## Language models and formal mathematics

Sean Welleck

University of Washington

Carnegie Mellon University (Jan 2024)

## Language models

General sequence generation

- Summarize documents
- Generate code from a description


## Language models

Sequence generation in mathematics

- Solve problems
- Open-ended dialogue
- Help generate verified proofs
- ...

[^0]
## "Informal" mathematics

Input:
Let $r$ be a root of $x^{2}+5 x+7=0$. Compute $(r-1)(r+2)(r+6)(r+3)$.
LLEmmA 34B solution:
We have

$$
\begin{aligned}
(r-1)(r+2)(r+6)(r+3) & =\left(r^{2}+r-2\right)\left(r^{2}+9 r+18\right) \\
& =r^{4}+10 r^{3}+25 r^{2}+19 r-36 \\
& =r^{4}+5 r^{3}+7 r^{2}+5 r+7+5 r^{3}+20 r^{2}+14 r-36 \\
& =r^{4}+5 r^{3}+25 r^{2}+19 r-29 \\
& =\left(r^{2}+5 r+7\right)\left(r^{2}+4 r-5\right) \\
& =(0)\left(r^{2}+4 r-5\right) \\
& =0 .
\end{aligned}
$$

Final Answer: The final answer is 0 .
Figure 1: A LLEMMA 34B solution to a MATH [4] problem.

## "Informal" mathematics

Input:
Let $r$ be a root of $x^{2}+5 x+7=0$. Compute $(r-1)(r+2)(r+6)(r+3)$.
LLEmmA 34B solution:
We have

$$
\begin{aligned}
(r-1)(r+2)(r+6)(r+3) & =\left(r^{2}+r-2\right)\left(r^{2}+9 r+18\right) \\
& =r^{4}+10 r^{3}+25 r^{2}+19 r-36 \\
& =r^{4}+5 r^{3}+7 r^{2}+5 r+7+5 r^{3}+20 r^{2}+14 r-36 \\
& =r^{4}+5 r^{3}+25 r^{2}+19 r-29 \\
& =\left(r^{2}+5 r+7\right)\left(r^{2}+4 r-5\right) \\
& =(0)\left(r^{2}+4 r-5\right) \\
& =0 .
\end{aligned}
$$

Final Answer: The final answer is 0 .
Figure 2: No correctness guarantees, errors can be difficult to detect.

## Formal mathematics



Figure 3: Mathematics as verifiable source code

## Formal mathematics (Demo)

## If $R \subseteq S$ and $S \subseteq T$ then $R \subseteq T$ <br> 

## Formal mathematics

## Lean Mathlib

- 1+ million lines of code
- > 300 contributors
- Algebra, Linear Algebra, Topology,
 Analysis, Probability, Geometry, Combinatorics, ...


## Formal mathematics

- Liquid tensor project: Lean formalization with Peter Scholze ${ }^{1}$
- Courses at CMU, Imperial College London, Fordham, JHU, ...2
- eXperimental Lean Lab at the University of Washington! ${ }^{3}$

[^1]
## Formal mathematics

佱
Terence Tao
@tao@mathstodon.xyz
Finished formalizing in \#Lean4 the proof of an actual new theorem (Theorem 1.3) in my recent paper arxiv.org/abs/2310.05328 :

Figure 4: Terence Tao's Lean formalization project (October 2023)

## Language models and formal mathematics

## Generative Language Modeling for Automated Theorem Proving

Stanislas Polu OpenAI spolu@openai.com

Ilya Sutskever
OpenAI
ilyasu@openai.com

Figure 5: gpt-f (2020)

## Language models and formal mathematics


metamath / set.mm

f- Merged nmegill merged 6 commits into metamath:develop from spolu:openai-shorten on Mar 27, 2020
"Any ML-based system is impressive if it can find many shorter proofs than the ones we already have. Nice work."

This Pr
models "The shorter proof is easier to translate. It's more symmetric in that it treats A and B identically. It's philosophically more concise in that it doesn't rely on the existence of a universal class of all sets."

Figure 6: gpt-f (2020)

## Language models and formal mathematics

Terence Tao
@tao@mathstodon.xyz
Finished formalizing in \#Lean4 the proof of an actual new theorem (Theorem 1.3) in my recent paper arxiv.org/abs/2310.05328 :
The ability of Github copilot to correctly anticipate multiple lines of code for various routine verifications, and inferring the direction I want to go in from clues such as the names I am giving the theorems, continues to be uncanny.

Figure 7: Terence Tao's Lean formalization project (October 2023)

## This talk: "build your own Lean copilot"

- Part 1: Small models trained to predict the next step of a proof
- Part 2: Llemma: foundation model for mathematics

[^2]
## This talk: "build your own Lean copilot"

- Part 1: Small models trained to predict the next step of a proof
- Part 2: Llemma: foundation model for mathematics

LLMSTEP: tool for receiving verified language model suggestions


Proof state

Vlimstep suggestions


Try this:

- \$ linarith
- rw [- sub_eq_zero] at ha
. apply eq_neg_of_add_eq_zero_left
- rw [- Int.negSucc_coe] at ha

Next-step suggestions

[^3]PART I:
Next-step prediction

## Next-step prediction

| Topic | Notebook |
| :--- | ---: |
| 0. Intro | notebook |
| 1. Data | notebook |
| 2. Learning | notebook |
| 3. Proof Search | notebook |
| 4. Evaluation | notebook |
| 5. llmsuggest | notebook |

Interactive notebooks and code: github.com/wellecks/ntptutorial ${ }^{4}$

[^4]
## Next-step prediction

- Language model suggests next-proof-steps
- Generate a full proof via tree search


[^5]
## Language models

- Model: $p_{\theta}(y \mid x ; \mathcal{D})$
- y : output sequence
- $x$ : input sequence
- $\theta$ : parameters (e.g., transformer)
- $\mathcal{D}$ : dataset


## Language models

- Model: $p_{\theta}(y \mid x ; \mathcal{D})$
- $y$ : output sequence
- x : input sequence
- $\theta$ : parameters (e.g., transformer)
- $\mathcal{D}$ : dataset
- Learning:
- $\arg \max _{\theta} \sum_{y \in \mathcal{D}} \log p_{\theta}(y)$


## Language models

- Model: $p_{\theta}(y \mid x ; \mathcal{D})$
- $y$ : output sequence
- x : input sequence
- $\theta$ : parameters (e.g., transformer)
- $\mathcal{D}$ : dataset
- Learning:
- $\arg \max _{\theta} \sum_{y \in \mathcal{D}} \log p_{\theta}(y)$


## Language models

- Model: $p_{\theta}(y \mid x ; \mathcal{D})$
- y : output sequence
- x : input sequence
- $\theta$ : parameters (e.g., transformer)
- $\mathcal{D}$ : dataset
- Learning:
- $\arg \max _{\theta} \sum_{y \in \mathcal{D}} \log p_{\theta}(y)$
- Inference:
- $y=f\left(p_{\theta}(\cdot \mid x)\right)$
- f: e.g., sampling


## Problem setup

Proof: sequence of (state, next step)

- $\left(x_{0}, y_{0}\right), \ldots,\left(x_{t}, y_{t}\right), \ldots,\left(x_{T}, y_{T}\right)$
- $x_{t}$ : proof state from Lean
- $y_{t}$ : proof step ("tactic")
- $x_{T}$ : proof complete


## Problem setup

Proof: sequence of (state, next step)

- $\left(x_{0}, y_{0}\right), \ldots,\left(x_{t}, y_{t}\right), \ldots,\left(x_{T}, y_{T}\right)$
- $x_{t}$ : proof state from Lean
- $y_{t}$ : proof step ("tactic")
- $x_{T}$ : proof complete

Idea:

- Collect a dataset $\mathcal{D}$ of (state, next step) examples
- Train a language model $p_{\theta}\left(y_{t} \mid x_{t}\right)$ using $\mathcal{D}$


## Data



- Extract (state, next step) pairs, e.g. from Mathlib
- Open-source tooling available ${ }^{5}$

[^6]
## Data



- Extract (state, next step) pairs, e.g. from Mathlib
- Open-source tooling available ${ }^{5}$
- Mathlib yields a dataset $\mathcal{D}$ with 170,000 pairs

[^7]
## Learning

- Standard supervised learning on $\mathcal{D}$ :

$$
\arg \max _{\theta} \sum_{\left(x_{t}, y_{t}\right) \in \mathcal{D}} \log p_{\theta}\left(y_{t} \mid x_{t}\right)
$$

## Learning

```
    Input:
    [GOAL]l : Type u_1
    It J+ : Box l
Xt x y : ı }->\mathbb{R
    I J : WithBot (Box r)
    \vdash \uparrowI = \uparrowJ ↔ I = J[PROOFSTEP]
    Output:
yt simp only [Subset.antisymm_iff, & le_antisymm_iff, withBotCoe_subset_iff]<|endoftext|>
```


## Proof search

- Use generator $p_{\theta}\left(y_{t} \mid x_{t}\right)$ to generate a full proof $y_{1}, \ldots, y_{T}$
- Standard approach: Best-first search


## Best-first search



Figure 8: Best-first search ${ }^{6}$

[^8]
## Evaluation

- Proof search on held-out theorems from the training distribution



## Evaluation



Figure 9: Proof search performance on held-out Mathlib theorems. ${ }^{\text { }}$

[^9]
## Evaluation (out of domain)

Benchmarks evaluate problems drawn from a different distribution:

- miniF2F [11]: competition problems (AMC, AIME, IMO)
- ProofNet [1]: undergraduate textbooks (e.g. Analysis, Topology)


## Evaluation

```
-- from mathlib:
theorem prod_mono
    {\mp@subsup{s}{1}{}}\quad\mp@subsup{s}{2}{}:\mathrm{ : Subsemiring R|} (hs : }\mp@subsup{s}{1}{}\leq\mp@subsup{s}{2}{}
    {t, t t 2 : Subsemiring S|} (ht : t 
    s
        intro x hx
        simp_rw [Subsemiring.mem_prod]
        cases' x with x_fst x_snd
        exact \langlehs hx.1, ht hx.2\rangle
-- from miniF2F:
theorem mathd_algebra_159 (b : \mathbb{R})(f:\mathbb{R}->\mathbb{R})
    (ho: : }\forall\textrm{x},\textrm{f}x=3*\mp@subsup{x}{}{\wedge}4-7*\mp@subsup{x}{}{\wedge}3+2*\mp@subsup{x}{}{\wedge}2-b*x+1
    (h): f 1 = 1) : b = -2 := by
        apply eq_neg_of_add_eq_zero_left
        rw [ho] at h1
        norm_num at h}\mp@subsup{h}{1}{
        linarith
```

Figure 10: Generated proofs on mathlib and miniF2F

## Verified Lean copilot: LLMSTEP



LLMSTEP: tool for receiving verified language model suggestions ${ }^{8}$

[^10]
## Verified Lean copilot: LLMstep

## LLMSTEP: tool for receiving verified language model suggestions

```
variable {\alpha : Type _} (R S T : Set \alpha)
example (h1: R\subseteqS) (h2: S \subseteq T) : R \subseteq T := by
llmstep "|
```

1 goal
$\alpha$ : Type ?u. 14
R S T : Set $\alpha$
h1: R $\subseteq 5$
h2: $\mathrm{S} \subseteq \mathrm{T}$
$\vdash \mathrm{R} \subseteq T$

حllmstep suggestions

## Try this:

-     * apply Set. Subset.trans h1 h2
- exact Set.Subset.trans h1 h2
- intros
- intro $\times$ h

Figure 11: DEMO

## Verified Lean copilot: LLMSTEP

Runtime Comparison


Figure 12: LLMSTEP suggestion latency on CPU and GPU

Part II:
LLEMMA: foundation model for mathematics

## Modeling the distribution of mathematics

- Previous: train a model specifically for tactic prediction


## Modeling the distribution of mathematics

- Previous: train a model specifically for tactic prediction
- Next: model a diverse distribution of math-related sequences


## Modeling the distribution of mathematics

- Previous: train a model specifically for tactic prediction
- Next: model a diverse distribution of math-related sequences
- Perform a task by prompting with a few (input, output) examples


## Modeling the distribution of mathematics

- Previous: train a model specifically for tactic prediction
- Next: model a diverse distribution of math-related sequences
- Perform a task by prompting with a few (input, output) examples
- "Foundation model" [3]: train on large quantity of data, adapt to tasks via prompting or further training


## Foundation model for mathematics

- Language model learning:

$$
\arg \max _{\theta} \sum_{y \in \mathcal{D}} \log p_{\theta}(y)
$$

- Equivalent to:

$$
\begin{gathered}
\arg \min _{\theta} \mathrm{KL}\left(p_{*}, p_{\theta}\right), \\
\quad \text { where } \mathcal{D} \sim p_{*}
\end{gathered}
$$

## Approach 1: train a good generalist



Figure 13: Increasing compute predictably improves language modeling. ${ }^{9}$

[^11]
## Approach 1: train a good generalist



Figure 13: Increasing compute predictably improves language modeling. ${ }^{9}$

- Idea: let $\mathcal{D}$ be as general as possible (with math as a subset), increase compute as much as possible ( $|\mathcal{D}|$ and $|\theta|)$.

[^12]
## Approach 1: train a good generalist

Example: Llama 2

- $\theta$ : 7B parameter transformer
- D : 2 trillion tokens
- General domain: CommonCrawl, Github, Wikipedia, Arxiv, ...


## Approach 1: train a good generalist



Figure 14: Training a generalist can be inefficient. ${ }^{10}$

[^13]
## Approach 2: specialize via transfer



Figure 15: Pretraining on $p_{1}$ can make transfer to $p_{2}$ more efficient ${ }^{11}$

[^14]
## Approach 2: LLEMMA

LLEMMA:
Collect high-quality mathematics data $\mathcal{D}^{\prime}$, transfer to $\mathcal{D}^{\prime} \sim p_{2}$

- Initialize with $\theta_{\text {codellama }}$
- Continue training on $\mathcal{D}^{\prime}: 55$ billion token Proofplle II


## Approach 2: LLEMMA

LLEMMA:
Collect high-quality mathematics data $\mathcal{D}^{\prime}$, transfer to $\mathcal{D}^{\prime} \sim p_{2}$

- Initialize with $\theta_{\text {codellama }}$
- Continue training on $\mathcal{D}^{\prime}$ : 55 billion token Proofpile II
- Mathematical code
- Mathematical web data
- Scientific papers


## LLEMMA



Figure 16: LLEMMA improves with a modest amount of math-specific compute

Data: Proofpile II

## Data: PROOFPILE II

Proofpile II: 55 billion tokens

- Code: 11B tokens
- Web: 15B tokens
- Papers: 29B tokens


## Code: AlgebraicStack

Code: AlgebraicStack

- 11 billion tokens of math-related code
- 17 languages, from the Stack [6], public GitHub repos, proof steps


## AlgebraicStack - Data Quality



Figure 17: AlgEBRAICSTACK pipeline

## AlgebraicStack - Formal mathematics

1.5B tokens of formal math: Agda, Coq, Idris, Isabelle, Lean

- Extracted Lean and Isabelle goal states

$$
\begin{aligned}
& \nabla \text { Tactic state } \\
& \text { example : } \mathrm{R} \subseteq \mathrm{~S} \rightarrow \mathrm{~S} \subseteq \mathrm{~T} \rightarrow \mathrm{R} \subseteq \mathrm{~T}:=\text { by } \\
& 1 \text { goal } \\
& \text { a: Type } \\
& \text { R S T : Set } \alpha \\
& h_{1}: R \subseteq S \\
& h_{2}: S \subseteq T \\
& \vdash \mathrm{R} \subseteq \mathrm{~T}
\end{aligned}
$$

Figure 18: Lean code (left) and goal state (right)

## Web: OpenWebMath

Web: OpenWebMath [Paster et al 2023] ${ }^{12}$

- 14.7 billion tokens of math-related web data
- CommonCrawl with math-specific filtering and extraction

[^15]
## Training



Figure 19: Llemma validation loss

## Formal theorem proving

Traditional proof search: $p_{\theta}$ (next-tactic|state) + best-first search.

- We implement a few-shot version by providing LLemma with 3 (state, next-tactic) examples in its prompt


## LLEMMA formal-to-formal theorem proving



Figure 20: Few-shot proving in Lean with LLemmA ${ }^{13}$

[^16]
## LLemMA as a Lean copilot

LLMSTEP + LLEMMA:

- Part I: send proof state to a tactic prediction model
- $y_{t} \sim p_{\theta}\left(\cdot \mid x_{t}\right)$
- Now: send file content and proof state to Llemma
- $y_{t} \sim p_{\theta}(\cdot \mid$ preceding file content $)$


## Verified Lean copilot: LLMstep

Key difference: use new definitions and theorems

```
structure my_object (\Omega : Type*)[Fintype \Omega] :=
    ( f : \Omega ->\mathbb{R})
    (cool_property : }\forall\textrm{x}:\Omega,0\leqf\times\mathrm{ ) 
theorem my_object_sum_nonneg (o1 o2: my_object \Omega) :
    o1.f + o2.f \geq 0 := by
    apply add_nonneg
    - apply o1.cool_property
    ~ llmstep ""
```

```
Vllmstep suggestions
    Try this:
    - exact o2.cool_property
    - apply o2.cool_property
    - intro hb
    - intro h
```

DEMO

## Verified Lean copilot: LLMstep

Key difference: use new definitions and theorems

| 37 | -- Probability of any outcome is at most one. | Vllmstep suggestions |
| :---: | :---: | :---: |
| 38 | theorem px_le_one ( $\mathrm{p}: \mathrm{pmf} \Omega$ ) ( $\mathrm{x}: \Omega)$ : $\mathrm{p} \times \mathrm{x} 1$ := by |  |
| 39 | refine ${ }^{\text {d }}$ hasSum_le _ (hasSum_ite_eq $\times$ ( $\mathrm{p} \times$ ) ) (hasSum | Try this: |
| 40 | intro x | - > exact p.my_non_neg |
| 41 | split_ifs with h | - * exact p.my_non_neg x |
| 42 | rw [h] | - apply p.my_non_neg |
| 43 | llmstep "' | - simp |
| 44 |  | - apply le_of_lt |

DEMO

## Beyond formal-to-formal proving

Problem: If $3 a+b+c=-3, a+3 b+c=9, a+b+3 c=19$, then find $a b c$. Show that it is -56 .
Informal Proof (Human-written): Summing all three equations yields that $5 a+5 b+5 c=-3+9+$ $19=25$. Thus, $a+b+c=5$. Subtracting this from each of the given equations, we obtain that $2 a=$ $-8,2 b=4,2 c=14$. Thus, $a=-4, b=2, c=7$, and their product is $a b c=-4 \times 2 \times 7=-56$.

## Formal Statement and Proof:

```
theorem mathd_algebra_338:
    fixes a b c :: real
    assumes "3* a + b + c = -3" and "a + 3* b + c = 9" and "a + b + 3* c=19"
    shows "a * b * c = -56"
proof -
        (* Summing all three equations yields that 5a + 5b + 5c = -3 +9 + 19 = 25.
        Thus, a + b + c = 5. *)
        have "5* a + 5 * b + 5 * c = -3 + 9 + 19" using assms <ATP>
        then have "5 * (a+b + c) = 25" <ATP>
        then have "a+b+c=5" <ATP>
        (* Subtracting this from each of the given equations, we obtain that 2a =
        -8, 2b=4, 2c=14. Thus, a = 4, b = 2, c =7, and their product is abc =
        -4 \times 2\times 7 = -56. *)
        then have "2 * a = -8" "2 * b = 4" "2 * c = 14" using assms <ATP>
        then have "a = -4" "b = 2" "c = 7" <ATP>
        then show ?thesis <ATP>
qed
```


## Part II conclusion

- Recipe for specializing a language model to mathematics
- Llemma: 7B and 34B CodeLLama further trained on Proofpile II
- Open platform for research:
- Code/Models/Data: https://github.com/EleutherAI/math-lm


## Conclusion

Neural theorem proving: language models $\cap$ formal proof assistants

- Next-step predictors: small/fast, narrow
- Foundation models: larger, flexible
- Both enable new practical tools


## Conclusion

Neural theorem proving: language models $\cap$ formal proof assistants

- Next-step predictors: small/fast, narrow
- Foundation models: larger, flexible
- Both enable new practical tools

Claim: mathematical foundation models (e.g. LLEMMA) open a new frontier of methods, capabilities, and applications to explore!

## Thank you!

- Zhangir Azerbayev (Princeton, Eleuther)
- Hailey Schoelkopf (Eleuther)
- Keiran Paster (Toronto, Vector)
- Marco Dos Santos (Cambridge)
- Stephen McAleer (CMU)
- Albert Jiang (Cambridge)
- Jia Deng (Princeton)
- Stella Biderman (Eleuther)
- Sean Welleck (Washington, CMU)
- Tutorial: github.com/wellecks/ntptutorial
- LLMstep: github.com/wellecks/llmstep, arXiv:2310.18457
- Llemma: github.com/EleutherAI/math-lm, arXiv:2310.10631


## LLEMMA.

## Appendix

## Web: OpenWebMath



Figure 21: OpenWebMath pipeline. ${ }^{14}$

[^17]
## Web: OpenWebMath

```
This paper concerns the quantity
<img src="https://s0.wp.com/
latex.php?latex=%7BM%28x%29..."
alt="{M(x)}" />, defined as the
length of the longest
subsequence of the numbers from
```

Image Equations

```
Suppose I have a smooth map
[tex]f\colon \mathbb{R}^3
\longrightarrow S^2[/tex]. If I
identify [tex]\mathbb{R}^3[/tex]
with [tex]U_S = S^3 - \
{(0,0,1)\}[/tex] via
stereographic projection
```

<math>
<semantics>
<annotation ...>
\\displaystyle \mathrm (MA)
\(=\left\{\backslash f r a c\left\{f \_\{0\}\right\}\left\{f^{2}\{\right.\right.\) E \(\left.\left.\left.\}\right\}\right\}\right\}\) </annotation>
</semantics>
</math>
Delimited Math
Special Tags
Figure 22: Extraction: OpenWebMath extracts Latex from MathJax and 6 other sources of embedded Latex. ${ }^{16}$

[^18]
## Web: OpenWebMath

Classifier Input


Figure 23: Filtering: the MathScore classifier predicts whether a document contains a popular Latex command given the surrounding words. ${ }^{18}$

[^19]
## Scientific Papers

ArXiv papers (29 billion tokens)

- From RedPajama, an open-source replication of Llama data


## LLEMMA as initialization for further fine-tuning



Figure 24: LLEMMA vs. Llama 2 as initialization for finetuning on MetaMathQA

Not the focus of our work! A lot more to explore with fine-tuning.

## Analysis: overlap

## LLEMMA's open dataset allows for studying the effects of train/test

 overlap: ${ }^{19}$| Proof-Pile-2 | Test | ProblemExample Docs |  | Solution <br> Example Docs |  | Same solution <br> Different solution, same answer | 149 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  |  |  |
| OpenWebMath | MATH | 348 | 717 | 34 | 46 | Different solution, different answer | 9 |
| AlgebraicStack | MATH | 3 | 3 | 1 | 1 | No solution | 41 |
| OpenWebMath | GSM8k | 2 | 3 | 0 | 0 | Different problem | 0 |
| AlgebraicStack | GSM8k | 0 | 0 | 0 | 0 | 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.

[^20]
## Analysis: memorization

Surprisingly, Llemma did not perform any better on MATH problems that are contained in its training set:


Figure 25: LLEMMA-34b's accuracy on hits and non-hits by MATH level.

## References i

嘈
Z．Azerbayev，B．Piotrowski，H．Schoelkopf，E．W．Ayers，D．Radev， and J．Avigad．
Proofnet：Autoformalizing and formally proving undergraduate－level mathematics， 2023.
围 Z．Azerbayev，H．Schoelkopf，K．Paster，M．D．Santos，S．McAleer， A．Q．Jiang，J．Deng，S．R．Biderman，and S．Welleck．
Llemma：An open language model for mathematics．
ArXiv，abs／2310．10631， 2023.
圊 R．B．et al．
On the opportunities and risks of foundation models， 2022.

## References ii

嗇 D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang,
D. Song, and J. Steinhardt.

Measuring mathematical problem solving with the math dataset.
NeurIPS, 2021.

- J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai,
E. Rutherford, D. de Las Casas, L. A. Hendricks, J. Welbl, A. Clark,
T. Hennigan, E. Noland, K. Millican, G. van den Driessche,
B. Damoc, A. Guy, S. Osindero, K. Simonyan, E. Elsen, O. Vinyals, J. W. Rae, and L. Sifre.

Training Compute-Optimal Large Language Models.
In Advances in Neural Information Processing Systems, 2022.

## References iii

冨 D. Kocetkov, R. Li, L. Ben Allal, J. Li, C. Mou, C. Muñoz Ferrandis, Y. Jernite, M. Mitchell, S. Hughes, T. Wolf, D. Bahdanau, L. von Werra, and H. de Vries.
The stack: 3 tb of permissively licensed source code.
Preprint, 2022.
A. Lewkowycz, A. J. Andreassen, D. Dohan, E. Dyer, H. Michalewski, V. V. Ramasesh, A. Slone, C. Anil, I. Schlag, T. Gutman-Solo, Y. Wu, B. Neyshabur, G. Gur-Ari, and V. Misra.

Solving quantitative reasoning problems with language models. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, Advances in Neural Information Processing Systems, 2022.

## References iv

R. Muennighoff, A. M. Rush, B. Barak, T. L. Scao, A. Piktus, N. Tazi, S. Pyysalo, T. Wolf, and C. Raffel.

Scaling data-constrained language models.
arXiv preprint arXiv:2305.16264, 2023.
目
S. Welleck and R. Saha.

Llmstep: Llm proofstep suggestions in lean.
ArXiv, abs/2310.18457, 2023.

- K. Yang, A. Swope, A. Gu, R. Chalamala, P. Song, S. Yu, S. Godil,
R. Prenger, and A. Anandkumar.

LeanDojo: Theorem proving with retrieval-augmented language models.
In Neural Information Processing Systems (NeurIPS), 2023.

## References v

固
K. Zheng, J. M. Han, and S. Polu.
minif2f: a cross-system benchmark for formal olympiad-level mathematics.
In International Conference on Learning Representations, 2022.


[^0]:    ${ }^{1}$ See, e.g. Harvard New Technologies in Mathematics: 1/2021 (Christian Szegedy), 3/2022 (Stan Polu), 10/2022 (Guy Gur-Ari), 10/2023 (Alex Gu), 12/2023 (Katherine Collins), ...;

[^1]:    ${ }^{1}$ https://www.nature.com/articles/d41586-021-01627-2
    ${ }^{2}$ https://leanprover-community.github.io/teaching/courses.html
    ${ }^{3} h t t p s: / / s i t e s . m a t h . w a s h i n g t o n . e d u / \sim j a r o d / x l l . h t m l$

[^2]:    ${ }^{1}$ Llemma [2], LLMstep [9]

[^3]:    ${ }^{1}$ Llemma [2], LLMstep [9]

[^4]:    ${ }^{4}$ From A tutorial on neural theorem proving, IJCAI 2023

[^5]:    ${ }^{1}$ E.g., [Polu \& Sutskever 2020], [Han et al 2021], [Jiang et al 2022], [Yang et al 2023]

[^6]:    ${ }^{5}$ E.g., Lean Dojo [10] and github.com/semorrison/lean-training-data

[^7]:    ${ }^{5}$ E.g., Lean Dojo [10] and github.com/semorrison/lean-training-data

[^8]:    ${ }^{6}$ Example scoring function $\frac{1}{2} \sum_{t} \log p_{\theta}\left(y_{t} \mid x_{t}\right)$

[^9]:    ${ }^{7}$ tidy and GPT-4 (Lean 3) from [10]. Tutorial model: Ilmstep Pythia 2.8b, best-first-search with beam size 32 and a 10 minute timeout.

[^10]:    ${ }^{8}$ https://arxiv.org/abs/2310.18457. github.com/wellecks/Ilmstep.

[^11]:    ${ }^{9}$ Image from [Kaplan et al 2020]. See [5, 8] for more recent scaling laws.

[^12]:    ${ }^{9}$ Image from [Kaplan et al 2020]. See [5, 8] for more recent scaling laws.

[^13]:    ${ }^{10}$ Minerva [7]: general language model finetuned on a large mathematical corpus

[^14]:    ${ }^{11}$ Image from [Hernandez et al 2020] Scaling laws for transfer.

[^15]:    ${ }^{12}$ OpenWebMath: An Open Dataset of High-Quality Mathematical Web Text. Keiran Paster, Marco Dos Santos, Zhangir Azerbayev, Jimmy Ba

[^16]:    ${ }^{13}$ CodeLlama and Llemma use best-first-search with beam size 32 and a 10 minute timeout. GPT-4, COPRA and Reprover (no retrieval) from [Thakur et al 2023] (Lean 3)

[^17]:    ${ }^{14}$ Image from [Paster et al 2023]

[^18]:    ${ }^{8}$ Image from [Paster et al 2023]

[^19]:    ${ }^{9}$ Image from [Paster et al 2023]

[^20]:    ${ }^{19}$ Overlap tool at https://github.com/wellecks/overlap

