

LLEMMA: AN OPEN LANGUAGE MODEL FOR MATHEMATICS

Zhangir Azerbayev^{1,2} Hailey Schoelkopf² Keiran Paster^{3,4}
 Marco Dos Santos⁵ Stephen McAleer⁶ Albert Q. Jiang⁵ Jia Deng¹
 Stella Biderman² Sean Welleck^{6,7}

¹ Princeton University ² EleutherAI ³ University of Toronto ⁴ Vector Institute
⁵ University of Cambridge ⁶ Carnegie Mellon University ⁷ University of Washington

ABSTRACT

We present LLEMMA, a large language model for mathematics. We continue pretraining Code Llama on Proof-Pile-2, a mixture of scientific papers, web data containing mathematics, and mathematical code, yielding LLEMMA. On the MATH benchmark LLEMMA outperforms all known open base models, as well as the unreleased Minerva model suite on an equi-parameter basis. Moreover, LLEMMA is capable of tool use and formal theorem proving without any further finetuning. We openly release all artifacts, including 7 billion and 34 billion parameter models, the Proof-Pile-2, and code to replicate our experiments.¹

1 INTRODUCTION

Language models trained on diverse mixtures of text display remarkably general language understanding and generation capabilities (Brown et al., 2020; Chowdhery et al., 2022), serving as base models that are adapted to a wide range of applications (Raffel et al., 2023). Applications such as open-ended dialogue (Thoppilan et al., 2022; Touvron et al., 2023) or instruction following (Ouyang et al., 2022; Wei et al., 2022) require balanced performance across the entire distribution of natural text, thus favoring *generalist models*. However, if we seek to maximize performance within one domain, such as medicine (Singhal et al., 2022; 2023), finance (Wu et al., 2023), or science (Taylor et al., 2022), a *domain-specific language model* may offer superior capabilities for a given computational cost, or lower computational cost for a given level of capability.

In this work, we train a domain-specific language model for mathematics. We have several motivations for doing so. First, solving mathematical problems requires pattern matching against a large body of specialized prior knowledge, thus serving as an ideal setting for domain adaptation. Second, mathematical reasoning is in itself a central AI task, its study dating back to at least Gelernter (1959) and Wang (1960) and continuing to today (Lu et al., 2023). Third, language models capable of strong mathematical reasoning are upstream of a number of research topics, such as reward modeling (Uesato et al., 2022; Lightman et al., 2023), reinforcement learning for reasoning (Polu et al., 2022; Lample et al., 2022), and algorithmic reasoning (Zhou et al., 2022; Zhang et al., 2023).




Figure 1: Continued pretraining on Proof-Pile-2 yields LLEMMA, a base model with improved mathematical capabilities.

¹<https://github.com/EleutherAI/math-lm>

Although domain-specific models for mathematics have been trained in the past, they have either been closed access (Lewkowycz et al., 2022), limiting their ability to become a platform for further research, or have lagged far behind the closed access state-of-the-art (Azerbayev et al., 2023).

We present a recipe for adapting a language model to mathematics through continued pretraining (Lewkowycz et al., 2022; Rozière et al., 2023) on Proof-Pile-2, a diverse mixture of math-related text and code. Applying the recipe to Code Llama (Rozière et al., 2023) yields LLEMMA: 7 billion and 34 billion parameter base language models with substantially improved mathematical capabilities.

Specifically, our contributions are as follows:

1. We train and release the LLEMMA models: 7B and 34B parameter language models specialized for mathematics. The LLEMMA models are a new state-of-the-art for publicly released base models on MATH (Lewkowycz et al., 2022).
2. We release the AlgebraicStack, a dataset of 11B tokens of code specifically related to mathematics.
3. We demonstrate that LLEMMA is capable of using computational tools to solve mathematical problems, namely, the Python interpreter and formal theorem provers.
4. Unlike prior mathematics language models such as Minerva (Lewkowycz et al., 2022), the LLEMMA models are open access and we open source our training data and code. This allows LLEMMA to serve as a platform for future research in mathematical reasoning.

Our work builds on findings in Minerva (Lewkowycz et al., 2022), but differs in several ways: (1) LLEMMA’s training and evaluation covers a wider range of data and tasks, notably code data (e.g., the AlgebraicStack), tool use, and formal mathematics; (2) our work only depends on publicly accessible tools and data; (3) we provide new analyses related to the continued training data mixture, memorization, and additional supervised finetuning; (4) we make all artifacts publicly available.

2 APPROACH

LLEMMA models are 7 billion and 34 billion parameter language models specialized for mathematics. Our approach is to continue pretraining Code Llama (Rozière et al., 2023) on the Proof-Pile-2.

Model	Adaptation tokens	Open	Dataset	Tokens	Open
Minerva-8b	164B	✗	Minerva Dataset	38.5B	✗
Minerva-62b	109B	✗	Proof-Pile-2 (ours)	55B	✓
LLEMMA-7b (ours)	200B	✓	Code (AlgebraicStack)	11B	✓
LLEMMA-34b (ours)	50B	✓	OpenWebMath (Paster et al., 2023))	15B	✓
			ArXiv (Computer, 2023))	29B	✓

Figure 2: Comparison of LLEMMA and Minerva training

2.1 DATA: Proof-Pile-2

We form the Proof-Pile-2, a 55B-token mixture of scientific papers, web data containing mathematics, and mathematical code. With the exception of the Lean proofsteps subset (see Appendix B), the Proof-Pile-2 has a knowledge cutoff of April 2023.

Code. Computational tools such as numerical simulations, computer algebra systems, and formal theorem provers are of ever increasing importance to mathematicians (Avigad, 2018). Motivated by this fact, we create AlgebraicStack, an 11B-token dataset of source code from 17 languages, spanning numerical, symbolic, and formal math. The dataset consists of filtered code from the Stack (Kocetkov et al., 2022), public GitHub repositories, and formal proofstep data. Table 9 shows the number of tokens by language in AlgebraicStack. See Appendix B.1 for further details on AlgebraicStack.

Web data. We use OpenWebMath (Paster et al., 2023), a 15B-token dataset of high-quality web pages filtered for mathematical content. OpenWebMath filters CommonCrawl web pages based

on math-related keywords and a classifier-based math score, preserves mathematical formatting (e.g., \LaTeX , AsciiMath), and includes additional quality filters (e.g., perplexity, domain, length) and near-deduplication. Refer to Paster et al. (2023) for a full description of OpenWebMath.

Scientific papers. We use the ArXiv subset of RedPajama (Computer, 2023), an open-access reproduction of the LLaMA training dataset. The ArXiv subset contains 29B tokens.

General natural language and code data. Following Lewkowycz et al. (2022), our training mixture consists of a small amount of general domain data, which functions as a form of regularization. Since the pretraining dataset for LLaMA 2 is undisclosed, we use the Pile (Gao et al., 2020; Biderman et al., 2022) as a surrogate training dataset. We set 95% of our training mixture to be the Proof-Pile-2, 2% to be from the Pile (with ArXiv removed, as it is separately in Proof-Pile-2), and 3% to be the GitHub subset of RedPajama (Computer, 2023).

Further information on dataset composition and a datasheet are in Appendix B and Appendix E, respectively. We publicly release Proof-Pile-2 at hf.co/datasets/EleutherAI/proof-pile-2.

2.2 MODEL AND TRAINING

Each model is initialized from Code Llama (Rozière et al., 2023). Code Llama models are decoder-only transformer language models initialized from Llama 2 (Touvron et al., 2023) and further trained on 500B tokens of code. We continue training the Code Llama models on Proof-Pile-2 using a standard autoregressive language modeling objective. We train the 7B model for 200B tokens, and the 34B model for 50B tokens.

We train all models in bfloat16 mixed precision using the GPT-NeoX library (Andonian et al., 2023) across 256 A100 40GB GPUs. We use Tensor Parallelism (Shoeybi et al., 2019) with a world size of 2 for LLEMMA-7B, and a world size of 8 for LLEMMA-34B, alongside ZeRO Stage 1 sharded optimizer states (Rajbhandari et al., 2020) across Data Parallel (Goyal et al., 2017) replicas. We use Flash Attention 2 (Dao, 2023) to improve throughput and further reduce memory requirements.

LLEMMA 7B is trained for 42,000 steps with a global batch size of 4 million tokens and a 4096 token context length. This corresponds to roughly 23,000 A100-hours. The learning rate is warmed up to $1 \cdot 10^{-4}$ over 500 steps, then set to cosine decay to 1/30th of the maximum learning rate over 48,000 steps. The reason for the discrepancy between the number of training steps and the scheduler length is that we planned to train for 48,000 steps, but encountered NaN losses after step 42,000, likely caused by unstable optimization or hardware failures (Elsen et al., 2023).

LLEMMA 34B is trained for 12,000 steps with a global batch size of 4 million tokens and a 4096 context length. This corresponds to roughly 47,000 A100-hours. The learning rate is warmed up to $5 \cdot 10^{-5}$ over 500 steps, then decayed to 1/30th the peak learning rate.

Before training LLEMMA 7B, we contract the RoPE (Su et al., 2022) base period of the Code Llama 7B initialization from $\theta = 1,000,000$ to $\theta = 10,000$. This is so that the long context finetuning procedure described in Peng et al. (2023) and Rozière et al. (2023) can be repeated on the trained LLEMMA 7B (we leave actually doing so to future work). Due to compute constraints, we were unable to verify that training LLEMMA 34B with a contracted RoPE base period did not come with a performance penalty, therefore for that model we preserved $\theta = 1,000,000$.

3 EVALUATION

Our goal is to evaluate LLEMMA as a base model for mathematical text. To this end, we compare LLEMMA models using few-shot evaluation (Brown et al., 2020), and primarily focus on state-of-the-art models that have not been finetuned on supervised examples for the task. First, we evaluate the model’s ability to solve mathematics problems using chain of thought reasoning (Wei et al., 2023) and majority voting (Wang et al., 2023). Our evaluations include MATH (Hendrycks et al., 2021b) and GSM8k (Cobbe et al., 2021), the de-facto standard benchmarks for evaluating quantitative reasoning in language models (Lewkowycz et al., 2022). Second, we explore few-shot tool use and formal theorem proving. Third, we study the effects of memorization and the data mixture. Appendix G contains a preliminary study of supervised finetuning with LLEMMA.

3.1 CHAIN-OF-THOUGHT MATHEMATICAL PROBLEM SOLVING

These tasks involve generating self-contained text solutions to problems expressed in \LaTeX or natural language, without using external tools (Lewkowycz et al., 2022). We use the following evaluation:

- **MATH** (Hendrycks et al., 2021b), a dataset with 12.5k problems (5k evaluation) from high-school math competitions. Given a problem statement, the model generates a \LaTeX solution and an answer that must match a reference answer. We follow a similar task implementation to Lewkowycz et al. (2022), using their four-example prompt and evaluating answers for exact string match or `SymPy` equivalence.
- **GSM8k** (Cobbe et al., 2021), a dataset of middle-school level math word problems. We use the 8-shot prompt from Wei et al. (2023), as Lewkowycz et al. (2022) do not specify their evaluation prompt or number of few-shot examples.
- **OCWCourses** (Lewkowycz et al., 2022), a collection of undergraduate-level STEM problems harvested from MIT’s OpenCourseWare. We use the four-example prompt provided by (Lewkowycz et al., 2022).
- **MMLU-STEM** (Hendrycks et al., 2021a), a subset of 18 out of 57 subjects in the MMLU benchmark. We follow Lewkowycz et al. (2022) and use their provided four-example chain-of-thought prompt.
- **SAT**, we create a dataset consisting of the 32 math questions that do not contain figures from the May 2023 College Board SAT examination, which is after our model’s knowledge cutoff.

Input:
 Let $f(r) = \sum_{j=2}^{2008} \frac{1}{j^r} = \frac{1}{2^r} + \frac{1}{3^r} + \dots + \frac{1}{2008^r}$. Find $\sum_{k=2}^{\infty} f(k)$.

LLEMMA 34B solution:
 We have

$$\begin{aligned} \sum_{k=2}^{\infty} f(k) &= \sum_{k=2}^{\infty} \sum_{j=2}^{2008} \frac{1}{j^k} = \sum_{j=2}^{2008} \sum_{k=2}^{\infty} \frac{1}{j^k} = \sum_{j=2}^{2008} \frac{1}{j^2} \sum_{k=0}^{\infty} \frac{1}{j^k} = \sum_{j=2}^{2008} \frac{1}{j^2} \frac{1}{1 - \frac{1}{j}} \\ &= \sum_{j=2}^{2008} \frac{1}{j(j-1)} = \sum_{j=2}^{2008} \left(\frac{1}{j-1} - \frac{1}{j} \right) \\ &= \left(\frac{1}{1} - \frac{1}{2} \right) + \left(\frac{1}{2} - \frac{1}{3} \right) + \dots + \left(\frac{1}{2007} - \frac{1}{2008} \right) \\ &= 1 - \frac{1}{2008} \\ &= \boxed{\frac{2007}{2008}}. \end{aligned}$$

Final Answer: The final answer is $\frac{2007}{2008}$.

Figure 3: Example of a LLEMMA 34B solution to a MATH (Hendrycks et al., 2021a) problem. This problem is tagged with difficulty level 5, the highest in MATH. The model was conditioned on the 4-shot prompt described in subsection 3.1, and the solution was produced by greedy decoding. The model had to apply two nontrivial steps to solve this problem: (1) noticing that swapping the order of summation simplifies the problem, and (2) noticing that the resulting sum telescopes.

We compare with Minerva (Lewkowycz et al., 2022), which continued pretraining the PaLM language model on a dataset of technical content; Code Llama, the initialization of LLEMMA’s continued pretraining; and Llama 2, the initialization of Code Llama’s continued pretraining on code. For open access models, we report scores computed using our evaluation suite, which is implemented as a fork of the Language Model Evaluation Harness (Gao et al., 2021). For Minerva models, we report benchmark scores from Lewkowycz et al. (2022).

Results. LLEMMA’s continued pretraining on Proof-Pile-2 improves few-shot performance on the five mathematical benchmarks. LLEMMA 34B improves over Code Llama by 20 percentage points on GSM8k and 13 points on MATH, and LLEMMA 7B outperforms the proprietary Minerva model. Our approach also outperforms all open-weight language models at the time of writing. We conclude that continued pretraining on Proof-Pile-2 is effective for improving a pretrained model’s ability to perform mathematical problem solving.

LLEMMA is pretrained on a diverse distribution of mathematics-related data, and is not tuned for a particular task. Therefore, we expect that LLEMMA can adapt to many other tasks via task-specific finetuning and few-shot prompting.

		GSM8k	OCW	MMLU-STEM	SAT	MATH
Llama 2	7B	11.8%	3.7%	29.9%	25.0%	3.2%
Code Llama	7B	10.5%	4.4%	25.1%	9.4%	4.5%
Minerva	8B	16.2%	7.7%	35.6%	-	14.1%
LLEMMA	7B	36.4%	7.7%	37.7%	53.1%	18.0%
Code Llama	34B	29.6%	7.0%	40.5%	40.6%	12.2%
LLEMMA	34B	51.5%	11.8%	49.0%	71.9%	25.0%
Minerva	62B	52.4%	12.0%	53.9%	-	27.6%
Minerva	540B	58.8%	17.6%	63.9%	-	33.6%

Table 1: Results on our five chain-of-thought reasoning tasks with samples generated via greedy decoding. Minerva results are quoted from Lewkowycz et al. (2022). Note that CodeLlama 7B performs worse than random guessing (25%) on MMLU and SAT, largely due to failing to conclude its chain of thought with a valid answer.

		GSM8k	OCW	MMLU-STEM	SAT	MATH
		maj@ <i>k</i>	maj@ <i>k</i>	maj@ <i>k</i>	maj@ <i>k</i>	maj@ <i>k</i>
Minerva	8B	28.4%	12.5%	43.4%	-	25.4%
LLEMMA	7B	54.0%	14.3%	49.9%	78.1%	33.5%
LLEMMA	34B	69.3%	18.4%	59.7%	81.3%	43.1%
Minerva	62B	68.5%	23.5%	63.5%	-	43.4%
Minerva	540B	78.5%	30.8%	75.0%	-	50.3%

Table 2: Majority voting results for LLEMMA and Minerva. Minerva results are quoted from Lewkowycz et al. (2022). Voting is done with $k = 256$ for MATH, $k = 100$ for GSM8k and OCW, and $k = 16$ for MMLU-STEM and SAT. We sample with temperature $T = 0.6$ for $k = 256$ and $k = 100$ and $T = 0.3$ for $k = 16$, and use nucleus sampling with $p = 0.95$ (Holtzman et al., 2020). Due to compute constraints, we do not calculate majority voting scores for Llama 2 and Code Llama.

3.2 MATHEMATICAL PROBLEM SOLVING WITH TOOL USE

These tasks involve solving problems with access to computational tools. We evaluate the following:

- **MATH+Python**, the model is prompted to alternately describe a solution step in natural language, then execute that step with code. The final answer is a program that executes to a numeric type or a `SymPy` object. Our few-shot prompt includes examples that use built-in numeric operations, the `math` module, and `SymPy`.
- **GSM8k+Python**, solving a GSM8k word problem by writing a Python program that executes to an integer answer. We use the prompt from Gao et al. (2023).

Results. As seen in Table 3, LLEMMA improves over Code Llama on both tasks. Its performance on MATH and GSM8k with tools is also higher than its performance on these datasets without tools.

		GSM8k+Python	MATH+Python
		pass@1	pass@1
Code Llama	7B	27.1%	17.2%
LLEMMA	7B	40.1%	21.5%
Code Llama	34B	52.7%	23.5%
LLEMMA	34B	62.6%	27.1%

Table 3: Mathematical problem solving with tool use.

3.3 FORMAL MATHEMATICS

Interactive proof assistants such as Lean (de Moura et al., 2015), Isabelle (Wenzel et al., 2008), and Coq (Paulin-Mohring, 1989a;b) express mathematics in programming languages that allow for verification. These languages are data scarce compared to mainstream languages, especially in the context of pretraining. For instance, the Stack dataset used to pretrain language models in the BigCode project (Allal et al., 2023) has over 700 gigabytes of Python, compared to 322 megabytes of Lean. Proof assistants also require models to leverage information that is not present in raw source code, such as goal states that contain information about each step of a proof.

Problem (MATH Number theory 185): When a number is divided by 5, the remainder is 3. What is the remainder when twice the number is divided by 5? Show that it is 1.

Human-written informal proof: If our number is n , then $n \equiv 3 \pmod{5}$. This tells us that

$$2n = n + n \equiv 3 + 3 \equiv 1 \pmod{5}.$$

The remainder is 1 when the number is divided by 5.

Informal-to-formal (Isabelle):
{Problem, human-written informal proof}

```

theorem mathd_numbertheory_185:
  fixes n ::nat
  assumes "n mod 5 = 3"
  shows "(2 * n) mod 5 = 1"

proof -
  have "2 * n = n + n" <ATP>
  also have "... mod 5 =
    (n mod 5 + n mod 5) mod 5" <ATP>
  also have "... = (3 + 3) mod 5"
  using assms <ATP>
  also have "... = 1" <ATP>
  finally show ?thesis <ATP>
qed
        
```

Formal-to-formal (Lean 4):

```

theorem mathd_numbertheory_185
  (n : ℕ) (h₀ : n % 5 = 3)
  : 2 * n % 5 = 1 := by

-- INPUT (step 1):
-- n: ℕ
-- h₀: n % 5 = 3
-- ⊢ 2 * n % 5 = 1
rw [mul_mod, h₀]

-- INPUT (step 2):
-- n: ℕ
-- h₀: n % 5 = 3
-- ⊢ 2 % 5 * 3 % 5 = 1
simp only [h₀, mul_one]
        
```

Figure 4: Example formal proofs from LLEMMA-7b. *Left:* The model is given a problem, informal proof, and formal statement, following Jiang et al. (2023). It generates a formal proof (starting with `proof -`) containing Isabelle code and calls to automation (shown as `<ATP>`). *Right:* The model is given a proof state, visualized as a grey comment, and generates the subsequent step (e.g. `rw [..]`).

Proof-Pile-2’s AlgebraicStack contains over 1.5 billion tokens of formal mathematics data, including proof states extracted from Lean and Isabelle formalizations. While a full investigation of formal math is outside the scope of this paper, we evaluate LLEMMA few-shot on two tasks:

- **Informal-to-formal proving** (Jiang et al., 2023), the task of generating a formal proof, given a formal statement, an informal \LaTeX statement, and an informal \LaTeX proof. The formal proof is checked by the proof assistant. We use the Isabelle proof assistant and evaluate on miniF2F (Zheng et al., 2021), a benchmark consisting of problem statements from Olympiads and undergraduate coursework. For the prompt, we use 11 (formal statement, informal statement, informal proof, formal proof) examples from Jiang et al. (2023), selecting 7 examples for number theory problems, and 6 examples for all others. We generate a single proof with greedy decoding.
- **Formal-to-formal proving** (e.g., Polu & Sutskever (2020)), the task of proving a formal statement by generating a sequence of proof steps (tactics). At each step, the input is a state x_t given by the proof assistant, and the language model’s task is to generate a proof step y_t (a sequence of code). The proof step is checked by the proof assistant, yielding a new state x_{t+1} or an error message. The process continues, stopping if a proof is completed or a timeout is reached. We prompt the model using three (x_t, y_t) examples. We evaluate on miniF2F (Zheng et al., 2021) using the Lean 4 proof assistant, and use a standard best first search. See Appendix D for more details.

Results. As seen in Table 4, LLEMMA’s continued pretraining on Proof-Pile-2 improved few-shot performance on the two formal theorem proving tasks.

Method	Informal-to-formal		Method	Formal-to-formal	
	miniF2F-valid	miniF2F-test		Search	miniF2F-test
Sledgehammer	14.72%	20.49%	ReProver (fine-tuned)	1×64	26.50%
Code Llama 7b	16.31%	17.62%	Code Llama 7b	1×32	20.49%
Code Llama 34b	18.45%	18.03%	Code Llama 34b	1×32	22.13%
LLEMMA-7b	20.60%	22.13%	COPRA (GPT-4)	- [†]	23.36%
LLEMMA-34b	21.03%	21.31%	LLEMMA-7b	1×32	26.23%
			LLEMMA-34b	1×32	25.82%

Table 4: Formal theorem proving tasks. *Left*: Informal-to-formal proving in Isabelle, showing the percentage of proven theorems with greedy decoding. *Right*: Formal-to-formal proving in Lean, showing the percentage of proven theorems with the given number of attempts \times generations-per-iteration of best first search, and a 10-minute timeout. Sledgehammer (Paulson & Nipkow, 2023) is built-in Isabelle automation. ReProver (Yang et al., 2023) is a supervised and retrieval-augmented model. COPRA (Thakur et al., 2023) is a retrieval-augmented GPT-4 based method. [†] COPRA does not use best first search, but instead samples from GPT-4 (OpenAI, 2023) a maximum of 60 times.

On informal-to-formal proving, LLEMMA-7b closes 22.1% of the theorems, improving upon its Code Llama initialization and the Sledgehammer prover. The theorems that LLEMMA proves are often complementary to those proved with Sledgehammer: taking the union of Sledgehammer and LLEMMA proofs results in 26 new validation proofs (an 11 percentage-point increase), and 17 new test proofs (a 7 point increase); see Appendix Table 11. Prior to our work, the only demonstration of few-shot proof autoformalization used the proprietary Codex model (Jiang et al., 2023).

On Lean 4 formal-to-formal proving, LLEMMA-7b improves upon its Code Llama initialization, and performs similar to ReProver (Yang et al., 2023), a retrieval-augmented language model finetuned for tactic prediction. LLEMMA adapts to the task using a 3 example prompt, which to our knowledge is the first demonstration of few-shot tactic prediction for theorem proving by an open model.

3.4 IMPACT OF DATA MIXTURE

When training a language model, it is common to upsample high-quality subsets of the training data according to mixture weights (Brown et al., 2020; Gao et al., 2020; Xie et al., 2023). We select mixture weights by doing short training runs on several hand-picked mixture weights, then choosing the one which minimizes perplexity on a set of high-quality held-out text (we use the MATH training set). Table 5 shows the MATH training set perplexity of models trained using different mixtures of arXiv to web to code. Based on these results, we trained LLEMMA with a ratio of 2 : 4 : 1. Note that our methodology uses the MATH training set to determine a training hyperparameter, though we expect that the effect is similar to that of related high-quality texts.

Mixture	MATH training set perplexity							
	Overall	Prealgebra	Algebra	Number Theory	Counting & Probability	Geometry	Intermediate Algebra	Precalculus
2:4:1	1.478	1.495	1.515	1.552	1.475	1.519	1.439	1.331
2:4:2	1.482	1.500	1.519	1.556	1.477	1.524	1.443	1.334
4:2:1	1.487	1.505	1.524	1.561	1.481	1.534	1.447	1.338
4:2:2	1.489	1.508	1.527	1.562	1.483	1.538	1.447	1.339
4:4:1	1.487	1.506	1.525	1.561	1.482	1.529	1.446	1.335
4:4:2	1.485	1.503	1.523	1.559	1.480	1.529	1.444	1.334

Table 5: MATH training set perplexity of Code Llama 7B models trained using different data mixtures for a reduced number of steps. Each mixture is represented by its arXiv:Web:Code ratio.

3.5 DATASET OVERLAP AND MEMORIZATION

Do test problems or solutions appear in the corpus? We check whether any 30-gram in a test sequence (either an input problem or an output solution) occurs in any OpenWebMath or AlgebraicStack document. If so, we say that a *hit* occurred between the sequence and the document. Table 6 shows hits between sequences from MATH and documents from Proof-Pile-2. Using our methodology, around 7% of MATH test problem statements and 0.6% of MATH test solutions have hits. Note that our methodology gives a lower bound on the number of semantically equivalent sequences (e.g., it does not account for alternative phrasing).

We manually inspected 100 uniformly sampled hits between a test problem statement and an OpenWebMath document. 41 of the cases had no solution, which included websites with a list of problems, discussions, or hints. 49 had an alternative solution to the MATH ground-truth solution, but with the same answer. These include solutions that solve the problem differently than the ground-truth, solutions with missing details, and discussions that include the answer. 9 cases had a missing or incorrect answer, and 1 had the same solution as in the ground-truth. In summary, we find that solutions can appear in a corpus derived from web documents, particularly alternative solutions to those in the evaluation set. We repeated our analysis with 20-gram hits and our findings were similar, though with false positives; see Appendix Figure 6 for examples.

Proof-Pile-2	Test	Problem		Solution			
		Example	Docs	Example	Docs		
OpenWebMath	MATH	348	717	34	46	Same solution	1
AlgebraicStack	MATH	3	3	1	1	Different solution, same answer	49
OpenWebMath	GSM8k	2	3	0	0	Different solution, different answer	9
AlgebraicStack	GSM8k	0	0	0	0	No solution	41
						Different problem	0

Table 6: *Left*: 30-gram hits between MATH test problems or solutions and Proof-Pile-2 documents. *Example* and *Docs* are the numbers of unique test examples and Proof-Pile-2 documents with a hit. *Right*: manual inspection of 100 hits between a problem statement and a Proof-Pile-2 document.

How do problems in the corpus impact performance?

Next, we evaluate LLEMMA-34b on the test examples with a 30-gram hit, and the test examples without a 30-gram hit. Table 7 shows the accuracy partitioned by MATH difficulty level. The model’s accuracy remains low on difficult problems (e.g., 6.08% on Level 5 problems with a hit, versus 6.39% on problems without a hit), and we observe no clear relationship between 30-gram hits and accuracy across difficulty levels. We conclude that a nontrivial match between a test example and a training document did not imply that the model generated a memorized correct answer. We repeated the analysis with 20-grams and with the 7b model, and our findings were analogous. Figure 7 shows an example.

MATH Level	Hit Accuracy	Nonhit Accuracy	# Hits
Level 1	72.73	61.50	11
Level 2	35.71	40.18	28
Level 3	30.36	26.88	56
Level 4	14.89	16.61	94
Level 5	6.08	6.39	181

Table 7: LLEMMA-34b’s accuracy on hits (a 30-gram overlap between a problem or solution and a training sequence) and non-hits by MATH difficulty level.

Finally, we check 30-gram hits between LLEMMA’s MATH generations and OpenWebMath. There were 13 hits, which occurred when the model generated a common sequence of numbers (e.g., a list of Fibonacci numbers), plus one instance of factoring a polynomial. Appendix Figure 6 shows an example. We find all of these observations worthy of further study. Using LLEMMA and Proof-Pile-2 to better understand data, memorization, and performance is an interesting future direction. We include the code for our analysis in the LLEMMA repository.

4 RELATED WORK

Large-scale language modeling. Recent progress in large language models involves two connected threads: the increasing scale of models and data (Hoffmann et al., 2022; Kaplan et al., 2020; Chowdhery et al., 2022), and a progression toward more generalist models (Radford et al., 2019; Brown et al., 2020) which are capable of solving diverse problems and adapting quickly to novel tasks. A third thread relates to enabling open access to language models with these capabilities (Black et al., 2022; Biderman et al., 2023; Touvron et al., 2023; Rozière et al., 2023). Our work provides a recipe for specializing these language models to the domain of mathematics, providing a platform for further research and applications.

Domain adaptation. Language model applications typically require a general-domain pretraining step, followed by a shorter fine-tuning step. The finetuning step is often aimed at imbuing instruction-following ability (Sanh et al., 2022; Wei et al., 2022) or aligning a model’s outputs with human preferences (Ziegler et al., 2019; Ouyang et al., 2022; Bai et al., 2022). Other work explores adapting pretrained models to novel domains by continued training (Rozière et al., 2023; Beltagy et al., 2019), parameter-efficient finetuning methods (Yong et al., 2023), retrieval augmentation (Min et al., 2023; Asai et al., 2023), and other techniques. We provide an adaptation recipe involving continued training and targeted data collection.

Language models for mathematics. Applying large language models to problems in mathematics is an active subfield of machine learning, including benchmarking mathematical knowledge and reasoning at varying levels (Hendrycks et al., 2021b; Zheng et al., 2021; Welleck et al., 2022; Azerbayev et al., 2023). Although achieving strong mathematical reasoning is an important target, it is difficult to assess the correctness of models’ answers and processes, especially as models become more capable (Bowman et al., 2022; Uesato et al., 2022; Lightman et al., 2023; Cobbe et al., 2021).

A number of recent works focus on supervised finetuning on task-relevant (input, output) pairs (e.g., Yu et al. (2023); Yue et al. (2023)). Doing so boosts performance on some common mathematical language modeling benchmarks, but trains the model for these specific tasks. In contrast, Lewkowycz et al. (2022) and our work seek to train a *base* language model as a platform for further development.

Language models for formal mathematics. An ongoing line of work explores integrating language models with interactive proof assistants in the context of mathematics. This includes synthesizing proofs via tactic prediction (Polu & Sutskever, 2020; Han et al., 2022; Lample et al., 2022; Jiang et al., 2022), autoformalization (Wu et al., 2022; Jiang et al., 2023), and integrated tools (Welleck & Saha, 2023). Due to high computational costs of search, language models applied to this domain have traditionally been small, but recent work has demonstrated promise in the use of larger models (First et al., 2023; Jiang et al., 2023). Our work provides a demonstration of few-shot proof autoformalization and tactic prediction, a large collection of formal mathematics data, along with an open access model for further exploring these directions.

5 CONCLUSION

We introduce LLEMMA and Proof-Pile-2, a novel base model and corpus for language modeling of mathematics. Our models, dataset, and code are openly available. We have shown that LLEMMA achieves state-of-the-art results for open-weights models on mathematical problem solving benchmarks, shown capabilities of using external tools via Python code, and demonstrated few-shot tactic prediction for theorem proving. We hope that LLEMMA and Proof-Pile-2 will be a useful base for future work on understanding language model generalization and dataset composition, investigating the limits of domain-specific language models, using language models as tools for mathematicians, and improving the mathematical capabilities of language models.

ACKNOWLEDGEMENTS

We would like to thank Dragomir Radev, Arman Cohan, Jesse Michael Han, and the Deepmind Blueshift team for valuable guidance. We thank Jonah Pillion for the model name. We thank Aviya Skowron for advising us on ethical considerations in the development and release of our models. We thank Jonathan Laurent and Leo Du for contributions to our open-source code.

We would also like to thank several parties for donating computing resources for this project: Stability AI (training the LLEMMA models), CoreWeave (evaluations and finetuning), the Province of Ontario and companies sponsoring the Vector Institute for Artificial Intelligence (www.vectorinstitute.ai/partners), and Brigham Young University (finetuning). KP is supported by an NSERC PGS-D award.

REFERENCES

- Loubna Ben Allal, Raymond Li, Denis Kocetkov, Chenghao Mou, Christopher Akiki, Carlos Munoz Ferrandis, Niklas Muennighoff, Mayank Mishra, Alex Gu, Manan Dey, Logesh Kumar Umapathi, Carolyn Jane Anderson, Yangtian Zi, Joel Lamy Poirier, Hailey Schoelkopf, Sergey Troshin, Dmitry Abulkhanov, Manuel Romero, Michael Lappert, Francesco De Toni, Bernardo García del Río, Qian Liu, Shamik Bose, Urvashi Bhattacharyya, Terry Yue Zhuo, Ian Yu, Paulo Villegas, Marco Zocca, Sourab Mangrulkar, David Lansky, Huu Nguyen, Danish Contractor, Luis Villa, Jia Li, Dzmitry Bahdanau, Yacine Jernite, Sean Hughes, Daniel Fried, Arjun Guha, Harm de Vries, and Leandro von Werra. Santacoder: don't reach for the stars! In *Deep Learning for Code (DLAC) Workshop*, 2023.
- Alex Andonian, Quentin Anthony, Stella Biderman, Sid Black, Preetham Gali, Leo Gao, Eric Hallahan, Josh Levy-Kramer, Connor Leahy, Lucas Nestler, Kip Parker, Michael Pieler, Jason Phang, Shivanshu Purohit, Hailey Schoelkopf, Dashiell Stander, Tri Songz, Curt Tigges, Benjamin Thérien, Phil Wang, and Samuel Weinbach. GPT-NeoX: Large scale autoregressive language modeling in PyTorch. GitHub Repo, 9 2023. URL <https://www.github.com/eleutherai/gpt-neox>.
- Akari Asai, Sewon Min, Zexuan Zhong, and Danqi Chen. Retrieval-based language models and applications. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 6: Tutorial Abstracts)*, pp. 41–46, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-tutorials.6. URL <https://aclanthology.org/2023.acl-tutorials.6>.
- Jeremy Avigad. The mechanization of mathematics. *Notices of the AMS*, 65(6):681–90, 2018.
- Zhangir Azerbayev, Bartosz Piotrowski, Hailey Schoelkopf, Edward W. Ayers, Dragomir R. Radev, and Jeremy Avigad. Proofnet: Autoformalizing and formally proving undergraduate-level mathematics. *ArXiv*, abs/2302.12433, 2023.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- Iz Beltagy, Kyle Lo, and Arman Cohan. SciBERT: A pretrained language model for scientific text. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 3615–3620, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1371. URL <https://aclanthology.org/D19-1371>.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pp. 2397–2430. PMLR, 2023.

- Stella Rose Biderman, Kieran Bicheno, and Leo Gao. Datasheet for the pile. *ArXiv*, abs/2201.07311, 2022.
- Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, et al. Gpt-neox-20b: An open-source autoregressive language model. In *Proceedings of BigScience Episode# 5–Workshop on Challenges & Perspectives in Creating Large Language Models*, pp. 95–136, 2022.
- Samuel R. Bowman, Jeeyoon Hyun, Ethan Perez, Edwin Chen, Craig Pettit, Scott Heiner, Kamilė Lukošiušė, Amanda Askell, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Christopher Olah, Daniela Amodei, Dario Amodei, Dawn Drain, Dustin Li, Eli Tran-Johnson, Jackson Kernion, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Noudou, Liane Lovitt, Nelson Elhage, Nicholas Schiefer, Nicholas Joseph, Noemí Mercado, Nova DasSarma, Robin Larson, Sam McCandlish, Sandipan Kundu, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Timothy Telleen-Lawton, Tom Brown, Tom Henighan, Tristan Hume, Yuntao Bai, Zac Hatfield-Dodds, Ben Mann, and Jared Kaplan. Measuring progress on scalable oversight for large language models. *arXiv preprint arXiv:2211.03540*, 2022.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *ArXiv*, abs/2005.14165, 2020.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Katherine M. Collins, Albert Q. Jiang, Simon Frieder, Lionel Wong, Miri Zilka, Umang Bhatt, Thomas Lukasiewicz, Yuhuai Wu, Joshua B. Tenenbaum, William Hart, Timothy Gowers, Wenda Li, Adrian Weller, and Mateja Jamnik. Evaluating language models for mathematics through interactions. *arXiv preprint arXiv:2306.01694*, 2023.
- Together Computer. Redpajama: An open source recipe to reproduce llama training dataset, April 2023. URL <https://github.com/togethercomputer/RedPajama-Data>.
- Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. *arXiv preprint arXiv:2307.08691*, 2023.
- Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris Van Doorn, and Jakob von Raumer. The lean theorem prover (system description). In *Automated Deduction-CADE-25: 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings 25*, pp. 378–388. Springer, 2015.
- Erich Elsen, Curtis Hawthorne, and Arushi Somani. The adventure of the errant hardware, 2023. URL <https://www.adept.ai/blog/sherlock-sdc>.
- Emily First, Markus N. Rabe, Talia Ringer, and Yuriy Brun. Baldur: Whole-proof generation and repair with large language models. *arXiv preprint arXiv:2303.04910*, 2023.
- Leo Gao, Stella Rose Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. The pile: An 800gb dataset of diverse text for language modeling. *ArXiv*, abs/2101.00027, 2020.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Jason Ociepa, Chris Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika,

- Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for few-shot language model evaluation, September 2021. URL <https://doi.org/10.5281/zenodo.5371628>.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. Pal: Program-aided language models. *arXiv preprint arXiv:2211.10435*, 2022.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. Pal: Program-aided language models. *arXiv preprint arXiv:2211.10435*, 2023.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III au2, and Kate Crawford. Datasheets for datasets, 2021.
- Herbert L. Gelernter. Realization of a geometry theorem proving machine. In *IFIP Congress*, 1959. URL <https://api.semanticscholar.org/CorpusID:18484295>.
- Priya Goyal, Piotr Dollár, Ross B. Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch SGD: training imagenet in 1 hour. *CoRR*, abs/1706.02677, 2017. URL <http://arxiv.org/abs/1706.02677>.
- Jesse Michael Han, Jason Rute, Yuhuai Wu, Edward Ayers, and Stanislas Polu. Proof artifact co-training for theorem proving with language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=rpxJc9j04U>.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2021a.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *NeurIPS*, 2021b.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and L. Sifre. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration, 2020.
- Albert Q. Jiang, Wenda Li, Jesse Michael Han, and Yuhuai Wu. Lisa: Language models of isabelle proofs. *6th Conference on Artificial Intelligence and Theorem Proving*, 2021.
- Albert Q. Jiang, Wenda Li, Szymon Tworkowski, Konrad Czechowski, Tomasz Odrzygóźdź, Piotr Miłoś, Yuhuai Wu, and Mateja Jamnik. Thor: Wielding hammers to integrate language models and automated theorem provers. *arXiv preprint arXiv:2205.10893*, 2022.
- Albert Qiaochu Jiang, Sean Welleck, Jin Peng Zhou, Timothee Lacroix, Jiacheng Liu, Wenda Li, Mateja Jamnik, Guillaume Lample, and Yuhuai Wu. Draft, sketch, and prove: Guiding formal theorem provers with informal proofs. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=SMa9EAovKMC>.
- Jared Kaplan, Sam McCandlish, T. J. Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeff Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- Denis Kocetkov, Raymond Li, Loubna Ben Allal, Jia Li, Chenghao Mou, Carlos Muñoz Ferrandis, Yacine Jernite, Margaret Mitchell, Sean Hughes, Thomas Wolf, Dzmitry Bahdanau, Leandro von Werra, and Harm de Vries. The stack: 3 tb of permissively licensed source code. *Preprint*, 2022.
- Guillaume Lample, Marie-Anne Lachaux, Thibaut Lavril, Xavier Martinet, Amaury Hayat, Gabriel Ebner, Aurélien Rodriguez, and Timothée Lacroix. Hypertree proof search for neural theorem proving. *arXiv preprint arXiv:2205.11491*, 2022.

- Aitor Lewkowycz, Anders Johan Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Venkatesh Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. *arXiv preprint arXiv:2305.20050*, 2023.
- Pan Lu, Liang Qiu, Wenhao Yu, Sean Welleck, and Kai-Wei Chang. A survey of deep learning for mathematical reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14605–14631, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.817. URL <https://aclanthology.org/2023.acl-long.817>.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*, 2023.
- The mathlib Community. The lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2020*, pp. 367–381, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450370974. doi: 10.1145/3372885.3373824. URL <https://doi.org/10.1145/3372885.3373824>.
- Sewon Min, Suchin Gururangan, Eric Wallace, Hannaneh Hajishirzi, Noah A. Smith, and Luke Zettlemoyer. Silo language models: Isolating legal risk in a nonparametric datastore, 2023.
- Scott Morrison. `lean-training-data`. <https://github.com/semorrison/lean-training-data>, 2023.
- OpenAI. Gpt-4 technical report, 2023.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*, 2022.
- Keiran Paster, Marco Dos Santos, Zhangir Azerbayev, and Jimmy Ba. Openwebmath: An open dataset of high-quality mathematical web text. *CoRR*, abs/2310.06786, 2023. doi: 10.48550/ARXIV.2310.06786. URL <https://doi.org/10.48550/arXiv.2310.06786>.
- Christine Paulin-Mohring. Extracting ω ’s programs from proofs in the calculus of constructions. In *Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pp. 89–104, 1989a.
- Christine Paulin-Mohring. *Extraction de programmes dans le Calcul des Constructions*. PhD thesis, Université Paris-Diderot-Paris VII, 1989b.
- Larry Paulson and Tobias Nipkow. The sledgehammer: Let automatic theorem provers write your isabelle scripts!, 2023. URL <https://isabelle.in.tum.de/website-Isabelle2009-1/sledgehammer.html>.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. Yarn: Efficient context window extension of large language models. *arXiv preprint arXiv:2309.00071*, 2023.
- Stanislas Polu and Ilya Sutskever. Generative language modeling for automated theorem proving. *arXiv preprint arXiv:2009.03393*, 2020.
- Stanislas Polu, Jesse Michael Han, Kunhao Zheng, Mantas Baksys, Igor Babuschkin, and Ilya Sutskever. Formal mathematics statement curriculum learning. *arXiv preprint arXiv:2202.01344*, 2022.

- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 2019.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer, 2023.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. ZeRO: Memory optimizations toward training trillion parameter models. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '20. IEEE Press, 2020. ISBN 9781728199986. doi: 10.5555/3433701.3433727. URL <https://dl.acm.org/doi/10.5555/3433701.3433727>.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*, 2023.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Tali Bers, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. Multitask prompted training enables zero-shot task generalization. *arXiv preprint arXiv:2110.08207*, 2022.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-LM: Training multi-billion parameter language models using model parallelism. *Computing Research Repository*, 2019. doi: 10.48550/arXiv.1909.08053. URL <https://arxiv.org/abs/1909.08053v4>. Version 4.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Nathaneal Scharli, Aakanksha Chowdhery, Philip Mansfield, Blaise Agueray Arcas, Dale Webster, Greg S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkomar, Joelle Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. Large language models encode clinical knowledge, 2022.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, Mike Schaekermann, Amy Wang, Mohamed Amin, Sami Lachgar, Philip Mansfield, Sushant Prakash, Bradley Green, Ewa Dominowska, Blaise Agueray Arcas, Nenad Tomasev, Yun Liu, Renee Wong, Christopher Semturs, S. Sara Mahdavi, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Shekoofeh Azizi, Alan Karthikesalingam, and Vivek Natarajan. Towards expert-level medical question answering with large language models, 2023.
- Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*, 2022.
- Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. Galactica: A large language model for science, 2022.
- Amitayush Thakur, Yeming Wen, and Swarat Chaudhuri. A language-agent approach to formal theorem-proving, 2023.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent

- Zhao, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh Srinivasan, Laichee Man, Kathleen Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc Le. Llama: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process- and outcome-based feedback, 2022.
- H. Wang. Toward mechanical mathematics. *IBM Journal of Research and Development*, 4(1):2–22, 1960. doi: 10.1147/rd.41.0002.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=1PL1NIMMrw>.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2022.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023.
- Sean Welleck. Neural theorem proving tutorial. <https://github.com/wellecks/ntptutorial>, 2023.
- Sean Welleck and Rahul Saha. llmstep: Llm proofstep suggestions in lean. <https://github.com/wellecks/llmstep>, 2023.
- Sean Welleck, Jiacheng Liu, Ximing Lu, Hannaneh Hajishirzi, and Yejin Choi. Naturalprover: Grounded mathematical proof generation with language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=rhdfTOiXBng>.
- Makarius Wenzel, Lawrence C Paulson, and Tobias Nipkow. The isabelle framework. In *Theorem Proving in Higher Order Logics: 21st International Conference, TPHOLs 2008, Montreal, Canada, August 18-21, 2008. Proceedings 21*, pp. 33–38. Springer, 2008.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhajan Kambadur, David Rosenberg, and Gideon Mann. Bloomberggpt: A large language model for finance, 2023.

- Yuhuai Wu, Albert Qiaochu Jiang, Wenda Li, Markus Norman Rabe, Charles E Staats, Mateja Jamnik, and Christian Szegedy. Autoformalization with large language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=IUikebJ1Bf0>.
- Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy Liang, Quoc V. Le, Tengyu Ma, and Adams Wei Yu. Doremi: Optimizing data mixtures speeds up language model pretraining. *arXiv preprint arXiv:2305.10429*, 2023.
- Kaiyu Yang, Aidan Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan Prenger, and Anima Anandkumar. LeanDojo: Theorem proving with retrieval-augmented language models. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- Zheng Xin Yong, Hailey Schoelkopf, Niklas Muennighoff, Alham Fikri Aji, David Ifeoluwa Adelani, Khalid Almubarak, M Saiful Bari, Lintang Sutawika, Jungo Kasai, Ahmed Baruwa, Genta Winata, Stella Biderman, Edward Raff, Dragomir Radev, and Vassilina Nikoulina. BLOOM+1: Adding language support to BLOOM for zero-shot prompting. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11682–11703, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.653. URL <https://aclanthology.org/2023.acl-long.653>.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T. Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*, 2023.
- Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhua Chen. Mammoth: Building math generalist models through hybrid instruction tuning. *CoRR*, abs/2309.05653, 2023. doi: 10.48550/arXiv.2309.05653. URL <https://doi.org/10.48550/arXiv.2309.05653>.
- Shizhuo Dylan Zhang, Curt Tigges, Stella Biderman, Maxim Raginsky, and Talia Ringer. Can transformers learn to solve problems recursively?, 2023.
- Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. Minif2f: a cross-system benchmark for formal olympiad-level mathematics. *arXiv preprint arXiv:2109.00110*, 2021.
- Hattie Zhou, Azade Nova, Hugo Larochelle, Aaron Courville, Behnam Neyshabur, and Hanie Sedghi. Teaching algorithmic reasoning via in-context learning, 2022.
- Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.

A AUTHOR CONTRIBUTIONS

Training Data. Zhangir Azerbayev, Keiran Paster, Marco Dos Santos, Sean Welleck.

Model training. Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster.

Evaluations. Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen McAleer, Albert Q. Jiang, Sean Welleck.

Formal math evaluations. Sean Welleck.

Memorization analysis. Sean Welleck, Keiran Paster.

Senior Authorship and Advising. Jia Deng, Stella Biderman, Sean Welleck.

B DATA: Proof-Pile-2

Data source	Tokens	Weight
Proof-Pile-2	55B	–
Code (AlgebraicStack)	11B	1.00
Web (OpenWebMath)	15B	4.00
Papers (ArXiv)	29B	2.00
General code (RedPajama)	59B	0.22
General language (Pile)	300B	0.15

Table 8: Proof-Pile-2 data sources (top), general language and code data included during training (bottom), and the mixture weights of each component during training.

B.1 MATHEMATICAL CODE: AlgebraicStack

AlgebraicStack contains roughly 11B tokens of code related to mathematics. We describe its sources, filtering, and content below. Table 9 shows the number of tokens per language in AlgebraicStack.

Language	AlgebraicStack tokens	Language	AlgebraicStack tokens
Agda	35.2 M	Julia	531.0 M
C	25.1 M	Jupyter	199.1 M
C++	954.1 M	Lean	285.6 M
Coq	281.9 M	Maple	2.0 M
Fortran	724.9 M	Matlab	65.8 M
GAP	3.6 M	Python	6,098.8 M
Haskell	9.1 M	R	71.3 M
Idris	10.9 M	Tex	567.7 M
Isabelle	1,089.7 M	Total	10,955.7 M

Table 9: Tokens in AlgebraicStack, computed with the Llama tokenizer.

B.1.1 GITHUB CODE

The following programming languages were either barely present in the Stack or consisted of largely incorrect filetypes, so we downloaded data for these languages directly via the Github Python API.

- **Coq** : We filter for files with the `.v` extension, and include Coq via including files that match a heuristic filter for the keywords "**Theorem**", "**Proof**", "**Qed**", "**Inductive**", "**Definition**", "**Fixpoint**" and exclude Verilog files via the keyword blacklist "**pragma**", "**endmodule**", "**posedge**", "**negedge**", "**wire**". We additionally exclude files noted as automatically generated.

- **Isabelle** : We filter for files with the `.thy` extension and include files matching the keyword whitelist `"theorem "`, `"lemma "`. We keep only `isabelle-prover/mirror-afp-devel` and discard all other older copies of the Archive of Formal Proofs. We further remove theorem statements and proofs that have a theorem name in the PISA (Jiang et al., 2021) test set.
- **Lean** : We filter for files with the `.lean` extension, using the keyword whitelist `"theorem "`, `"lemma "`, `"example "`. We remove all dependency files, and in order to avoid known benchmark contamination, we blacklist the `ProofNet` and `MiniF2F` repositories. We further remove theorems or lemmas that share a theorem name with the `LeanDojo` (Yang et al., 2023) val or test sets.
- **MATLAB** : We filter for files with the `.m` extension, using the keyword whitelist `"#import"`, `"#interface"`, `"#implementation"`, `"#property"`, and blacklist C files via the keywords `"#include"` and the regex `r' main\(. *{\$'`

We implemented a cutoff date for our Github API downloads, and used a cutoff date of April 1, 2023.

For all languages, unless otherwise stated, we additionally filtered out files with a filesize greater than 1048575 bytes or with a numerical density (ratio of digit characters to non-digit characters) of 0.5. We additionally perform document-level exact deduplication by removing documents which contain an overlapping 2048-character chunk as another document.

B.1.2 LEAN PROOFSTEPS

We extract a dataset of (tactic state, next tactic) pairs from Mathlib 4 (mathlib Community, 2020) using the `lean-training-data` (Morrison, 2023) tool. We use Mathlib 4 commit `c779bd5`, which was created on August 20th 2023.

B.1.3 ISABELLE PROOFSTEPS

We construct a dataset of Isabelle proofs, building upon the PISA dataset Jiang et al. (2021). Isabelle Proofsteps comprises proofs from the Archive of Formal Proofs and Isabelle Standard Library, scraped with PISA Jiang et al. (2021). Each entry in the dataset includes the theorem statement, the proof states and the proof steps, separated by specific tags. To maintain the integrity of evaluations using the PISA test set, we decontaminate Isabelle Proofsteps by removing theorems whose names overlap with those in the PISA test set. Although this approach results in a strict filtering – removing more than 10,000 theorems although there are only 3600 in the PISA test set – we consider it acceptable in order to mitigate data contamination. After filtering, Isabelle Proofsteps contains 251,000 theorems.

B.1.4 STACK FILTERING

We source the following programming languages from the Stack (Kocetkov et al., 2022) dataset, and describe our filtering process and quality issues we chose to mitigate beyond our default quality heuristics:

- **Agda**: Only standard filters applied.
- **C** : We include documents based on a keyword whitelist, namely: `"#include <fftw.h>"`, `"#include <fftw3.h>"`, `"#include <rfftw.h>"`, `"#include <gsl>"`, `"#include <blas.h>"`, `"#include <blas.h>"`, `"#include <lapacke.h>"`, `"#include <nlopt.h>"`, `"#include <petsc.h>"`.
- **C++** : We include documents based on a keyword whitelist, namely: `"#include <adept_arrays.h>"`, `"#include <adept.h>"`, `"#include <alglib>`, `"#include <boost"`, `"#include <armadillo"`, `"#include <blitz"`, `"#include <Eigen"`, `"#include <deal.II"`, `"#include <dlib"`, `"#include <NTL"`, `"#include <ntl"`.
- **Fortran** : Only standard filters applied.
- **GAP** : Only standard filters applied.
- **Haskell** : We filtered the data to only contain files with the following imports: `Numeric.LinearAlgebra`, `Numeric.SpecFunctions`, `Numeric.Vector`, `Statistics`, `Data.Complex`.

- **Idris** : Only standard filters applied.
- **Julia** : We filtered out mislabeled JSON lines files. We removed files larger than 10,000 characters long which both were not files containing tests and which had a lower numerical density than 0.5, and otherwise ignored numerical density. We additionally only accepted files within a specific keyword whitelist, to attempt to control relevance to scientific computing, namely: "**LinearAlgebra**", "**DifferentialEquations**", "**Symbolics**", "**Distributions**", "**DataFrames**", "**DynamicalSystems**", "**Turing**", "**Gen**", "**JuMP**", "**sqrt**", "**abs**", "**zeros**", "**ones**", "**sin**", "**cos**", "**tan**", "**log**", "**exp**", "**integrate**", "**likelihood**", "**Matrix**", " π ", "**pi**", "**rand**", "**grad**".
- **Jupyter** : We found that many Jupyter notebook files were large due to containing long cell outputs, such as base64 images, long tracebacks, or other extra JSON cell metadata. We use `nbconvert` to convert notebooks to a markdown format, removing metadata.
- **Maple** : We filtered out files with a size greater than 100,000 bytes, and found that some files were XML. We filtered all files beginning with an XML declaration.
- **Python** : We filtered notebooks and JSON files out by excluding documents with beginning "{" characters, and included only files importing from a fixed list of libraries.
- **R** : We excluded all files beginning with an XML declaration. We additionally filtered out all notebooks, and filtered all files containing MacOS "Resource Fork" files.
- **Tex** : We used a max file size of 10,000,000 bytes. We excluded tex files found in directories named "**latex**" because these were often auto-generated files, and excluded documents using **gnuplot**. We included only documents containing one of the keywords "**\chapter{**", "**\chapter*{**", "**\section{**", "**\section*{**", "**\subsection{**", "**\subsection*{**", "**\subsubsection{**", "**\subsubsection*{**", "**\paragraph{**", "**\subparagraph{**", and additionally only included documents identified as English by a classifier from the `langid` package.

For all languages we used within the Stack, unless otherwise stated, we additionally filtered out files with a filesize greater than 1048575 bytes or with a numerical density (ratio of digit characters to non-digit characters) of 0.5.

We used v1.2 of the near-deduplicated Stack as a base for processing.

B.2 PAPERS: ARXIV

We use the entirety of ArXiv, as accessed by Computer (2023) in April 2023. For further information on preprocessing applied to ArXiv, see Computer (2023).

B.3 WEB: OPENWEBMATH

For the web portion of our training dataset, we use OpenWebMath (Paster et al., 2023).

C EVALUATION HARNESS

We implement a variety of math-related tasks and evaluation protocols into a public fork of the Language Model Evaluation Harness (Gao et al., 2021). The Harness provides a model-agnostic framework for standardized, reproducible evaluation of language models.

We add the following tasks for the evaluations in this paper:

- `hendrycks_math_ppl`: Perplexity evaluation on MATH (Hendrycks et al., 2021a) sub-tasks.
- `minif2f_isabelle`: Proof autoformalization in Isabelle on the miniF2F benchmark based on Jiang et al. (2023), with a Portal-to-Isabelle (Jiang et al., 2021) proof checker.
- `minerva_math`: The MATH benchmark with the prompt and Sympy evaluation from Minerva (Lewkowycz et al., 2022).
- `minerva-hendrycksTest`: MMLU-STEM tasks following Lewkowycz et al. (2022).

- `ocw_courses`: The OCW Courses task from Lewkowycz et al. (2022).
- `python_gsm8k`: GSM8k with Python, based on Gao et al. (2022).
- `sympy_math`: MATH with Sympy evaluation.

We include a link to the implementations for these tasks, including full prompts, in our public codebase.

D EVALUATION: EXPERIMENT DETAILS

D.1 ISABELLE INFORMAL-TO-FORMAL THEOREM PROVING

We follow Jiang et al. (2023), allowing the model to issue a call to built-in Isabelle automation in the output proof by generating `sledgehammer`. This calls Sledgehammer (Paulson & Nipkow, 2023) and the list of heuristics listed in Jiang et al. (2023). Following Jiang et al. (2023), as a baseline we use Sledgehammer and the heuristics executed at the beginning of the proof (referred to as Sledgehammer in the main text for brevity). We use a 30-second timeout for Sledgehammer and implement proof checking via Portal-to-Isabelle (Jiang et al., 2021). Refer to the implementation in the Evaluation Harness for further details.

D.2 LEAN THEOREM PROVING

Theorem proving via tactic prediction involves interacting with a proof assistant after each step of a proof. Implementing these interactions within the evaluation harness is outside the scope of this work. Therefore, for the Lean theorem proving task we use a separate evaluation setup based on an open-source implementation (Welleck, 2023). We include our evaluation code in our public codebase.

Setup. We evaluate on miniF2F (Zheng et al., 2021), which consists of 488 formalized statements from math competitions and undergraduate coursework. Given a formalized statement, the task is to generate a formal proof that is checked by Lean.

We use best first search, commonly used for neural tactic prediction models (e.g., Polu & Sutskever (2020)). Best first search is parameterized by the number of attempts (N), generated tactics per iteration (S), and maximum iterations (T). We define the *search budget* to be the maximum number of generated tactics, $N \times S \times T$. We set our search budget to $N = 1$, $S = 32$, and $T = 100$, less than that of the baseline model. Following Yang et al. (2023), we generate tactics with beam search and use a 10 minute timeout. We adapt the proof search implementation from Welleck (2023), which uses LeanDojo v.1.1.2 (Yang et al., 2023) for interaction. We use Lean 4 miniF2F, using <https://github.com/rah4927/lean-dojo-mew> commit `d00c776260c77de7e70125ef0cd119de6c0ff1de`. Note that the ReProver baseline from (Yang et al., 2023) reports performance with Lean 3.

Prompt. We prompt the model with three (state, tactic) examples, shown in Figure 5.

```

"""Given the Lean 4 tactic state, suggest a next tactic.
Here are some examples:

Tactic state:
---
 $\alpha$  : Type u_1
r :  $\alpha \rightarrow \alpha \rightarrow \text{Prop}$ 
inst1 : DecidableEq  $\alpha$ 
inst : IsIrrefl  $\alpha$  r
 $\vdash \text{CutExpand } r \leq \text{InvImage } (\text{Finsupp.Lex } (\text{ $\lambda$ r } \Pi \text{ fun } x \text{ } x_1 \Rightarrow x \neq x_1)$ 
  fun x x_1 => x < x_1)  $\uparrow$ toFinsupp
---
Next tactic:
---
rintro s t <u, a, hr, he>
---

Tactic state:
---
 $\iota$  : Type u_1
I J : Box  $\iota$ 
x y :  $\iota \rightarrow \mathbb{R}$ 
I J : WithBot (Box  $\iota$ )
 $\vdash \uparrow I = \uparrow J \leftrightarrow I = J$ 
---
Next tactic:
---
simp only [Subset.antisymm_iff,  $\leftarrow$  le_antisymm_iff,
  withBotCoe_subset_iff]
---

Tactic state:
---
m n :  $\mathbb{N}$ 
h : Nat.coprime m n
 $\vdash \text{Nat.gcd } m \text{ } n = 1$ 
---
Next tactic:
---
rw [ $\leftarrow$  h.gcd_eq_one]
---

Tactic state:
---
%s
---
Next tactic:
----"""

```

Figure 5: Prompt for the Lean theorem proving experiments.

E DATASHEET

We provide a datasheet for Proof-Pile-2, following the framework in Gebru et al. (2021).

MOTIVATION	
For what purpose was the dataset created?	Proof-Pile-2 was created for the training or finetuning of domain-specific large language models for general mathematics tasks.
Who created the dataset and on behalf of which entity?	The dataset was created by the authors of this paper for the purposes of this research project.
Who funded the creation of the dataset?	The creation of the dataset was funded by the coauthors' grants and employers, as further described in section 5.
Any other comment?	
COMPOSITION	
What do the instances that comprise the dataset represent?	Instances are text-only documents.
How many instances are there in total?	We detail fine-grained token counts elsewhere in this paper.
Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?	Our dataset is filtered based on our assessments of quality for the language modeling task. More detail on methodology can be found in Appendix B.
What data does each instance consist of?	Each instance is a text-only document, alongside metadata about its originating split and filename or location.
Is there a label or target associated with each instance?	No.
Is any information missing from individual instances?	Yes, we filter undesired noise, such as base64-encoded images, from some documents.
Are relationships between individual instances made explicit?	No.
Are there recommended data splits?	Yes, we release a canonical train, validation, and test split of the dataset, which we follow in this work.
Are there any errors, sources of noise, or redundancies in the dataset?	We make our best efforts to remove errors or sources of noise, but our dataset will naturally contain documents with errors or noise, and may contain near-duplicate documents.
Is the dataset self-contained, or does it link to or otherwise rely on external resources?	The dataset is self-contained, but can also be reconstructed based on external publicly available data sources and datasets following our instructions.
Does the dataset contain data that might be considered confidential?	All documents in Proof-Pile-2 are publicly available online.

Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety?	We estimate toxic content to be less prevalent in our dataset than other more general web-based datasets, due to its technical focus. However, it is likely to contain such content.
COLLECTION	
How was the data associated with each instance acquired?	Data was largely sourced from existing public subsets, such as the RedPajama dataset (Computer, 2023), OpenWebMath dataset (Paster et al., 2023), and via filtering the Stack (Kocetkov et al., 2022). Some data was collected using the Github API.
What mechanisms or procedures were used to collect the data?	See above.
If the dataset is a sample from a larger set, what was the sampling strategy?	We release the entirety of the dataset following the application of our quality filters. We randomly held out validation and test splits from the dataset.
Who was involved in the data collection process and how were they compensated?	The authors of this paper participated in locating, retrieving, and filtering the dataset.
Over what timeframe was the data collected?	This data was collected in 2023, with a cut-off date of April 2023 for all subsets with the exception of our Lean proofstep data.
Were any ethical review processes conducted?	Yes, the authors conducted an informal ethical review internally.
PREPROCESSING	
Was any preprocessing/cleaning/labeling of the data done?	Yes, the authors extensively filtered the dataset subsets in keeping with our expectations for high-quality language modeling data in our domain. See Appendix B for further detail on filtering steps taken.
Was the “raw” data saved in addition to the preprocessed/cleaned/labeled data?	Raw data can be accessed via reuse of our provided codebase.
Is the software that was used to preprocess/clean/label the data available?	Yes. We release our codebase, which can be used to reproduce our dataset and its construction process, at https://github.com/EleutherAI/math-lm .
USES	
Has the dataset been used for any tasks already?	Yes, this dataset has been used to train the LLEMMA language models as a domain adaptation and continued pretraining corpus.
Is there a repository that links to any or all papers or systems that use the dataset?	No.
What (other) tasks could the dataset be used for?	The dataset was specifically targeted as a high quality language modeling corpus for the mathematics domain, but may be useful for general-purpose language modeling or unforeseen other downstream uses.

Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses?	We filtered the dataset with the intent of creating a model useful for mathematical tasks with solely English text.
Are there tasks for which the dataset should not be used?	The dataset should not be used with the intent to cause harm or for models intended for the purposes of harm.
DISTRIBUTION	
Will the dataset be distributed to third parties outside of the entity on behalf of which the dataset was created?	We make the dataset publicly available for reproducibility, analysis, and other further downstream uses.
How will the dataset will be distributed?	We provide code to replicate the dataset, and release it via the Huggingface Hub.
When will the dataset be distributed?	The dataset is available immediately.
Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?	We do not relicense the dataset’s components, and do not impose our own use restrictions.
Have any third parties imposed IP-based or other restrictions on the data associated with the instances?	Not to our knowledge.
Do any export controls or other regulatory restrictions apply to the dataset or to individual instances?	Not to our knowledge.
MAINTENANCE	
Who will be supporting/hosting/maintaining the dataset?	The dataset will be hosted on the HuggingFace Hub and able to be recreated via code at https://github.com/EleutherAI/math-1m . The dataset will not be updated post-release.
How can the owner/curator/manager of the dataset be contacted?	Via email at za2514@princeton.edu
Is there an erratum?	No.
Will the dataset be updated?	No.
If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?	No.

Table 10: **Datasheet for Proof-Pile-2**, following the framework introduced by Gebru et al. (2021).

F ADDITIONAL RESULTS

F.1 PROOF AUTOFORMALIZATION

Table 11 shows additional results on Isabelle proof autoformalization, including the union of theorems closed by Sledgehammer and the given language model.

Method	Autoformalization pass@1	
	miniF2F-valid*	miniF2F-test
Sledgehammer	14.72%	20.49%
Code Llama 7b	16.31%	17.62%
LLEMMA-7b	20.60%	22.13%
Code Llama 7b \cup Sledgehammer	20.17%	25.00%
LLEMMA-7b \cup Sledgehammer	25.97%	27.46%

Table 11: **Isabelle autoformalization.** *We exclude the 11 examples used in the few-shot prompts. Pass@1 with greedy decoding.

G SUPERVISED FINETUNING

A full exploration of finetuning applications for LLEMMA, such as instruction following (Ouyang et al., 2022; Wei et al., 2022), dialogue modeling (Thoppilan et al., 2022; Touvron et al., 2023; Collins et al., 2023), and reward modeling (Cobbe et al., 2021; Lightman et al., 2023) are outside the scope of this work. However, to establish that LLEMMA retains its advantage over other open models when finetuned, we conduct preliminary experiments finetuning LLEMMA-7B on MetaMathQA (Yu et al., 2023), a supervised dataset targeted at the MATH and GSM8k benchmarks. Results are shown in Table 12.

Initialization	Finetune Dataset	MATH	GSM8k
Llama 2 7B	WizardMath (Proprietary)	10.7%	54.9%
Llama 2 7B	MetaMathQA	19.4%	66.4%
LLEMMA 7B	MetaMathQA	25.2%	66.5%
Llama 2 70B	WizardMath (Proprietary)	22.7%	81.6%
Llama 2 70B	MetaMathQA	26.6%	82.3%

Table 12: Finetuning of various 7B base models on supervised mathematics datasets. All results with a Llama 2 initialization are copied from the literature (Luo et al., 2023; Yu et al., 2023). The LLEMMA 7B finetune is trained with identical hyperparameters to the models in Yu et al. (2023)

H QUALITATIVE EXAMPLES

Dataset overlap. Figure 6 shows example false positives when checking n -gram overlap with OpenWebMath documents for various n . Figure 7 shows an example OpenWebMath document that has 30-gram overlap with a MATH problem, and LLEMMA-7b’s generated solution.

Task outputs. Figure 8 shows a generated proof in the informal2formal theorem proving task.

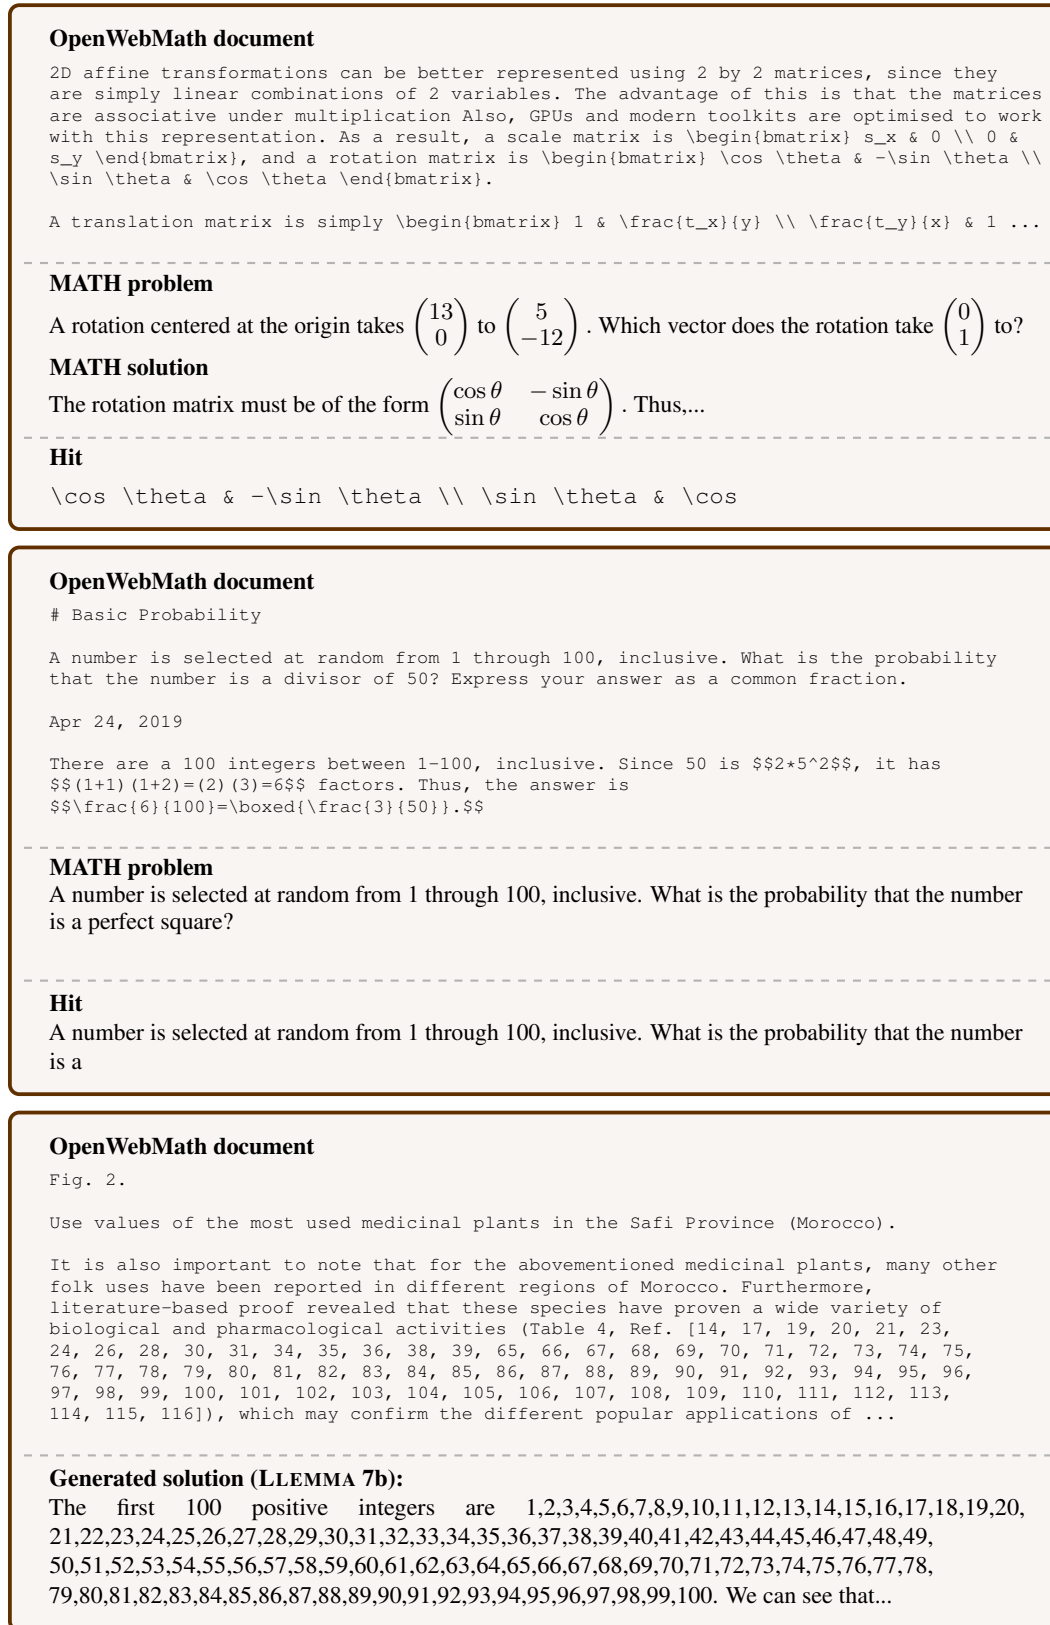


Figure 6: **Data overlap:** Example false positives using 10-gram match between MATH solutions and OpenWebMath documents (*top*), 20-gram match between MATH problems and OpenWebMath documents (*middle*), and 30-gram match between LLEMMA-7b’s generated solutions and OpenWebMath documents (*bottom*).

OpenWebMath document

A triangle is formed with edges along the line $y = \frac{2}{3}x + 5$, the x -axis, and the line $x = k$. If the area of the triangle is less than 20, find the sum of all possible integral values of k .

Feb 28, 2018

Look at the graph, here...two triangles are possible :
<https://www.desmos.com/calculator/m6wnjppgldq>
 The height of the triangles at any point will be formed by

$$\left[\frac{2}{3}x + 5 \right]$$

And the bases will be $[x - (-7.5)] = [x + 7.5]$
 So...we want to solve this
 $(1/2) [\frac{2}{3}x + 5] [x + 7.5] = 20$
 $[\frac{2}{3}x + 5] [x + 7.5] = 40$
 $(2/3)x^2 + 5x + 5x + 37.5 = 0$
 $(2/3)x^2 + 10x - 2.5 = 0$

Using a little technology....the max x value for the triangle formed above the x axis will be = .246

And the min x value for the triangle formed below the x axis will be = -15.246

With the given boundaries, the integer sums of all possible x values of k giving triangles with an area < 20 units² =

$$[(-15) + (-14) + (-13) + \dots + (-2) + (-1) + 0] =$$

$$- (15)(16) / 2 =$$

$$-120$$

Feb 28, 2018

MATH Problem: A triangle is formed with edges along the line $y = \frac{2}{3}x + 5$, the x -axis, and the line $x = k$. If the area of the triangle is less than 20, find the sum of all possible integral values of k .
Answer: -120

Generated solution (LLEMMA 7b):
 The triangle has vertices $(k, 0)$, $(0, 5)$, and $(k, \frac{2}{3}k + 5)$. The base of the triangle has length k , and the height of the triangle has length $\frac{2}{3}k + 5$. The area of the triangle is therefore

$$A = \frac{1}{2}bh = \frac{1}{2}k\left(\frac{2}{3}k + 5\right) = \frac{1}{3}k^2 + \frac{5}{2}k.$$

Since the area of the triangle is less than 20, we have

$$\frac{1}{3}k^2 + \frac{5}{2}k < 20 \Rightarrow \frac{1}{6}k^2 + \frac{5}{4}k - 10 < 0.$$

This inequality is true if and only if $\frac{1}{6}k^2 + \frac{5}{4}k - 10$ is positive. Since $\frac{1}{6}k^2 + \frac{5}{4}k - 10$ is a quadratic with positive leading coefficient, it is positive if and only if its vertex is less than 0. The vertex of $\frac{1}{6}k^2 + \frac{5}{4}k - 10$ is $(-\frac{10}{3}, -\frac{125}{18})$. Since $-\frac{10}{3} < 0$, we have that $\frac{1}{6}k^2 + \frac{5}{4}k - 10 > 0$ if and only if $k < -\frac{10}{3}$. Therefore, the only possible integral values of k are $k = 0, 1, 2, 3, 4$. The sum of these values is $\boxed{10}$. Final Answer: The final answer is 10.

Figure 7: **Data overlap:** Example OpenWebMath document that has a 30-gram overlap with the given MATH problem, and LLEMMA-7b's generated solution.

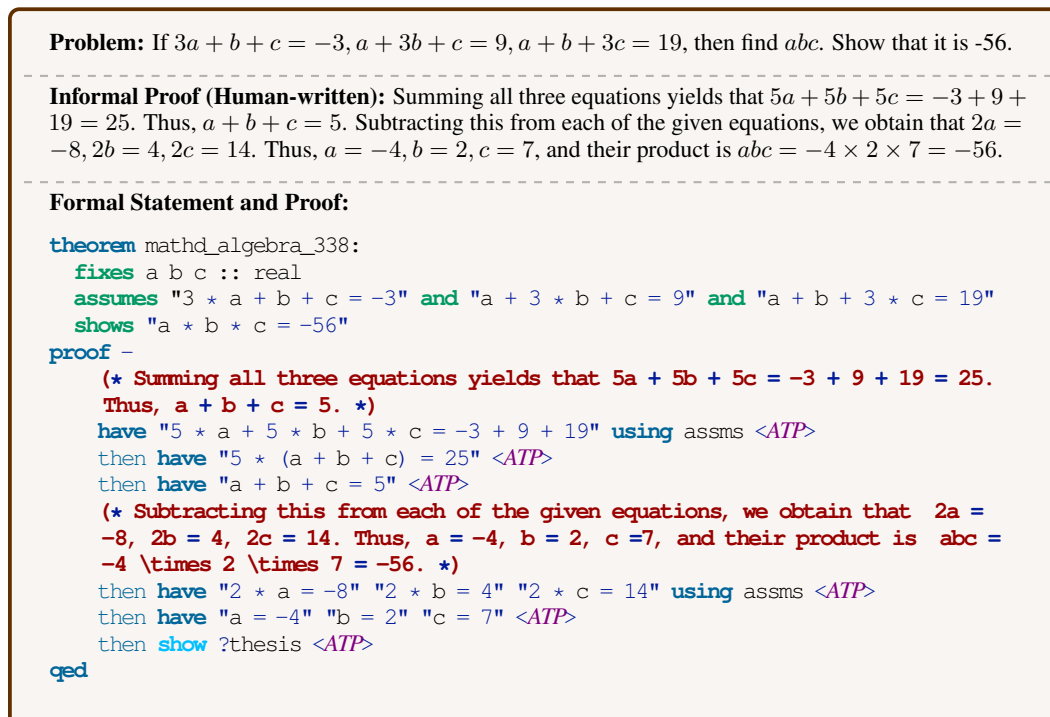


Figure 8: **Informal-to-formal proving.** The model is given the problem, informal proof, and formal statement, following Jiang et al. (2023). It generates a formal proof (starting with `proof -`) containing Isabelle code, comments (`(* . . . *)`) that align the informal and formal proofs, and calls to an automated prover (shown as `<ATP>`). The proof is from LLEMMA-7b with greedy decoding.