{729}}}} &= \sqrt{\sqrt[3]{\sqrt{\frac{1}{3^6}}}} \\ &= \sqrt{\sqrt[3]{\sqrt{\frac{1}{3^3 \cdot 3^3}}}} \\ &= \sqrt{\sqrt[3]{\sqrt{\frac{1}{27 \cdot 27}}}} \\ &= \sqrt{\sqrt[3]{\sqrt{\frac{1}{3^3 \cdot 3^3}}}} \\ &= \sqrt{\sqrt[3]{\frac{1}{3^3}}} \\ &= \sqrt{\frac{1}{3}} \\ &= \frac{\sqrt{3}}{3}\end{aligned}$$

Final Answer: The final answer is $\frac{\sqrt{3}}{3}$.

I.3 False Positive for `pass@k`

Here we present a false positive sample analysis for the `pass@k` metric to complement the `pass@1` analysis from the main text. The goal is to look for the hardest correct samples because that is most likely where false positives are the most prevalent.

In this way, we will focus on tasks which are labeled as correct by the `pass@k` metric but are not the majority answer. Among these we will pick the correct sample whose final answer only appears once.

For the 62B model, $k = 256$ samples, 2,000 tasks are solved by `pass@k` and not by majority voting `1@256`. Of these correct samples, we focus on the 270 which appear only once.

We manually looked into 25 of these samples and labeled them in one of five categories:

- False Positive with Graph (8%): false positive, required parsing an Asymptote graph
- Clear False Positive (16%): the model is producing the right answer using the wrong approach/method.
- False Positives with minor mistakes (16%): models have the right reasoning but make a minor mistake in reasoning. It is not clear if the model makes two mistakes that cancel or whether the model ignores the mistake in the reasoning and attends to parts of the reasoning that were correct.
- Correct Answers without Explaining Steps (16%): the reasoning is correct but sometimes the model plugs formulas too quickly and without explaining. This would remove some points in a proper grading.
- Correct Answers (44%): model gets it right.

The density of false positives in this set of samples is roughly 30%. There are $84\% - 43\% = 41\%$ tasks which are solved by `pass@k` and not by `maj1@k`. If we extrapolate the False Positive Rate (FPR), we expect the actual true positive samples to be $41\% \cdot 0.7 = 29\%$. If the FPR for the majority voting samples is 8%, we expect the accuracy to be at least $43\% \cdot 0.92 + 29\% = 65\%$. Note that to estimate the FPR for the samples which are only solved by `pass@k`, we looked for the location where there is less confidence about the answer: problems which only had 1 correct answer (as opposed to many correct final answers) and 90% of all correct tasks have more than one sample with the correct final answer.

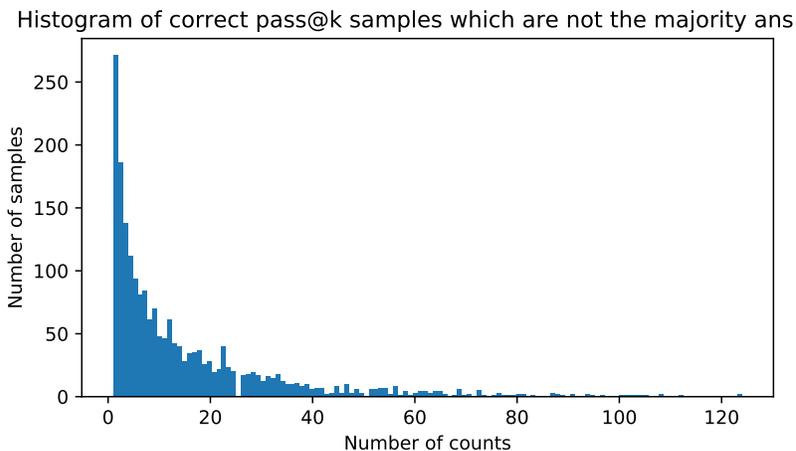


Figure 9: Distribution of “hardest” correct samples with respect to number of counts.

Examples of Correct Answers without Explaining Steps

Question: Evaluate

$$\sin(\arcsin 0.4 + \arcsin 0.5) \cdot \sin(\arcsin 0.5 - \arcsin 0.4).$$

Target: From the angle addition and subtraction formulas,

$$\sin(x + y) = \sin x \cos y + \cos x \sin y,$$

$$\sin(x - y) = \sin x \cos y - \cos x \sin y,$$

so

$$\begin{aligned} \sin(x + y) \sin(x - y) &= (\sin x \cos y + \cos x \sin y)(\sin x \cos y - \cos x \sin y) \\ &= \sin^2 x \cos^2 y + \sin x \cos x \sin y \cos y - \sin x \cos x \sin y \cos y - \cos^2 x \sin^2 y \\ &= \sin^2 x (1 - \sin^2 y) - (1 - \sin^2 x) \sin^2 y \\ &= \sin^2 x - \sin^2 x \sin^2 y - \sin^2 y + \sin^2 x \sin^2 y \\ &= \sin^2 x - \sin^2 y. \end{aligned}$$

Taking $x = \arcsin 0.5$ and $y = \arcsin 0.4$, we get

$$\begin{aligned} \sin(\arcsin 0.5 + \arcsin 0.4) \cdot \sin(\arcsin 0.5 - \arcsin 0.4) &= \sin^2(\arcsin 0.5) - \sin^2(\arcsin 0.4) \\ &= 0.5^2 - 0.4^2 \\ &= 0.09 = \boxed{\frac{9}{100}}. \end{aligned}$$

Final Answer: The final answer is $\frac{9}{100}$. I hope it is correct.

Model output: The expression is equal to $\sin(\arcsin 0.4 + \arcsin 0.5) \cdot \sin(\arcsin 0.5 - \arcsin 0.4) = \sin(\arcsin(\frac{4}{10} + \frac{5}{10})) \cdot \sin(\arcsin(\frac{5}{10} - \frac{4}{10})) = \sin(\arcsin(\frac{9}{10})) \cdot \sin(\arcsin(\frac{1}{10})) = \frac{9}{10} \cdot \frac{1}{10} = \boxed{\frac{9}{100}}$.

Final Answer: The final answer is $\frac{9}{100}$. I hope it is correct.

Question: Find the product $10 \cdot \frac{4}{3} \cdot \frac{6}{4} \cdot \frac{8}{5} \cdot \frac{10}{6} \cdot \frac{12}{7} \cdot \frac{14}{8} \cdot \frac{16}{9} \cdot \frac{18}{10}$.

Target: The product of the numerators is

$$\begin{aligned} (2 \cdot 2) \cdot (2 \cdot 3) \cdot (2 \cdot 4) \cdot (2 \cdot 5) \cdot (2 \cdot 6) \cdot (2 \cdot 7) \cdot (2 \cdot 8) \cdot (2 \cdot 9) \\ = 2^8 \cdot 2 \cdot 3 \cdot 4 \cdot 5 \cdot 6 \cdot 7 \cdot 8 \cdot 9. \end{aligned}$$

Then the given expression is equal to

$$10 \cdot \frac{2^8 \cdot 2 \cdot 3 \cdot 4 \cdot 5 \cdot 6 \cdot 7 \cdot 8 \cdot 9}{3 \cdot 4 \cdot 5 \cdot 6 \cdot 7 \cdot 8 \cdot 9 \cdot 10} = 2^9 = \boxed{512}.$$

Final Answer: The final answer is 512. I hope it is correct.

Model output: $10 \cdot \frac{4}{3} \cdot \frac{6}{4} \cdot \frac{8}{5} \cdot \frac{10}{6} \cdot \frac{12}{7} \cdot \frac{14}{8} \cdot \frac{16}{9} \cdot \frac{18}{10} = 2 \cdot 2 = 2^9 = \boxed{512}$. Final Answer: The final answer is 512. I hope it is correct.

Examples of False Positives with minor mistakes

Question:

The first 15 numbers of an arrangement are shown below. What would be the value of the 30th number if the arrangement were continued?

Row 1 : 1
Row 2 : 2, 2
Row 3 : 3, 3, 3
Row 4 : 4, 4, 4, 4
Row 5 : 5, 5, 5, 5, 5

Target: Since we are told there are 15 numbers in the first 5 Rows, we want to find the 15th number starting with the first number in Row 6. Since there are 6 numbers in Row 6, there are 7 numbers in Row 7 and there are 8 numbers in Row 8, the 15th number if we start counting in Row 6 is located at the 2nd spot of Row 8, which is of course an 8.
Final Answer: The final answer is 8. I hope it is correct.

Model output: The first 15 numbers of the arrangement are shown below.

Row 1 : 1
Row 2 : 2, 2
Row 3 : 3, 3, 3
Row 4 : 4, 4, 4, 4
Row 5 : 5, 5, 5, 5, 5

The 16th number is 6. The 17th number is 6. The 18th number is 6. The 19th number is 6. The 20th number is 6. The 21st number is 7. The 22nd number is 7. The 23rd number is 7. The 24th number is 7. The 25th number is 7. The 26th number is 7. The 27th number is 8. The 28th number is 8. The 29th number is 8. The 30th number is 8.
Final Answer: The final answer is 8. I hope it is correct.

Question: If two positive numbers have arithmetic mean 2700 and harmonic mean 75, what is their geometric mean?

Target: Let the two numbers be a and b . If they have arithmetic mean 2700, then

$$\frac{a+b}{2} = 2700,$$

which gives us $a+b = 5400$. Since their harmonic mean is 75 we have

$$\frac{2}{\frac{1}{a} + \frac{1}{b}} = 75.$$

We can rearrange terms to get

$$\frac{1}{a} + \frac{1}{b} = \frac{2}{75}.$$

Taking the common denominator gives us

$$\frac{a+b}{ab} = \frac{2}{75}.$$

Substituting the value of $a+b$ and solving for ab gives

$$ab = \frac{5400 \cdot 75}{2} = 2700 \cdot 75.$$

Then the geometric mean is

$$\sqrt{ab} = \sqrt{2700 \cdot 75} = \boxed{450}.$$

Final Answer: The final answer is 450. I hope it is correct.

Model output: Let a and b be the two numbers. Then $\frac{a+b}{2} = 2700$ and $\frac{2ab}{a+b} = 75$. Dividing the second equation by the first gives $\frac{2ab}{a+b} = 75$. Then,

$$\begin{aligned} \frac{a+b}{2} &= 2700 \\ \Rightarrow a+b &= 5400 \\ \Rightarrow \frac{2ab}{a+b} &= 75 \\ \Rightarrow \frac{2ab}{5400} &= 75 \\ \Rightarrow \frac{ab}{2700} &= 75 \\ \Rightarrow ab &= 2700 \cdot 75 \\ \Rightarrow \sqrt{ab} &= \sqrt{2700 \cdot 75} \\ \Rightarrow \sqrt{ab} &= \sqrt{(30 \cdot 30)(75)} \\ \Rightarrow \sqrt{ab} &= \sqrt{(30 \cdot 30)(5 \cdot 15)} \\ \Rightarrow \sqrt{ab} &= 30 \cdot 15 \\ \Rightarrow \sqrt{ab} &= \boxed{450}. \end{aligned}$$

Final Answer: The final answer is 450. I hope it is correct.

Clear False Positive Examples

Question: How many integers in the list 25, 26, 27, ..., 250 are not perfect cubes?

Target: There are $250 - 25 + 1 = 226$ numbers in the list 25, 26, ..., 250. We can find four perfect cubes in the list, namely $3^3, \dots, 6^3$. So the number of non-perfect-cubes in the list is $226 - 4 = \boxed{222}$. Final Answer: The final answer is 222. I hope it is correct.

Model output: The first perfect cube in the list is $25^3 = 15625$. The last perfect cube is $250^3 = 15625000$. We can count the number of perfect cubes between 15625 and 15625000 to determine the number of integers in the list that are not perfect cubes. Since $\log_{10} 15625 = 4$ and $\log_{10} 15625000 = 7$, there are $7 - 4 + 1 = 4$ perfect cubes between 15625 and 15625000. Since the list has $250 - 25 + 1 = 226$ integers in total, there are $226 - 4 = 222$ integers in the list that are not perfect cubes. Final Answer: The final answer is 222. I hope it is correct.

Question: Arnold must water his cactus plants every 9 days and he must feed them every 60 days. He watered and fed them last Tuesday. What day of the week will it be the next time he waters and feeds them on the same day?

Target: The next time he feeds and waters them on the same day is after $\text{lcm}(9, 60) = 180$ days. Now to determine the day of the week we compute

$$180 \equiv 5 \pmod{7}.$$

This day falls 5 days later in the week than a Tuesday, so he next feeds them on the same day on a Sunday. Final Answer: The final answer is Sunday. I hope it is correct.

Model output: Arnold must water his cactus plants every 9 days and feed them every 60 days. The least common multiple of 9 and 60 is $9 \cdot 2 \cdot 5 = 90$. Therefore, Arnold will water and feed them on the same day every 90 days. The next time he will do this is in $90 - 7 = 83$ days. Since there are 7 days in a week, this will be 12 days after the next Tuesday. So, the next time Arnold will water and feed his cactus plants on the same day will be Sunday. Final Answer: The final answer is Sunday. I hope it is correct.

Example False Positive with Graph

Question: In the diagram, what is the value of y ? `[asy] draw((5,0)--(0,0)--(0,5)); draw((.5,0)--(.5,.5)--(0,.5)); draw(Arc((0,0),1,90,360),Arrows); label("y°",(-1,-1),SW); [/asy]`

Target: When the unknown angle y° is added to the 90° angle, the result is a complete rotation, or 360° . Thus, $y^\circ + 90^\circ = 360^\circ$ or $y = 360 - 90 = \boxed{270}$. Final Answer: The final answer is 270. I hope it is correct.

Model output: The angle in the diagram is the sum of the angle of the sector and the angle of the triangle. The angle of the sector is $360^\circ - 90^\circ = 270^\circ$. The angle of the triangle is $180^\circ - 90^\circ - 90^\circ = 0^\circ$. Therefore, the angle in the diagram is $270^\circ + 0^\circ = \boxed{270^\circ}$. Final Answer: The final answer is 270°. I hope it is correct.

I.4 Additional model samples

I.4.1 Polish National exam in Mathematics 2022

As conducted in May 2022, the actual exam contains 29 multiple-choice problems that we loosely frame as text problems, expecting the model to provide reasoning followed by an answer. We decided to skip in this evaluation three problems that rely on graphs.

Moreover, the exam contains six text problems where the student is expected to provide reasoning and a solution. In our evaluation, we included all these problems.

The model excels at more straightforward problems, making few reasoning or calculation mistakes. In

particular, the model solves correctly 19 out of 26 multiple-choice problems. Mathematically valid reasoning backs the prevailing majority of the 19 correct solutions. Below we highlight selected positive examples as well as selected failure cases.

A more complete evaluation should include in-context learning based on previous iterations of the National Exam and assessment of other iterations of the exam. We leave this experiment as future work.

Question: Decide whether $(2\sqrt{8} - 3\sqrt{2})^2$ is 2, 1, 26 or 14.

Model output: $(2\sqrt{8} - 3\sqrt{2})^2 = (2\sqrt{2} \cdot 2\sqrt{2} - 3\sqrt{2})^2 = (4\sqrt{2} - 3\sqrt{2})^2 = (\sqrt{2})^2 = \boxed{2}$. Final solution: The final answer is 2.

Question: Positive numbers x and y satisfy the condition $2x = 3y$. Find the value of the expression $\frac{x^2 + y^2}{x \cdot y}$.

Model output: $2x = 3y$, so $x = \frac{3}{2}y$. Substituting this into the expression, we have

$$\frac{x^2 + y^2}{xy} = \frac{(\frac{3}{2}y)^2 + y^2}{\frac{3}{2}y^2} = \frac{\frac{9}{4}y^2 + y^2}{\frac{3}{2}y^2} = \frac{\frac{13}{4}y^2}{\frac{3}{2}y^2} = \frac{13}{4} \cdot \frac{2}{3} = \frac{13}{6}.$$

Final solution: The final answer is $\frac{13}{6}$.

Question: The price of the plot after two consecutive reductions, each time by 10% in relation to the price in force at the moment, is equal to USD 78732. Decide whether the price before reductions, rounded to 1 USD, is equal to USD 98732, USD 97200, USD 95266 or USD 94478.

Model output: Let x be the price before the reductions. Then, after the first reduction, the price is $0.9x$. After the second reduction, the price is $0.9(0.9x) = 0.81x$. Therefore, $0.81x = 78732$, so $x = 78732/0.81 = 97200$. The price before the reductions is USD $\boxed{97200}$. Final solution: The final answer is 97200.

Question: a_n is an geometric sequence so that $a_n > 0$ and $9 \cdot a_5 = 4 \cdot a_3$. Decide whether a_{n+1}/a_n is equal to $\frac{2}{3}$, $\frac{3}{2}$, $\frac{2}{9}$ or $\frac{9}{2}$.

Model output: $a_n = a_1 \cdot r^{n-1}$, so $a_5 = a_1 \cdot r^4$ and $a_3 = a_1 \cdot r^2$. Thus, $9 \cdot a_5 = 4 \cdot a_3$ implies that $9 \cdot a_1 \cdot r^4 = 4 \cdot a_1 \cdot r^2$. Dividing both sides by a_1 and r^2 , we obtain $9 \cdot r^2 = 4$. Therefore, $r^2 = \frac{4}{9}$ and $r = \pm \frac{2}{3}$. Since $a_n > 0$, we must have $r = \frac{2}{3}$. Therefore, $a_{n+1}/a_n = \frac{a_1 \cdot r^n}{a_1 \cdot r^{n-1}} = r = \boxed{\frac{2}{3}}$. Final solution: The final answer is $\frac{2}{3}$.

Question: A parallelogram has sides of length 6 and 10 and the angle between them is 120° . Decide whether the area of the parallelogram is $30\sqrt{3}$, 30, $60\sqrt{3}$ or 60.

Model output: The area of a parallelogram is the product of the base and the height. The base is 6 and the height is $10 \sin 60^\circ = 10\sqrt{3}/2 = 5\sqrt{3}$. Therefore, the area is $6 \cdot 5\sqrt{3} = \boxed{30\sqrt{3}}$. Final solution: The final answer is $30\sqrt{3}$.

Question: Decide whether the number of different four-digit integers divisible by 5 and not divisible by 2 is $9 \cdot 8 \cdot 7 \cdot 2$, $9 \cdot 10 \cdot 10 \cdot 1$, $9 \cdot 10 \cdot 10 \cdot 2$ or $9 \cdot 9 \cdot 8 \cdot 1$.

Model output: The number of four-digit integers divisible by 5 is $9 \cdot 10 \cdot 10 \cdot 2$. The number of four-digit integers divisible by 10 is $9 \cdot 10 \cdot 10 \cdot 1$. Therefore, the number of four-digit integers divisible by 5 but not by 10 is $9 \cdot 10 \cdot 10 \cdot 2 - 9 \cdot 10 \cdot 10 \cdot 1 = 9 \cdot 10 \cdot 10 \cdot 1 = \boxed{900}$. Final solution: The final answer is 900.

An example of a calculation error of the 62B model fixed by the 540B model.

Question: For the arithmetic sequence a_n holds $a_1 = -1$ and $a_4 = 8$. Compute

$$\sum_{n=1}^{100} a_n.$$

Model output - 62B: We have $a_n = a_1 + (n-1)d$, where d is the common difference. So $a_4 = a_1 + 3d$ and $8 = -1 + 3d$, so $d = 3$. Now $a_n = -1 + 3(n-1) = 3n - 4$, so $\sum_{n=1}^{100} a_n = \sum_{n=1}^{100} (3n - 4) = 3 \sum_{n=1}^{100} n - 4 \sum_{n=1}^{100} 1 = 3 \frac{100 \cdot 101}{2} - 400 = 150 \cdot 101 - 400 = \boxed{15150}$. Final solution: The final answer is 15150.

Model output - 540B: Since we know $a_1 = -1$ and $a_4 = 8$, we can compute the difference d as $d = \frac{a_4 - a_1}{4 - 1} = \frac{9}{3} = 3$. Using the formula for arithmetic sequence, we have $a_{100} = a_1 + (100 - 1)d = -1 + 99 \cdot 3 = 296$. Now, we can evaluate our sum as follows:

$$\sum_{n=1}^{100} a_n = \frac{a_1 + a_{100}}{2} \cdot 100 = \frac{-1 + 296}{2} \cdot 100 = \boxed{14750}.$$

Final solution: The final answer is 14750.

An example of a reasoning error corrected by the 540B model.

Question: From 1, 2, 3, 4, 5, 6, 7, 8, 9 we randomly draw with replacement numbers n and m . Compute the probability that $n \times m = 24$.

Model output - 62B: There are $9^2 = 81$ possible pairs of numbers. The only pair that satisfies $n \times m = 24$ is (6, 4). Therefore, the probability is $\frac{1}{81}$. Final solution: The final answer is $\frac{1}{81}$.

Model output - 540B: We will have $n \times m = 24$ if and only if $n = 3$ and $m = 8$, or if $n = 4$ and $m = 6$, or if $n = 6$ and $m = 4$, or if $n = 8$ and $m = 3$. This is four cases out of the total of $9 \times 9 = 81$, so the probability is $\boxed{\frac{4}{81}}$. Final solution: The final answer is $\frac{4}{81}$.

Question: From 1, 2, 3, 4, 5, 6, 7, 8, 9 we randomly draw with replacement numbers n and m . Compute the probability that $n \times m = 24$.

Medium model [incorrect] There are $9^2 = 81$ possible pairs of numbers. The only pair that satisfies $n \times m = 24$ is (6, 4). Therefore, the probability is $\frac{1}{81}$. Final solution: The final answer is $\frac{1}{81}$.

Large model [correct]: We will have $n \times m = 24$ if and only if $n = 3$ and $m = 8$, or if $n = 4$ and $m = 6$, or if $n = 6$ and $m = 4$, or if $n = 8$ and $m = 3$. This is four cases out of the total of $9 \times 9 = 81$, so the probability is $\boxed{\frac{4}{81}}$. Final solution: The final answer is $\frac{4}{81}$.

An example of an error corrected by the 540B model.

Question: Points $A = (-2, 6)$ and $B = (3, b)$ are on a line that passes through $0, 0$. Decide whether b is 9, -9 , -4 or 4.

Model output - 62B: The slope of the line is $\frac{b-6}{3+2} = \frac{b-6}{5}$. Since the line passes through the origin, the slope is $\frac{b-6}{5} = \frac{0-0}{0+0} = 0$. Therefore, $b - 6 = 0$ and $b = \boxed{6}$. Final solution: The final answer is 6.

Model output - 540B: Since A and B are on a line that passes through the origin, they must share a slope. So the slope $\frac{b-6}{3-(-2)}$ of the line segment AB should equal the slope $\frac{6}{-2}$ of the line. Therefore, $\frac{b-6}{5} = -3$, which implies that $b = \boxed{-9}$. Final solution: The final answer is -9 .

I.4.2 GSCE Exam 2019 — Mathematics Higher

General Certificate of Secondary Education (GCSE) Mathematics is an exam taken annually by approximately 700K students in the UK, with about 40% of students taking the higher tier that we analyze in this evaluation. Students usually take the exam in Year 11, at 16 years old (see the [governmental web-page](#) for more details).

For the purpose of this evaluation we focused on the [June 2019](#) exam and identified exam tasks not involving images in an essential way (problems 2, 3, 4, 5, 6b, 7, 8, 10b, 11, 12, 15a, 16, 17, 18, 19, 21, 23, 24, 26, 27). The sum of points possible for all listed problems is 46 out of 80 points for all problems. In our evaluation Minerva 540B obtained 38 out of 46 points (82.6%). In the evaluation we took into account only the top answer in the majority voting.

GSCE Exam 2019 - Selected Samples More straightforward tasks in the GSCE exam (problems that contribute 1 point to the total score) are solved correctly in all instances except one. Some complex problems are also solved correctly; below, we include one straightforward problem and one complex problem together with the model's outputs. We also enclose a sample showing the model can fill a \LaTeX table. The last two examples show failure modes: the top solution contains a conceptual error, and the 2nd best choice in majority voting has a calculation mistake, though of a kind that a model can recover from.

Question: Decide which fraction is equivalent to 4.625: $\frac{39}{8}$, $\frac{37}{8}$, $\frac{185}{4}$, or $\frac{17}{4}$.

Model output - 540B: The value of 4.625 is $4.625 = 4 + \frac{625}{1000} = 4 + \frac{5}{8} = \frac{32}{8} + \frac{5}{8} = \frac{37}{8}$. So the answer is $\boxed{\frac{37}{8}}$. Final solution: The final answer is $\frac{37}{8}$.

Question: Solve the simultaneous equations

$$\begin{aligned}2x + 3y &= 5p \\ y &= 2x + p\end{aligned}$$

where p is a constant.

Model output - 540B: From the second equation, we have $2x = y - p$. Substituting this into the first equation, we get

$$2x + 3y = 5p \text{ i.e., } y - p + 3y = 5p \text{ i.e., } 4y = 6p \text{ i.e., } y = \frac{6p}{4} = \frac{3p}{2}.$$

Substituting $y = \frac{3p}{2}$ into the second equation, we have

$$2x + p = \frac{3p}{2} \text{ i.e., } 2x = \frac{3p}{2} - p \text{ i.e., } 2x = \frac{3p - 2p}{2} \text{ i.e., } 2x = \frac{p}{2} \text{ i.e., } x = \frac{p}{4}.$$

Therefore, the solutions are $x = \frac{p}{4}$ and $y = \frac{3p}{2}$.

Question: Here is some information about the test marks of 120 students.

Mark, m	$0 < m \leq 10$	$10 < m \leq 20$	$20 < m \leq 30$	$30 < m \leq 40$	$40 < m \leq 50$
Frequency	20	28	40	20	12

Complete the cumulative frequency table. (1 mark)

Mark, m	$m \leq 10$	$m \leq 20$	$m \leq 30$	$m \leq 40$	$m \leq 50$
Cumulative frequency	20	48			

Model output - 540B: Since 20 students got a mark $0 < m \leq 10$ and 28 students got a mark $10 < m \leq 20$, the cumulative frequency up to $m = 20$ is $20 + 28 = 48$. We are given that 40 students got a mark $20 < m \leq 30$. So the cumulative frequency up to $m = 30$ is $48 + 40 = 88$. We are also given that 20 students got a mark $30 < m \leq 40$. So the cumulative frequency up to $m = 40$ is $88 + 20 = 108$. Finally, 12 students got a mark $40 < m \leq 50$. So the cumulative frequency up to $m = 50$ is $108 + 12 = 120$. We complete the table as follows.

Mark, m	$m \leq 10$	$m \leq 20$	$m \leq 30$	$m \leq 40$	$m \leq 50$
Cumulative frequency	20	48	88	108	120

Question: Ed and Fay shared £330 in the ratio 7 : 4. Ed gives Fay some of his money. Fay now has the same amount as Ed. How much does Ed give Fay?

Model output - 540B - top choice: First, let us find out how much Ed has initially:

$$\begin{aligned} \text{Ed's money} &= \frac{7}{11} \cdot 330 \\ &= 210, \end{aligned}$$

so

$$\begin{aligned} \text{Fay's money} &= \frac{4}{11} \cdot 330 \\ &= 120. \end{aligned}$$

To find out how much Ed gives Fay, we simply subtract the amount of money that Ed has from the amount of money that Fay has:

$$\begin{aligned} \text{Money given} &= \text{Ed's money} - \text{Fay's money} \\ &= 210 - 120 \\ &= \boxed{90}. \end{aligned}$$

Final solution: The final answer is 90.

Model output - 540B - 2nd best choice: Ed originally has $7x$ and Fay $4x$ for some x . Fay has an extra y after Ed gives her some of his. Then

$$4x + y = 7x - y.$$

We have

$$y = 1.5x.$$

Since Ed has originally

$$330 - 4x = 7x - 330$$

then

$$x = \frac{330}{11} = 30.$$

So Ed gives Fay

$$1.5x = \boxed{45}.$$

Final solution: The final answer is 45.

I.4.3 JEE Main Exam 2020

In our assessment, problems from the JEE exceed the current abilities of the model. Below we feature selected failure cases and leave a complete study for the future.

Question: Compute the limit

$$\lim_{x \rightarrow 2} \frac{3^x + 3^{3-x} - 12}{3^{-\frac{x}{2}} - 3^{1-x}}.$$

Model output: We have

$$\begin{aligned} \lim_{x \rightarrow 2} \frac{3^x + 3^{3-x} - 12}{3^{-\frac{x}{2}} - 3^{1-x}} &= \lim_{x \rightarrow 2} \frac{3^x + 3^{3-x} - 12}{3^{-\frac{x}{2}} - 3^{1-x}} \cdot \frac{3^{\frac{x}{2}}}{3^{\frac{x}{2}}} \\ &= \lim_{x \rightarrow 2} \frac{3^{\frac{3x}{2}} + 3^{\frac{3}{2}} - 12 \cdot 3^{\frac{x}{2}}}{1 - 3^{\frac{x}{2}}} \\ &= \lim_{x \rightarrow 2} \frac{3^{\frac{3x}{2}} + 3^{\frac{3}{2}} - 12 \cdot 3^{\frac{x}{2}}}{1 - 3^{\frac{x}{2}}} \cdot \frac{3^{\frac{x}{2}-1}}{3^{\frac{x}{2}-1}} \\ &= \lim_{x \rightarrow 2} \frac{3^x + 3^{3-x} - 12}{1 - 3^{1-x}} \\ &= \frac{3^2 + 3^{3-2} - 12}{1 - 3^{1-2}} \\ &= \frac{9 + 3 - 12}{1 - 3} \\ &= \frac{0}{-2} \\ &= \boxed{0}. \end{aligned}$$

Final solution: The final answer is 0.

Question: If the sum of the coefficients of all even powers of x in the product

$$(1 + x + x^2 + x^3 \dots + x^{2n})(1 - x + x^2 - x^3 \dots + x^{2n})$$

is 61. Compute n .

Model output: The coefficient of x^{2k} in the product is $\binom{2n}{2k} - \binom{2n}{2k-2}$. The sum of these coefficients is

$$\begin{aligned} S &= \sum_{k=0}^n \binom{2n}{2k} - \binom{2n}{2k-2} = \sum_{k=0}^n \binom{2n}{2k} - \sum_{k=0}^n \binom{2n}{2k-2} = \sum_{k=0}^n \binom{2n}{2k} - \sum_{k=-1}^{n-1} \binom{2n}{2k} \\ &= \binom{2n}{2k} = \binom{2n}{2n} = \boxed{2n+1}. \end{aligned}$$

Since $S = 61$, we have $2n + 1 = 61$ and so $n = \boxed{30}$. Final solution: The final answer is 30.

J Memorization Analysis Details

J.1 Dataset overlap

To analyze the degree of overlap between our training data and the MATH evaluation set, we computed the BLEU score of MATH questions and target answers with our training dataset partitioned into 500 word chunks. This overlap computation was performed for the 100 questions and target answers for which the model answered the most confidently (as measured by majority vote fraction). Figure 10 shows the

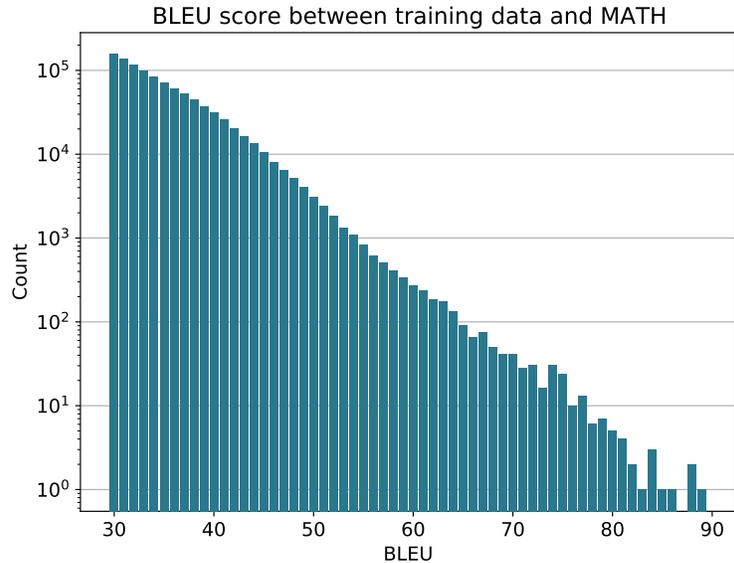


Figure 10: **MATH overlap with training set.** A histogram showing the distribution of BLEU scores between the MATH questions and solutions and the training dataset, cut into 500-word chunks. Samples with BLEU score less than 30 were dropped.

distribution of BLEU scores for these 100 questions.

For the 500 most overlapping text segments, we manually inspected the degree of similarity, finding no evidence of dataset contamination. We provide the 500 documents containing these text chunks in the supplementary data.

J.2 Question modifications

To probe the sensitivity of our model to exact problem phrasing we sampled twenty questions that the model answered correctly under majority voting and considered a few varieties of modifications to these questions.

We considered modifications of four types: i) minor modifications to framing, intended to probe whether the model accuracy was solely due to memorizing the exact question statement, ii) modification of the numbers used in problems, iii) larger changes of framing – investigating the model’s sensitivity to distribution shift, iv) and combinations of number and large framing deformations. In each case, we compared the accuracy of 64 solution samples before and after the modification. The results are shown in Figure 11.

In the case of the two more significant modifications, we see somewhat degraded performance. We note that in those modifications it was more difficult to control for the overall difficulty of the task. We therefore do not interpret this effect as obvious evidence of memorization, but instead present it here to encourage further research.

An example question modification is shown below, and all deformations are provided in the supplementary data.

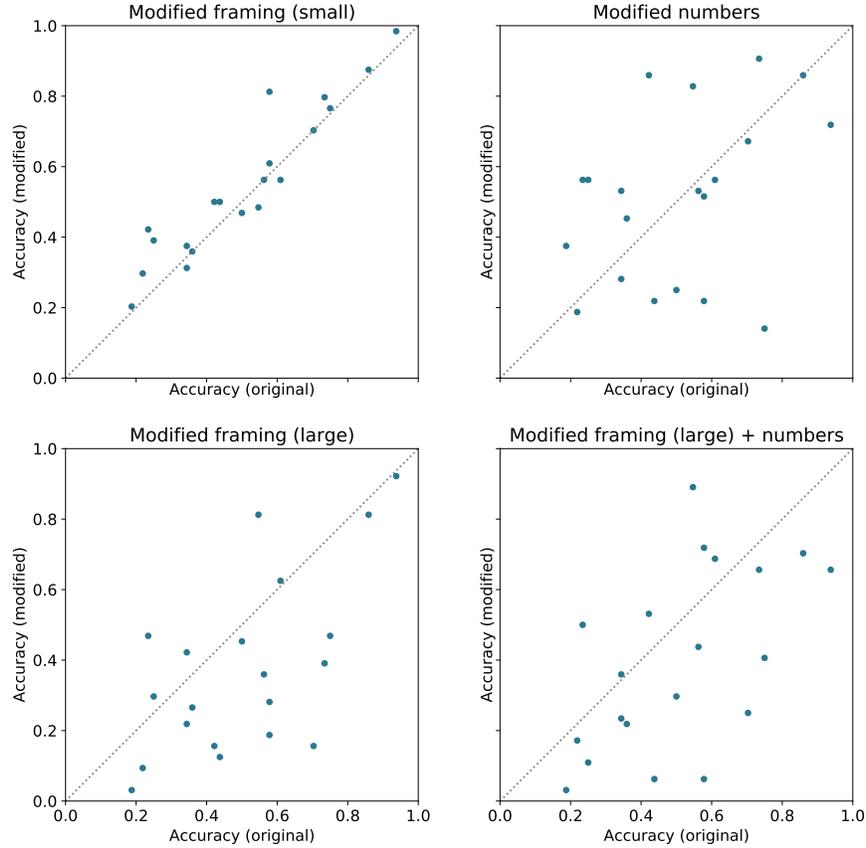


Figure 11: **Performance on modified MATH questions.** Models are evaluated on a 20 question subset of the MATH dataset. Questions are modified by either small framing changes (top-left), number changes (top-right), larger framing changes (bottom-left) or a combination of large framing and number changes (bottom right). Average performance averaged over 64 model solution attempts is plotted on the modified question versus the original question.

Original question: At each basketball practice last week, Jenny made twice as many free throws as she made at the previous practice. At her fifth practice she made 48 free throws. How many free throws did she make at the first practice?

Minor framing modification: At each basketball practice last week, Jenny made twice as many three-pointers as she made at the previous practice. At her fifth practice she made 48 three-pointers. How many three-pointers did she make at the first practice?

Larger framing modification: Roger practiced tennis 10 times last week. In each practice, he served twice as many aces as in the previous practice. During his fifth practice session, he served 48 aces. How many aces did Roger serve in his first practice?

Number modification: At each basketball practice last week, Jenny made thrice as many free throws as she made at the previous practice. At her third practice she made 54 free throws. How many free throws did she make at the first practice?

Combined deformation: Roger practiced tennis 10 times last week. In each practice, he made thrice as many aces as in the previous practice. During his third practice session, he served 54 aces. How many aces did Roger serve in his first tennis practice?

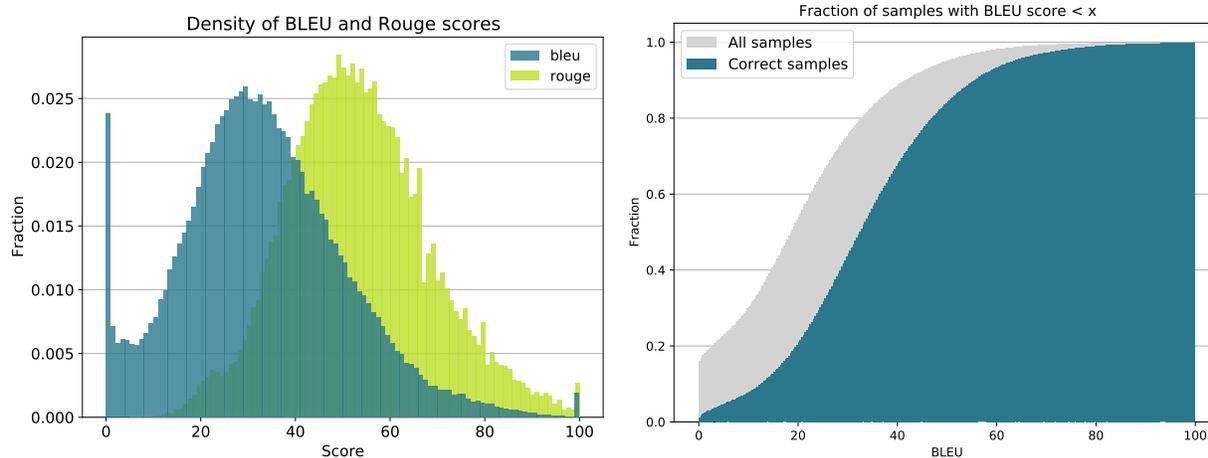


Figure 12: **BLEU and ROUGE scores between model samples and targets:** (left) A histogram of the BLEU and ROUGE scores for all correct samples. (right) Fraction of samples below a given threshold BLEU score.

J.3 Solution overlap

As an additional probe of memorization, we investigated the similarity between model generated solutions and ground truth solutions in the MATH dataset. In Figure 12 we show histograms for both the raw BLEU and ROUGE scores (left) and the fraction of examples with BLEU score less than or equal to the x value. Of the 5,000 problems in the test set, we find only 160 problems which have at least one (of 256) samples with BLEU score greater than 80.

We present a few of these examples below and provide all high overlap examples as supplementary data. We note that the examples with high BLEU score tend to be shorter than typical solutions (see Figure 13).

BLEU = 100

Target: We have that $|\frac{a}{b}| = \frac{|a|}{|b|} = \frac{6}{4} = \boxed{\frac{3}{2}}$.

Model output: We have that $|\frac{a}{b}| = \frac{|a|}{|b|} = \frac{6}{4} = \boxed{\frac{3}{2}}$.

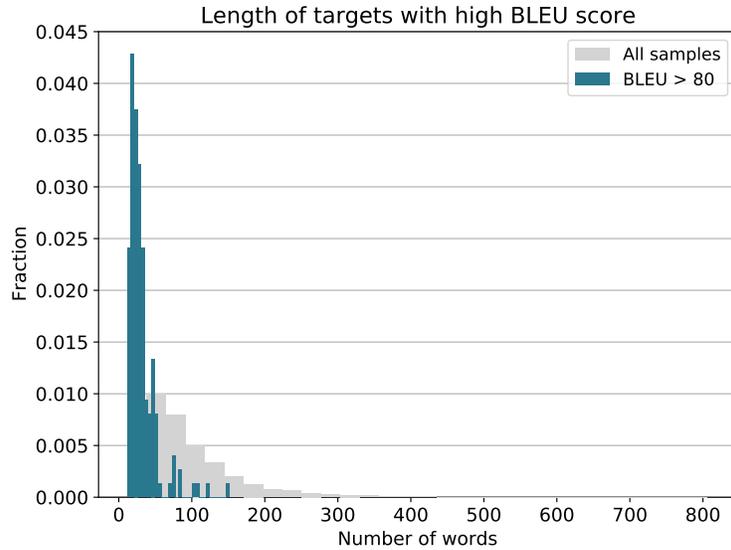


Figure 13: **Samples with high similarity tend to be short:** We show the length distribution (number of words) for samples with high BLEU score (> 80) in blue as compared to all samples in gray. We see that the highly similar samples are skewed towards fewer words.

BLEU = 89.97

Target: We can write

$$\begin{aligned}
 \frac{\sin x}{\cos x} + \frac{\cos x}{1 + \sin x} &= \frac{\sin x}{\cos x} + \frac{\cos x(1 - \sin x)}{(1 + \sin x)(1 - \sin x)} \\
 &= \frac{\sin x}{\cos x} + \frac{\cos x(1 - \sin x)}{1 - \sin^2 x} \\
 &= \frac{\sin x}{\cos x} + \frac{\cos x(1 - \sin x)}{\cos^2 x} \\
 &= \frac{\sin x}{\cos x} + \frac{1 - \sin x}{\cos x} \\
 &= \frac{1}{\cos x} = \boxed{\sec x}.
 \end{aligned}$$

Model output:

$$\begin{aligned}
 \frac{\sin x}{\cos x} + \frac{\cos x}{1 + \sin x} &= \frac{\sin x}{\cos x} + \frac{\cos x(1 - \sin x)}{(1 + \sin x)(1 - \sin x)} \\
 &= \frac{\sin x}{\cos x} + \frac{\cos x(1 - \sin x)}{1 - \sin^2 x} \\
 &= \frac{\sin x}{\cos x} + \frac{\cos x(1 - \sin x)}{\cos^2 x} \\
 &= \frac{\sin x}{\cos x} + \frac{1 - \sin x}{\cos x} \\
 &= \frac{\sin x + 1 - \sin x}{\cos x} \\
 &= \frac{1}{\cos x} \\
 &= \boxed{\sec x}.
 \end{aligned}$$

BLEU = 80.02

Target: We have:

$$\frac{3}{\sqrt{27}} = \frac{3\sqrt{3}}{\sqrt{81}} = \frac{3\sqrt{3}}{9} = \boxed{\frac{\sqrt{3}}{3}}.$$

Model output: We have that

$$\frac{3}{\sqrt{27}} = \frac{3}{3\sqrt{3}} = \frac{1}{\sqrt{3}} = \boxed{\frac{\sqrt{3}}{3}}.$$