

NATURALPROVER: Grounded Mathematical Proof Generation with Language Models

Sean Welleck^{1,2*}, Jiacheng Liu^{1*}, Ximing Lu², Hannaneh Hajishirzi^{1,2}, Yejin Choi^{1,2}

¹Paul G. Allen School of Computer Science & Engineering, University of Washington

²Allen Institute for Artificial Intelligence, *Equal contribution
wellecks@uw.edu

Abstract

Theorem proving in natural mathematical language – the mixture of symbolic and natural language used by humans – plays a central role in mathematical advances and education, and tests aspects of reasoning that are core to intelligence. Yet it has remained underexplored with modern generative models. We study large-scale language models on two new generation tasks: suggesting the next step in a mathematical proof, and full proof generation. We develop NATURALPROVER, a language model that generates proofs by conditioning on background references (e.g. theorems and definitions that are either retrieved or human-provided), and optionally enforces their presence with constrained decoding. On theorems from the NATURALPROOFS benchmark, NATURALPROVER improves the quality of next-step suggestions and generated proofs over fine-tuned GPT-3, according to human evaluations from university-level mathematics students. NATURALPROVER is capable of proving some theorems that require short (2-6 step) proofs, and providing next-step suggestions that are rated as correct and useful over 40% of the time, which is to our knowledge the first demonstration of these capabilities using neural language models.¹

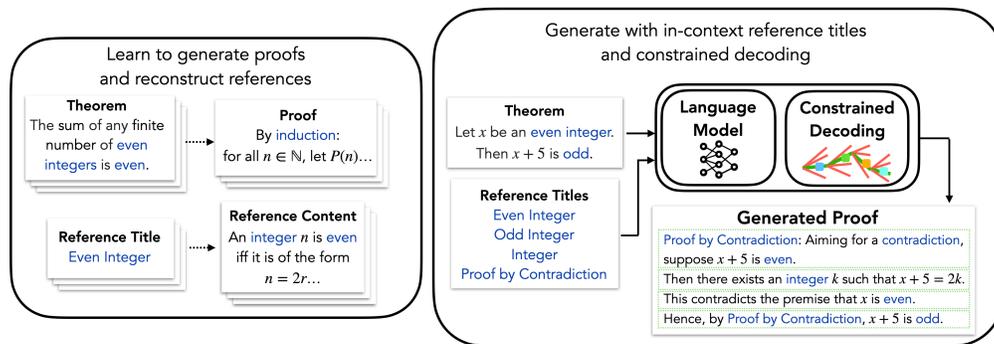


Figure 1: NATURALPROVER proves Even Integer Plus 5 is Odd. At training time, NATURALPROVER obtains background knowledge about references (e.g. theorems or definitions) via *reference reconstruction*: learning to map a reference’s title to its content. At test time, NATURALPROVER grounds its generations through in-context reference constraints that are retrieved or human-provided, and optionally enforced with *stepwise constrained decoding*. This theorem’s human-written proof in ProofWiki contains an error and differs substantially from NATURALPROVER’s correct proof.

¹Code and data available at <https://github.com/wellecks/naturalprover>.

1 Introduction

Constructing a rational argument that justifies a claim is a key aspect of explaining, verifying, and communicating ideas in situations ranging from everyday interactions, to legal and political discourse, to science and mathematics [10, 42, 24]. Within the latter context, a *mathematical proof* – a sequence of logical arguments expressed in a mixture of symbolic and natural language – assumes this role by providing justification and insight into why a claim is true [12]. Proofs operate on a relatively explicit and objective set of ground knowledge, isolating a subset of reasoning that is desirable for models that form the foundation of machine learning systems [3]. Moreover, we envision assistive systems that provide suggested proofs or next-steps, analogous to language-model-based code suggestions (e.g. GitHub CoPilot [6]) or formal proof assistants (e.g. GPT-*f* [20]), which could make learning or using mathematics more productive and accessible.

To this end, we study the capabilities of large-scale language models (e.g. GPT-3 [5]) on two new theorem proving tasks in natural mathematical language: *next-step suggestion*, in which a model suggests the next step of a proof, and *full-proof generation*, in which a model fully proves a claim. As proofs are grounded in knowledge from past results (e.g. theorems, definitions), analogous to facts deployed in a conversation [13], prior rulings used in a legal opinion [16], or articles used to justify an answer [30], we develop a methodology for obtaining and using background knowledge to prove theorems with a generic language model.

We develop NATURALPROVER, a language model that generates proofs by conditioning on background references (e.g. theorems and definitions that are either retrieved or human-provided), and optionally enforces their presence with a constrained decoding algorithm that leverages the multi-step structure of proofs. On a collection of theorems from the NATURALPROOFS benchmark [45], NATURALPROVER improves the quality of next-step suggestions and generated proofs over fine-tuned GPT-3, according to human evaluations from university-level mathematics students. NATURALPROVER is capable of proving some theorems that require short (2-6 step) proofs, and providing next-step suggestions that are rated as correct and useful more than 40% of the time, which is to our knowledge the first demonstration of these capabilities using neural language models.

Along with these successes, we study deficiencies in our current models. We find that models can struggle with logical coherence on longer proofs, with providing valid justifications, and with performing multi-step symbolic derivations. Taken together, our tasks, methodology, and evaluation show the feasibility of language models as interactive aids in mathematics, along with open challenges.

2 NATURALPROOFS-GEN Dataset and Tasks

We create a NATURALPROOFS-GEN dataset adapted from NATURALPROOFS [45], and use the dataset for two tasks: suggesting the next step of a proof, and fully proving a theorem.

NATURALPROOFS-GEN. NATURALPROOFS-GEN adapts data from NATURALPROOFS, which contains theorem statements, proofs, definitions, and additional pages (e.g. axioms, corollaries) sourced from ProofWiki, an online compendium of community-contributed mathematical proofs. In NATURALPROOFS-GEN, each example $(\mathbf{x}, \mathbf{y}) \in \mathcal{D}$ pairs a theorem \mathbf{x} with a gold proof \mathbf{y} , both of which are a mixture of text and \LaTeX . [45] split the examples and reference sets into training, dev, and test splits to ensure that no theorem in the dev or test splits was mentioned in the training split. We adopt these splits of roughly 12.5k training, 1k validation, and 1k test examples, and sampled *core evaluation sets* with 100 dev and 100 test theorems that are used for human evaluation. The proofs contain additional structure, discussed next.

Multi-step proof structure. Each proof has a *multi-step* structure, meaning that a proof $\mathbf{y} = (y_1, \dots, y_{|y|})$ is a variable-length token sequence that is segmented into *proof steps*, where each step y_t is itself a variable-length sequence of tokens (either text or LaTeX). The segmentation is largely determined by ProofWiki’s formatting and community standards for structuring proofs, and we additionally merge steps to ensure that each step contains non-trivial semantic content. For example, Figure 1 shows a 4-step (generated) proof with each step highlighted in green.

References. Each proof mentions a variable-number of *references* $\{\mathbf{r}_1, \dots, \mathbf{r}_{R_y}\}$ from a set \mathcal{R} of roughly 33k theorems and definitions, analogous to how Wikipedia articles reference other pages. For example, Figure 1 shows a proof with reference mentions in blue. Each mention identifies a

reference by its title and provides a natural language surface form. For instance, in Figure 1, the first proof step mentions the definition of even integer as `even`, which is formatted in the proof as `[[Definition:Even_Integer|even]]` and tokenized along with the rest of the proof.

Tasks. We consider two tasks that are motivated by an assistive system that provides suggested proofs or next-steps to a user. The **full proof generation** task is to generate a proof \mathbf{y} given a theorem \mathbf{x} . The **next-step suggestion** task is to generate a set of next steps $\{y_t^k\}_{k=1}^K$ given theorem \mathbf{x} and proof history $y_{<t}$ from a gold proof. In each case, we consider an additional **provided reference** setting where the model is also given the set of references $\{\mathbf{r}_1^*, \dots, \mathbf{r}_{R_y}^*\}$ from a gold proof of the theorem. The next-step task simulates a human correctly proving the theorem up to a point, then querying a system for suggested next-steps when stuck, while the provided reference setting simulates a human specifying a plan for a system that writes a proof.

3 NATURALPROVER: Grounded Proof Generation via Language Modeling

We describe NATURALPROVER, a language model which generates grounded proofs by conditioning on references and optionally enforcing their presence with constrained decoding.

Setup. Our objective is to generate correct proofs, $\hat{\mathbf{y}} = \arg \max_{\mathbf{y}} \text{correct}(\mathbf{x}, \mathbf{y})$. Unfortunately, evaluating proof correctness is costly, and is only done once at test time. A naive approach is to approximate the objective, $\hat{\mathbf{y}} \approx \arg \max_{\mathbf{y}} \log p_{\theta}(\mathbf{y}|\mathbf{x})$, by fine-tuning a language model p_{θ} on (\mathbf{x}, \mathbf{y}) examples and using a decoding algorithm (e.g. greedy decoding). We instead investigate conditioning on background knowledge in the form of reference documents, $p_{\theta}(\mathbf{y}|\mathbf{x}, R)$, which is beneficial in related generation settings (e.g. [38]), and offers control over the generated proof. To do so, NATURALPROVER uses in-context references and a reference reconstruction objective.

In-context references. Language models have a limited context window that prevents conditioning on full documents. Instead, NATURALPROVER conditions on a set of reference titles, $p_{\theta}(\mathbf{y}|\mathbf{x}, R_{\text{title}})$. Concretely, we fine-tune on (theorem, reference titles, proof) sequences of the form,

$$\begin{aligned} &\langle \text{theorem} \rangle \langle \text{title} \rangle \{ \text{theorem-title} \} \langle / \text{title} \rangle \langle \text{content} \rangle \{ \text{theorem-content} \} \langle / \text{content} \rangle \langle / \text{theorem} \rangle \\ &\langle \text{ref} \rangle \{ \text{ref-title-1} \} \langle / \text{ref} \rangle \dots \langle \text{ref} \rangle \{ \text{ref-title-R} \} \langle / \text{ref} \rangle \langle \text{proof} \rangle \{ \text{proof} \} \langle / \text{proof} \rangle \end{aligned} \quad (1)$$

with new-lines and `{}` tokens omitted, relevant strings inserted, and loss only on tokens after `<proof>`.

Reference reconstruction. Reference titles do not capture all of the information contained in the reference documents. We learn a mapping between each reference title and its underlying document with a reference reconstruction objective, $p_{\theta}(\mathbf{r}|\mathbf{r}_{\text{title}})$ for references \mathbf{r} in the training reference set. Concretely, we fine-tune on additional (title, content) pairs of the form,

$$\langle \{ \text{type} \} \rangle \langle \text{title} \rangle \{ \text{title} \} \langle / \text{title} \rangle \langle \text{content} \rangle \{ \text{content} \} \langle / \text{content} \rangle \langle / \{ \text{type} \} \rangle, \quad (2)$$

where the `{type}` is theorem/definition/other, and the loss is only on tokens after `<content>`. Intuitively, this lets the model associate each reference title with the reference’s underlying content.

The joint objective. For training, we minimize the joint loss,

$$\mathcal{L}(\theta) = \frac{1}{|\mathcal{D}^{\text{train}}| + |\mathcal{R}^{\text{train}}|} \left[\sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}^{\text{train}}} -\log p_{\theta}(\mathbf{y}|\mathbf{x}, R_{\text{title}}) + \sum_{\mathbf{r} \in \mathcal{R}^{\text{train}}} -\log p_{\theta}(\mathbf{r}|\mathbf{r}_{\text{title}}) \right]. \quad (3)$$

Evaluation-time references. We consider two settings for evaluation-time references: (i) *retrieved* references, from a retrieval model $f(\mathbf{x}) \rightarrow \{\mathbf{r}_1, \dots, \mathbf{r}_k\}$, and (ii) *human-provided* references from the ground-truth proof. The retrieval setting simulates a fully automated proof assistant, while the second simulates a human specifying a plan for an assistant that writes a proof, and acts as an upper bound for a retrieval system optimized to predict references in a ground-truth proof.

3.1 Stepwise constrained decoding

In the provided-reference setting, the conditioned references are known to be relevant to a correct proof. We hypothesize that explicitly encouraging generated proofs to contain the references will improve correctness, by placing lexical constraints on the reference-titles at decoding time,

$$\hat{\mathbf{y}} \approx \arg \max_{\mathbf{y}} \log p_{\theta}(\mathbf{y}|\mathbf{x}, R_{\text{title}}), \text{ subject to } \sum_{\mathbf{r}_{\text{title}} \in R_{\text{title}}} \mathbb{I}[\mathbf{r}_{\text{title}} \in \mathbf{y}] = |R_{\text{title}}|, \quad (4)$$

where $\mathbb{I}[\cdot]$ is an indicator function. To approximate this objective, we generate step-by-step by sampling multiple proof-step candidates, retaining those with high value (reference coverage and log-probability) in a beam, and continuing to the next step, which we call stepwise beam search.

Value function. The search supports any function of the proof-so-far, $v(y_{\leq t}) \rightarrow \mathbb{R}$. We use a value function that is a weighted combination of constraint satisfaction and log-probability,

$$v_{\alpha}(y_{\leq t}) = \alpha v_{\text{constraint}}(y_{\leq t}) + (1 - \alpha) v_{\text{LM}}(y_{\leq t}), \quad (5)$$

where $v_{\text{constraint}}(y_{\leq t})$ is the number of unique in-context reference-titles in $y_{\leq t}$, and $v_{\text{LM}}(y_{\leq t})$ is $\log p_{\theta}(y_{\leq t})$. We normalize each term by dividing by the maximum absolute value among candidates.

Stepwise beam search. The procedure generates a proof $\mathbf{y} = (y_1, \dots, y_T)$ by iteratively sampling and pruning next-proof-step candidates y_t . Each iteration expands a size- K beam of proofs-so-far, $S_{t-1} = \{y_{<t}^k\}_{k=1}^K$, by generating N next-step candidates,

$$S'_t = \cup_{y_{<t} \in S_{t-1}} \{(y_{<t} \circ y_t^n) \mid y_t^n \sim q(\cdot \mid y_{<t}, \mathbf{x}, R_{\text{title}})\}_{n=1}^N, \quad (6)$$

where q is a decoding algorithm (e.g. temperature sampling) and \circ is concatenation. The next iteration’s beam is formed by selecting the top scoring candidates, $S_t = \arg \text{top-K}_{y_{\leq t} \in S'_t} v_{\alpha}(y_{\leq t})$. When a proof in the beam terminates, it is not expanded further. The search ends when the beam consists of K terminated proofs. The highest value proof is returned as the final output.

Stepwise++. We add two mechanisms for promoting exploration at each step. First, we expand each prefix in the beam (Eqn. 6) by sampling with multiple temperatures, $\{y_t^n \sim q_{\tau}(\cdot \mid y_{<t}, \mathbf{x}, R_{\text{title}}) \mid \tau \in \{\tau_i\}_{i=1}^m\}$, where q_{τ} is sampling with temperature τ . This relaxes the commitment to a single temperature for all proof steps, balancing exploration (higher τ) with exploitation (lower τ).

Second, rather than selecting the top- K candidates, we select clusters based on different value weights: $S_t = \cup_{\alpha \in \{\alpha_j\}_{j=1}^{\ell}} \text{top}_{K'}(S_t^{\alpha})$, where S_t^{α} is the set of candidates scored with v_{α} , and $K' = K/\ell$. This interpolates between selecting steps based on likelihood (low α) and constraint satisfaction (high α).

Full proof sampling and greedy decoding. An alternative is to sample full proofs and select the best one according to the value function. This can be viewed as expansion (Eqn. 6) done at the full proof, rather than the step level. Moreover, greedy decoding corresponds to expanding only 1 candidate with temperature $\rightarrow 0$. We formalize this in §D as a segment-level search that contains stepwise++, full proof sampling, and greedy decoding as special cases.

4 Proof Evaluation

A proof’s correctness is contingent on a variety of factors, including reasoning with past results, performing symbolic derivations, and altogether providing sufficient evidence that the claim is true. We design a human-evaluation schema that isolates these aspects at the proof-step level, along with a full-proof summary. Table 1 summarizes the schema, which we overview below.

References. First, proofs involve deploying statements from references, such as applying a definition or adapting it to fit the context. Deployments should be consistent with the reference, e.g. deploying the definition of even integer as ‘...by definition, $\exists k \in \mathbb{Z} : x = 2k$...’, rather than ‘... $\exists k \in \mathbb{Z} : x = 2k + 1$ ’, and are a common source of errors in student proofs [15].

Second, proofs use references as justification for steps of reasoning; for instance, Real Addition is Commutative provides justification for the statement $x + y = y + x$ where $x, y \in \mathbb{R}$, but not for $xy = yx$. This aspect is analogous to using an article to justify a claim (e.g. [30]). Finally, proofs should not hallucinate references, or ‘beg the question’ by self-referencing the current theorem.

Equations. Proofs contain a variety of multi-step derivations, ranging from simple arithmetic to more sophisticated derivations (e.g. see Table 17). A derivation should start with a valid equation given the surrounding context (e.g. $x + x = 2x$ in Table 1 versus $x + x = 3x$). Each subsequent step should be a valid derivation from the previous step, e.g. stating $= (2k + 6) - 1$ after $y = 2k + 5$.

Other reasoning, language, & symbolic errors. A proof should provide sufficient evidence that a claim is true to a human reader; it should not skip steps. Proof steps should make progress towards proving the goal; in particular, they should not repeat known conditions in the theorem or conclusions made in a prior step. Finally, our schema leaves room for any other reasoning errors, as well as symbol errors (e.g. undefined symbols) and language errors (e.g. incomplete statements).

Error Type	Example
Reasoning: Reference	
Invalid Deployment	Since x is an even integer, $\exists k \in \mathbb{Z} : x = 2k + 1$.
Invalid Justification	$\mathbb{E}(X^2) = \sum_{k=1}^n k^2 \Pr(X = k)$ <u>Power Series for Exponential Function</u>
Hallucinated Ref.	From <u>Power of Number are Irrational</u> , $\sqrt[3]{2}$ is irrational.
Self Loop	(Proving Pythagoras’s Theorem:) From <u>Pythagoras’s Theorem</u> , $c^2 = a^2 + b^2$.
Reasoning: Equation	
Invalid Equation	$\forall x \in \mathbb{R}, x + x = 3x$.
Invalid Derivation	(Since x is an even integer, $x + 1 = 2r + 1$) $= 2(r + 1)$
Reasoning: Other	
Skips Steps	($x \in \mathbb{Z}$ is not a multiple of 3.) Therefore, $x^3 \equiv 1$ or $8(\text{mod } 9)$
Repetition	(Let $\triangle ABC$ be a right triangle.) Then $\triangle ABC$ is a right triangle.
Invalid (Other)	(x is an even integer.) So, $x + 1$ is an even integer.
Language	Let $c = \sqrt{a^2 \text{add } b^2}$ be the (<i>incomplete statement ; unknown symbol \add</i>)
Symbolic	(Let $x \in \mathbb{R}$.) Let $y = x \circ x^{-1}$. (<i>undefined operator \circ for real numbers</i>)

Table 1: Overview of human evaluation error schema. See Table 24 for full schema. Reference. Hallucinated reference. The necessary context (e.g. known conditions, prior steps).

Usefulness and correctness. To judge the potential utility of language models as assistive systems in natural mathematics, we measure whether generated next-steps and full proofs are potentially useful hints for proving the theorem on one’s own. Additionally, we measure a summary judgment of correctness. Note that an incorrect statement can still be helpful; for instance, it could give a hint for the type of reference to use, derivation to perform, argument to make, etc.

Human evaluation protocol. We measure these aspects through human annotation at a *step-wise* and an *overall* level. For a step-wise annotation, an annotator is presented with the theorem, proof-so-far, and a generated next-step. The annotator labels the $\{0, 1\}$ correctness, usefulness, and presence of fine-grained errors outlined above. After labeling each step of a proof, the annotator rates the full proof’s overall correctness and usefulness on a 0-5 scale. A rating of 4 or 5 is needed to be considered as correct, and a rating of 3 or above is needed to be considered as useful.

Automatic metrics: lexical content. As automatic proxies for quality, we compare each generated proof against its ground-truth counterpart using the sentence-level n -gram matching metric GLEU [29], and following work in knowledge-grounded dialogue [38] we use F1 overlap between generated and ground-truth tokens. Prior to computing the metrics, we normalize the generated and ground-truth proofs by only keeping the surface form of references, removing formatting characters with a MediaWiki parser, and collapsing any consecutive whitespace into a single space.

Automatic metrics: knowledge grounding. We define knowledge grounding as meaning that a generated proof contains the same references as those found in the ground-truth proof. To measure this, we use precision, recall, and F1-score between the reference sets contained in the generated and ground-truth proofs; i.e. $m(\{\hat{r}_1, \dots, \hat{r}_R\}, \{r_1^*, \dots, r_{R^*}^*\})$, where $m(\cdot)$ is precision, recall, or F1. We also use Knowledge Token-F1 (kF1) ([38]), the overlap of the generated proof’s tokens with tokens contained in the references mentioned in the ground-truth proof.

5 Experiments

We use the training and dev splits of NATURALPROOFS-GEN during fine-tuning, and the *core evaluation sets* consisting of 100 theorems from the validation set and 100 from the test set for evaluation (see §2). These theorems were selected by the authors such that by looking at the theorem title each author could recall its content and sketch a proof. While this may shift the evaluation towards an easier slice of the dataset, it was necessary to make human evaluation at a meaningful scale feasible. We also use the core sets for explorations and ablations.

We finetune three GPT-3 [5] (Curie) models, using the OpenAI API (see Appendix E for details):

	Reasoning Errs (↓)			Lexical Errs (↓)		Per-Step (↑)		Full Proof (↑)	
	Ref.	Eqn.	Other	Lang.	Sym.	Useful	Correct	Useful	Correct
GPT-3	30.92	32.54	40.15	5.61	5.24	25.69	28.18	20%	13%
NATURALPROVER _{RETRIEVE}	23.52	37.55	23.66	4.54	6.19	41.54	33.56	32%	24%
NATURALPROVER	25.84	35.93	25.23	8.41	5.35	39.60	26.30	35%	24%
NATURALPROVER ₊₊	23.61	28.54	18.45	5.58	3.65	46.57	35.41	45%	32%
Next-step (NATURALPROVER)	19.70	26.32	19.10	8.57	5.86	51.43	42.86	–	–

Table 2: Human evaluation results on the core test set for full proof generation and next-step suggestion (bottom row). All models are fine-tuned on NATURALPROOFS-GEN. Knowledge – either retrieved or human provided – and constrained decoding improve proof generation, with 46% of proof steps rated as useful and 35% correct according to university-level mathematics students.

1. **Baseline GPT-3.** We finetune a baseline GPT-3 model, $p_\theta(\mathbf{y}|\mathbf{x})$, on theorem-proof examples $\{(\mathbf{x}, \mathbf{y})\}$ from the training split. At test time, we condition the model on a test theorem.
2. **NATURALPROVER_{RETRIEVE}.** We finetune GPT-3 with retrieved references, $p_\theta(\mathbf{y}|\mathbf{x}, \hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_{20})$. We use a pretrained joint retrieval model $f(\mathbf{x}) \rightarrow (\mathbf{r}_1, \dots, \mathbf{r}_{|\mathcal{R}|})$ from [45], which was trained to retrieve an input theorem’s ground truth references. At test time, the model receives a theorem and the top-20 reference titles that are retrieved given the theorem.
3. **NATURALPROVER.** We finetune GPT-3 with human-provided references, $p_\theta(\mathbf{y}|\mathbf{x}, \mathbf{r}_1^*, \dots, \mathbf{r}_{R_y}^*)$, where $\{\mathbf{r}_1^*, \dots, \mathbf{r}_{R_y}^*\}$ is the set of reference-titles in the ground-truth proof. We use reference-title conditioned examples (Eqn. 1) and reference-reconstruction (Eqn. 2) on the training split/reference set. At test time, the model receives a theorem and reference titles from its ground-truth proof.

For **next-step suggestion** we use the human-provided knowledge model (NATURALPROVER).

Decoding. For full proof generation, we use stepwise++ decoding with the provided knowledge model, which we refer to as **NATURALPROVER₊₊**, and otherwise use greedy decoding. We do not use stepwise constrained decoding with retrieved references since these references introduce noisy constraints, nor for next-step prediction since the algorithm is designed for multi-step proofs. See §E for additional experimental details.

Human evaluation setup. To evaluate the proofs generated by NATURALPROVER, we recruited 15 students from the Department of Mathematics and Applied Mathematics at the University of Washington, including undergraduate, masters, and Ph.D. students. The annotators were trained on how to evaluate proof correctness and compensated according to IRB requirements; see §F.2. For each task, we first reveal the theorem and its gold proof to the annotator. If they cannot understand a theorem or its gold proof, they may skip evaluating it. Otherwise, they may proceed to see the model-generated proof, one step at a time, and annotate each step under the step-wise evaluation schema (outlined in §4). After all the steps are shown and evaluated, for the full-proof generation task, the annotator is asked to annotate the entire proof under the overall evaluation schema.

5.1 Main Results

Our best method is capable of generating correct and useful proofs. According to human evaluation results (Table 2), our best method is NATURALPROVER with human-provided references and stepwise++ inference. 32% of the proofs generated by this method are rated as **correct** and 45% are rated as **useful** as an aid for human proof writers. On the per-step level, 35% of the proof steps are **correct** and 47% are **useful**. Taking a more granular view, our best method makes significantly less **reference**, **equation**, and **other reasoning errors** than other baselines. It makes very few **language** and **symbolic errors**, meaning it produces mostly complete, well-formatted mathematical statements, and defines and uses symbols accordingly. It mostly avoids hallucinating references or creating circular proofs. On the other hand, despite improving over the other methods, the model often struggles with correctly deploying and using references (23.6% reference error rate), as well as symbolic computations (28.5% equation error rate), especially multi-step derivations (21.9% invalid).

What do the model’s correct proofs look like? We inspected the proofs labeled as correct and found three main categories: (1) *reference-assembly* proofs whose correctness is heavily determined by reference statements (e.g. Table 18, Table 20); (2) *template-adaptation* proofs in which the model adapts the structure and content of a training theorem’s proof to prove the unseen evaluation theorem

Theorem	Singleton Set is not Dense-in-itself	
	Let $T = (S, \tau)$ be a topological space. Let $x \in S$. Then the singleton set $\{x\}$ is not dense-in-itself.	
GPT-3	Let $x \in S$. From Clopen Points in Topological Space we have that $\{x\}$ is open in T . From Point is Open iff it is not Dense-in-itself, it follows that $\{x\}$ is not dense-in-itself.	✓ ✗ ✗
NATURALPROVER ₊₊	From Singleton Point is Isolated, $\{x\}$ has an isolated point. Hence the result by definition of dense-in-itself.	✓ ✓

Singleton Point is Isolated
Let $T = (S, \tau)$ be a topological space.
Let $x \in S$.
Then x is an isolated point of the singleton set $\{x\}$, but not necessarily an isolated point of T .

Dense-in-itself
Let $T = S, \tau$ be a topological space.
Let $H \subseteq S$.
Then H is dense-in-itself iff it contains no isolated points.

Table 3: GPT-3 hallucinates references, while the knowledge-grounded NATURALPROVER₊₊ with constrained decoding correctly uses references, resulting in a correct and useful proof.

(e.g. Table 21, Table 22); (3) *complex* proofs that are not fully determined by reference statements and differ significantly from training proofs (e.g. Figure 1, Table 3). In terms of techniques, our method demonstrates some ability to produce direct proofs (Table 19), proofs by cases (Table 22), proofs by induction (Table 23), utilize references (Table 20) and do symbolic computations (Table 21).

Vanilla fine-tuned GPT-3 struggles with proof generation. The vanilla fine-tuned GPT-3 model yielded fewer useful and correct proofs, with more reference-based and other reasoning errors than all three knowledge-grounded settings. The model showed severe reference hallucination (18%) and repetition (23%). It also makes significantly more reasoning errors related to reference usage. Language and symbolic error rates roughly stay the same. Overall, naively fine-tuning GPT-3 on theorem-proof examples alone is suboptimal for proof generation.

Human-provided knowledge improves proof generation. Grounding the generations with human-provided references significantly raises correctness and usefulness of the proofs in both full-proof and per-step evaluation. It most substantially reduces reference errors, especially invalid deployments and hallucinated references. For example, Table 3 shows the model grounding a proof with information from the theorem Singleton Point is Isolated and the definition of Dense-in-itself, in contrast to the vanilla GPT-3 model which hallucinates references.

Retrieved knowledge also improves proof generation. Retrieved knowledge also turns out to be very helpful, and even comparable to human-provided knowledge in some metrics. Although the retrieval model is far from perfect, the proof generation model is capable of narrowing down the retrieved reference titles provided in its context, assembling proofs that are useful and correct more often than the no-knowledge model. Qualitatively, we found examples where grounding in retrieved references eliminates repetition, enables multi-step derivations justified by references (Table 21), and assembles references into a correct proof (Table 20). This paves a promising path towards fully automated mathematical proof generation in natural mathematical language.

Constrained decoding further improves proof generation. Table 4 confirms that stepwise₊₊ decoding approximates the constrained objective (Eqn. 4) better than greedy search, yielding proofs with lower perplexity and higher constraint satisfaction (Ref-F1). This translates to generations that are correct and useful more often according to the annotators. Intuitively, the constraints encourage the model to include references that help prove the claim (e.g. Table 18).

Next-step suggestion. The next-step suggestion task characterizes a model’s performance on making a single proof step given a correct proof-so-far. In Table 2 we use the provided-knowledge model with greedy decoding for next-step suggestion, and find that reasoning errors decrease and per-step usefulness and correctness improve compared to the full proof setting, with 51% of the proof steps rated as useful and 43% correct. Although we used a single suggestion in our human evaluation study, in Table 5 we simulate a user choosing from among multiple suggestions by sampling 10 next-steps from our model and computing automatic metrics on the sample with the best sum of metrics. Us-

In-context	Stepwise ₊₊	PPL (↓)	Ref-F1 (↑)
✗	✗	1.0639	26.33
✗	✓	1.0549	30.07
✓	✗	1.0644	89.43
✓	✓	1.0549	94.25

Table 4: Stepwise₊₊ decoding approximates the constrained objective better than greedy decoding, resulting in both lower perplexity and better reference coverage, regardless of whether knowledge is provided in-context.

		Lexical		Grounding				
		GLEU	Token F1	kF1	Ref-P	Ref-R	Ref-F1	Halluc (\downarrow)
GPT-3		24.40	49.96	49.30	29.93	24.73	23.69	17.92
NATURALPROVER _{RETRIEVE}		26.58	53.02	55.88	38.17	28.48	27.10	2.25
NATURALPROVER		35.27	66.00	90.07	93.05	86.05	87.08	1.60
NATURALPROVER ₊₊		34.49	65.61	96.39	94.66	95.00	93.92	1.71
human metric	Correctness [full]	0.93	0.91	0.86	0.83	0.85	0.85	0.94
	Usefulness [full]	0.90	0.87	0.82	0.78	0.80	0.80	0.97
	Correctness [step]	0.81	0.80	0.74	0.69	0.73	0.72	0.97
	Usefulness [step]	0.65	0.61	0.53	0.47	0.52	0.51	0.98
	Reasoning Errors: Ref.	0.71	0.64	0.52	0.48	0.50	0.50	0.95
	Reasoning Errors: Eqn.	0.70	0.74	0.75	0.69	0.74	0.73	0.78
	Reasoning Errors: Other	0.65	0.61	0.53	0.47	0.52	0.51	0.98
	Language Errors	0.99	1.00	0.99	0.98	0.99	0.99	0.73
	Symbolic Errors	-0.72	-0.80	-0.88	-0.89	-0.89	-0.88	-0.21

Table 6: Automatic metrics on the core test set for full-proof generation, and correlation between human metrics and automatic metrics on the core validation set.

ing 10 samples instead of greedily decoding a single sequence substantially improves each metric, suggesting that utility might be increased further by presenting multiple suggestions.

How good are Automatic Metrics? We study how well the automatic lexical and grounding metrics introduced in (§4) can reflect the real quality of proofs, as a guide for using them as a proxy evaluation protocol for NATURALPROOFS-GEN. We compute the Pearson correlation coefficient between each pair of human and automatic metrics, with data from the four experiment settings for full-proof generation. Results are shown in the lower part of Table 6, with error metrics negated, meaning positive correlation is desired.

The lexical and grounding metrics positively correlate with full proof correctness and usefulness (≥ 0.8). At the step-level, the metrics show (i) high correlation with step-level correctness and language errors ; (ii) varied, but positive, correlations with aggregate reasoning errors; (iii) negative correlation with symbolic errors (though symbolic errors are relatively low for all models). The results suggest that optimizing for automatic metrics may be a viable strategy, albeit without guarantees on how finer-grained reasoning aspects vary across proofs.

5.2 Ablations and error analysis.

Reference reconstruction. We fine-tune an additional GPT-3 model that is provided with in-context reference titles, but without reference reconstruction. As seen in Table 7, reference reconstruction improves content and reference usage.

Constrained decoding. First, Table 9 compares the step-level search in stepwise++ with searching at the full-proof level through sampling multiple proofs and selecting the best with the NATURALPROVER value function ($rerank(n)$). Reranking 60 samples matches the cost of stepwise++ in terms of number of decoded tokens. Full-proof reranking yields the best Gleu, though with lower reference-F1. Second, Table 8 shows that the expansion and selection mechanisms together result in the best reference matching, while holding Gleu at a similar level. Finally, Table 13 shows that both terms in the NATURALPROVER value function $\alpha v_{constraints} + (1 - \alpha)v_{LM}$ are needed: increasing the constraint weight α increases reference-matching, with a tradeoff in Gleu at high values.

Language model comparison. Table 10 varies the language model used to parameterize NATURALPROVER. The content and reference usage metrics improve with larger models. Separately, we find that increasing inference-time compute closes the gap in reference-matching between GPT-2 and the larger GPT-3 model (Table 11): sampling 10 full-proofs from GPT-2 and selecting the best

	Decoding	GLEU	Ref-F1
Greedy		47.87	65.50
Temp (t=.6)		60.60	84.44
Temp (t=.8)		61.89	86.74
Temp (t=1.0)		62.12	86.87

Table 5: *Next-step suggestion:* Sampling 10 suggestions improves over a single greedy suggestion.

	Recon.	Gleu	Ref-F1	Halluc.
\times		33.03	82.85	3.32
\checkmark		35.93	84.15	2.68

Table 7: Effect of reference reconstruction in NATURALPROVER (greedy decoding, full validation set).

Expand	Select	GLEU	Ref-F1	Decoding	Gleu	Ref-F1
✗	✗	40.62 (.84)	91.78 (.49)	Greedy	41.12 (-)	89.30 (-)
✓	✗	41.12 (.58)	92.61 (.63)	Rerank (10)	43.88 (.29)	91.72 (.28)
✗	✓	39.14 (.55)	93.11 (.34)	Rerank (60)	42.23 (.80)	93.16 (.27)
✓	✓	40.11 (1.55)	94.13 (.45)	Stepwise++	40.11 (1.55)	94.13 (.45)

Table 8: Ablation of the stepwise++ expansion Table 9: Stepwise versus full-proof search. Mean and selection mechanisms. Mean (std) over 3 (std) over 3 runs on the core dev set. runs shown on the core dev set.

using the NATURALPROVER value function achieves the same reference-F1 as GPT-3 with a single greedily-decoded proof. However, Gleu remains much higher with the larger GPT-3 model.

Challenge: Reasoning with references. Although reference reasoning errors were decreased through knowledge-grounding and constrained decoding, NATURALPROVER still commits a reference error on 23.6% of test steps (27% dev), with 15% of steps containing invalid deployments and 10% invalid justifications. For next-step prediction, the reference error rate remains nontrivial (19.7% test, 13% dev), meaning that the model can struggle to correctly deploy references or use them as justification even in the absence of compounding errors from previous steps. Table 15 shows example invalid deployments and justifications; the errors are at times subtle, and require reasoning about the theorem statement, reference content, and proof context.

Challenge: Equations and derivations. NATURALPROVER commits an equation-related error on 28.5% of test steps (22.8% dev), including invalid equations (9.4%) and derivations (21.9%). Though an improvement over vanilla fine-tuned GPT-3 (32.5%), the errors occur frequently and remain high for next-step prediction (26%). Table 17 shows representative errors, which range from simple ‘commonsense’ mistakes (e.g. $24 = 2^3$) to making invalid steps with false justification within more sophisticated multi-step proofs. Investigating the role of pretraining, in-context techniques [31], and autoformalization [39] is interesting future work.

Challenge: Proof length. Although NATURALPROVER demonstrates some ability to write long proofs (e.g. Table 23), the 42% next-step correctness suggests that compounding errors are likely as proof length increases. Indeed, our best model’s full-proof correctness is 48% on 1-4 step proofs ($n = 102$), decreasing to 15.6% on proofs with 5 or more steps ($n = 64$), with lower per-step usefulness and correctness at later steps (Figure 2). Our findings are analogous to recent work on language modeling for formal theorem proving [32], where current models are typically limited to chaining 2 or 3 non-trivial steps of mathematical reasoning.

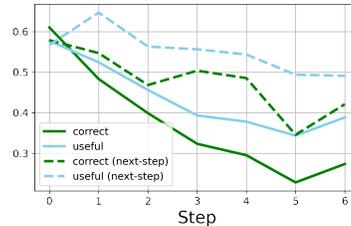


Figure 2: Per-step correctness and usefulness as a function of step number, for full-proof generation with NATURALPROVER++ and next-step prediction with NATURALPROVER.

5.3 Additional discussion

Finally, we provide higher-level comments on future work related to interactive systems, mathematical assistants, and generating proofs in informal versus formal mathematics.

Interactive & improving systems. Currently, our tasks are at two ends of a spectrum: in next-step generation, we always assume previous steps are from a human-written proof, while in full proof generation they are always from the model. Our results with multiple next-step suggestions suggest that users might find *some* suggestion among the multiple returned useful at a high rate, pointing to a middle ground: a human-in-the-loop NATURALPROVER, in which a human picks the next step from among the returned suggestions, or writes one based on the suggestions. The selected or written next-step could then be used as feedback to improve the system, enabling an iteratively improving NATURALPROVER. This notion of a continuously improving, teachable system is an emerging (e.g. [9]) and interesting future direction.

Assistants for mathematics. Our tasks were motivated by an assistant that helps a user write a proof, either from scratch or when stuck part of the way through. Our study here focuses on *capability*:

investigating whether neural language models are capable of performing the underlying mathematics that would be expected from such an assistant. A further challenge is to also ensure *reliability* – a user should have confidence that the model is not deceptive or incorrect, and is robust to changes in domain, on nearby problems, and on alternative ways of expressing a problem. Even further, we would like *flexibility* – human teachers can interact with a student flexibly through dialogue, natural language, and diagrams, rather than the strict input-output format defined by a dataset. Our work provides an initial step towards this larger vision.

Informal and formalized mathematics. Our work investigates theorem proving entirely in natural mathematical language (i.e. ‘informal’ mathematics), as it reflects an interface that a student typically uses when working with mathematics. An alternative is proving theorems in a formalized system, in which proof steps are expressed in a programming language (e.g. Lean [11]). Operating purely in a formalized system allows for verifying correctness – unlike our setting which must be verified by a human – arguably at the cost of flexibility and interpretability, as the mathematics is no longer expressed in natural language and must adhere to constraints of the formal system. Investigating combinations of the two – e.g. expressing a theorem in natural language, receiving a verified formal proof, then providing an interpretation in natural language – presents a wide range of interesting directions for future work.

6 Related Work

Formalized mathematics with neural language models. A large portion of work on machine learning for mathematics focuses on formalized mathematics. Language models have been used for interactive theorem proving, including in GPT-*f* [33, 32], PACT [20], and in [41]. In these settings proof steps are expressed in a programming language (e.g. Lean [11]) and there is access to a verifier, which differs from our setting of theorem proving in natural mathematical language.

Informal mathematics with neural language models. Previous work on theorem proving in natural mathematical language focuses on retrieving relevant premises (e.g. theorems, definitions) [17, 18, 45, 21], or informal-to-formal translation [43], which differ from our setting of generating next-steps or full proofs. Outside of theorem proving, various works use sequence models for problem solving, including benchmarking language models on arithmetic [37] or competition problems [22], symbolic mathematics [25, 46], augmenting LMs with verifiers [7] or in-context rationales [44] for math word problems, or using language models for math-related program synthesis [2, 14] and competitive programming [26]. These settings focus on generating executable programs or a numerical answer, which differ from our theorem proving setting, where the goal is to generate sound and convincing arguments on a range of topics in natural mathematical language.

Related areas in NLP. Systematic reasoning in natural language (outside of math) has been studied with synthetic proofs [36, 40], single-step deductions [4], or entailment trees [8], which differ from proving real-world mathematical theorems. Augmenting LMs with knowledge reduces hallucinations in dialogue [38] which has an analogous step-wise structure, while [30] use references within long-form answers; these and related NLP findings differ from improving the utility of mathematical proofs. Lexically-constrained decoding algorithms include variants of (token-level) beam search (e.g. [1, 23, 28, 27]) which assume access to per-token logits, and gradient-based decoding [34]; our segment-level decoding only assumes a sampler that returns text and its log-probability, making it compatible with recent language model API interfaces (e.g. the GPT-3 API).

7 Conclusion

We described NATURALPROVER, a knowledge-grounded language model that generates mathematical proofs by conditioning on background theorems and definitions, and optionally enforces their presence with constrained decoding. Our system improves the quality of next-step suggestions and generated proofs over fine-tuned GPT-3, demonstrating an ability to correctly prove theorems and provide useful suggestions to human proof writers.

Acknowledgments and Disclosure of Funding

This work was funded in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) (funding reference number 401233309), DARPA MCS program through NIWC Pacific (N66001-19-2-4031), and the Allen Institute for AI. We also thank Google Cloud Compute, as well as OpenAI.

The authors would like to thank Alisa Liu, Julian Michael, Yuren (Rock) Pang, and Kaiming Cheng for dogfooding and providing valuable feedback to our human evaluation system. We would also like to thank James McGivern for developing an interactive demo for NaturalProver.

References

- [1] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Guided open vocabulary image captioning with constrained beam search. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 936–945, Copenhagen, Denmark, September 2017. Association for Computational Linguistics.
- [2] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. Program synthesis with large language models, 2021.
- [3] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, S. Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen Creel, Jared Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren E. Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas F. Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, O. Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir P. Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avani Narayan, Deepak Narayanan, Benjamin Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, J. F. Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Robert Reich, Hongyu Ren, Frieda Rong, Yusuf H. Roohani, Camilo Ruiz, Jack Ryan, Christopher R’e, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishna Parasuram Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei A. Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models. *ArXiv*, abs/2108.07258, 2021.
- [4] Kaj Bostrom, Xinyu Zhao, Swarat Chaudhuri, and Greg Durrett. Flexible generation of natural language deductions. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6266–6278, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [5] 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.
- [6] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

- [7] 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.
- [8] Bhavana Dalvi, Peter Jansen, Oyvind Tafjord, Zhengnan Xie, Hannah Smith, Leighanna Pipatanangkura, and Peter Clark. Explaining answers with entailment trees. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7358–7370, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [9] Bhavana Dalvi, Oyvind Tafjord, and Peter Clark. Towards teachable reasoning systems. *ArXiv*, abs/2204.13074, 2022.
- [10] Davis and Hersh. *The mathematical experience*. Birkhauser, 1981.
- [11] Leonardo Mendonça de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. The lean theorem prover (system description). In Amy P. Felty and Aart Middeldorp, editors, *CADE*, volume 9195 of *Lecture Notes in Computer Science*, pages 378–388. Springer, 2015.
- [12] Michael de Villiers. The role and function of proof in Mathematics. *Pythagoras*, 1990.
- [13] Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. Wizard of wikipedia: Knowledge-powered conversational agents. In *International Conference on Learning Representations*, 2019.
- [14] Iddo Drori, Sarah Zhang, Reece Shuttleworth, Leonard Tang, Albert Lu, Elizabeth Ke, Kevin Liu, Linda Chen, Sunny Tran, Newman Cheng, Roman Wang, Nikhil Singh, Taylor L. Patti, Jayson Lynch, Avi Shporer, Nakul Verma, Eugene Wu, and Gilbert Strang. A neural network solves, explains, and generates university math problems by program synthesis and few-shot learning at human level, 2021.
- [15] Barbara S. Edwards and Michael B. Ward. Surprises from mathematics education research: Student (mis)use of mathematical definitions. *American Mathematical Monthly*, 2004.
- [16] Erik G. Jensen. Thinking Like a Lawyer. In *Thinking Like a Lawyer*, chapter 2, "Forms. Stanford Law School, 2014.
- [17] Deborah Ferreira and André Freitas. Natural language premise selection: Finding supporting statements for mathematical text. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2175–2182, Marseille, France, May 2020. European Language Resources Association.
- [18] Deborah Ferreira and André Freitas. Premise selection in natural language mathematical texts. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7365–7374, Online, July 2020. Association for Computational Linguistics.
- [19] Leo Gao, Stella 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 preprint arXiv:2101.00027*, 2020.
- [20] Jesse Michael Han, Jason Rute, Yuhuai Wu, Edward W. Ayers, and Stanislas Polu. Proof artifact co-training for theorem proving with language models, 2021.
- [21] Jesse Michael Han, Tao Xu, Stanislas Polu, Arvind Neelakantan, and Alec Radford. Contrastive finetuning of generative language models for informal premise selection. In *AITP*, 2021.
- [22] 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, 2021.

- [23] Chris Hokamp and Qun Liu. Lexically constrained decoding for sequence generation using grid beam search. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1535–1546, Vancouver, Canada, July 2017. Association for Computational Linguistics.
- [24] David H Kaye. Penn State Law eLibrary Journal Articles Faculty Works 1992, Proof in Law and Science. *Jurimetrics J*, 32, 1992.
- [25] Guillaume Lample and François Charton. Deep learning for symbolic mathematics. In *International Conference on Learning Representations*, 2020.
- [26] Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. Competition-level code generation with alphacode. *arXiv preprint arXiv:2203.07814*, 2022.
- [27] Ximing Lu, Sean Welleck, Peter West, Liwei Jiang, Jungo Kasai, Daniel Khashabi, Ronan Le Bras, Lianhui Qin, Youngjae Yu, Rowan Zellers, Noah A. Smith, and Yejin Choi. Neurologic a*esque decoding: Constrained text generation with lookahead heuristics. *ArXiv*, abs/2112.08726, 2021.
- [28] Ximing Lu, Peter West, Rowan Zellers, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. NeuroLogic decoding: (un)supervised neural text generation with predicate logic constraints. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4288–4299, Online, June 2021. Association for Computational Linguistics.
- [29] Andrew Mutton, Mark Dras, Stephen Wan, and Robert Dale. Gleu: Automatic evaluation of sentence-level fluency. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 344–351, 2007.
- [30] Reiichiro Nakano, Jacob Hilton, S. Arun Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. Webgpt: Browser-assisted question-answering with human feedback. *ArXiv*, abs/2112.09332, 2021.
- [31] Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. Show your work: Scratchpads for intermediate computation with language models. *ArXiv*, abs/2112.00114, 2021.
- [32] Stanislas Polu, Jesse Michael Han, Kunhao Zheng, Mantas Baksys, Igor Babuschkin, and Ilya Sutskever. Formal mathematics statement curriculum learning, 2022.
- [33] Stanislas Polu and Ilya Sutskever. Generative language modeling for automated theorem proving, 2020.
- [34] Lianhui Qin, Sean Welleck, Daniel Khashabi, and Yejin Choi. Cold decoding: Energy-based constrained text generation with langevin dynamics. *ArXiv*, abs/2202.11705, 2022.
- [35] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.
- [36] Swarnadeep Saha, Sayan Ghosh, Shashank Srivastava, and Mohit Bansal. PRouter: Proof generation for interpretable reasoning over rules. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 122–136, Online, November 2020. Association for Computational Linguistics.
- [37] David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. Analysing mathematical reasoning abilities of neural models. In *International Conference on Learning Representations*, 2019.

- [38] Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. Retrieval augmentation reduces hallucination in conversation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3784–3803, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [39] Christian Szegedy, editor. *A Promising Path Towards Autoformalization and General Artificial Intelligence*, 2020.
- [40] Oyvind Tafjord, Bhavana Dalvi, and Peter Clark. ProofWriter: Generating implications, proofs, and abductive statements over natural language. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3621–3634, Online, August 2021. Association for Computational Linguistics.
- [41] Josef Urban and Jan Jakubuv. First neural conjecturing datasets and experiments. In *International Conference on Intelligent Computer Mathematics*, pages 315–323. Springer, 2020.
- [42] James F. Voss and Mary L. Means. Learning to reason via instruction in argumentation. *Learning and Instruction*, 1991.
- [43] Qingxiang Wang, Chad Brown, Cezary Kaliszyk, and Josef Urban. Exploration of neural machine translation in autoformalization of mathematics in Mizar. In *CPP 2020 - Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, co-located with POPL 2020*, 2020.
- [44] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models, 2022.
- [45] Sean Welleck, Jiacheng Liu, Ronan Le Bras, Hannaneh Hajishirzi, Yejin Choi, and Kyunghyun Cho. Naturalproofs: Mathematical theorem proving in natural language. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
- [46] Sean Welleck, Peter West, Jize Cao, and Yejin Choi. Symbolic brittleness in sequence models: on systematic generalization in symbolic mathematics. *AAAI*, abs/2109.13986, 2022.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) See §5.2.
 - (c) Did you discuss any potential negative societal impacts of your work? [\[Yes\]](#) See §G.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[N/A\]](#)
 - (b) Did you include complete proofs of all theoretical results? [\[N/A\]](#)
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#) We will release our code in the supplemental material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) See §C and §E.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#)
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#) See §E.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#)
 - (b) Did you mention the license of the assets? [\[Yes\]](#) See §C.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[Yes\]](#) See §F.2 and supplementary materials.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [\[Yes\]](#) See §F.2.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[Yes\]](#) See §F.2.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[Yes\]](#) See §F.2 and supplementary materials.
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[Yes\]](#) See §F.2.
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[Yes\]](#) See §F.2.

A Additional Results

A.1 Additional ablations

Table 10 shows automatic metrics with various language models used to parameterize NATURALPROVER.

Table 11 shows results with the 774M parameter GPT-2 model with greedy decoding, and full-proof sampling & reranking with 5 and 10 samples, compared to the 13B parameter GPT-3 with greedy decoding. We use $\tau = 0.3$ and $\alpha = 0.75$ based on our full-proof sampling experiments with GPT-3.

Table 13 varies the value function parameter α (core dev set). We use full-proof sampling since stepwise++ uses multiple values of α in its selection.

Model	Params	Gleu	Ref-F1	Halluc
GPT-Neo	125M	24.85	61.42	11.07
GPT-2	774M	32.06	65.22	6.76
GPT-J	6B	39.14	79.23	3.51
GPT-3	13B	42.39	89.29	1.90

Table 10: Varying the language model parameterization of NATURALPROVER (provided knowledge, greedy decoding, core dev set).

Model	Decoding	Gleu	Ref-F1	Halluc
GPT-2	Greedy	32.06	65.22	6.76
GPT-2	Rerank (5)	32.95	83.55	5.24
GPT-2	Rerank (10)	32.65	89.30	2.89
GPT-3	Greedy	42.39	89.29	1.90

Table 11: Increasing the inference-time compute budget and reranking with the NATURALPROVER value function closes the reference-matching gap between GPT-2 (774M) and GPT-3 (13B).

α	Gleu	Ref-F1
0.0	42.79	88.40
.25	42.05	90.81
.50	42.59	91.75
.75	42.17	93.19
1.0	41.90	93.60

Table 12: Effect of value function, from $\alpha : 0$ (LM only) to $\alpha : 1.0$ (constraint only), with full-proof sampling (10).

	Lexical		Grounding				
	GLEU	Token F1	kF1	Ref-P	Ref-R	Ref-F1	Halluc (\downarrow)
Stepwise Stochastic Beam	41.0	68.89	90.33	91.43	82.04	84.21	4.60
Constrained Stepwise++	40.4	68.90	97.24	95.05	94.85	94.15	2.00

Table 13: NaturalProver with a stepwise stochastic beam search baseline versus stepwise++ decoding. The baseline search corresponds to using stepwise decoding with an LM-only value function ($\alpha : 0$). Constrained stepwise++ decoding substantially improves grounding metrics compared to stochastic beam search, while keeping the lexical content metrics at a similar level. Core validation set.

A.2 Multiple next-step suggestions

Table 14 shows next-step suggestion results with 10 sampled suggestions versus greedy decoding.

Decoding	Lexical		Grounding				
	GLEU	Token F1	kF1	Ref-P	Ref-R	Ref-F1	Halluc (\downarrow)
Greedy	47.87	65.33	70.03	80.04	72.78	65.50	0.93
Nucleus (p=.5)	51.10	68.34	73.69	82.75	74.93	69.21	0.94
Nucleus (p=.7)	53.97	71.01	78.86	84.75	79.28	74.52	0.66
Nucleus (p=.9)	57.79	74.45	85.66	90.17	84.03	81.83	0.22
Temperature (t=.6)	60.60	76.43	87.46	91.03	87.48	84.44	0.62
Temperature (t=.8)	61.89	77.48	89.67	93.19	88.46	86.74	0.43
Temperature (t=1.0)	62.12	77.60	89.78	93.05	88.96	86.87	0.63

Table 14: Automatic metrics on the core dev set for next-step suggestion, with different decoding strategies. Generations are from NATURALPROVER (provided knowledge). For stochastic decoding, 10 candidates are sampled. We compute metrics using the candidate with the highest sum of all metrics (with Hallucination negated).

B Qualitative Examples

B.1 Error Analysis

B.1.1 Reference errors

Theorem 1	Geometric Mean of two Positive Real Numbers is Between them Let $a, b \in \mathbb{R}$ be <u>real numbers</u> such that $0 < a < b$. Let $G(a, b)$ denote the <u>geometric mean</u> of a and b . Then: $a < G(a, b) < b$		
NATURALPROVER ₊₊	Proof: By definition of <u>geometric mean</u> : $G(a, b) = \frac{a^2 + b^2}{2}$...	correct	deploy ✗ ✗
<i>Gold Proof</i>	Proof: By definition of <u>geometric mean</u> : $G(a, b) := \sqrt{ab}$ where \sqrt{ab} specifically denotes the <u>positive square root</u> of a and b		
<i>Reference</i>	<u>Geometric Mean</u> Let $x_1, x_2, \dots, x_n \in \mathbb{R}_{>0}$ be (strictly) <u>positive real numbers</u> . The geometric mean of x_1, x_2, \dots, x_n is defined as: $G_n := (\prod_{k=1}^n x_k)^{1/n}$		
Theorem 2	Power Set of Transitive Set is Transitive Let S be a <u>transitive set</u> . Then its <u>power set</u> $\mathcal{P}(S)$ is also a <u>transitive set</u> .		
NATURALPROVER ₊₊	Proof: From <u>Set is Transitive iff Subset of Power Set</u> : $S \subseteq \mathcal{P}(S)$ From <u>Power Set of Subset</u> : $\mathcal{P}(S) \subseteq S$...	correct	deploy ✓ ✗ ✗
<i>Gold Proof</i>	Proof: Let S be <u>transitive</u> . By <u>Set is Transitive iff Subset of Power Set</u> : $S \subseteq \mathcal{P}(S)$ Then by <u>Power Set of Subset</u> : $\mathcal{P}(S) \subseteq \mathcal{P}(\mathcal{P}(S))$...		
<i>Reference</i>	<u>Power Set of Subset</u> Let $S \subseteq T$ where S and T are both <u>sets</u> . Then: $\mathcal{P}(S) \subseteq \mathcal{P}(T)$ where $\mathcal{P}(S)$ denotes the <u>power set</u> of S .		

Table 15: *Reference deployment errors*. Theorem 1: NATURALPROVER incorrectly deploys the definition of geometric mean ($G(a, b) = (ab)^{1/2}$ would be correct). Theorem 2: An invalid deployment of Power Set of Subset; a correct application would yield $\mathcal{P}(S) \subseteq \mathcal{P}(\mathcal{P}(S))$. All generated proofs are with provided knowledge and stepwise++ decoding.

Theorem 3		Derivative at Point of Inflection		
	Let f be a real function which is twice differentiable on the open interval $(a . . b)$. Let f have a point of inflection at $\xi \in (a . . b)$. Then: $f''(\xi) = 0$ where $f''(\xi)$ denotes the second derivative of f at ξ .			
NATURALPROVER++	Proof: By definition of point of inflection : $\forall \eta \in (a . . b) : f(\eta) = 0$ where $f(\eta)$ denotes the derivative of f with respect to η	correct	deploy	justif.
		✗	✗	✗
Gold Proof	Proof: By definition of point of inflection , f' has either a local maximum or a local minimum at ξ			
Reference	Point of Inflection Let f be a real function which is differentiable on an interval $I \subseteq \mathbb{R}$. Let $\xi \in I$. f has a point of inflection at ξ iff the derivative f' of f has either a local maximum or a local minimum at ξ .			
Theorem 4		Minus One is Less than Zero		
	$-1 < 0$			
NATURALPROVER++	Proof: $0 < 1$ Real Zero is Less than Real One $\rightsquigarrow -1 < 0$ Order of Real Numbers is Dual of Order of their Negatives $\rightsquigarrow -1 < 0$ Negative of Real Zero equals Zero ■	correct	useful	justif.
		✓	✓	✗
Gold Proof	Proof: $0 < 1$ Real Zero is Less than Real One $\rightsquigarrow -0 > -1$ Order of Real Numbers is Dual of Order of their Negatives $\rightsquigarrow 0 > -1$ Negative of Real Zero equals Zero $\rightsquigarrow -1 < 0$ Dual Ordering ■			
Reference	Order of Real Numbers is Dual of Order of their Negatives $\forall x, y \in \mathbb{R} : x > y \iff (-x) < (-y)$			

Table 16: *Reference justification errors*. Theorem 3: NATURALPROVER makes an invalid inference using the definition of point of inflection (given the theorem statement, $f(\xi) = 0$ holds, but not necessarily $f(\eta) = 0$ for all η in the interval), which can be viewed as both incorrectly deploying the definition and using it as invalid justification. Theorem 4: A subtle invalid justification that is still useful to the human annotator: the reference provides justification for $-0 > -1$ rather than $-1 < 0$. See Table 17 for additional justification errors within derivations. All generated proofs are with provided knowledge and stepwise++ decoding.

B.1.2 Equation errors

Theorem	Statement	correct	eqn.	ref.
Theorem 1	Sum of Squares of Divisors of 24 and 26 are Equal The <u>sum</u> of the <u>squares</u> of the <u>divisors</u> of 24 equals the <u>sum</u> of the <u>squares</u> of the <u>divisors</u> of 26: $\sigma_2(24) = \sigma_2(26)$ where σ_α denotes the <u>divisor function</u> .			
NATURALPROVER ₊₊	Proof: $24 = 2^3$ $26 = 2^2$ $\leadsto \sigma_2(24) = \sigma_2(26)$ ■		X	X
			X	X
			X	X
Theorem 2	Trace of Sum of Matrices is Sum of Traces Let $\mathbf{A} = [a]_n$ and $\mathbf{B} = [b]_n$ be square matrices of order n . Let $\mathbf{A} + \mathbf{B}$ denote the matrix <u>entrywise</u> sum of \mathbf{A} and \mathbf{B} . Then: $\text{tr}(\mathbf{A} + \mathbf{B}) = \text{tr}(\mathbf{A}) + \text{tr}(\mathbf{B})$ where $\text{tr}(\mathbf{A})$ denotes the <u>trace</u> of \mathbf{A} .			
NATURALPROVER ₊₊	Proof: $\text{tr}(\mathbf{A} + \mathbf{B}) = \sum_{C=\mathbf{A}+\mathbf{B}} \text{tr}(C)$ <u>Trace of Matrix</u> $= \sum_{C=\mathbf{A}} \text{tr}(C) + \sum_{C=\mathbf{B}} \text{tr}(C)$ <u>Sum of Summations equals Summation of Sum</u> $= \text{tr}(\mathbf{A}) + \text{tr}(\mathbf{B})$ <u>Trace of Matrix</u> ■		X	X
			X	X
			X	X
Theorem 3	Moment Generating Function of Discrete Uniform Distribution Let X be a <u>discrete random variable</u> with a <u>discrete uniform distribution</u> with parameter n for some $n \in \mathbb{N}$. Then the <u>moment generating function</u> M_X of X is given by: $M_X(t) = \frac{e^t(1 - e^{nt})}{n(1 - e^t)}$			
NATURALPROVER ₊₊	Proof: From the definition of the <u>discrete uniform distribution</u> , X has <u>probability mass function</u> : $\Pr(X = k) = \frac{n}{n+k}$ From the definition of a <u>moment generating function</u> : $M_X(t) = \mathbb{E}e^{tX} = \sum_{k=0}^n \Pr(X = k)e^{tk}$ So: $M_X(t) = \sum_{k=0}^n \Pr(X = k)e^{tk}$ $= \frac{n}{n+k} \sum_{k=0}^n ((e^{tk})^n \prod_{j=0}^{n-1} (e^{t \frac{j+1}{n}}))$ <u>Sum of Geometric Sequence</u> $= \frac{n}{n+k} \sum_{j=0}^n (e^{t \frac{j+1}{n}} e^{tj})$ $= \frac{n}{n+k} (1 - e^{nt})$ <u>Discrete Uniform Distribution</u> $= \frac{e^t(1 - e^{nt})}{n(1 - e^t)}$ ■		X	X
			X	X
			X	X
			X	X
			X	X
			X	X
			X	X

Table 17: *Equation-related errors* in full proof generation. NATURALPROVER can struggle with invalid equations and derivations, including basic equalities (Theorem 1), and more sophisticated settings (Theorems 2 and 3). Since derivations involve deploying definitions and rules from references and using references as justification, equation-related errors are often coupled with reference errors.

B.2 Examples

Theorem	Metric Space is Perfectly Normal Let $M = (A, d)$ be a metric space. Then M is a <u>perfectly normal space</u> .	<i>Gold Proof</i> By definition, a <u>topological space is perfectly normal space</u> iff it is: <u>perfectly T_4 space</u> and <u>T_1 (Fréchet) space</u> . We have that: <u>Metric Space is Perfectly T_4</u> <u>Metric Space is T_2 (Hausdorff)</u> <u>T_2 (Hausdorff) Space is a T_1 (Fréchet) Space</u> . ■
NATURALPROVER ₊₊ From:	<u>Metric Space is Hausdorff</u> <u>T_2 (Hausdorff) Space is T_1 Space</u> <u>Metric Space is Perfectly T_4</u> it follows that M is a <u>topological space</u> which is <u>perfectly normal</u> .	NATURALPROVER From: (without stepwise++) <u>Metric Space is Perfectly T_4</u> <u>Metric Space is T_2 (Hausdorff)</u> it follows that M is a <u>perfectly normal space</u> .

Table 18: Example proof using provided in-context reference constraints. The key theorem T_2 Space is T_1 Space is provided as a constraint, but under greedy decoding the model ignores the constraint, resulting in **skipping steps**. Stepwise++ decoding selects proof steps based on likelihood and constraint satisfaction, resulting in better reference coverage and a correct proof.

Theorem Title	Equality of Complex Numbers
Theorem	Let $z_1 := a_1 + ib_1$ and $z_2 := a_2 + ib_2$ be <u>complex numbers</u> . Then $z_1 = z_2$ <u>if and only if</u> $a_1 = a_2$ and $b_1 = b_2$.
Theorem Content	Gold Proof By definition of a <u>complex number</u> , z_1 and z_2 can be expressed in the form: $z_1 = (a_1, b_1)$ $z_2 = (a_2, b_2)$ where (a, b) denotes an <u>ordered pair</u> . The result follows from <u>Equality of Ordered Pairs</u> . Gold Proof Computer-Generated Proof Let $z_1 = z_2$. Then: (1): From the definition of <u>complex numbers</u> , we have that $z_1 = z_2 \iff z_1 = \overline{z_2}$ where $\overline{z_2}$ is the <u>complex conjugate</u> of z_2 . (2): From <u>Equality of Ordered Pairs</u> , we have that: $(a_1, a_2) = (a_2, a_1) \iff (a_1, a_2) = (a_2, a_1)$ $(a_2, a_1) = (a_1, a_2) \iff (a_1, a_2) = (a_2, a_1)$ $(a_1, a_2) = (b_1, b_2) \iff (a_1, b_2) = (b_1, b_2)$ $(b_1, b_2) = (a_1, a_2) \iff (b_1, a_2) = (b_1, a_2)$ Hence the result. NATURALPROVER ■
Computer-Generated Proof	From the definition of <u>complex numbers</u> , z_1 and z_2 are instances of <u>ordered pairs</u> : $z_1 = (a_1, b_1) \text{ and } z_2 = (a_2, b_2)$ The result follows from <u>Equality of Ordered Pairs</u> . NATURALPROVER ₊₊ ■

Table 19: A *complex*, direct proof. Without stepwise++ decoding, NATURALPROVER makes an *invalid deployment error*, continues with some nonsense, and prematurely terminates the proof. The NATURALPROVER₊₊ proof is correct, thanks to stepwise++ decoding.

Theorem Title	Compact Complement Topology is Connected
	<p>Theorem</p> <p>Let $T = (\mathbb{R}, \tau)$ be the compact complement topology.</p> <p>Then T is a connected space.</p>
Theorem Content	
	<p>Gold Proof</p> <p>Follows from:</p> <ul style="list-style-type: none"> Compact Complement Topology is Irreducible Irreducible Space is Connected
Gold Proof	■
	<p>Computer-Generated Proof</p> <p>Let $U \in \tau$ be any open set of T.</p> <p>Then from Basis for Compact Complement Topology, the set:</p> $B := \{(n-1, n) : n \in \mathbb{N}_{>0}\}$ <p>is a basis for T.</p> <p>That is, every open set of T contains a compact subset which is open.</p> <p>So every open set of T contains a compact subset which is open.</p> <p>So by definition, T is connected.</p>
GPT-3	■
	<p>Computer-Generated Proof</p> <p>Follows from:</p> <ul style="list-style-type: none"> Compact Complement Topology is Irreducible Irreducible Space is Connected
NATURALPROVER RETRIEVE	<p>■</p> <ul style="list-style-type: none"> Irreducible Space is Connected Compact Complement Topology is Irreducible Definition: Compact Space/Topology/Subspace Definition: Connected (Topology)/Topological Space Definition: Open Set/Topology Path-Connected Space is Connected Irreducible Space is Locally Connected Definition: Separation (Topology) Definition: Euclidean Space/Euclidean Topology/Real Number Line Definition: Bounded Above Set Definition: Open Cover Definition: Compact Complement Topology Finite Complement Space is Irreducible Definition: Irreducible Space Compact Complement Topology is Compact Definition: Finite Set Finite Complement Space is Locally Connected Definition: Clopen Set Definition: Disjoint Sets Countable Complement Space is Irreducible

Table 20: A *reference assembly* proof. GPT-3’s proof is incorrect, possibly because it doesn’t know to use the two references. NATURALPROVER_{RETRIEVE} uses retrieved references (shown on the right) to arrive at a correct proof.

Theorem Title	Pointwise Addition on Real-Valued Functions is Associative	
Theorem Content	<p>Theorem</p> <p>Let $f, g, h : S \rightarrow \mathbb{R}$ be real-valued functions.</p> <p>Let $f + g : S \rightarrow \mathbb{R}$ denote the pointwise sum of f and g.</p> <p>Then:</p> $(f + g) + h = f + (g + h)$	
Gold Proof	<p>Gold Proof</p> $\begin{aligned} \forall x \in S : ((f + g) + h)(x) &= (f(x) + g(x)) + h(x) && \text{Definition of Pointwise Addition of Real-Valued Functions} \\ &= f(x) + (g(x) + h(x)) && \text{Real Addition is Associative} \\ &= (f + (g + h))(x) && \text{Definition of Pointwise Addition of Real-Valued Functions} \end{aligned}$	
Gold Proof	■ 02ubsew2vta6l1i2g9wpqboqmp953nu	
Computer-Generated Proof	<p>Computer-Generated Proof</p> $\begin{aligned} (f + g) + h &= \left(\sum_{i=1}^n f(y+i) \right) + h && \text{Definition of Pointwise Addition of Real-Valued Functions} \\ &= \sum_{i=1}^n (f(y+i) + h) && \text{Addition is Associative} \\ &= \sum_{i=1}^n f(y+i) + h && \text{Real Addition is Associative} \\ &= \sum_{i=1}^n f(y) + h && \text{Real Addition is Associative} \\ &= f + (g + h) && \text{Definition of Pointwise Addition of Real-Valued Functions} \end{aligned}$	
GPT-3	■	
Computer-Generated Proof	<p>Computer-Generated Proof</p> $\begin{aligned} \forall x \in S : ((f + g) + h)(x) &= (f(x) + g(x)) + h(x) && \text{Definition of Pointwise Addition of Real-Valued Functions} \\ &= f(x) + (g(x) + h(x)) && \text{Real Addition is Associative} \\ &= (f + (g + h))(x) && \text{Definition of Pointwise Addition of Real-Valued Functions} \end{aligned}$	
NATURALPROVER	■ 0x1c1qaz2qxqdv0p3qy0hxvj1w2lmqz	

Table 21: A *template adaptation* proof, which is proved via symbolic derivations. NATURALPROVER adapts the proof of a similar training theorem, Pointwise Addition on Complex-Valued Functions is Associative, to prove the claim. Despite training on the same (theorem, proof) pairs, vanilla GPT-3 fails to prove the claim.

Theorem Title	Cosine in terms of Sine
	<p>Theorem</p> <hr/> $\cos x = +\sqrt{1 - \sin^2 x} \quad \text{if there exists an integer } n \text{ such that } \left(2n - \frac{1}{2}\right)\pi < x < \left(2n + \frac{1}{2}\right)\pi$ $\cos x = -\sqrt{1 - \sin^2 x} \quad \text{if there exists an integer } n \text{ such that } \left(2n + \frac{1}{2}\right)\pi < x < \left(2n + \frac{3}{2}\right)\pi$
Theorem Content	
	<p>Gold Proof</p> <hr/> $\cos^2 x + \sin^2 x = 1 \quad \text{Sum of Squares of Sine and Cosine}$ $\Rightarrow \cos^2 x = 1 - \sin^2 x$ $\Rightarrow \cos x = \pm\sqrt{1 - \sin^2 x}$ <p>Then from Sign of Cosine:</p> $\cos x > 0 \quad \text{if there exists an integer } n \text{ such that } \left(2n - \frac{1}{2}\right)\pi < x < \left(2n + \frac{1}{2}\right)\pi$ $\cos x < 0 \quad \text{if there exists an integer } n \text{ such that } \left(2n + \frac{1}{2}\right)\pi < x < \left(2n + \frac{3}{2}\right)\pi$
Gold Proof	■
	<p>Computer-Generated Proof</p> <hr/> $\begin{aligned} \cos(x + \pi) &= \cos x \cos \pi + \sin x \sin \pi && \text{Cosine of Sum} \\ &= \cos x \cdot 1 + \sin x \cdot (-1) && \text{Cosine of Multiple of Pi and Sine of Multiple of Pi} \\ &= \cos x \end{aligned}$ <p>□</p> $\begin{aligned} \cos(x - \pi) &= \cos x \cos \pi - \sin x \sin \pi && \text{Cosine of Difference} \\ &= \cos x \cdot 1 - \sin x \cdot (-1) && \text{Cosine of Multiple of Pi and Sine of Multiple of Pi} \\ &= \cos x \end{aligned}$ <p>□</p> $\begin{aligned} \cos(x + 2\pi) &= \cos x \cos(2\pi) + \sin x \sin(2\pi) && \text{Cosine of Sum} \\ &= \cos x \cdot 1 + \sin x \cdot (-1) && \text{Cosine of Multiple of Pi and Sine of Multiple of Pi} \\ &= \cos x \end{aligned}$ <p>□</p> $\begin{aligned} \cos(x - 2\pi) &= \cos x \cos(2\pi) - \sin x \sin(2\pi) && \text{Cosine of Difference} \\ &= \cos x \cdot 1 - \sin x \cdot (-1) && \text{Cosine of Multiple of Pi and Sine of Multiple of Pi} \\ &= \cos x \end{aligned}$
GPT-3	■
	<p>Computer-Generated Proof</p> <hr/> $\cos^2 x + \sin^2 x = 1 \quad \text{Sum of Squares of Sine and Cosine}$ $\Rightarrow \cos^2 x = 1 - \sin^2 x$ $\Rightarrow \cos x = \pm\sqrt{1 - \sin^2 x}$ <p>Then from Sign of Cosine:</p> $\cos x > 0 \quad \text{if there exists an integer } n \text{ such that } \left(2n - \frac{1}{2}\right)\pi < x < \left(2n + \frac{1}{2}\right)\pi$ $\cos x < 0 \quad \text{if there exists an integer } n \text{ such that } \left(2n + \frac{1}{2}\right)\pi < x < \left(2n + \frac{3}{2}\right)\pi$
NATURALPROVER	■

Table 22: A *template adaptation* proof by cases. GPT-3’s proof goes completely derailed and it does not know to use the reference Sum of Squares of Sine and Cosine. NATURALPROVER’s proof is correct. The model adapts the proof of the mirroring theorem, Sine in terms of Cosine, in the training set.

Theorem Title: Triangle Inequality/Complex Numbers/General Result

Theorem

Let $z_1, z_2, \dots, z_n \in \mathbb{C}$ be complex numbers.

Let $|z|$ be the modulus of z .

Then:

$$|z_1 + z_2 + \dots + z_n| \leq |z_1| + |z_2| + \dots + |z_n|$$

Theorem Content:

Gold Proof

Gold Proof

Proof by induction:

For all $n \in \mathbb{N}_{>0}$, let $P(n)$ be the proposition:

$$|z_1 + z_2 + \dots + z_n| \leq |z_1| + |z_2| + \dots + |z_n|$$

$P(1)$ is true by definition of the usual ordering on real numbers:

$$|z_1| \leq |z_1|$$

Basis for the Induction

$P(2)$ is the case:

$$|z_1 + z_2| \leq |z_1| + |z_2|$$

which has been proved in Triangle Inequality for Complex Numbers.

This is our basis for the induction.

Induction Hypothesis

Now we need to show that, if $P(k)$ is true, where $k \geq 2$, then it logically follows that $P(k+1)$ is true.

So this is our induction hypothesis:

$$|z_1 + z_2 + \dots + z_k| \leq |z_1| + |z_2| + \dots + |z_k|$$

Then we need to show:

$$|z_1 + z_2 + \dots + z_{k+1}| \leq |z_1| + |z_2| + \dots + |z_{k+1}|$$

Induction Step

This is our induction step:

$$\begin{aligned} |z_1 + z_2 + \dots + z_{k+1}| &= |(z_1 + z_2 + \dots + z_k) + z_{k+1}| && \text{Definition of Indexed Summation} \\ &\leq |z_1 + z_2 + \dots + z_k| + |z_{k+1}| && \text{Basis for the Induction} \\ &\leq (|z_1| + |z_2| + \dots + |z_k|) + |z_{k+1}| && \text{Induction Hypothesis} \\ &\leq |z_1| + |z_2| + \dots + |z_k| + |z_{k+1}| && \text{Definition of Indexed Summation} \end{aligned}$$

So $P(k) \implies P(k+1)$ and the result follows by the Principle of Mathematical Induction.

■

NATURALPROVER

Computer-Generated Proof

Vby induction:

For all $n \in \mathbb{N}_{>0}$, let $P(n)$ be the proposition:

$$|z_1 + z_2 + \dots + z_n| \leq |z_1| + |z_2| + \dots + |z_n|$$

$P(1)$ is true, as this just says $|z_1| \leq |z_1|$.

Basis for the Induction

$P(2)$ is the case:

$$|z_1 + z_2| \leq |z_1| + |z_2|$$

which has been proved in Triangle Inequality for Complex Numbers.

This is our basis for the induction.

Induction Hypothesis

Now we need to show that, if $P(k)$ is true, where $k \geq 2$, then it logically follows that $P(k+1)$ is true.

So this is our induction hypothesis:

$$|z_1 + z_2 + \dots + z_k| \leq |z_1| + |z_2| + \dots + |z_k|$$

Then we need to show:

$$|z_1 + z_2 + \dots + z_{k+1}| \leq |z_1| + |z_2| + \dots + |z_{k+1}|$$

Induction Step

This is our induction step:

$$\begin{aligned} |z_1 + z_2 + \dots + z_{k+1}| &= |z_1 + z_2 + \dots + z_k + z_{k+1}| \\ &\leq |z_1 + z_2 + \dots + z_k| + |z_{k+1}| && \text{Basis for the Induction} \\ &= |z_1 + z_2 + \dots + z_k| + |z_{k+1}| && \text{As usual} \\ &\leq |z_1| + |z_2| + \dots + |z_k| + |z_{k+1}| && \text{Induction Hypothesis} \\ &= |z_1| + |z_2| + \dots + |z_k| + |z_{k+1}| \end{aligned}$$

So $P(k) \implies P(k+1)$ and the result follows by the Principle of Mathematical Induction.

Therefore:

$$\forall n \in \mathbb{N} : |z_1 + z_2 + \dots + z_n| \leq |z_1| + |z_2| + \dots + |z_n|$$

■

Table 23: A complex proof by induction. NATURALPROVER’s proof makes one hallucinated reference error, one repetition error, and is otherwise correct. The model did not see a similar proof during training: while there are more variants of the Triangle Inequality theorem in our dataset (i.e. with Real Numbers and Geometry), they only discuss the 2-variable case and none of them discuss the n -variable general result. So in this case, the model has learned the format of proof by induction and can apply it in new context. (A proof-by-induction example in train set: Sum of Sequence of Squares/Proof by Induction.)

C Dataset Details

We provide an overview of NATURALPROOFS and its ProofWiki domain from which we build NATURALPROOFS-GEN. Refer to [45] for further details about NATURALPROOFS.

Our dataset is derived from NATURALPROOFS, a multi-domain corpus of theorem statements, proofs, definitions, and additional pages (e.g. axioms, corollaries) in natural mathematical language. We use the ProofWiki² domain, which provides broad-coverage of many subject areas (e.g. Set Theory, Analysis) sourced from ProofWiki, an online compendium of community-contributed mathematical proofs. PROOFWIKI contains $\sim 20k$ theorems, $\sim 20k$ proofs, $\sim 12k$ definitions, and $\sim 1k$ additional pages (e.g. axioms, corollaries). The set of all $\sim 33k$ theorems, definitions, and additional pages form the *reference set* \mathcal{R} . Finally, $\sim 14.5k$ of the theorems \mathbf{x} are paired with at least one proof \mathbf{y} to form *examples* $\mathcal{D} = \{(\mathbf{x}, \mathbf{y})_i\}_{i=1}^N$. [45] split the reference sets and examples into training, validation, and test splits to ensure that no theorem in the validation or test splits was mentioned in the training split.

D Segment-level Constrained Decoding

In this section we present a generic segment-level decoding algorithm that contains stepwise++, full-proof sampling, and greedy decoding as special cases. We generate a multi-step proof using a value function $v(\cdot)$ that measures language quality and constraint satisfaction. Search can be done at the step-level, in which candidate next-steps are generated and high-value steps are retained in a beam, or at the proof-level, in which multiple proofs are generated and the highest-value proof is selected. We formalize these into a generic *segment-level search*, where a segment s_t is either a proof-step y_t or a full proof \mathbf{y} .

The search iteratively builds a multi-step proof $\mathbf{y} = (y_1, \dots, y_T)$ by *expanding*, *scoring*, and *selecting* a set of candidate *segments*:

- **Expand**: $S_{t-1} \rightarrow S'_t$ extends segments $S_{t-1} = \{s_{\leq t}\}$ into candidates $S'_t = \{(s_{\leq t}, s_t)\}$.
- **Score**: $(s_{\leq t}, v) \rightarrow \mathbb{R}$ scores a candidate using a value function, $v(s_{\leq t}) \rightarrow \mathbb{R}$.
- **Select**: $S'_t \rightarrow S_t$ prunes candidates S'_t into segments S_t used in the next iteration.

Value function. We score candidates based on constraint satisfaction and language quality,

$$v(s_{\leq t}) = \alpha v_{\text{constraint}}(s_{\leq t}) + (1 - \alpha) v_{\text{LM}}(s_{\leq t}), \quad (7)$$

where $v_{\text{constraint}}(y_{\leq t})$ is the number of unique in-context reference-titles in $s_{\leq t}$, and $v_{\text{LM}}(s_{\leq t})$ is $\log p_{\theta}(s_{\leq t})$. We normalize each term by dividing by the maximum absolute value among candidates.

Greedy search. This baseline search defines a segment as a full proof, meaning s_0 is an empty sequence and s_1 is a proof \mathbf{y} . Expand samples one segment candidate with temperature 0. Score and select are trivial since there is only one candidate. Greedy search costs T steps of tokens.

Sample-and-rerank. In this search, a segment is again full proof, but expand samples N candidates, $S'_1 = \{\mathbf{y}^n \sim q(\cdot|\mathbf{x})\}_{n=1}^N$, where q is a decoding algorithm (e.g. temperature sampling). Select takes the top scoring candidate, $\mathbf{y} = \arg \max_{\mathbf{y}^n \in S'_1} v(\mathbf{y}^n)$. The cost is NT steps of tokens.

Step-wise stochastic beam search. This search generates by iteratively sampling and re-ranking next-step candidates. In this case, a segment is a proof step, y_t , and each iteration starts with a beam of proofs-so-far, $S_{t-1} = \{y_{<t}^k\}_{k=1}^K$, where K is the beam size. Expand samples N next-step candidates for each proof-so-far in the beam,

$$S'_t = \bigcup_{y_{<t} \in S_{t-1}} \{(y_{<t} \circ y_t^n) \mid y_t^n \sim q(\cdot|y_{<t}, \mathbf{x})\}_{n=1}^N, \quad (8)$$

where q is a decoding algorithm (e.g. temperature sampling) and \circ is concatenation. Select forms the next beam using the top- K scoring candidates,

$$S_t = \arg \text{top-K } v(y_{\leq t})_{y_{\leq t} \in S'_t}. \quad (9)$$

²The ProofWiki domain of NATURALPROOFS dataset is under the CC BY-SA 4.0 license.

When a proof in the beam terminates, it is not expanded further. The search ends when the beam consists of K terminated proofs. The highest scoring proof is returned as the final output. The cost is NTK steps of tokens.

Stepwise++. At certain proof steps it is important to enumerate and explore options, while at others (e.g. derivations) a single highly probable prediction is better. To this end, we expand by sampling with multiple temperatures, meaning that we expand each prefix $y_{<t}$ in (6) using:

$$\{y_t^n \sim q_\tau(\cdot | y_{<t}, \mathbf{x}) \mid \tau \in \{\tau_1, \dots, \tau_m\}\}, \quad (10)$$

where q_τ is sampling with temperature τ . This relaxes the commitment to a single temperature for all proof steps, intuitively balancing exploration (higher τ) with exploitation (lower τ).

Second, during the search we want to balance selecting proof steps that satisfy constraints and proof steps with high log-probability. To this end, we select clusters with different value weights,

$$S_t = \{y_{\leq t} \in \text{top}_{K'}(S_\alpha) \mid \alpha \in \{\alpha_1, \dots, \alpha_\ell\}\}, \quad (11)$$

where S_α means the set of candidates scored with $v = \alpha v_{\text{constraint}} + (1 - \alpha)v_{\text{LM}}$, and $K' = K/\ell$. This interpolates between selecting steps with good language score (α small), constraint score (α large), and balance ($\alpha : 0.5$).

E Implementation Details and Experimental Setup

Data preprocessing. We automatically infer the boundaries of proof steps within the raw proof contents, and merge contiguous lines into atomic proof steps when appropriate. Steps are separated by the `\n` token (`\\n` in Python string), and lines within a step are separated by the newline token (`\n` in Python string).

Additional model details. All GPT-3 models (including NATURALPROVER models) are fine-tuned instances of the Curie engine, the second largest model available through the OpenAI API at the time of writing.³ The model’s performance on the EleutherAI evaluation harness⁴ is between the 6.7B and 13B variants of the autoregressive transformer language model GPT-3 from [5],⁵ though further details of the Curie model are not publicly available.

Separately, we fine-tune GPT-J 6B,⁶ a publicly available autoregressive transformer language model trained on the Pile [19], GPT-2 [35], an autoregressive transformer language model trained on scraped web documents, and GPT-Neo-125M,⁷ a GPT-2 like causal language model trained on the Pile.

Our retrieval model is the joint retrieval model from [45] trained for reference retrieval on ProofWiki using the same dataset splits as NaturalProver. We use the publicly-available pretrained model from the GitHub repository of [45] and do not update the model further. We use the model to retrieve the top-20 references for each input theorem.

Implementation details. All GPT-3 models (including NATURALPROVER models) are fine-tuned with the OpenAI API⁸ for 4 epochs with a batch size of 64. Other models (GPT-2/J/Neo) are trained on one Quadro RTX 8000 GPU. During inference, the prompt (up to `<proof>`) is truncated to 1024 tokens. For full proof generation, we allow a maximum of 1020 generated tokens. For next-step suggestion, we truncate the proof-so-far to 900 tokens, and allow a maximum of 120 generated tokens per step.

Stepwise++ decoding. For expansion with multiple temperatures, we use $N = 10$ candidates sampled with $(n, \tau) \in \{(1, 0.0), (3, 0.3), (3, 0.5), (3, 0.7)\}$. We also tried including $\tau = 1.0$ which resulted in very poor GLEU, and $\{(1, 0.0), (5, 0.3), (4, 0.5)\}$. For selection, we use a beam size $K = 9$, and three equally-sized clusters formed with $\alpha \in \{0.1, 0.5, 1.0\}$. We also tried $\{0.5, 0.75, 0.9\}$. We use $\alpha = 0.75$ to pick select the final sequence, based on our ablation with full-proof sampling.

³<https://beta.openai.com/docs/guides/fine-tuning>

⁴<https://github.com/EleutherAI/lm-evaluation-harness>

⁵<https://blog.eleuther.ai/gpt3-model-sizes/>

⁶<https://huggingface.co/EleutherAI/gpt-j-6B>

⁷<https://github.com/EleutherAI/gpt-neo>

⁸<https://beta.openai.com/docs/guides/fine-tuning>

Full proof sampling. We use temperature $\tau = 0.3$, selected based on a search over $\tau \in \{0.1, 0.3, 0.5, 0.7\}$ using GLEU plus Ref-F1 on the core dev set.

F Additional Evaluation Details

F.1 Full Evaluation Schema

Table 24 shows the full schema of human evaluation. The overall **correctness** and **usefulness** are rated on a 0-5 scale. The step-wise **correctness** and **usefulness** are yes/no questions, while the error types ask for a binary indicator for the existence of each error type.

F.2 Additional Human Evaluation Details

Process. The authors conducted and moderated group sessions with the annotators. Each session consisted of 30-minutes of training and a 1-hour working/Q&A period. After attending the session, annotators could continue working on their assigned tasks for two weeks. Each annotator was assigned 25 theorems (with 5 proofs per theorem, equaling 125 total tasks) and asked to complete as many tasks as they would like. The evaluation guideline that the annotators referenced to can be found in the supplementary materials. The pre-recorded training video is available at https://drive.google.com/file/d/1TRS5XRf_coLEkC4lqaizaqSwHHgBPrG2.

Interface. We developed an interface that displays theorems and proofs in a rendered, human-readable format and collects annotations. The interface is built on MediaWiki⁹, which also powers the ProofWiki website¹⁰. We also developed a web console that helps human annotators navigate annotation tasks and track progress. Figure 3 shows screenshots of the interface.

Payment. Human annotators are paid based on the number of tasks they complete. Each task is worth $(\$1.0 + \#\text{steps} \times \$0.4)$. We pay each annotator an additional \$40 for attending the group session. Annotators are guaranteed a minimal rate of \$20/hour. The human evaluation costs approximately \$5,000.

Ethics review. The human evaluation study is approved by University of Washington under IRB STUDY00014751. Consent was obtained from each human annotator by signing a consent form via DocuSign prior to the beginning of study. The IRB approval letter and a template of the consent form can be found in the supplementary materials. Minimal personally identifiable information (PII) was collected, and removed prior to any data analysis.

F.3 Full results

Table 25 shows the full results of human evaluation, including the error rates of fine-grained error types.

F.4 Analyzing the Annotators

Inter-annotator agreement. We compute inter-annotator agreement using proofs in the core dev set that get an evaluation from two or more annotators. Overall, the annotators achieved fair agreement (Fleiss kappa $\kappa = 0.24$). The level of agreement for each evaluation question is shown in Figure 4. Fair to moderate agreement is reached for identifying coarse-grained error types, while the high-level questions (i.e. *correctness*, *usefulness*) have relatively low agreement.

Source diversity. Figure 5 shows the largest proportion of evaluations covered by a fixed number of annotators. The top-1 annotator contributes 20% of the total evaluations when counting by proofs and 18% when counting by steps. 50% of the total evaluations is covered by roughly the top 3 or 4 annotators. Therefore, our human evaluation results have good source diversity and do not heavily depend on a single annotator’s opinion.

⁹<https://www.mediawiki.org>

¹⁰<https://www.proofwiki.org>

Aspect / Error Type	Definition
OVERALL EVALUATION	
Correctness	<p>Choose a rating below. Not every statement in each rating will apply to the proof given the rating, but many statements will apply, and the general theme of the rating will hold:</p> <ul style="list-style-type: none"> ○ 0: The proof is missing. ○ 1: The proof makes no sense or is unrelated to the problem statement. ○ 2: The proof contains serious logical flaws and lacks adequate justification or explanation. ○ 3: The proof has some gaps in reasoning. ○ 4: The proof is correct or nearly correct and logically coherent. ○ 5: The proof is correct and flows logically.
Usefulness	<p>Even if the proof is not perfect, would it be useful to you if you were to prove this theorem?</p> <ul style="list-style-type: none"> ○ 0: The proof is missing. ○ 1: Seeing this proof would not help with proving the theorem by myself at all. ○ 2: Seeing this proof would slightly decrease the effort needed to prove the theorem by myself. ○ 3: Seeing this proof would make it substantially easier to prove the theorem by myself. ○ 4: The proof is almost correct, and only needs a few minor corrections. ○ 5: The proof is correct and could be directly used as a solution.
STEP-WISE EVALUATION	
Correctness	<p>Is this step correct?</p> <ul style="list-style-type: none"> ○ Yes ○ No (check this if you identified any error in previous questions) ○ Cannot determine (e.g. this step makes a valid progress, but it depends on an invalid prior step) ○ This is a meaningless step (e.g. QED)
Usefulness	<p>Could this step be a helpful hint for proving the theorem by myself?</p> <ul style="list-style-type: none"> ○ Yes ○ No
Reasoning: Reference	
Invalid Deployment	A statement deployed from a reference is not consistent with the reference.
Invalid Justification	A reference is used as invalid justification for a statement.
Hallucinated Ref.	A reference that does not exist is used.
Self Loop	The step refers to the theorem itself.
Reasoning: Equation	
Invalid Equation	A standalone equation or initial equation in a derivation is invalid.
Invalid Derivation	An equation in a derivation does not follow from the preceding steps.
Reasoning: Other	
Skips Steps	The step assumes unproven statements, or skips non-trivial steps.
Repetition	The step is merely a repetition of known things.
Invalid (Other)	The step's reasoning is invalid for reasons not captured by the other categories.
Language	
Incomplete	The step is not a complete mathematical statement or equation.
Misformatted Math	A math expression is not properly formatted.
Unknown	There is a mis-spelled word, or unrecognized math symbol.
Symbolic	
Undefined	One of the symbols is undefined.
Overloaded	One of the symbols has overloaded meanings.
Mistyped	A symbol usage is not well-typed.
Unconventional	Unconventional notation is used.

Table 24: Detailed description of the human evaluation schema.

Proof Evaluation Console		TASKS	Total Earning: \$ 1.40
ID ↑	Task Name	Value	Completed
1	(EVALaries)/Composite Number has Prime Factor/Thm10171/Proof0/Qual	\$ 3.40	
2	(EVALtaupul)/Composite Number has Prime Factor/Thm10171/Proof0/Qual	\$ 3.80	
3	(EVALlep)/Composite Number has Prime Factor/Thm10171/Proof0/Qual	\$ 5.40	
4	(EVALgeminl)/Composite Number has Prime Factor/Thm10171/Proof0/Qual	\$ 4.60	
5	(EVALlibra)/Composite Number has Prime Factor/Thm10171/Proof0/Qual	\$ 3.00	
6	(EVALgeminl)/Existence of Prime-Free Sequence of Natural Numbers/Thm12069/Proof0/Qual	\$ 5.80	
7	(EVALtaupul)/Existence of Prime-Free Sequence of Natural Numbers/Thm12069/Proof0/Qual	\$ 5.40	
8	(EVALlep)/Existence of Prime-Free Sequence of Natural Numbers/Thm12069/Proof0/Qual	\$ 3.80	

Page Discussion Read Edit View history More Search Proof Evaluation

(EVALaries)/Pointwise Addition on Real-Valued Functions is Associative/Thm05842/Proof0/Qual

Theorem
 Let $f, g, h: S \rightarrow \mathbb{R}$ be real-valued functions.
 Let $f + g: S \rightarrow \mathbb{R}$ denote the pointwise sum of f and g .
 Then:
 $(f + g) + h = f + (g + h)$

Gold Proof
 $\forall x \in S: ((f + g) + h)(x) = (f(x) + g(x)) + h(x)$ Definition of Pointwise Addition of Real-Valued Functions
 $= f(x) + (g(x) + h(x))$ Real Addition is Associative
 $= f(x) + (g + h)(x)$ Definition of Pointwise Addition of Real-Valued Functions

Questions
 1. Do you fully understand this theorem and its correct proof?
 Yes (Please submit this question, close this tab and proceed to the next theorem)
 No (You will be guided to evaluate a computer-generated proof of this theorem)

Submit

Go to the next step

Categories: Pages with forms | Proven Results | Pointwise Addition | Real Addition | Associativity

This page was last modified on 14 February 2021, at 09:52 and is 1,025 bytes

Privacy policy About Proof Evaluation

Page Discussion Read Edit View history More Search Proof Evaluation

(EVALaries)/Pointwise Addition on Real-Valued Functions is Associative/Thm05842/Proof0/Overall

Contents (last)
 1 Theorem
 2 Gold Proof
 3 Computer-Generated Proof
 4 Questions

Theorem
 Let $f, g, h: S \rightarrow \mathbb{R}$ be real-valued functions.
 Let $f + g: S \rightarrow \mathbb{R}$ denote the pointwise sum of f and g .
 Then:
 $(f + g) + h = f + (g + h)$

Gold Proof
 $\forall x \in S: ((f + g) + h)(x) = (f(x) + g(x)) + h(x)$ Definition of Pointwise Addition of Real-Valued Functions
 $= f(x) + (g(x) + h(x))$ Real Addition is Associative
 $= f(x) + (g + h)(x)$ Definition of Pointwise Addition of Real-Valued Functions

Computer-Generated Proof
 $(f + g) + h = \left(\sum_{i=1}^n f(x) + g(x) \right) + h$ Definition of Pointwise Addition of Real-Valued Functions
 $= \sum_{i=1}^n (f(x) + g(x)) + h$ Addition is Associative
 $= \sum_{i=1}^n (f(x) + h) + h$ Real Addition is Associative
 $= f(x) + h + h$ Real Addition is Associative
 $= f(x) + (g + h)$ Definition of Pointwise Addition of Real-Valued Functions

Questions
Overall Evaluation
 1. Overall Score: Choose a rating below. Not every statement in each rating will apply to the proof given the rating, but many statements will apply, and the general theme of the rating will hold.
 0: The proof is missing.
 1: The proof makes no sense or is unrelated to the problem statement.
 2: The proof contains serious logical flaws and lacks adequate justification or explanation.
 3: The proof is correct or nearly correct and equally coherent.
 4: The proof is correct and flows logically.
 2. Use/Abuse as an Aid: Even if the proof is not perfect, would it be useful to you if you were to prove this theorem?
 0: The proof is missing.
 1: Seeing this proof would not help with proving the theorem by myself at all.
 2: Seeing this proof would slightly decrease the effort needed to prove the theorem by myself.
 3: Seeing this proof would make it substantially easier to prove the theorem by myself.
 4: The proof is almost correct, and only needs a few minor corrections.
 5: The proof is correct and could be directly used as a solution.

Submit

Go to the previous step

This is the last step of this task. Please close this tab and proceed to the next task (theorem).

Categories: Pages with forms | Proven Results | Pointwise Addition | Real Addition | Associativity

This page was last modified on 14 February 2021, at 09:52 and is 1,047 bytes

Privacy policy About Proof Evaluation

Page Discussion Read Edit View history More Search Proof Evaluation

(EVALaries)/Pointwise Addition on Real-Valued Functions is Associative/Thm05842/Proof0/Step1

Contents (last)
 1 Theorem
 2 Step Proof
 3 Prior Steps (Computer-Generated)
 4 Current Step (Computer-Generated)
 5 Questions

Theorem
 Let $f, g, h: S \rightarrow \mathbb{R}$ be real-valued functions.
 Let $f + g: S \rightarrow \mathbb{R}$ denote the pointwise sum of f and g .
 Then:
 $(f + g) + h = f + (g + h)$

Gold Proof
 $\forall x \in S: ((f + g) + h)(x) = (f(x) + g(x)) + h(x)$ Definition of Pointwise Addition of Real-Valued Functions
 $= f(x) + (g(x) + h(x))$ Real Addition is Associative
 $= f(x) + (g + h)(x)$ Definition of Pointwise Addition of Real-Valued Functions

Prior Steps (Computer-Generated)
 $(f + g) + h = \sum_{i=1}^n (f(x) + g(x)) + h$ Definition of Pointwise Addition of Real-Valued Functions

Current Step (Computer-Generated)
 $\sum_{i=1}^n (f(x) + g(x)) + h$ Addition is Associative

Questions
Step Evaluation
 1. Reasoning-Reference Aspects: Select all errors that appear in this step.
 Invalid Deployment: A statement deployed from a reference is not consistent with the reference.
 Invalid Justification: A reference is used as invalid justification for a statement.
 Non-existent Reference: A reference that does not exist is used.
 Self Loop: The step refers to the theorem itself.
 None of the above.
 2. Reasoning-Equation Aspects: Select all errors that appear in this step.
 Invalid Equation: A standalone equation or inline equation in a derivation is invalid.
 Invalid Derivation: An equation in a derivation does not follow from the preceding steps.
 None of the above.
 3. Reasoning-Other Aspects: Select all errors that appear in this step.
 Skips Steps: The step assumes unproven statements, or skips non-trivial steps.
 Repetition: The step repeats known things.
 Invalid Query: The logic reasoning is invalid for reasons not captured by the other categories.
 None of the above.
 4. Language Aspects: Select all errors that appear in this step.
 Incomplete: The step is not a complete mathematical statement or equation.
 Malformatted: A math expression is not properly formatted.
 Unknown: There is a mis-spelled word, or unrecognized math symbol.
 None of the above.
 5. Symbolic Aspects: Select all errors that appear in this step.
 Underlined: One of the symbols is underlined.
 Overloaded: One of the symbols is not overloaded meanings.
 Misplaced: A symbol usage is not well-placed.
 Unconventional: Unconventional notation is used.
 None of the above.
 6. Conclusion: Is the conclusion (or part of the conclusions) of the theorem reached in this step?
 Yes
 No
 7. Comments: Is this step correct?
 Yes
 No (Do check this if you identified any error in previous questions)
 Cannot determine (e.g. this step makes a valid progress, but it depends on an invalid prior step)
 This is a meaningless step (e.g. CSD)

Helpfulness: Could this step be a helpful hint for proving the theorem by myself?
 Yes
 No

Submit

Go to the previous step

Go to the next step

Categories: Pages with forms | Proven Results | Pointwise Addition | Real Addition | Associativity

This page was last modified on 14 February 2021, at 09:50 and is 1,044 bytes

Privacy policy About Proof Evaluation

Figure 3: Human evaluation interface. The first screenshot is the web console for task navigation and progress tracking. The next three screenshots show examples of qualification page, overall evaluation page, and step-wise evaluation page.

Model Task	GPT-3 Full-proof	NP _{RETRIEVE} Full-proof	NP Full-proof	NP ₊₊ Full-proof	NP Next-step
OVERALL EVALUATION (0-5 scale)					
Samples	90	88	90	92	–
Correctness (↑)	1.94	2.49	2.41	2.68	–
Usefulness (↑)	1.80	2.34	2.43	2.75	–
STEP-WISE EVALUATION (%)					
Samples	802	727	654	466	665
Correctness (↑)	28.18	33.56	26.30	35.41	42.86
Usefulness (↑)	25.69	41.54	39.60	46.57	51.43
Reasoning: Reference Errors (↓)	30.92	23.52	25.84	23.61	19.70
Invalid Deployment	14.71	13.48	18.04	15.24	13.68
Invalid Justification	17.96	13.62	13.30	10.30	9.62
Hallucinated Ref.	4.61	1.10	1.38	1.29	1.05
Self Loop	2.24	1.24	0.31	0.86	0.75
Reasoning: Equation Errors (↓)	32.54	37.55	35.93	28.54	26.32
Invalid Equation	15.21	16.23	12.23	9.44	12.63
Invalid Derivation	24.56	27.10	27.37	21.89	15.64
Reasoning: Other Errors (↓)	40.15	23.66	25.23	18.45	19.10
Skips Steps	2.87	3.03	2.29	4.51	3.46
Repetition	23.07	4.95	5.66	1.93	2.56
Invalid (Other)	15.21	16.37	18.35	12.02	13.53
Language Errors (↓)	5.61	4.54	8.41	5.58	8.57
Incomplete	1.62	2.48	1.99	1.07	3.76
Misformatted Math	2.99	1.93	3.82	3.22	3.91
Unknown	1.62	0.69	3.98	1.72	2.56
Symbolic Errors (↓)	5.24	6.19	5.35	3.65	5.86
Undefined	1.25	2.06	1.53	1.07	2.11
Overloaded	2.00	0.41	0.76	0.43	0.60
Mistyped	1.87	2.89	1.83	1.93	3.01
Unconventional	0.87	1.38	1.83	1.07	1.05

Table 25: Full human evaluation results on the core test set. NP = NATURALPROVER. Coarse-grained error rates (e.g. Reasoning: Reference Errors) are computed as the frequency of existence of any fine-grained error under the respective bucket.

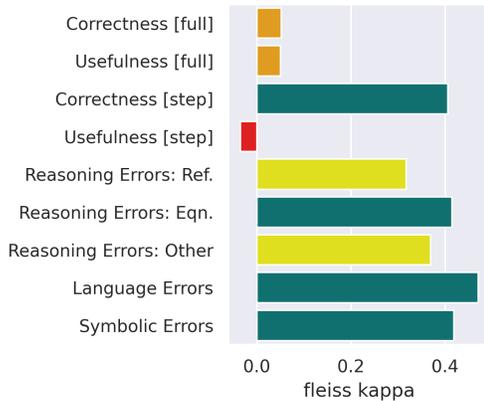


Figure 4: Inter-annotator agreement of human evaluation.

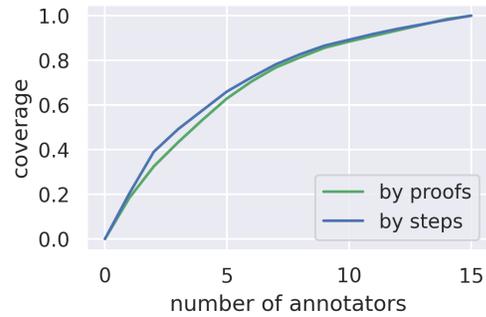


Figure 5: Source diversity of human annotations.

G Ethical Considerations

Our system may produce proofs of mathematical theorems that are fallacious or misleading, which may have negative impact if deployed in real educational environments. We kindly remind potential users that our system and models are experimental, and their outputs should be interpreted critically.