

Reliable Evaluation and Benchmarks for Statement Autoformalization

Auguste Poiroux Gail Weiss
Viktor Kunčak Antoine Bosselut

School of Computer and Communication Sciences
EPFL, Switzerland

{auguste.poiroux, gail.weiss, viktor.kuncak, antoine.bosselut}@epfl.ch

Abstract

Evaluating statement autoformalization, translating natural language mathematics into formal languages like Lean 4, remains a significant challenge, with few metrics, datasets, and standards to robustly measure progress. In this work, we present a comprehensive approach combining improved metrics, robust benchmarks, and systematic evaluation, to fill this gap. First, we introduce **BEq+**, an automated metric that correlates strongly with human judgment, along with **ProofNetVerif**, a new dataset for assessing the quality of evaluation metrics, containing 3,752 annotated examples. Second, we develop two new autoformalization benchmarks: **ProofNet#**, a corrected version of ProofNet, and **RLM25**, with 619 new pairs of research-level mathematics from six formalization projects. Through systematic experimentation across these benchmarks, we find that current techniques can achieve up to 45.1% accuracy on undergraduate mathematics but struggle with research-level content without proper context. Our work establishes a reliable foundation for evaluating and advancing autoformalization systems.

1 Introduction

Automatic verification of logical reasoning holds promise for formal verification of mathematical proofs, software verification, and artificial intelligence. Proof assistants allow users to rigorously express mathematical statements and mechanically check their proofs (Mahboubi and Tassi, 2022; Paulson, 2023; Avigad, 2024), but require inputs to be *formalized*: translated from informally stated mathematical statements into a formal language, such as Lean 4. New research explores methods that automate this process, a task referred to as *autoformalization* (Szegedy, 2020; Wang et al., 2020), but there exists few reliable resources for evaluating statement autoformalization methods (Liu et al., 2024; Lu et al., 2024b).

To enable scalable experimentation and robust evaluation, we propose new metrics, benchmarks, and standards for evaluating statement autoformalization methods. First, we introduce **BEq+**, a reference-based metric inspired by BEq (Liu et al., 2024). **BEq+** is deterministic, runs efficiently on CPU alone, and does not require a 20B LLM prover. We measure its strong correlation with human judgments, demonstrating its utility as a metric for comparing autoformalization models. Recognizing that progress in autoformalization will require continued improvements in metrics, we also release **ProofNetVerif**, a new benchmark of 3752 formal-informal pairs with human-annotated binary semantic equivalence labels that can be used to benchmark the faithfulness of new metrics.

Next, we propose two benchmarks as evaluation testbeds of statement autoformalization methods. First, we identify numerous formalization errors in existing Lean 4 ports (31.8% of the entries), and present **ProofNet#**, a meticulously corrected version of the leading ProofNet benchmark (Azerbaiyev et al., 2023a). At the same time, we observe that ProofNet focuses on undergraduate-level and self-contained statement formalization. To assess autoformalization in more realistic, context-dependent, research-oriented scenarios, we curate **RLM25**, a novel benchmark comprising 619 pairs from real-world, research-level formalization projects, the first of its kind, to our knowledge.

Finally, we set new standards for evaluating autoformalization. Current state-of-the-art methods, whether based on large language models (Wu et al., 2022; Azerbaiyev et al., 2023a), distilled back-translation (Jiang et al., 2023; Azerbaiyev et al., 2023a), or retrieval augmented generation (Azerbaiyev et al., 2023a; Liu et al., 2024), show limited success, with reported accuracies typically below 20% on benchmarks such as ProofNet (Azerbaiyev et al., 2023a; Liu et al., 2024). A common failure of these methods is the inability to generate for-

malizations that *type-check* in Lean 4 (Moura and Ullrich, 2021), a crucial precursor to correctness. Our results show that sampling multiple formalizations and discarding the ones that do not type-check already substantially increases downstream correctness.

Given that statement autoformalization is more likely to be applied in limited data contexts (i.e., mathematics), increasing test-time compute is a promising factor. Consequently, we propose that autoformalization studies should explore a variety of inference-time compute budgets to broadly test method capabilities. While prior work explored sampling multiple formalizations (up to $n = 20$) (Li et al., 2024b; Agrawal et al., 2022), we systematically study how performance scales with a significantly larger number of samples (up to $n = 1000$). We also investigate various strategies for filtering ill-typed samples and selecting the final prediction.

Our experiments demonstrate that scaling test-time compute by increasing candidate sampling, combined with type-check filtering and selection heuristics, can substantially improve autoformalization accuracy. Notably, the performance of GPT-4o on ProofNet# improves from 31.0% (single generation) to 45.1% (using 50 samples and a selection heuristic). These findings, alongside our contributions in metrics and benchmarks, establish a more reliable foundation for evaluating and advancing statement autoformalization systems, particularly for complex undergraduate and research-level mathematics.

We summarize our contributions as follows:

- We develop **BEq+**, a reference-based metric for statement autoformalization based on deterministic symbolic computation and running exclusively on a CPU. We demonstrate that **BEq+ correlates strongly with human annotations at the benchmark level**, making it reliable to compare different autoformalization models.
- We release **ProofNetVerif**, a benchmark of 3 752 formal-informal pairs annotated with binary semantic equivalence labels to evaluate statement autoformalization metrics.
- We curate **ProofNet#**, a revised version of ProofNet with many corrections.
- We release **RLM25**, a new statement autoformalization benchmark based on several research-level natural language-aligned formalization projects. To the best of our knowledge, this is the first autoformalization benchmark on research-level mathematical topics.

- Using BEq+ and manual evaluation, we study the statement autoformalization performance of various leading models.

We release **BEq+**, **ProofNetVerif**, and **ProofNet#** with an MIT license. **RLM25** is released under an Apache 2.0 license, similar to the license of the underlying projects.

2 Related Work

Interactive Theorem Proving. Autoformalization in mathematics depends on formal systems, such as Coq (Castéran and Bertot, 2004), Lean (Moura and Ullrich, 2021), Isabelle (Nipkow et al., 2002), and their math libraries. In this work, we focus on Lean (specifically, its current version, Lean 4): a powerful interactive theorem prover with a growing formal library of definitions and proven statements known as Mathlib (mathlib Community, 2020). de Moura et al. (2015) and Moura and Ullrich (2021) provide insights into the inner workings of Lean type-checking.

Autoformalization. Classical programmatic tools can be used to translate *constrained* natural language statements into formal systems (Pathak, 2024). In contrast, we are interested in translating *unconstrained* natural language statements. In Wu et al. (2022), the authors find LLMs to be a promising approach, capable of autoformalization through the use of in-context learning. In Azerbayev et al. (2023a) and Jiang et al. (2023), the authors demonstrate that distilled back-translation improves performance of some base models. Agrawal et al. (2022) use an advanced post-processing step to automatically fix type errors in LLM predictions. They find that, for a given problem, keeping only well-typed predictions when generating several formalization attempts is a strong filter. In a more recent work, (Li et al., 2024b) proposes a self-consistency approach specifically designed for autoformalization. Their method clusters logically equivalent formalizations using automated theorem-proving techniques. They evaluate this approach up to $n = 10$ samples per problem.

Metrics. Evaluating the accuracy of models on the statement autoformalization task is a non-trivial problem. In previous works (Wu et al., 2022; Agrawal et al., 2022; Azerbayev et al., 2023a; Jiang et al., 2023), manual evaluation is the standard practice to report statement autoformalization performance. Manual evaluation, though comprehensive and methodical, is a bottleneck for evaluation, mo-

tivating automatic metrics as a proxy for manual accuracy, such as BLEU (Wu et al., 2022; Azerbayev et al., 2023a; Ying et al., 2024b). However, prior work (Azerbayev et al., 2023a) showed that the correlation between BLEU and formalization accuracy is low. Other works have also proposed type-check rate (Azerbayev et al., 2023a; Agrawal et al., 2022) and symbolic equivalence metrics (Liu et al., 2024; Li et al., 2024b) as proxies for autoformalization accuracy. In our work, we further discuss shortcomings of these metrics, and propose new metrics to overcome them.

Benchmarks. ProofNet (Azerbayev et al., 2023a) is a benchmark specifically designed for autoformalization. It consists of 371 undergraduate mathematical exercises, making it an essential benchmark for evaluating the performance of autoformalization models. In a recent work on neural theorem proving (Hu et al., 2024), the authors evaluated their automated theorem prover method on research-level formal projects. Similarly, in Liu et al. (2024), the authors evaluate their method on Con-NF (Wilshaw, 2025) using LLM-generated natural language statements.

LLM sampling-based methods. For our evaluation, we study a method using self-consistency algorithms such as majority voting (Wang et al., 2023) and Self-BLEU (Zhu et al., 2018). Such methods have empirically proven to be effective across a wide range of NLP tasks (Li et al., 2024a). In particular, Lewkowycz et al. (2022) demonstrated the effectiveness of the combination of sampling and majority voting on the MATH benchmark (Hendrycks et al., 2021). Further works in this direction improve over majority voting by using trained verifiers (Hosseini et al., 2024).

3 Manual and Symbolic Metrics

Currently, the most reliable evaluation for autoformalization is a manual evaluation by persons with sufficient understanding of the formal proof assistant and its library. As in prior work, we define a formalization as **correct** if it is semantically equivalent to the provided natural language statement. Throughout this paper, **accuracy** exclusively refers to the proportion of statements evaluated as correct by *manual annotation*.

Symbolic Computation Metrics BEq_L and BEq^+ . However, as manual evaluation is time-consuming, we propose novel automatic evaluation metrics BEq_L and BEq^+ that compare a can-

didate formalization to a reference formalization by checking equivalence between two formulas using symbolic algorithms inside the proof assistant. We invoke proof scripts in Lean that try to prove each formula from the other. These metrics are motivated by the prior BEq metric (Liu et al., 2024) which leverages a 20B LLM trained on the theorem proving task to check equivalence. Instead of an LLM, our metrics employ two purely symbolic, CPU-only equivalence checks that are not only more computationally efficient and interpretable, but also eliminate the need for specialized hardware, allowing us to scale our experiments to $n = 1000$ samples per query.

BEq_L . This metric is based on proving formula equivalence using the `exact?` tactic (tactics can be seen as proof steps). As noted by the authors (Liu et al., 2024), this tactic can prove equivalence for a wide variety of syntactic variations of equivalent formal statements. `exact?` can, in general, use other theorems from the libraries to prove the current formalization, resulting in a situation where unrelated true theorems are considered equivalent. We therefore restrict the `exact?` tactic to only use the candidate equivalent formula. Even though this metric has a relatively high false negative rate, its simplicity makes it easier to reproduce results across works.

BEq^+ . To improve recall compared to BEq_L , we introduce a new symbolic metric: BEq^+ . BEq^+ explores several alternative proof strategies for each direction of the implication, as summarized in Algorithm 1. A description of the tactics used can be found in subsection A.1. We apply this algorithm once for each of the two implication directions.

3.1 ProofNetVerif: Benchmarking Metrics

When proposing new automated metrics (section 3), we must validate whether these metrics are accurate enough to approximate human evaluation. In related work (Liu et al., 2024), the authors conduct a study on 200 sampled formalization attempts and show that BEq is accurate at the instance level, but with a relatively high amount of false negatives.

Using manual annotations from this paper, we curate a new benchmark: ProofNetVerif. Achieving high performance on this benchmark serves two key purposes: (1) establishing robust evaluation metrics for autoformalization, and (2) improving selection strategies by filtering out incorrect formalizations, thereby narrowing the search space in sampling-based methods.

Algorithm 1 BEq+ - Unidirectional

Input: Theorem formalizations t_1 and t_2
Output: Whether t_2 can be derived from t_1

1. *Run BEq_L*
if `exact`? closes proof by using t_1 **then**
return TRUE
2. *Leverage conclusion matching*
if `apply` t_1 or convert t_1 succeeds **then**
Proving t_1 assumptions can be derived from t_2
if repeated applications of `tauto`, `simp_all_arith!`, `noncomm_ring`, or `exact`? close proof **then**
return TRUE
3. *Attempt direct assumption of t_1*
if `have` : `goal(t_1) := by apply_rules [t_1]` succeeds **then**
if repeated applications of `tauto`, `simp_all_arith!`, `noncomm_ring`, or `exact`? close the subproof introduced by `have` **then**
if repeated applications of `tauto`, `simp_all_arith!`, or `exact`? **using this** close the main proof **then**
return TRUE
return FALSE

ProofNetVerif contains aligned examples of a natural language statement, reference formalization, predicted formalization, and a boolean label indicating semantic equivalence between the predicted formalization and the natural language statement, making it suitable to evaluate both reference-free and reference-based metrics. The benchmark contains 3752 examples, with 1142 entries of predicted formalizations equivalent to a natural language input (the remainder being non-equivalent). With this scale, ProofNetVerif provides a challenging benchmark for designing automated metrics to evaluate autoformalization mistakes.

We report the results of our new metrics on ProofNetVerif in [Table 1](#). In particular, we find that BEq+ largely improves over BEq_L and captures non-trivial semantical equivalence. We measured the performance of BEq+ and BEq_L only once on this benchmark and did not tune their design to fit this specific dataset. Additional details and discussions can be found in [subsection A.5](#).

Table 1: Binary performance, in percentage, of BEq_L and BEq+ metrics on ProofNetVerif.

Metric	BEq _L	BEq+
Precision	100.0	98.0
Recall	30.9	48.3
F1 Score	47.2	64.7

3.2 Benchmark-Level Agreement

ProofNetVerif measures instance-level agreement between human evaluation and automated metrics.

Table 2: Correlation between human-reported accuracy and different automated metrics on ProofNet# using data from all models evaluated in this paper.

Metric	Pearson	Kendall
Type-Check	0.655	0.560
BEq _L	0.966	0.846
BEq+	0.974	0.872

However, current metrics still show relatively low recall relative to near-perfect precision, underestimating true autoformalization accuracy of models.

To complement this instance-level view, we also study benchmark-level agreement, where we no longer look at whether a single formalization attempt is correct, but rather whether aggregated scores across an entire benchmark (i.e., across all problems solved by a given model under a given method) align with human judgments (i.e., accuracy). We show that BEq+ has strong correlation with accuracy: leveraging all our manual annotations on ProofNet# conducted for this work, we measure the correlation of symbolic metrics and human evaluation at the benchmark level, amounting to 65 data points (different models and method choices, see [subsection A.17](#)), and report the results in [Table 2](#). Correlation factors between BEq+ and human evaluation are fairly high, with a Kendall coefficient of 0.872. While symbolic metrics are conservative in comparison to human-measured accuracy, this high Kendall correlation coefficient shows that BEq+ can be reliably used to compare model performance for autoformalization.

We also find that the type-check rate correlation with human evaluation is relatively low, experimentally confirming that this metric can not be used as a substitute for human evaluation. However, our empirical results in [Figure 3](#) suggest that type-check rates can probably be used to approximate accuracy when comparing different experiments for a *single* model. This can be particularly useful in reference-free setups outside of benchmarks where no reference formalizations are available.

4 New Benchmarks

In addition to metrics, we contribute several new benchmarks for autoformalization, including RLM25, a new benchmark for research-Level mathematics, and ProofNet#, a manual re-annotation of ProofNet that identifies and corrects a large percentage of mislabeled pairs.

Table 3: Lean Blueprint projects used to build RLM25

Project	#Thms	Lean	First Commit
Carleson	111	4.14.0-rc2	20 Oct 2023
	56	4.14.0-rc2	19 Nov 2023
	84	4.7.0-rc2	22 Mar 2024
	145	4.14.0-rc3	13 Nov 2023
PrimeNumberTheoremAnd testing-lower-bounds	99	4.14.0-rc2	9 Jan 2024
	124	4.13.0-rc3	22 Feb 2024

4.1 RLM25: Research-Level Mathematics

To better evaluate the use of autoformalization for formalizing new mathematical results, we introduce a new benchmark, RLM25, comprised of 619 pairs natural language statement and their Lean formalizations with context¹ from six formalization projects (Table 3) that use the Lean blueprint framework (Massot, 2025). Because these projects contain natural language-aligned formalizations of research-level mathematics, they are suitable for evaluating statement autoformalization methods. We believe RLM25 is more representative of the intended use of autoformalization in mathematics research. To the best of our knowledge, we are the first to conduct such a study on real formalization projects. With 619 pairs, RLM25 is slightly larger than existing benchmarks MiniF2F (488 pairs) and ProofNet (371 pairs).

Curating this benchmark required notable engineering and analysis efforts. We use plasTeX (plasTeX Development Team, 2024) to extract natural language statements from blueprint latex files along with the Lean labels. We then use LeanDojo (Yang et al., 2023) to extract formal statements along with their context from the Lean files. Finally, we align the natural language statements with their formal counterparts using the Lean labels in the latex files. For evaluation, we use Lean REPL (Morrison, 2023), which we backported to previous Lean versions to make the latest features and bug fixes available for all the projects in RLM25.

The projects included in RLM25 have been selected as follows: from 11 Lean blueprint project candidates available at the time of our study, we selected 6 projects that (1) began after October 2023 to avoid data contamination, (2) contained at least 50 extractable formal-informal pairs, and (3) for which we got use agreement from their primary authors to conduct evaluations of AI systems.

¹We obtained agreement from the primary authors of these projects to evaluate our models on them.

4.2 ProofNet#: A Corrected Lean 4 ProofNet

Past autoformalization approaches (Azerbayev et al., 2023a; Jiang et al., 2023; Liu et al., 2024) often use the ProofNet benchmark (Azerbayev et al., 2023a) for evaluation, containing 371 pairs (185 validation, 186 test) of informal statements in undergraduate mathematics and corresponding formalizations in Lean 3. As we focus on Lean 4, we start from two Lean 4 ports of ProofNet (Vishwakarma, 2024; Xin et al., 2024). However, our analysis revealed that these ports, direct translations from Lean 3 to Lean 4, contained 118 entries with formalization errors (31.8% of the total entries), biasing downstream evaluation. We corrected these errors, leading to a new dataset, ProofNet# (see subsection A.2), that is compatible and well-typed for Lean versions between 4.7.0 to 4.16.0-rc2. ProofNet# remains very close to ProofNet, as only the reference formalizations are updated. Hence, the results reported in other works using reference-free metrics (e.g., human evaluation and type-check rate) remain the same.

5 Autoformalization Methods

In this section, we describe the leading autoformalization approaches evaluated using our new metrics and benchmarks.

5.1 Adaptation Methods

In-context learning. On ProofNet#, we use the 12-shot prompt from Azerbayev et al. (2023a) updated to Lean 4, which we share in subsection A.15. On research-level formalization projects, similar to Hu et al. (2024) for neural theorem proving, we also consider in-file context, i.e., the content preceding the official formalizations in the project files.

Fine-tuning on MMA. Empirically, it has been shown that LLMs are better at informalization, i.e., translating formal statements to informal mathematical statements, than autoformalization (Wu et al., 2022; Azerbayev et al., 2023a). Using this fact, Jiang et al. (2023) informalized the Lean 4 Mathlib library with GPT-4 (OpenAI et al., 2024a) to create a dataset, MMA, of formal-informal pairs that is often used for finetuning autoformalization methods.

Fine-tuning on Lean Workbook. In a recent work (Ying et al., 2024a), the authors release a synthetically generated training set for statement autoformalization. They train a model on MiniF2F and ProofNet benchmarks, and then, through active

learning, they curate a train set of $\sim 57K$ formal-informal pairs. Because of this training on the benchmarks, we only evaluate models fine-tuned on this dataset on the RLM25 benchmark.

5.2 Sampling-Based Methods

We consider sampling methods designed as a plug-and-play improvement that can be seamlessly integrated with existing techniques to enhance their performance.

Sampling. In our experiments, unless otherwise stated, we employ temperature sampling with $T = 0.7$ and generate $n = 50$ autoformalization attempts per informal statement.

Filtering. We use Lean REPL (Morrison, 2023) to implement our filtering step, which assesses if provided statements are well-typed. de Moura et al. (2015) and Moura and Ullrich (2021) provide detailed insights into the inner workings of the Lean type system and type-checking.

Selection. In our selection process, we employ and compare four distinct heuristics to refine and choose the best outputs generated by the models: random selection, majority voting (Wang et al., 2023), Self-BLEU (Zhu et al., 2018), and the symbolic equivalence method presented in Li et al. (2024b). A more detailed description of each component is presented in subsection A.3.

6 Experiments

We evaluate statement autoformalization across models and methods using BEq+ and human annotations. Our experiments test (i) optimal temperature, (ii) various selection methods, (iii) how performance scales with the number of attempts, and (iv) the effect of contextual information on research-level tasks.

6.1 Models Used in Experimental Setup

Llemma-7B & 34B (Azerbayev et al., 2023b). These open models are based on CodeLlama 7B and 34B (Rozière et al., 2024) and have been further pre-trained on the ProofPile-2 collection of mathematical data (explicitly excluding ProofNet), which was introduced along with these models. Due to their training data, these math models are particularly suited for formal-related tasks.

Llama3-8B-Instruct (Grattafiori et al., 2024). This is a state-of-the art open 8B model from the Llama3 family.

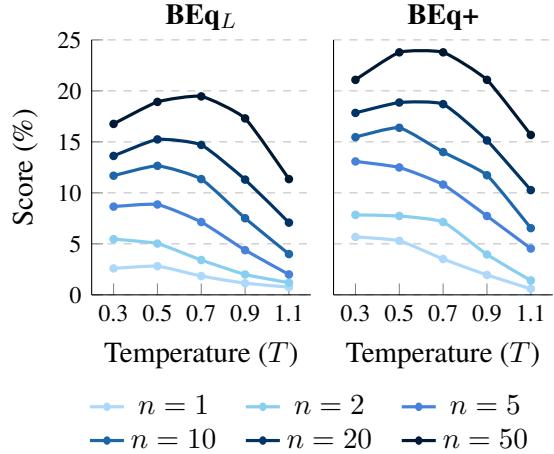


Figure 1: **Evolution of BEq+ and BEq_L pass@n scores for different temperatures on top of type-check filtering.** We evaluate Llemma 7B on ProofNet# validation split.

GPT-4o (OpenAI et al., 2024b). This is a state-of-the-art general LLM. We use version gpt-4o-2024-05-13 for reproducibility.

To address data contamination concerns, Table 3 includes the first commit date of each project in the RLM25 benchmark.² For all these projects, this commit date comes after the announced knowledge cut-off date of the models used in this work: October 2023 for GPT-4o, March 2023 for Llama3 8B, and August 2023 as the release date of Llemma 7B.

6.2 Analysis of Sampling-Based Methods

In this section, we conduct an ablation study of various parameters involved in sampling-based methods. By default, we use a temperature of $T = 0.7$, $n = 50$ samples, and 12-shot prompting in our experiments. All the results in this section are conducted on the **validation** split of the ProofNet# benchmark.

Optimal temperature In Figure 1, we report the evolution of BEq+ pass@n metric³ using the Llemma 7B model. We find that the optimal temperature for balancing exploration and coherent outputs depends on the number of samples, and that this optimal temperature increases with the number of samples. For the rest of this study, we continue with the value of $T = 0.7$ for our sampling-based experiments.

²The authors of the projects confirmed that the first commit dates correspond to the first public appearance of these projects.

³pass@n measures the percentage of tasks where at least one of the n outputs is correct, as measured by BEq+ here.

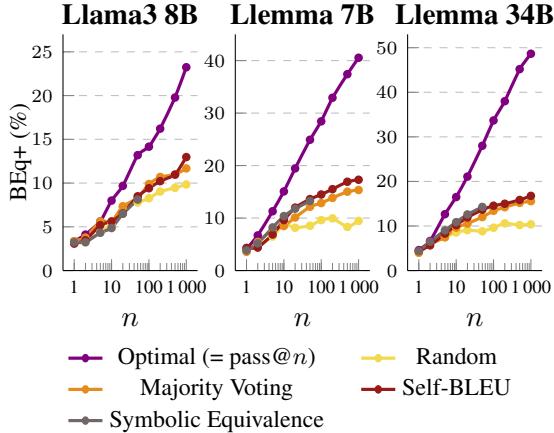


Figure 2: **Evolution of BEq+ metric for different selection methods** on top of type-check filtering. We evaluate Llama3 8B, Llemma 7B, and Llemma 34B on ProofNet# validation split with a number of candidate samples up to $n = 1000$. Given its quadratic scaling with n and high computational cost, the symbolic equivalence method is limited to $n \leq 50$ candidate samples.

Optimal selection method In Figure 2, we report BEq+ scores of the different selection methods, along with the optimal score that can be achieved with a perfect selection method, which is represented by pass@ n . First, BEq+ pass@ n scores steadily increase with the number of samples, going, for example, from 4.11% at $n = 1$, to 40.54% at $n = 1000$ for Llemma 7B, meaning that there is great potential in enhancing autoformalization performance through sampling. Additional studies on BEq+ pass@ n can be found in subsection A.7. Second, we find that the studied self-consistency methods performance monotonically increases with the number of samples. We confirm these findings with more models and manual evaluation in Figure 3. Finally, we find that the symbolic equivalence method does not yield better empirical results for selection compared to simpler and less compute-intensive methods such as Self-BLEU or majority voting. We therefore continue with Self-BLEU and majority voting in subsequent experiments.

Type-Check Filtering In Figure 4, we empirically study the contribution of the filtering component in sampling-based methods by evaluating with and without filtering, as well as with different selection heuristics. While majority voting (*No filter + Majority voting*) and Self-BLEU selection (*No filter + Self-BLEU*) generally improve the accuracy of random sampling, both struggle to increase the performance of random sampling beyond that of the

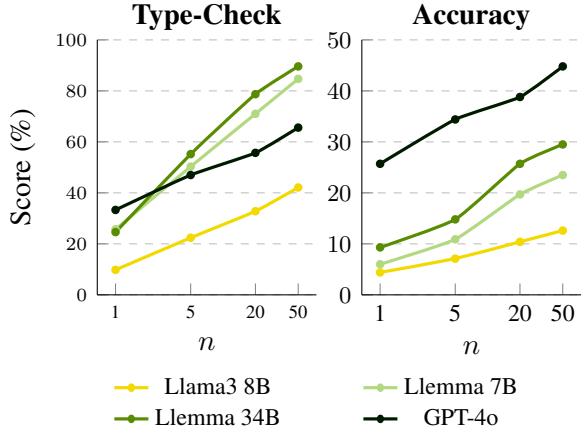


Figure 3: **Type-Check rate and Accuracy scaling trends with respect to the number of samples** on ProofNet# validation split using 12-shot prompting and a sampling-based method (type-check filter + Self-BLEU). The number of samples varies from $n = 1$ to 50 (where Llemma 34B has top type-check rate and GPT-4o top accuracy). Numbers are in Table 15.

greedy decoding baseline. Meanwhile, adding type-check filtering substantially outperforms the greedy decoding baseline even without any final selection heuristic (*Filter + Random selection*). We conclude that the type-check filter is a critical component in sampling-based methods and that it should be applied before selection.

6.3 Empirical Analysis on RLM25

Table 4: Results, in percentage, on RLM25 for different methods and models using only the natural language statement as input (i.e., no in-file context is provided).

Model	Method	Type-Check	BEq _L	BEq+
Llama3 8B	12-shot	2.80	0.20	0.40
	MMA	3.18	0.00	0.54
Llemma 7B	12-shot	5.06	2.78	2.78
	MMA	8.65	0.79	1.16
	Lean Workbook	23.06	0.17	1.13
Llemma 34B	12-shot	4.94	0.79	0.79
GPT-4o	12-shot	10.55	0.57	1.50

To further confirm our findings and evaluate current autoformalization methods in more realistic settings, we conduct studies on RLM25 in this section. We start by conducting an initial study using the current approach in the literature for statement autoformalization: only the natural language statement is provided as input to the methods. We report our results in Table 4. These results are very similar to the ones presented in Liu et al. (2024) on

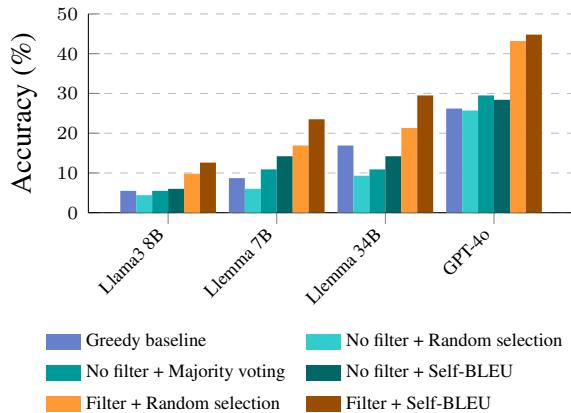


Figure 4: **Type-Check Filtering Ablation Study.** Accuracy scores are reported on ProofNet# validation split for various ablations of sampling-based methods. More details and exact numbers are reported in Table 14.

their semi-synthetic Con-NF benchmark, with BEq results close to 0% for all methods.

A manual analysis of the predictions quickly reveals a key issue: without in-file context, autoformalization methods lack access to crucial information. While some missing information from a Lean file, such as local definitions, can be tackled through current retrieval-based methods, others, such as opened namespaces and local variables are still missing. To assess the importance of different file components, we conducted an ablation study to identify which contextual elements help LLMs generate accurate formalizations. As shown in Table 5, the best performance across models is achieved when proofs are removed while retaining all other contextual elements. In contrast, removing both proofs and theorems significantly degrades performance, particularly for less capable models.

Table 5: **Ablation study on prompt content** using various models on RLM25. We report BEq+ (%) performance. In the prompt column, ‘-’ represents removal from the context.

Prompt	Llama3 8B	Llemma 7B	Llemma 34B	GPT-4o
12-shot	0.40	2.78	0.79	1.50
Full file context	18.67	22.15	25.99	20.64
- theorems & proofs	6.60	13.82	15.63	17.20
- proofs	20.29	24.16	30.33	24.56

In Table 5, we find that in-file context prompting substantially improves performance across all methods and models compared to Table 4. Furthermore, fine-tuning on existing autoformalization datasets does not yield improvement over base models on RLM25, suggesting that context-aware

formalization datasets are needed for tackling autoformalization for research-level projects.

6.4 Final Results

We conducted a detailed baseline study along with human evaluation of the different autoformalization methods on ProofNet#, which we report in the appendix in Table 8. Confirming results from prior works (Jiang et al., 2023; Azerbayev et al., 2023a), we note the large proportion of errors for all methods due to type-check failures. In Table 6, we report side-by-side final performance results between classical greedy decoding and a sampling-based method using Self-BLEU on both ProofNet# and RLM25. These results confirm the consistent performance improvement of sampling for all tested models on the two benchmarks, ProofNet# and RLM25. We also report the effect of sampling on MMA and Lean Workbook fine-tuned models in Table 9. However, we find that applying the sampling strategy on base models with few-shot learning achieves better absolute accuracy. We report additional results on other benchmarks such as PDA (Lu et al., 2024a) and MiniF2F (Zheng et al., 2022) in subsection A.11.

Table 6: **Performance comparison between greedy decoding and a sampling-based method on ProofNet# and the new RLM25 benchmark.** 12-shot is used for ProofNet#, and in-file context with proofs removed is used for RLM25. Sampled with $T = 0.7$ and $n = 50$.

Model	Method	ProofNet#		RLM25
		Accuracy	BEq+	BEq+
Llama3 8B	Greedy	3.3	3.3	20.3
	Filter + Self-BLEU	<u>12.0</u>	<u>9.2</u>	<u>23.9</u>
Llemma 7B	Greedy	10.9	6.5	24.2
	Filter + Self-BLEU	<u>29.3</u>	<u>17.9</u>	<u>28.8</u>
Llemma 34B	Greedy	12.5	7.1	30.3
	Filter + Self-BLEU	<u>28.3</u>	<u>14.7</u>	<u>33.4</u>
GPT-4o	Greedy	31.0	18.5	24.6
	Filter + Self-BLEU	45.1	23.4	<u>31.6</u>

7 Conclusion

Our work advances and standardizes the evaluation of statement autoformalization. We introduced several key resources: **BEq+**, an automated equivalence-checking metric demonstrating a strong correlation with human judgments, and **ProofNetVerif**, a dataset of 3752 annotated examples to validate new metrics. Alongside these, we published **ProofNet#**, a corrected version of the

popular ProofNet benchmark, and **RLM25**, the first benchmark for research-level mathematics formalization across six projects. Our experiments using these new resources reveal critical insights. Sampling strategies, combined with incorporating type-checking and selection heuristics, substantially boost performance, achieving up to 45.1% accuracy on undergraduate mathematics (ProofNet#), suggesting autoformalization evaluations should integrate test-time compute budgets and simple heuristic correctors for recording more realistic performance measures. These benchmarks and our metric provide the community with more reliable tools to measure progress. We believe these contributions will foster the development of more capable and practical autoformalization systems.

Limitations

We identify key limitations of our work. First, our newly introduced RLM25 benchmark is focused on statement autoformalization, while the underlying projects from which we construct the benchmark would support extension to proof autoformalization, which is not tackled in our work. Second, we observe that BEq+ and BEqL face challenges in proving equivalence for long formal statements (as shown in [subsection A.5](#)). The recall performance of BEq+ drops significantly from 62.5% on short statements to 29.6% on long statements, indicating that it over-penalizes longer formal expressions. These metrics also require accurate ground truth formalizations, which are not available when formalizing completely new mathematics.

Data contamination. Our in-depth study in [subsection A.16](#) suggests that data contamination is unlikely among the models we evaluated.

Acknowledgements

We thank the Lean community for their support and feedback, in particular the authors of the Lean blueprint projects included in RLM25. We thank Simon Sorg for finding and sharing an exploit of the BEq+ metric we address in [subsection A.6](#). We also gratefully acknowledge the support of the IC school of computer and communication sciences, the Swiss National Science Foundation (No. 215390), Innosuisse (PFFS-21-29), the EPFL Center for Imaging, Sony Group Corporation, and a Meta LLM Evaluation Research Grant.

References

Ayush Agrawal, Siddhartha Gadgil, Navin Goyal, Ashvni Narayanan, and Anand Tadipatri. 2022. [Towards a Mathematics Formalisation Assistant using Large Language Models](#). *arXiv preprint*. ArXiv:2211.07524 [cs].

Jeremy Avigad. 2024. Automated reasoning for mathematics. In *IJCAR (1)*, volume 14739 of *Lecture Notes in Computer Science*, pages 3–20. Springer.

Zhangir Azerbayev, Bartosz Piotrowski, Hailey Schoelkopf, Edward W. Ayers, Dragomir Radev, and Jeremy Avigad. 2023a. [ProofNet: Autoformalizing and Formally Proving Undergraduate-Level Mathematics](#). *arXiv preprint*. ArXiv:2302.12433 [cs].

Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. 2023b. [Llemma: An Open Language Model For Mathematics](#). *arXiv preprint*. ArXiv:2310.10631 [cs].

Pierre Castérán and Yves Bertot. 2004. *Interactive theorem proving and program development. Coq’Art: The Calculus of inductive constructions*. Texts in Theoretical Computer Science. Springer Verlag. Traduction en chinois parue en 2010. Tsinghua University Press. ISBN 9787302208136.

Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. 2015. The lean theorem prover (system description). In *Automated Deduction - CADE-25*, pages 378–388, Cham. Springer International Publishing.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The Llama 3 Herd of Models](#). *arXiv preprint*. ArXiv:2407.21783 [cs].

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring Mathematical Problem Solving With the MATH Dataset](#). *arXiv preprint*. ArXiv:2103.03874 [cs].

Arian Hosseini, Xingdi Yuan, Nikolay Malkin, Aaron Courville, Alessandro Sordoni, and Rishabh Agarwal. 2024. [V-STaR: Training Verifiers for Self-Taught Reasoners](#). *arXiv preprint*. ArXiv:2402.06457 [cs].

Jiewen Hu, Thomas Zhu, and Sean Welleck. 2024. [miniCTX: Neural Theorem Proving with \(Long\)Contexts](#). *arXiv preprint*. ArXiv:2408.03350 [cs].

Albert Q. Jiang, Wenda Li, and Mateja Jamnik. 2023. [Multilingual Mathematical Autoformalization](#). *arXiv preprint*. ArXiv:2311.03755 [cs].

Ashvini Jindal. 2023. [Arithmo-mistral-7b: Mathematical reasoning model](#). Hugging Face.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). *Preprint*, arXiv:2309.06180.

Aitor Lewkowycz, Anders Johan Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Venkatesh Ramasesh, Ambrose Sloane, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. 2022. [Solving Quantitative Reasoning Problems with Language Models](#). In *NeurIPS*.

Junyou Li, Qin Zhang, Yangbin Yu, Qiang Fu, and Deheng Ye. 2024a. [More agents is all you need](#). *Preprint*, arXiv:2402.05120.

Zenan Li, Yifan Wu, Zhaoyu Li, Xinming Wei, Xian Zhang, Fan Yang, and Xiaoxing Ma. 2024b. [Autoformalize Mathematical Statements by Symbolic Equivalence and Semantic Consistency](#). *arXiv preprint*. ArXiv:2410.20936.

Qi Liu, Xinhao Zheng, Xudong Lu, Qinxiang Cao, and Junchi Yan. 2024. [Rethinking and Improving Autoformalization: Towards a Faithful Metric and a Dependency Retrieval-based Approach](#).

Jianqiao Lu, Zhengying Liu, Yingjia Wan, Yinya Huang, Haiming Wang, Zhicheng Yang, Jing Tang, and Zhi-jiang Guo. 2024a. [Process-Driven Autoformalization in Lean 4](#). *arXiv preprint*. ArXiv:2406.01940 [cs].

Jianqiao Lu, Yingjia Wan, Yinya Huang, Jing Xiong, Zhengying Liu, and Zhi-jiang Guo. 2024b. [FormalAlign: Automated Alignment Evaluation for Autoformalization](#). *arXiv preprint*. ArXiv:2410.10135 [cs].

Assia Mahboubi and Enrico Tassi. 2022. [Mathematical Components](#). Zenodo.

Patrick Massot. 2025. [leanblueprint](#).

The mathlib Community. 2020. [The lean mathematical library](#). In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP)*. ACM.

Kim Morrison. 2023. [leanprover-community/repl](#). Original-date: 2023-03-30T23:12:19Z.

Leonardo de Moura and Sebastian Ullrich. 2021. The lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, pages 625–635, Cham. Springer International Publishing.

Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. 2002. *Isabelle/Hol a Proof Assistant for Higher-Order Logic*. Springer, Berlin and New York.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024a. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.

OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, A. J. Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, and 400 others. 2024b. [GPT-4o System Card](#). *arXiv preprint*. ArXiv:2410.21276 [cs].

Shashank Pathak. 2024. [Gflean: An autoformalisation framework for lean via gf](#). *Preprint*, arXiv:2404.01234.

Lawrence C. Paulson. 2023. Large-scale formal proof for the working mathematician - lessons learnt from the ALEXANDRIA project. In *CICM*, volume 14101 of *Lecture Notes in Computer Science*, pages 3–15. Springer.

plasTeX Development Team. 2024. [plastex/plastex](#). Original-date: 2014-03-06T16:10:23Z.

Baptiste Rozière, Jonas Gehrking, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémie Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, and 7 others. 2024. [Code llama: Open foundation models for code](#). *Preprint*, arXiv:2308.12950.

Christian Szegedy. 2020. [A Promising Path Towards Autoformalization and General Artificial Intelligence](#). In *Intelligent Computer Mathematics*, volume 12236, pages 3–20, Cham. Springer International Publishing.

Rahul Vishwakarma. 2024. [rahul3613/ProofNet-lean4](#). Original-date: 2024-03-04T14:12:29Z.

Qingxiang Wang, Chad Brown, Cezary Kaliszyk, and Josef Urban. 2020. [Exploration of Neural Machine Translation in Autoformalization of Mathematics in Mizar](#). In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 85–98. ArXiv:1912.02636 [cs].

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). *Preprint*, arXiv:2203.11171.

Sky Wilshaw. 2025. [leanprover-community/con-nf](#). Original-date: 2022-03-31T05:04:06Z.

Yuhuai Wu, Albert Q. Jiang, Wenda Li, Markus N. Rabe, Charles Staats, Mateja Jamnik, and Christian Szegedy. 2022. [Autoformalization with Large Language Models](#). *arXiv preprint*. ArXiv:2205.12615 [cs].

Huajian Xin, Z. Z. Ren, Junxiao Song, Zhihong Shao, Wanjia Zhao, Haocheng Wang, Bo Liu, Liyue Zhang, Xuan Lu, Qiushi Du, Wenjun Gao, Qihao Zhu, Dejian Yang, Zhibin Gou, Z. F. Wu, Fuli Luo, and Chong Ruan. 2024. [DeepSeek-Prover-V1.5: Harnessing Proof Assistant Feedback for Reinforcement Learning and Monte-Carlo Tree Search](#). *arXiv preprint*. ArXiv:2408.08152 [cs].

Kaiyu Yang, Aidan M. Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan Prenger, and Anima Anandkumar. 2023. [LeanDojo: Theorem Proving with Retrieval-Augmented Language Models](#). *arXiv preprint*. ArXiv:2306.15626 [cs, stat].

Huaiyuan Ying, Zijian Wu, Yihan Geng, Jiayu Wang, Dahua Lin, and Kai Chen. 2024a. [Lean Workbook: A large-scale Lean problem set formalized from natural language math problems](#). *arXiv preprint*. ArXiv:2406.03847 [cs].

Huaiyuan Ying, Shuo Zhang, Linyang Li, Zhejian Zhou, Yunfan Shao, Zhaoye Fei, Yichuan Ma, Jiawei Hong, Kuikun Liu, Ziyi Wang, Yudong Wang, Zijian Wu, Shuaibin Li, Fengzhe Zhou, Hongwei Liu, Songyang Zhang, Wenwei Zhang, Hang Yan, Xipeng Qiu, and 3 others. 2024b. [InternLM-Math: Open Math Large Language Models Toward Verifiable Reasoning](#). *arXiv preprint*. ArXiv:2402.06332 [cs].

Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. 2022. [MiniF2F: a cross-system benchmark for formal Olympiad-level mathematics](#). *arXiv preprint*. ArXiv:2109.00110 [cs] version: 2.

Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Texxygen: A benchmarking platform for text generation models](#). *Preprint*, arXiv:1802.01886.

A Appendix

A.1 Tactics used in BEq+

To develop BEq+, we checked existing tactics in Lean and Mathlib using the list provided at <https://github.com/haruhsa-enomoto/mathlib4-all-tactics/blob/main/all-tactics.md>. We provide a brief description of the tactics we ended up using in BEq+ using the official Mathlib documentation ([mathlib Community, 2020](#)):

- `exact?`: Searches environment for definitions or theorems that can solve the goal using `exact` with conditions resolved by `solve_by_elim`. While we are not directly interested in searching the library, this tactic is also capable of

handling a few transformations on the conclusion, but also variable/hypothesis assignment. For both `BEqL` and `BEq+`, if `exact?` succeeds we check that it is using the other formalization to close the goal. Otherwise it could lead to false positives.

- `apply` `e` tries to match the current goal against the conclusion of `e`'s type. If it succeeds, then the tactic returns as many subgoals as the number of premises that have not been fixed by type inference or type class resolution.
- `convert`: The `exact e` and `refine e` tactics require a term `e` whose type is definitionally equal to the goal. `convert e` is similar to `refine e`, but the type of `e` is not required to exactly match the goal. Instead, new goals are created for differences between the type of `e` and the goal using the same strategies as the `congr!` tactic. We use this tactic to try partial matching with the conclusion of the other theorem. In particular we use the `convert using n` variation, where `n` determines the matching depth. We vary `n` between 0 and 5.
- `tauto` breaks down assumptions of the form `_ ∧ _ → _`, `_ ∨ _`, `_ ↔ _` and `∃ _ , _` and splits a goal of the form `_ ∧ _ → _`, `_ ∨ _` or `∃ _ , _` until it can be discharged using `reflexivity` or `solve_by_elim`.
- `simp_all_arith!`: simplifies multiple times target and all (propositional) hypotheses using the other hypotheses. Additionally, it uses normalization by linear arithmetic.
- `noncomm_ring`: A tactic for simplifying identities in not-necessarily-commutative rings. It is pretty general and works on all types having a ring structure.
- `have : t := ...` adds the hypothesis `this : t` to the current goal.
- `apply_rules [l1, l2, ...]` tries to solve the main goal by iteratively applying the list of lemmas `[l1, l2, ...]` or by applying a local hypothesis. If `apply` generates new goals, `apply_rules` iteratively tries to solve those goals. `apply_rules` will also use `rfl`, `trivial`, `congrFun` and `congrArg`.

A.2 Annotation process

For all the evaluations presented in this paper, along with the curation of `ProofNetVerif` and `ProofNet#`, we adopted the following systematic annotation process for each formalized statement to determine its validity:

- Ill-typed statements and statements with

counter-examples found by `plausible` / `slim_check` were immediately marked as invalid.

- **Manual pass through all the statements.** Hypotheses and conclusions were carefully inspected for completeness and semantic correctness, correct handling of default Lean types, correct use of definitions/instances, ... A non-exhaustive list of issues that were looked for includes:
 - Missing/extra invalid hypotheses
 - Invalid implicit types (ex: $3/2 = 1$ in Lean because 3 and 2 are natural numbers by default)
 - Natural numbers (they start at 0 in Lean)
 - Parentheses, especially in quantified propositional formulas
 - Instances/definitions used: they must have the same semantic meaning as the ones in the informal statement to formalize
 - Semantically invalid hypotheses
 - Stronger / more restrictive hypotheses
 - In case of doubts, the prediction was marked as invalid
- (ProofNet# only) Running DeepSeek-Prover-V1.5 (Xin et al., 2024) to find proofs on the current iteration of ProofNet. We then manually analyze these proofs to check if flaws in the formalizations not detected by the manual pass have been exploited. This step helped finding a total of 2 additional formalization mistakes.

To reduce annotation mistakes and biases:

- All predictions were anonymized with respect to the model and method that produced them. This way, the annotator was not biased in favor of some models.
- All predictions related to a specific input problem were evaluated together, meaning that the annotator was annotating a few dozen predictions at the same time, giving them a good overview and leading to a more objective evaluation.
- Each sample has been annotated twice for model evaluations and the curation of ProofNetVerif, and three times for ProofNet#.
- ProofNet# and ProofNetVerif have been publicly shared with the Lean community several months before the reviewing process, including a website to explore ProofNet#. No errors have been reported during this period.

This annotation process was conducted by a sin-

gle annotator and required a total of ~ 200 human-hours for the results and datasets presented in this paper.

When curating ProofNet#, this annotation process led to the discovery of mistakes in 118 entries out of the 371 of the initial ProofNet benchmark.

A.3 Sampling-Based Method

The method is composed of three steps: (1) sampling, (2) type-check filtering, and then (3) selecting. We present them in this section.

A.3.1 Sampling

In our experiments, unless otherwise stated, we employ temperature sampling with $T = 0.7$ and generate $n = 50$ autoformalization attempts per informal statement. Depending on the models, we use the vLLM library (Kwon et al., 2023) or the OpenAI API to generate predictions.

Cleaning: Certain models often try to provide proofs after generating formal statements. Furthermore, we find that generated names for theorems sometimes clash with names in the Mathlib library. To avoid being considered invalid by the Lean type-checker, we trim proofs, substitute theorem names for dummy identifiers, and normalize whitespace when parsing the generated theorems. Additionally, the Lean proof assistant requires theorems to be accompanied by proofs. To address this, we append a safe dummy `sorry` proof to each theorem (which indicates to Lean that the proof will be provided later).

A.3.2 Filtering

We use Lean REPL (Morrison, 2023) to implement our filtering step. For any formal statement, if the statement is valid, REPL will return `declaration uses 'sorry'`, which means that the statement is well-typed and that we should provide an actual proof instead of `sorry`. Otherwise, the tool will return error messages explaining why the formal statement is ill-formed, which we use as an indicator to filter out such statements. de Moura et al. (2015) and Moura and Ullrich (2021) provide detailed insights into the inner workings of the Lean type system and type-checking.

A.3.3 Selection

In our selection process, we employ and compare four distinct heuristics to refine and choose the best outputs generated by the models: random selection, majority voting (Wang et al., 2023), Self-BLEU

(Zhu et al., 2018), and the symbolic equivalence method presented in Li et al. (2024b).

Random: As a baseline strategy, we randomly choose an output from the set of generated candidates.

Majority voting. (Wang et al., 2023): We aggregate multiple outputs and select the most frequently occurring candidate as the final choice, relying on consensus to mitigate the impact of any single erroneous output. Our cleaning process after the sampling step normalize the generated outputs, increasing the chance of exact string match between the predictions.

Self-BLEU. (Zhu et al., 2018): We evaluate the similarity of the generated outputs by calculating the BLEU score between all pairs of candidates. We then select the generated candidate with the highest aggregated BLEU score.

Symbolic Equivalence. (Li et al., 2024b): the core idea of this method is to compute equivalence classes of the generated predictions, using logical equivalence. A prediction from the largest equivalence class is then selected as final prediction. As the original work has been conducted in the Isabelle formal language (Nipkow et al., 2002), no implementation of this method is available in Lean. We therefore implemented a Lean version relying on our BEq+ method to compute the equivalence between statements.

A.4 Compute usage & cost

All experiments were run on 1xH100 for 7-8B models and 2xH100 for Llemma-34B. The most compute-intensive experiments, i.e., generating 1000 candidates per problem in ProofNet#, required less than 10 hours (wall-clock time). For cost reasons, we did not sample $n = 1000$ candidates with GPT-4o, and the largest GPT-4o results in the paper use $n = 50$, costing $\sim \$25$ per evaluation on ProofNet#.

A.5 BEq_L & BEq+ Performance on ProofNetVerif

In Table 7, we report the performance of our metrics at the instance level using binary metrics. While BEq+ outperforms BEq_L, we find that overall both struggle with a low recall. Empirically, we find that both BEq implementations show better recall on short statements than on long statements. Intuitively, long statements involve more logical clauses, generally making equivalence proving harder.

These results are not directly comparable to the ones presented in Liu et al. (2024) as we do not use the same samples to evaluate. For instance, regarding their results without LLM use, which correspond to BEq_L, they get a recall of 67.14% on their 200 sampled predictions, vs 30.9% on ProofNetVerif in our case. However, given the large improvement of BEq+ over BEq_L showcased in Table 7, it is very likely that BEq+ outperforms the original LLM-based BEq implementation. In fact, in their most compute-intensive setup, BEq achieves a recall of 72.86% on their samples, which is only slightly better than the 67.14% baseline. On ProofNetVerif, BEq+ outperforms BEq_L with a larger relative improvement for the recall: 48.3% vs. 30.9%.

A.6 BEq+ failure cases

A.6.1 False positives

A typical example of two non-semantically equivalent statements that are considered equivalent by BEq+:

```
theorem ground_truth (a b : ℤ) :  
  (ofInt a : GaussianInt) | ofInt b → a | b :=  
  sorry  
  
theorem prediction (a b : ℤ) (ha : a | b) : a |  
  (b : ℤ) := sorry
```

The main issue here lies from the fact that it is trivial to prove one formalization assuming the other. A generalization of this issue has been found by Simon Sorg:

```
theorem ground_truth (n : Nat) : n + n = 2 * n  
  := sorry  
  
theorem prediction (p : Prop) (h : p) : p :=  
  sorry
```

A fix with minimal performance loss is available in our [implementation](#) and is now the default.

A.6.2 False negatives

False negatives with BEq+ are caused by the relatively weak power of the static proof search we use. A typical example is:

```
theorem ground_truth : Infinite {p : Nat.Primes  
  // p ≡ -1 [ZMOD 6]} :=  
  
theorem prediction : Set.Infinite {p : ℙ | Nat.  
  Prime p ∧ p % 6 = 5} :=
```

Proving the equivalence between these two statements require a consequential work as it requires to (1) prove that $p \% 6 = 5$ is equivalent to $p \equiv -1 [ZMOD 6]$, (2) prove that the set of natural numbers is equivalent to the set of natural numbers with the

Table 7: Binary performance, in percentage, of BEq_L and $\text{BEq}+$ metrics on ProofNetVerif. The evaluation is performed on 3 splits based on the reference formalizations length. These splits contain roughly 1250 entries each.

Reference formalization length	Binary Metric	BEq_L	$\text{BEq}+$
Less than 115 characters	Precision	100.0	97.1
	Recall	39.9	62.5
	F1 Score	57.0	76.1
Between 115 and 165 characters	Precision	100.0	99.2
	Recall	28.2	41.0
	F1 Score	44.0	58.0
More than 165 characters	Precision	100.0	100.0
	Recall	17.2	29.6
	F1 Score	29.3	45.6
All	Precision	100.0	98.0
	Recall	30.9	48.3
	F1 Score	47.2	64.7

prime predicate, and (3) unstructuring the two sets to prove the equivalence between the two quantified statements. This is currently out of reach of our $\text{BEq}+$ implementation and solving these cases will likely require using powerful general-purpose proof search methods. Finding a way to develop such methods while keeping efficiency and alignment with human judgments is an interesting direction for future work.

A.7 $\text{BEq}+$ pass@n Additional Results

In Figure 5, we conduct a study on $\text{BEq}+$ pass@n with 4 models: Llama3 8B, Llemma 7B, Llemma 34B, and GPT-4o on both ProofNet# and RLM25. For ProofNet#, we use 12-shot prompting. For RLM25, we prompt models with in-file context with proofs removed as described in subsection 6.3. Llemma 34B model has not been run on the RLM25 benchmark.

A.8 ProofNet#: Baseline study

We report a baseline performance in Table 8. We evaluated the models described in subsection 6.1 using greedy decoding, coupled with either 12-shot learning or a fine-tuning on either MMA (Jiang et al., 2023) or Lean Workbook (Ying et al., 2024a).

A.9 Detailed Results of Sampling-Based Methods on ProofNet#

We present detailed results about the use of sampling-based methods on different models and autoformalization methods in Table 9.

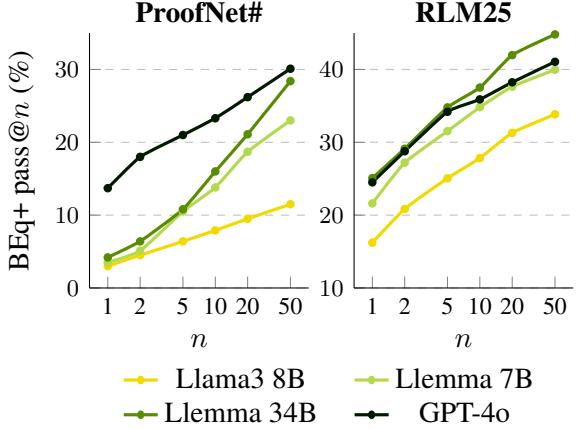


Figure 5: **BEq+ pass@n scaling trends with respect to the number of samples** on ProofNet# validation split and RLM25 using in-context learning prompting. We vary the number of candidate samples from $n = 1$ to 50.

In Figure 6, we report results on the ProofNet# test dataset by supplementing tested models using our self-consistency method described in subsection 5.2. Overall, we observe a consistent and significant improvement over the greedy baseline across all selection methods and all models evaluated. Interestingly, even using random selection over filtered generated statements is enough to outperform the greedy decoding baseline substantially. This demonstrates the practical efficacy of extensively leveraging the type-checker in the context of statement autoformalization. We find that the overall best strategy is to use type-check filtering

Table 8: **Baseline performance on ProofNet# using greedy decoding.** Except for Codex, which has been evaluated on Lean 3 in [Azerbayev et al. \(2023a\)](#) (indicated with an asterisk *, only the results on the test split are available), all models are evaluated on Lean 4.

Method	Model	Validation				Test			
		Type-Check	Accuracy \uparrow	BEq $_L\uparrow$	BEq+ \uparrow	Type-Check	Accuracy \uparrow	BEq $_L\uparrow$	BEq+ \uparrow
12-shot	Codex	-	-	-	-	23.7*	13.4*	-	-
Prompt retrieval	Codex	-	-	-	-	45.2*	16.1*	-	-
MMA	Llama3-8B	12.6	4.9	2.2	3.3	4.3	-	0.5	0.5
MMA	Llemma-7B	14.2	6.0	2.7	4.4	9.7	-	0.0	1.1
Lean Workbook ¹	Llemma-7B	39.5	-	5.4	7.0	39.3	-	5.4	6.5
12-shot	Llama3-8B	13.7	5.5	3.3	4.4	13.0	3.3	1.6	3.3
12-shot	Llemma-7B	26.8	8.7	3.3	6.0	29.9	10.9	5.4	6.5
12-shot	Llemma-34B	33.3	16.9	5.5	9.3	29.9	12.5	5.4	7.1
12-shot	GPT-4-turbo	24.6	19.7	8.7	12.6	27.7	22.8	13.0	16.8
12-shot	GPT-4o	33.3	26.2	10.9	16.4	42.9	31.0	13.6	18.5

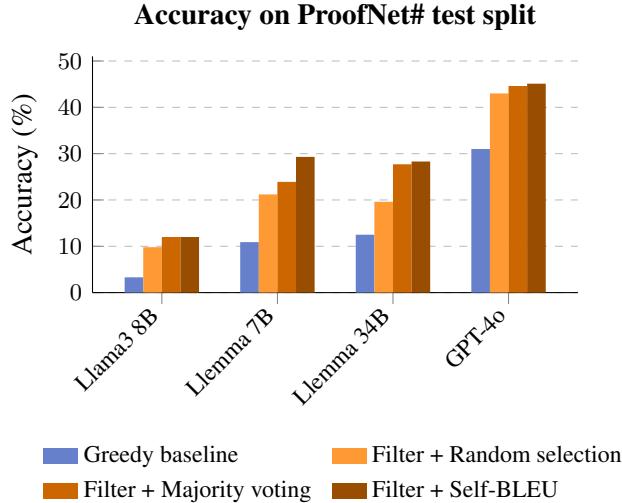


Figure 6: **Autoformalization accuracy on the ProofNet# test set.** All methods use 12-shot prompting in this figure. Detailed results are reported in [Table 8](#) and [Table 9](#).

and Self-BLEU for the selection step. One can notice that the absolute improvement in accuracy achieved by this method ranges from +8.7% to +18.4% over greedy decoding, with relative improvements between 1.4x and 3x. We also report results of sampling-based methods on other classical benchmarks, PDA and MiniF2F, in [subsection A.11](#).

A.10 Detailed Results on RLM25

We present detailed results about the use of sampling-based methods on different models and autoformalization methods in [Table 10](#).

A.11 Other Benchmarks

While we focused our work on the ProofNet# and RLM25 benchmarks, sampling-based methods are not specific to this benchmark, and we believe they should yield improvements out of the box on other statement autoformalization benchmarks.

In this section, we report results on two other benchmarks: Process-Driven Autoformalization (PDA) ([Lu et al., 2024a](#)), and MiniF2F ([Zheng et al., 2022](#)). The PDA benchmark from [Lu et al. \(2024a\)](#) has been designed to evaluate statement and proof autoformalization. Given the size of this benchmark, 3 test splits of 1000 theorems, and the fact that not all of these subsets contain reference formalizations, we sample 50 random problems for each test split and conduct a manual evaluation. We report our results in [Table 11](#).

On all test splits, we observe that sampling-based methods largely improve over greedy decoding. We have found the real test split to be challenging as problems are sampled from Arithmo dataset ([Jindal, 2023](#)) without solutions, therefore requiring first solving the problem before providing an accurate formalization.

The authors of the PDA benchmark reported compilation results for statement and proof autoformalization at once. However, they do not report results for statement autoformalization alone and do not report accuracy results either. This lack of data in [Lu et al. \(2024a\)](#) means that we cannot compare directly to their results.

We report results on the MiniF2F benchmark in [Table 12](#). We find the Llemma 7B model fine-tuned on Lean Workbook ([Ying et al., 2024a](#)) to perform particularly well on it. However, since

Table 9: Evaluation results (in percentage) of sampling-based methods on ProofNet#. For all these results, for each informal statement in the benchmark, we sampled 50 formalization attempts per model and filtered type-checking ones before applying a selection method. We observe some performance differences between the two splits which are caused by the small size of the ProofNet# benchmark (2x 185 statements).

Model	Selection method	Validation				Test			
		Type-Check	Accuracy \uparrow	BEq $_L$ \uparrow	BEq+ \uparrow	Type-Check	Accuracy \uparrow	BEq $_L$ \uparrow	BEq+ \uparrow
MMA fine-tune									
Llama3-8B	Random	33.3	9.3	3.3	4.9	29.0	-	2.2	4.8
	Majority	33.3	8.7	3.3	4.4	29.0	-	1.1	4.8
	Self-BLEU	33.3	8.7	3.3	4.4	29.0	-	1.1	4.3
Llemma-7B	Random	61.2	9.3	4.4	5.5	52.2	-	3.2	5.9
	Majority	61.2	10.9	3.8	5.5	52.2	-	3.8	7.0
	Self-BLEU	61.2	13.7	4.9	6.6	52.2	-	5.4	8.1
Lean Workbook fine-tune									
Llemma-7B	Random	86.0	-	7.6	9.2	86.6	-	7.0	8.6
	Majority	86.0	-	9.2	10.8	86.6	-	5.9	8.6
	Self-BLEU	86.0	-	9.7	12.4	86.6	-	6.5	10.2
12-shot									
Llama3-8B	Random	42.1	9.8	4.4	6.6	45.7	13.6	4.9	10.3
	Majority	42.1	12.0	4.9	7.1	45.7	14.7	4.9	10.9
	Self-BLEU	42.1	12.6	6.0	8.2	45.7	12.0	4.9	9.2
Llemma-7B	Random	84.7	16.9	4.9	7.7	88.6	21.2	9.8	11.4
	Majority	84.7	23.0	8.2	9.8	88.6	23.9	10.3	12.0
	Self-BLEU	84.7	23.5	7.7	11.5	88.6	29.3	11.4	17.9
Llemma-34B	Random	89.6	21.3	4.9	9.8	84.2	19.6	5.4	11.4
	Majority	89.6	27.3	8.7	12.6	84.2	27.7	10.9	14.1
	Self-BLEU	89.6	29.5	8.7	13.1	84.2	28.3	9.8	14.7
GPT-4o	Random	65.6	43.2	16.4	23.5	70.1	42.9	15.8	22.8
	Majority	65.6	45.4	15.8	21.3	70.1	44.6	15.8	22.8
	Self-BLEU	65.6	44.8	15.3	21.3	70.1	45.1	16.3	23.4

the Lean Workbook dataset has been synthetically generated by a model finetuned on the MiniF2F and ProofNet benchmarks, some data leakage might have happened.

A.12 Low-correction Effort Formalizations

One goal of autoformalization is the development of AI-assisted tools for formalization. In this setting, producing close-to-correct formal statements can already help users by providing hints and potential directions. Using the same setup as in the previous section, we report our results on the ProofNet# test split in Figure 6. Note that in this setup, contrary to Jiang et al. (2023), we are considering only well-typed statements.

We find that, by using sampling-based methods, open-source models Llemma-7B and Llemma-34B can autoformalize $\sim 50\%$ of the mathematical statements from the ProofNet# test benchmark in a *close-to-correct* way⁴. This makes these open models good fit for local autoformalization assistant, especially Llemma-7B for its relatively small size.

⁴We define *close-to-correct* formalizations as those with one slightly diverging hypothesis or conclusion, typically fixable in a matter of seconds.

In this section, we present several examples of autoformalizations on ProofNet# validation split that are evaluated as incorrect yet fixable with low effort. Evaluation results on ProofNet# test split are presented in Table 13.

Low-correction effort examples
Rudin exercise_1_1a
Informal statement: If r is rational ($r \neq 0$) and x is irrational, prove that $r + x$ is irrational.
Generated formalization:
<code>theorem exercise_1_1a (hr : r ≠ 0) (hi : Irrational x) : Irrational (r + x) :=</code>
Issue: r has not been declared as a rational number.
Herstein exercise_4_3_25
Informal statement: Let R be the ring of 2×2 matrices over the real numbers; suppose that I is an ideal of R . Show that $I = (0)$ or $I = R$.
Generated formalization:
<code>theorem exercise_4_3_25 {R : Type*} [</code>

Table 10: Results on RLM25 for different methods and models. In this setup, all models are prompted with in-file context and proofs are removed. Greedy decoding is used for generation.

Model	Method	Type-Check	BEq _L	BEq+
Llama3 8B	No fine-tuning	37.19	<u>18.90</u>	20.29
	MMA	41.73	17.55	19.24
Llemma 7B	No fine-tuning	55.96	22.49	<u>24.16</u>
	MMA	51.63	21.37	22.63
	Lean Workbook	53.89	21.84	23.55
GPT-4o	No fine-tuning	51.37	<u>21.18</u>	24.56

Table 11: Evaluation results (in percentage) of sampling-based methods on the **Process-Driven Autoformalization** benchmark. We evaluate on 50 samples on each of the "Basic", "Random", and "Real" splits of this benchmark. Greedy decoding is used for methods without sampling, i.e., when filtering is not mentioned. For sampling-based methods, we sample $n = 50$ predictions with temperature $T = 0.7$. We separate greedy decoding methods from sampling-based methods by a gray line for each model.

Model	Method	Basic		Random		Real	
		Type-Check	Accuracy↑	Type-Check	Accuracy↑	Type-Check	Accuracy↑
Llama3-8B	12-shot	18.0	16.0	20.0	16.0	20.0	4.0
	12-shot + Filter + Self-BLEU	48.0	28.0	46.0	32.0	72.0	8.0
Llemma-7B	12-shot	20.0	14.0	30.0	26.0	62.0	0.0
	12-shot + Filter + Self-BLEU	76.0	58.0	72.0	54.0	100.0	8.0
GPT-4o	12-shot	30.0	28.0	42.0	38.0	4.0	0.0
	12-shot + Filter + Self-BLEU	64.0	54.0	66.0	56.0	26.0	12.0

```
CommRing R] (I : Ideal (Matrix (Fin 2) (Fin 2) ℝ)) : I = ⊥ ∨ I = ⊤ :=
```

Issue: Superfluous declaration of `CommRing R`.

A.13 Ablation Study: Detailed Results

We present detailed results of our ablation study on the type-check filtering step in Table 14.

A.14 Sampling Scaling

Detailed results about our sampling study can be found in Table 15.

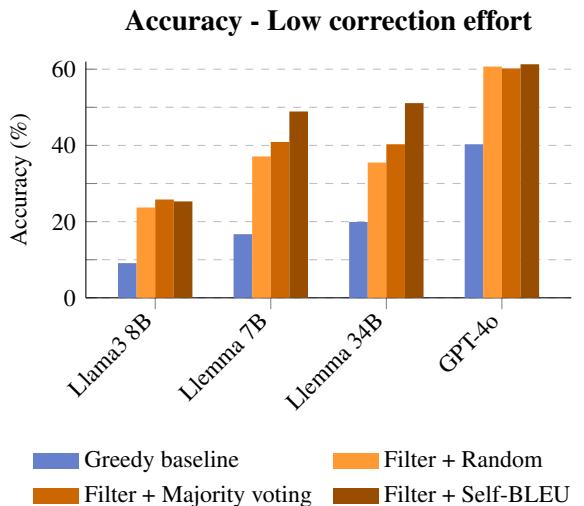


Figure 7: **Proportion of formalized statements evaluated as correct or as fixable with a low amount of effort (i.e., *close-to-correct*) on the ProofNet# test set.** All models are prompted with 12-shot examples. Detailed results are reported in Table 13.

Table 12: Performance of different methods on **MiniF2F**. Greedy decoding is used for methods without sampling, i.e., when filtering is not mentioned. For sampling-based methods, we sample $n = 50$ predictions with temperature $T = 0.7$. We separate greedy decoding methods from sampling-based methods by a gray line for each model. ¹Lean Workbook dataset has been curated with a model trained on the MiniF2F benchmark, thus data leakage concerns apply.

Model	Method	Type-Check	$\text{BEq}_L \uparrow$	$\text{BEq} \uparrow$
Llama3 8B	12-shot	45.9	6.8	14.6
	MMA	36.7	3.1	8.8
	12-shot + Filter + Self-BLEU	93.0	10.3	21.3
Llemma 7B	12-shot	67.0	7.6	16.0
	MMA	32.2	2.1	5.3
	Lean Workbook ¹	89.6	14.3	28.5
	12-shot + Filter + Self-BLEU	99.8	10.3	21.3
	Lean Workbook + Filter + Self-BLEU ¹	99.0	14.1	28.1
GPT-4o	12-shot	24.4	8.0	13.5

Table 13: Models performance (in percentage) on ProofNet# test split when accounting for formalizations that can be **corrected with a low amount of efforts**.

Model	Method	Accuracy \uparrow
Llama3-8B	Greedy	9.1
	Filter + Random	23.7
	Filter + Majority	25.8
	Filter + Self-BLEU	25.3
Llemma-7B	Greedy	16.7
	Filter + Random	37.1
	Filter + Majority	40.9
	Filter + Self-BLEU	48.9
Llemma-34B	Greedy	19.9
	Filter + Random	35.5
	Filter + Majority	40.3
	Filter + Self-BLEU	51.1
GPT-4o	Greedy	40.3
	Filter + Random	60.7
	Filter + Majority	60.2
	Filter + Self-BLEU	61.3

Table 14: Models performance (in percentage) on ProofNet# validation split removing different aspects of sampling-based methods. We also report Greedy baseline and the Filter + Self-BLEU results as reference.

Model	Method	Type-Check	Accuracy \uparrow	BEq $_L\uparrow$	BEq+ \uparrow
Llama3-8B	Greedy	13.7	5.5	3.3	4.4
	No filter + Random	9.8	4.4	2.7	2.7
	No filter + Majority	13.1	5.5	3.3	3.8
	No filter + Self-BLEU	14.2	6.0	2.7	3.8
	Filter + Random	42.1	9.8	4.4	6.6
	Filter + Self-BLEU	42.1	12.6	6.0	8.2
Llemma-7B	Greedy	26.8	8.7	3.3	6.0
	No filter + Random	25.7	6.0	3.3	3.8
	No filter + Majority	25.1	10.9	4.4	4.9
	No filter + Self-BLEU	32.2	14.2	6.6	9.3
	Filter + Random	84.7	16.9	4.9	7.7
	Filter + Self-BLEU	84.7	23.5	7.7	11.5
Llemma-34B	Greedy	33.3	16.9	5.5	9.3
	No filter + Random	24.6	9.3	2.7	4.4
	No filter + Majority	24.6	10.9	4.9	7.1
	No filter + Self-BLEU	32.8	14.2	3.8	6.0
	Filter + Random	89.6	21.3	4.9	9.8
	Filter + Self-BLEU	89.6	29.5	8.7	13.1
GPT-4o	Greedy	33.3	26.2	10.9	16.4
	No filter + Random	33.3	25.7	9.8	14.8
	No filter + Majority	35.5	29.5	12.0	16.4
	No filter + Self-BLEU	36.1	28.4	12.0	16.9
	Filter + Random	65.6	43.2	16.4	23.5
	Filter + Self-BLEU	65.6	44.8	15.3	21.3

Table 15: Evaluation results (in percentage) of sampling-based methods on ProofNet# validation split for different numbers of formalizations sampled during the sampling phase (represented by the number n in this table). We used a 12-shot prompt with the filter+Self-BLEU variant and a temperature of 0.7.

Model	Type-Check				Accuracy \uparrow				BEq $_L\uparrow$				BEq+ \uparrow			
	n=1	n=5	n=20	n=50	n=1	n=5	n=20	n=50	n=1	n=5	n=20	n=50	n=1	n=5	n=20	n=50
Llama3-8B	9.8	22.4	32.8	42.1	4.4	7.1	10.4	12.6	2.7	3.3	4.4	6.0	2.7	4.9	6.0	8.2
Llemma-7B	25.7	50.3	71.0	84.7	6.0	10.9	19.7	23.5	3.3	3.3	7.1	7.7	3.8	6.0	10.4	11.5
Llemma-34B	24.6	55.2	78.7	89.6	9.3	14.8	25.7	29.5	2.7	5.5	8.2	8.7	4.4	8.7	12.0	13.1
GPT-4o	33.3	47.0	55.7	65.6	25.7	34.4	38.8	44.8	9.8	14.2	13.7	15.3	14.8	20.2	19.7	21.3

A.15 12-shot Prompt

Note: We translated the 12-shot prompt from ProofNet# to Lean 4, with as minimal changes as possible, for the accuracy comparison with previous results to be as fair as possible. In particular, we did not remove/change the statements leaked from the benchmark and did not correct potential formalization mistakes in this prompt to make our results comparable with the results in [Azerbayev et al. \(2023a\)](#).

12-shot examples

Natural language version:

Let P be a p -subgroup of G . Then P is contained in a Sylow p -subgroup of G .

Translate the natural language version to a Lean 4 version:

```
theorem exists_le_sylow [Group G] {P : Subgroup G} (hP : IsPGroup p P) : ∃ Q : Sylow p G, P ≤ Q :=
```

Natural language version:

Let E and F be complex normed spaces and let $f : E \rightarrow F$. If f is differentiable and bounded, then f is constant. Translate the natural language version to a Lean 4 version:

```
theorem exists_eq_const_of_bounded {E : Type u} [NormedAddCommGroup E] [NormedSpace C E] {F : Type v} [NormedAddCommGroup F] [NormedSpace C F] {f : E → F} (hf : Differentiable C f) (hb : IsBounded (range f)) : ∃ c, f = const E c :=
```

Natural language version:

Let X be a topological space; let A be a subset of X . Suppose that for each $x \in A$ there is an open set U containing x such that $U \subset A$. Then A is open in X .

Translate the natural language version to a Lean 4 version:

```
theorem subset_of_open_subset_is_open (X : Type*) [TopologicalSpace X] (A : Set X) (hA : ∀ x ∈ A, ∃ U : Set X, IsOpen U ∧ x ∈ U ∧ U ⊆ A) : IsOpen A :=
```

Natural language version:

Two multiplicative functions $f, g : \mathbb{N} \rightarrow R$ are equal if and only if $f(p^i) = g(p^i)$ for all primes p .

Translate the natural language version to a Lean 4 version:

```
theorem eq_iff_eq_on_prime_powers [CommMonoidWithZero R] (f : ArithmeticFunction R) (hf : f.IsMultiplicative) (g : ArithmeticFunction R) (hg : g.IsMultiplicative) : f = g ↔ ∀ p i : ℕ, Nat.Prime p → f (p ^ i) = g (p ^ i) :=
```

Natural language version:

If z_1, \dots, z_n are complex, then $|z_1 + z_2 + \dots + z_n| \leq |z_1| + |z_2| + \dots + |z_n|$.

Translate the natural language version to a Lean 4 version:

```
theorem abs_sum_leq_sum_abs (n : ℕ) (f : ℕ → ℂ) : abs (Σ i in Finset.range n, f i) ≤ Σ i in Finset.range n, abs (f i) :=
```

Natural language version:

If x and y are in \mathbb{R}^n , then $|x + y|^2 + |x - y|^2 = 2|x|^2 + 2|y|^2$.

Translate the natural language version to a Lean 4 version:

```
theorem sum_add_square_sub_square_eq_sum_square (n : ℕ) (x y : EuclideanSpace ℝ (Fin n)) : ∥x + y∥^2 + ∥x - y∥^2 = 2*∥x∥^2 + 2*∥y∥^2 :=
```

Natural language version:

If x is an element of infinite order in G , prove that the elements $x^n, n \in \mathbb{Z}$ are all distinct.

Translate the natural language version to a Lean 4 version:

```
theorem distinct_powers_of_infinite_order_element (G : Type*) [Group G] (x : G) (hx_inf : ∀ n : ℕ, x ^ n ≠ 1) : ∀ m n : ℤ, m ≠ n → x ^ m ≠ x ^ n :=
```

Natural language version:

A set of vectors $\{v_i\}_{i \in I}$ orthogonal with respect to some bilinear form $B : V \times V \rightarrow K$ is linearly independent if for all $i \in I, B(v_i, v_i) \neq 0$.

Translate the natural language version to a Lean 4 version:

```
theorem linear_independent_of_is_Ortho {V K : Type*} [Field K] [AddCommGroup V] [Module K V] {n : Type*} {B : BilinForm K V} {v : n → V} (hv1 : B.iIsOrtho v) (hv2 : ∀ (i : n), ¬B.IsOrtho (v i) (v i)) : LinearIndependent K v :=
```

Natural language version:

Suppose that V is an n -dimensional vector space. Then for some set of vectors $\{v_i\}_{i=1}^k$, if $k > n$ then there exist scalars f_1, \dots, f_k such that $\sum_{i=1}^k f_i v_i = 0$.

Translate the natural language version to a Lean 4 version:

```
theorem
  exists_nontrivial_relation_sum_zero_of_
  dim_succ_lt_card {K V : Type*}
  [DivisionRing K] [AddCommGroup V] [Module K
  V] [FiniteDimensional K V]
  {t : Finset V} (h : FiniteDimensional.
  finrank K V + 1 < t.card) :
  ∃ (f : V → K), t.sum (λ (e : V) => f e · e)
  = 0 ∧ t.sum (λ (e : V) => f e) = 0
  ∧ ∃ (x : V) (H : x ∈ t), f x ≠ 0 :=
```

Natural language version:

A group is commutative if the quotient by the center is cyclic.

Translate the natural language version to a Lean 4 version:

```
theorem comm_group_of_cycle_center_quotient
  {G H : Type*} [Group G] [Group H]
  [IsCyclic H] (f : G →* H) (hf : f.ker ≤ (
  center G : Subgroup G)) :
  CommGroup G :=
```

Natural language version:

If H is a p -subgroup of G , then the index of H inside its normalizer is congruent modulo p to the index of H .

Translate the natural language version to a Lean 4 version:

```
theorem card_quotient_normalizer_modEq_
  card_quotient {G : Type*} [Group G] [
  Fintype G] {p : ℕ} {n : ℕ} [hp : Fact p
  .Prime]
  {H : Subgroup G} (hH : Fintype.card H = p ^
  n) :
  Fintype.card (normalizer H / Subgroup.comap
  ((normalizer H).subtype : normalizer H
  →* G) H) ≡
  Fintype.card (G / H) [MOD p] :=
```

Natural language version:

Suppose X, Y, Z are metric spaces, and Y is compact. Let f map X into Y , let g be a continuous one-to-one mapping of Y into Z , and put $h(x) = g(f(x))$ for $x \in X$. Prove that f is uniformly continuous if h is uniformly continuous.

Translate the natural language version to a Lean 4 version:

```
theorem uniform_continuous_of_continuous_
  injective_uniform_continuous_comp
  {X Y Z : Type*} [MetricSpace X] [MetricSpace
  Y] [MetricSpace Z]
  (hY : CompactSpace Y) (f : X → Y) (g : Y →
  Z) (hgc : Continuous g)
  (hgi : Function.Injective g)
  (h : UniformContinuous (g ∘ f)) :
  UniformContinuous f :=
```

A.16 Data Contamination

Data contamination is a serious issue in today's LLM benchmarks. In fact, large language models are trained on large-scale training data, thus, despite the filtering efforts, data leakage might happen. For the new dataset RLM25 we introduce, as stated in subsection 4.1, all projects selected for evaluation were made available after the knowledge cutoff dates of the evaluated models. In particular, Llemma 7B, performing almost on par with GPT-4o on this benchmark, is open-weight and has been released in August 2023, thus before the first commit of any of these projects.

ProofNet 3 was released in February 2023, and an unofficial port to Lean 4 has been publicly available since March 2024. Since the cutoff training dates for all models used in these experiments are before March 2024, Lean 4 data contamination due to training is not possible. However, it remains theoretically possible that some models were trained on the Lean 3 version and weakly generalized to Lean 4. Such data leakage for the Llemma models family (Azerbayev et al., 2023b) seems unlikely as the authors claim they have specifically excluded ProofNet from their training data.

For our data contamination study, we use an unofficial Lean 4 port (Vishwakarma, 2024) of the ProofNet benchmark made by an independent research team. This port shows minimal differences from the original Lean 3 ProofNet benchmark, preserving the order of hypotheses and terms. Upon analyzing the raw predictions of all models, we did not find any exact matches with the Lean 4 ground truths. This is primarily because the theorems in the benchmark follow an exercise_number naming scheme, which the models do not produce. Consequently, we employed fuzzy matching for our data contamination checks. This involved normalizing whitespaces and removing comments and theorem names. We found a maximum of 2.2% matches (4 statements out of 185/186) for each model independently on the validation split, including the

2 statements leaked by the prompt. Given that the space of correct formal statements is heavily constrained, this hit rate is quite reasonable. Below, we provide a list of all unique hits found across all models and experiments. Most of these hits are very short and almost unavoidable. Considering these results, it seems unlikely that significant data leakage occurred during the training of these models.

Nonetheless, during our data contamination study, we found that 4 examples from the 12-shot prompt in [Azerbayev et al. \(2023a\)](#), which we intended to compare to, were also present in the benchmark (2 in the validation set and 2 in the test set). Fortunately, this affects the results only negligibly (at most $\sim 1.1\%$). We report all our results with these statements removed.

List of all the hits found (using fuzzy matching) across all our experiments on the ProofNet validation split

Munkres|exercise_29_1:

```
theorem exercise_29_1 :  $\neg$ 
LocallyCompactSpace  $\mathbb{Q}$  :=
```

Dummit-Foote|exercise_1_1_22a:

```
theorem exercise_1_1_22a {G : Type*} [Group
G] (x g : G) :
orderOf x = orderOf (g-1 * x * g) :=
```

Herstein|exercise_2_1_27:

```
theorem exercise_2_1_27 {G : Type*} [Group G
]
[Fintype G] :  $\exists$  (m :  $\mathbb{N}$ ),  $\forall$  (a : G), am
= 1 :=
```

Munkres|exercise_17_4:

```
theorem exercise_17_4 {X : Type*} [
TopologicalSpace X]
(U A : Set X) (hU : IsOpen U) (hA :
IsClosed A) :
IsOpen (U \ A)  $\wedge$  IsClosed (A \ U) :=
```

Herstein|exercise_5_5_2:

```
theorem exercise_5_5_2 : Irreducible (X3 -
3*X - 1 : Polynomial  $\mathbb{Q}$ ) :=
```

Munkres|exercise_32_3:

```
theorem exercise_32_3 {X : Type*} [
TopologicalSpace X]
```

```
(hX : LocallyCompactSpace X) (hX' :
T2Space X) :
RegularSpace X :=
```

Herstein|exercise_4_3_25:

```
theorem exercise_4_3_25 (I : Ideal (Matrix (
Fin 2) (Fin 2)  $\mathbb{R}$ )) :
I =  $\perp$   $\vee$  I =  $\top$  :=
```

Statements from the validation split leaked by the ProofNet prompt

Munkres|exercise_13_1

```
theorem subset_of_open_subset_is_open (X :
Type*) [TopologicalSpace X]
(A : Set X) (hA :  $\forall$  x  $\in$  A,  $\exists$  U : Set X,
IsOpen U  $\wedge$  x  $\in$  U  $\wedge$  U  $\subseteq$  A) :
IsOpen A :=
```

Dummit-Foote|exercise_1_1_34

```
theorem
distinct_powers_of_infinite_order_element
(G : Type*) [Group G] (x : G)
(hx_inf :  $\forall$  n :  $\mathbb{N}$ , xn  $\neq$  1) :
 $\forall$  m n :  $\mathbb{Z}$ , m  $\neq$  n  $\rightarrow$  xm  $\neq$  xn :=
```

A.17 All results on ProofNet#

We report in the table below metric results for all autoformalization methods and models on which we conducted manual evaluation. Such manual evaluation has been conducted on both ProofNet# validation and test splits.

Model	Strategy	Split	Type Check	Accuracy	BEq_L	$BEq+$	Hyp. Rej.
ensemble-12shot-top50	top50 type-checked majority voting	test	96.2	44.6	17.4	23.4	3.8
ensemble-12shot-top50	top50 type-checked random	test	96.2	33.2	9.8	14.7	2.2
ensemble-12shot-top50	top50 type-checked self consistent	test	96.2	48.4	17.9	26.6	3.3
gpt-4-turbo-2024-04-09	greedy type-checked	test	27.7	22.8	13.0	16.8	1.1
gpt-4-turbo-2024-04-09	greedy type-checked	valid	24.6	19.7	8.7	12.6	0.0
gpt-4o-2024-05-13	greedy type-checked	test	42.9	31.0	13.6	18.5	1.6
gpt-4o-2024-05-13	greedy type-checked	valid	33.3	26.2	10.9	16.4	0.0
gpt-4o-2024-05-13-top50	top50 type-checked majority voting	test	70.1	44.6	15.8	22.8	2.2
gpt-4o-2024-05-13-top50	top50 type-checked random	test	70.1	42.9	15.8	22.8	1.1
gpt-4o-2024-05-13-top50	top50 type-checked self consistent	test	70.1	45.1	16.3	23.4	2.2
gpt-4o-2024-05-13-top50	top20 type-checked self consistent	valid	55.7	38.8	13.7	19.7	0.5
gpt-4o-2024-05-13-top50	top50 majority voting type-checked	valid	35.5	29.5	12.0	16.4	0.0
gpt-4o-2024-05-13-top50	top50 random type-checked	valid	33.3	25.7	9.8	14.8	0.0
gpt-4o-2024-05-13-top50	top50 self consistent type-checked	valid	36.1	28.4	12.0	16.9	0.0
gpt-4o-2024-05-13-top50	top50 type-checked majority voting	valid	65.6	45.4	15.8	21.3	0.0
gpt-4o-2024-05-13-top50	top50 type-checked random	valid	65.6	43.2	16.4	23.5	0.0
gpt-4o-2024-05-13-top50	top50 type-checked self consistent	valid	65.6	44.8	15.3	21.3	0.0
gpt-4o-2024-05-13-top50	top5 type-checked self consistent	valid	47.0	34.4	14.2	20.2	0.0
llama3-8b	greedy type-checked	test	13.0	3.3	1.6	3.3	1.1
llama3-8b	greedy type-checked	valid	13.7	5.5	3.3	4.4	0.5
llama3-8b-mma	valid	valid	12.6	4.9	2.2	3.3	2.2
llama3-8b-mma-top50	top50 type-checked most frequent	valid	33.3	8.7	3.3	4.4	2.2
llama3-8b-mma-top50	top50 type-checked random	valid	33.3	9.3	3.3	4.9	2.7
llama3-8b-mma-top50	top50 type-checked self consistent	valid	33.3	8.7	3.3	4.4	2.7
llama3-8b-top50	top50 type-checked majority voting	test	45.7	14.7	4.9	10.9	2.2
llama3-8b-top50	top50 type-checked random	test	45.7	13.6	4.9	10.3	1.1
llama3-8b-top50	top50 type-checked self consistent	test	45.7	12.0	4.9	9.2	1.6
llama3-8b-top50	top20 type-checked self consistent	valid	32.8	10.4	4.4	6.0	0.5
llama3-8b-top50	top50 majority voting type-checked	valid	13.1	5.5	3.3	3.8	0.5
llama3-8b-top50	top50 random type-checked	valid	9.8	4.4	2.7	2.7	0.5
llama3-8b-top50	top50 self consistent type-checked	valid	14.2	6.0	2.7	3.8	0.5
llama3-8b-top50	top50 type-checked most frequent	valid	42.1	12.0	4.9	7.1	1.1
llama3-8b-top50	top50 type-checked random	valid	42.1	9.8	4.4	6.6	1.1
llama3-8b-top50	top50 type-checked self consistent	valid	42.1	12.6	6.0	8.2	1.1
llama3-8b-top50	top5 type-checked self consistent	valid	22.4	7.1	3.3	4.9	1.1
llemma-34b	greedy type-checked	test	29.9	12.5	5.4	7.1	1.1
llemma-34b	greedy type-checked	valid	33.3	16.9	5.5	9.3	0.0
llemma-34b-top50	top50 type-checked majority voting	test	84.2	27.7	10.9	14.1	3.3
llemma-34b-top50	top50 type-checked random	test	84.2	19.6	5.4	11.4	2.2
llemma-34b-top50	top50 type-checked self consistent	test	84.2	28.3	9.8	14.7	2.2
llemma-34b-top50	top20 type-checked self consistent	valid	78.7	25.7	8.2	12.0	1.1
llemma-34b-top50	top50 majority voting type-checked	valid	24.6	10.9	4.9	7.1	0.0
llemma-34b-top50	top50 random type-checked	valid	24.6	9.3	2.7	4.4	0.5
llemma-34b-top50	top50 self consistent type-checked	valid	32.8	14.2	3.8	6.0	0.0
llemma-34b-top50	top50 type-checked majority voting	valid	89.6	27.3	8.7	12.6	1.1
llemma-34b-top50	top50 type-checked random	valid	89.6	21.3	4.9	9.8	2.2
llemma-34b-top50	top50 type-checked self consistent	valid	89.6	29.5	8.7	13.1	1.6
llemma-34b-top50	top5 type-checked self consistent	valid	55.2	14.8	5.5	8.7	2.2
llemma-7b	greedy type-checked	test	29.9	10.9	5.4	6.5	2.2
llemma-7b	greedy type-checked	valid	26.8	8.7	3.3	6.0	0.0
llemma-7b-mma	valid	valid	14.2	6.0	2.7	4.4	0.0
llemma-7b-mma-top50	top50 type-checked most frequent	valid	61.2	10.9	3.8	5.5	1.6
llemma-7b-mma-top50	top50 type-checked random	valid	61.2	9.3	4.4	5.5	2.7
llemma-7b-mma-top50	top50 type-checked self consistent	valid	61.2	13.7	4.9	6.6	2.7
llemma-7b-top50	top50 type-checked majority voting	test	88.6	23.9	10.3	12.0	2.7
llemma-7b-top50	top50 type-checked random	test	88.6	21.2	9.8	11.4	4.9
llemma-7b-top50	top50 type-checked self consistent	test	88.6	29.3	11.4	17.9	4.3
llemma-7b-top50	top20 type-checked self consistent	valid	71.0	19.7	7.1	10.4	1.6
llemma-7b-top50	top50 majority voting type-checked	valid	25.1	10.9	4.4	4.9	0.5
llemma-7b-top50	top50 random type-checked	valid	25.7	6.0	3.3	3.8	1.1
llemma-7b-top50	top50 self consistent type-checked	valid	32.2	14.2	6.6	9.3	0.5
llemma-7b-top50	top50 type-checked most frequent	valid	84.7	23.0	8.2	9.8	2.2
llemma-7b-top50	top50 type-checked random	valid	84.7	16.9	4.9	7.7	2.7
llemma-7b-top50	top50 type-checked self consistent	valid	84.7	23.5	7.7	11.5	1.6
llemma-7b-top50	top5 type-checked self consistent	valid	50.3	10.9	3.3	6.0	1.6