

# DON'T TRUST: VERIFY – GROUNDING LLM QUANTITATIVE REASONING WITH AUTOFORMALIZATION

**Jin Peng Zhou\***  
Cornell University

**Charles Staats†**

**Wenda Li**  
University of Edinburgh

**Christian Szegedy†**  
xAI

**Kilian Q. Weinberger**  
Cornell University

**Yuhuai Wu†**  
xAI

## ABSTRACT

Large language models (LLM), such as Google’s Minerva and OpenAI’s GPT families, are becoming increasingly capable of solving mathematical quantitative reasoning problems. However, they still make unjustified logical and computational errors in their reasoning steps and answers. In this paper, we leverage the fact that if the training corpus of LLMs contained sufficiently many examples of formal mathematics (e.g. in Isabelle, a formal theorem proving environment), they can be prompted to translate i.e. autoformalize informal mathematical statements into formal Isabelle code — which can be verified automatically for internal consistency. This provides a mechanism to automatically reject solutions whose formalized versions are inconsistent within themselves or with the formalized problem statement. We evaluate our method on GSM8K, MATH and MultiArith datasets and demonstrate that our approach provides a consistently better heuristic than vanilla majority voting — the previously best method to identify correct answers, by more than 12% on GSM8K. In our experiments it improves results consistently across all datasets and LLM model sizes. The code can be found at <https://github.com/jinpz/dtv>.

## 1 INTRODUCTION

Recently, language models (Devlin et al., 2018; Brown et al., 2020; Chowdhery et al., 2022) have advanced significantly in many natural language processing tasks such as machine translation, question answering, summarization, etc. More recent large language models (LLMs) such as Minerva (Lewkowycz et al., 2022), GPT3.5 (OpenAI) and GPT4 (OpenAI, 2023) have also become increasingly capable of solving quantitative reasoning problems, ranging from middle school math word problems (Cobbe et al., 2021) to challenging high school mathematical competition problems (Hendrycks et al., 2021). By training or finetuning the model on high-quality natural language mathematical and scientific text, these LLMs can generate self-contained step-by-step solutions to quantitative reasoning problems without relying on external tools. However, just like human beings, the solutions LLMs generate are prone to simple calculation errors and unjustified logical leaps.

Following Wang et al. (2022); Lewkowycz et al. (2022), one can sample many proposed solutions, extract the final answer from each, and select the most common answer. While aggregating answers like this improves performance at the problem level, the most common answer is sometimes wrong. Ideally, we would like a better heuristic to identify the correct answer. In fact, there are well-known techniques for computers to verify mathematical reasoning using formalization (Wiedijk, 2008). Those methods involve translating the problem and sometimes the solution into a formal language. Unfortunately this translation is difficult and time-intensive enough that formalization methods are less frequently leveraged.

Recently it was discovered that large language models, with few-shot prompting, can automatically formalize natural mathematical language into formal languages (Agrawal et al., 2022; Azerbayev

\*Work done while interning at Google Research. Correspondence to: [jpzhou@cs.cornell.edu](mailto:jpzhou@cs.cornell.edu).

†Work done while at Google Research.

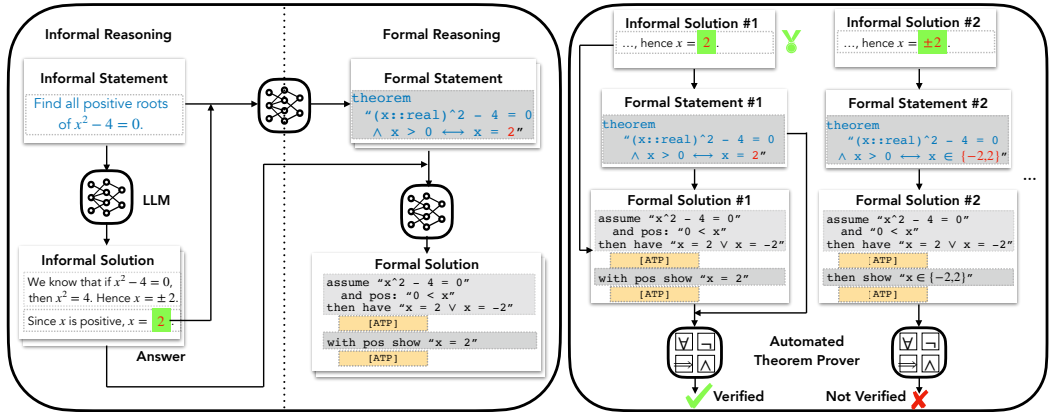


Figure 1: A pictorial illustration of **Don't Trust: Verify**. **Left**: starting from an informal statement, an informal solution is generated by a large language model. The informal statement and solution are then translated into their formal counterparts. Multiple informal solutions for each problem are generated with temperature sampling. **Right**: An automated theorem prover in the formal environment is used to verify formal solutions against formal statements step by step (indicated by [ATP]). The final answer is chosen using majority voting over only the verified solutions. In this example, Formal Solution #1 successfully proves Formal Statement #1. Formal Solution #2, however, fails to prove the "only if" direction of Formal Statement #2.

et al., 2023; Wu et al., 2022; Jiang et al., 2023; Wang et al., 2020). It is conjectured this capability arises because the training data for these LLMs contains sufficiently many examples of computer code and/or formal mathematics. While the translation is far from perfect and can fail on many complex statements, the correctness of the statements produced can be checked using formal theorem proving systems such as Isabelle (Nipkow et al., 2002) and Lean (de Moura et al., 2015).

In this paper, we demonstrate that in spite of its issues, autoformalization is already capable enough to identify correct answers for many quantitative reasoning problems. We call our method *Don't Trust: Verify* (DTV) and provide an overview in Figure 1. **Intuitively, instead of naively taking majority voting of all generated natural language solutions, we only aggregate solutions that can be verified with autoformalization.** To carry out verification, both a formal statement and solution are needed. Therefore, we attempt to translate the plain text problem and solution into formal language using large language models. However, because language models can erroneously translate incorrect statements into correct ones, we develop both **symbolic** and **neural filters** to improve the statement translation reliability. Additionally, even with a correct formal statement, a formal solution directly translated from informal solution can fail to prove it. In particular, correct informal solutions (whether generated by humans or LLMs) skip low-level steps that are necessary for formal reasoning. To this end, we instead generate a formal solution sketch following Wiedijk (2004); Jiang et al. (2023) and employ an automated theorem prover (ATP) to fill in the low-level gaps.

We evaluate DTV on GSM8K (Cobbe et al., 2021), three subsets of MATH (Hendrycks et al., 2021) following prior work (Zheng et al., 2021), and MultiArith (Roy & Roth, 2016) datasets. The results show that our method consistently outperforms vanilla majority voting (Wang et al., 2022; Lewkowycz et al., 2022), with a more than 12% improvement on GSM8K. We demonstrate that DTV improves performance across various model sizes and categories. Additionally, we provide case studies on our method identifying correct answers as well as informal solutions. Finally, we discuss the limitations of our approach inherited from LLM and theorem proving environments.

## 2 BACKGROUND AND RELATED WORK

**Informal quantitative reasoning with language models.** A number of large language models such as PaLM (Chowdhery et al., 2022), Minerva (Lewkowycz et al., 2022), GPT3.5 (OpenAI) and GPT4 (OpenAI, 2023) have demonstrated their impressive quantitative reasoning abilities by pretraining or finetuning on mathematical and science data. To improve informal reasoning

performance, chain-of-thought prompting (Wei et al., 2022) is typically used to encourage language models outputting intermediate reasoning steps before arriving at an answer. A diverse set of prompting methods have since been proposed for informal reasoning that tries to find better examples for prompting (Fu et al., 2023; Zhang et al., 2022) or better strategies for decoding from the model such as conditioning on references, multi-step decoding (Creswell et al., 2022; Zheng et al., 2023; Welleck et al., 2022; Khot et al., 2023). Cobbe et al. (2021) explores training informal verifiers to judge the correctness of reasoning. Additionally, to alleviate erroneous reasoning with language models, techniques that explore multi-sample consistency in the informal reasoning (Wang et al., 2022; Jung et al., 2022) have also been proposed. Our approach does not require training and is complementary to these methods since we rely on the consistency of a formal theorem proving system to identify correct reasoning paths and improve the reliability of large language model reasoning.

**Augmented language models.** Besides elucidating rationales via chain-of-thought prompting, recent research has also been devoted to augmenting LLMs with external tools such as web search engine (Nakano et al., 2021; Lazaridou et al., 2022; Schick et al., 2023), external memory retrieval (Shi et al., 2023; Izacard et al., 2022) and programming-based calculators (Chen et al., 2022; Gao et al., 2022; Imani et al., 2023) to bolster their downstream performance and reduce hallucination. Our work augments language models with a formal theorem proving environment that goes beyond simple arithmetic (Chen et al., 2022; Imani et al., 2023) and Boolean logic (Jung et al., 2022).

**Interactive theorem proving and autoformalization.** Modern interactive theorem provers such as Isabelle (Nipkow et al., 2002), Coq (Barras et al., 1997), and Lean (de Moura et al., 2015), provide an interactive environment to encode and mechanically verify mathematical proofs. Success stories with handcrafted proofs include the verification of industrial software systems (Klein et al., 2009) and research-level mathematics (Castelvecchi et al., 2021). Szegedy (2020) argues for automatically obtaining formal mathematics from their informal counterparts by translation i.e. autoformalization. Since then, researchers have shown the feasibility of autoformalization using neural networks in particular large language models (Agrawal et al., 2022; Azerbayev et al., 2023; Wu et al., 2022; Jiang et al., 2023; Wang et al., 2020). Our work which aims to ground LLM reasoning with formal theorem proving environments is distinct from Jiang et al. (2023) that improves theorem proving performance with LLM. DTV assumes a much more difficult but realistic setting where only natural language data is available. We are the first to demonstrate it is possible to automatically verify correct answer and informal solutions with autoformalization.

### 3 DON’T TRUST: VERIFY

We describe our approach *Don’t Trust: Verify* (DTV). Given a mathematical question and several solutions to it phrased in natural language, we assume the question has a well-defined answer and each solution contains an answer that may or may not be correct. Our goal is to understand which informal solutions are more likely to be correct. An overview of our method can be found in Figure 1.

#### 3.1 STATEMENT FORMALIZATION

We begin by finding a formal statement that corresponds to the description of the informal problem statement. In many mainstream formal theorem proving environments such as Isabelle (Nipkow et al., 2002), Lean (de Moura et al., 2015), and Coq (Barras et al., 1997), a formal statement needs to be an assertion rather than a question. For example, a formal statement cannot be in the form of *Is  $\pi$  irrational?* but rather *Show that  $\pi$  is irrational.* Since quantitative reasoning problems typically ask for an answer rather than a proof, we first extract the answer from each proposed informal solution. The formal statement is then generated conditioned on both the plain text statement and the extracted answer. Following Wu et al. (2022), we leverage the few-shot learning capability of large language models to generate formal representation of the informal statements. Specifically, we provide a few informal-formal statement translation pairs and prompt the language model to complete the subsequent problem statement formalization.

#### 3.2 SOLUTION FORMALIZATION AND VERIFICATION

The formalized statement itself does not tell us its correctness without a formal solution or proof. Because of this, we seek to generate a piece of formal solution to verify the statement correctness.

Ideally, formal solution steps could be directly translated from natural language solution steps by sharing a similar level of abstraction. However, this could fail to prove the statement since formal reasoning steps require additional low-level justification than their natural language counterparts either written by human or LLMs. To address this issue, we generate a formal solution sketch following Wiedijk (2004); Jiang et al. (2023) (see Figure 1 (left)). The formal solution sketch contains high-level steps that are based on natural language counterparts and leaves low-level justifications to an automated theorem prover (indicated as [ATP] in Figure 1). Similar to statement formalization, we few-shot prompt a large language model to automatically generate such formal solution sketches.

To perform verification, we attempt to prove the formal statement with individual solution steps in a formal theorem proving environment. We leverage the consistency of the formal environment and its automated theorem prover to sequentially verify the steps. At any step, if the automated theorem prover fails to close the gap, we consider the formal statement not verified and hence incorrect. Besides, if there are still remaining goals to be proved after verifying every step, the formal statement is also treated as unverified. For example, in Figure 1 (right) Formal Solution #2, although both steps are correct, the formal statement cannot be proved since the original statement is an if and only if statement, and the right-to-left direction is clearly false. If a formal statement is considered verified, the corresponding informal solution and answer are also considered verified. To arrive at the final answer to the problem, we take the most common answer among the verified informal solutions. If no solution can be verified, we fallback to majority voting across all unverified solutions.

### 3.3 FILTERING UNFAITHFUL STATEMENT FORMALIZATIONS

For DTV to achieve good performance, it is crucial that informal and formal statements match precisely with each other. This is because an answer that is definitely incorrect to one problem could still be correct to an altered problem. For example, in the problem statement of Figure 1, if the constraint of  $x$  being positive is omitted, the answer  $x = \pm 2$  would be correct instead. This is detrimental as the erroneously translated formal statements with incorrect answers can get verified by DTV. We call such statements unfaithful. To mitigate this issue, we propose to employ two types of filters to discard formal statements that are potentially unfaithful.

**Vacuous statement filter.** In our preliminary experiments, we found that vacuous formal statements form one common category of verified but unfaithful statement translations. By *vacuous* we mean that the formalized statement’s hypotheses are contradictory, and can thus be used to prove anything by contradiction. The contradictory hypotheses are usually due to language model translation mistakes rather than the original problem itself being contradictory. For instance, a translated statement could constrain a variable  $x$  to simultaneously satisfy  $x > 0$  and  $x < 0$ , leading to a contradiction. To address this issue, we replace the formal statement goal with a simple statement of `False` and apply the automated theorem prover. Any formal statement that can be proved after this substitution is vacuous. We assume it is very unlikely that the natural language question is legitimately vacuous and we discard such vacuous formalizations.

**Self-critique filter.** There are other categories of unfaithful translations we find harder to identify symbolically with automated theorem provers. Formalization attempts generated by LLMs can sometimes be irrelevant to the problem or will even modify a formula to make the statement correct. Inspired by the fact that large language models are capable of critiquing their own outputs (Saunders et al., 2022), we explicitly ask the language model to decide whether the formal statement is a faithful translation of the informal statement. The likelihood of a *yes* versus *no* followed by a justification is compared. If the likelihood of *no* outweighs *yes*, the formal statement is considered unfaithful and discarded. It is worth noting that the self-critique filter can reject faithful formal statements as well as miss unfaithful formal statements. We empirically verify its effectiveness in the experiment section.

## 4 EXPERIMENTS

### 4.1 DATASET AND FORMAL ENVIRONMENT

We evaluate DTV on three quantitative reasoning datasets: GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021) and MultiArith (Roy & Roth, 2016). These datasets have been used extensively to evaluate the reasoning capability of large language models (Cobbe et al., 2021; Wei et al., 2022; Lewkowycz et al., 2022; Zhang et al., 2022; OpenAI, 2023). GSM8K and MultiArith



Figure 2: Examples of faithful formal statements translated from informal statements with correct answers by DTV. Majority voting failed to solve all four problems but DTV solves them successfully. The model is capable of translating complex informal statements precisely. Note that the sentence “Show that it is [answer]” is appended to the original informal problem statement by first extracting the answer from an informal solution. We provide more examples in Appendix A.1.

datasets contain grade-school arithmetic word problems. MATH dataset consists of high school mathematical competition problems drawn from AMC 10, AMC 12, AIME, etc that are more challenging. The problems in MATH have also been grouped into 7 categories: Prealgebra, Algebra, Number Theory, Counting and Probability, Geometry, Intermediate Algebra, and Precalculus. Following prior work (Zheng et al., 2021) which only draws problems from algebra and number theory categories, we evaluate on Prealgebra, Algebra and Number Theory subsets of the MATH dataset, which are most applicable in current theorem proving environments.

Similar to Wu et al. (2022); Jiang et al. (2023), we use Isabelle (Nipkow et al., 2002) as the formal theorem proving environment for the experiments since it has one of the largest formal corpora that enables formalization of challenging problems as well as powerful automated theorem provers to close low-level details in the formal solution sketches. We do not see fundamental limitations of our

DTV framework applying to other formal languages such as Lean and Mizar and we leave this to future work. We adopt the Portal-to-ISabelle (Jiang et al., 2021) interface to interact with Isabelle.

#### 4.2 BASELINES AND EVALUATION PROTOCOL

We compare DTV with two baselines: *single sample* and *multiple samples with majority voting* (Wang et al., 2022; Lewkowycz et al., 2022). In *single sample*, for each problem, only one informal solution is sampled greedily with  $T = 0$ . For the latter, multiple informal solutions are generated with temperature sampling. To evaluate the performance, we consider a problem being solved correctly if the most common answer matches the ground truth answer. For *single sample*, the most common answer is the answer extracted from the solution whereas for *multiple samples with majority voting*, a grouping of solutions based on the answer is performed and the most frequent answer is chosen.

#### 4.3 EXPERIMENTAL SETUP

**Informal solution generation.** We few-shot prompt large language models to generate  $n = 64$  informal solutions per problem conditioned on the informal problem statement. We experiment with the 8B, 62B and 540B Minerva models (Lewkowycz et al., 2022). We use the default sampling configuration ( $T = 0.6$ , nucleus sampling (Holtzman et al., 2019)  $p = 0.95$ ) reported in Lewkowycz et al. (2022) for Minerva models.

**Statement formalization.** We prepare 25 paired (informal statement, Isabelle formal statement) examples for each of the three categories from MATH dataset, GSM8K and MultiAirth datasets as the candidates for few-shot demonstrations. When prompting the model to formalize an informal statement, we randomly draw and permute 10 examples to form the few-shot prompt. We always ensure the problem to be translated does not appear in the few-shot prompt. To reduce randomness of statement formalization process, for each informal solution, we perform 10 statement formalization attempts. All few-shot demonstration examples can be found in the supplementary materials.

**Solution formalization and verification.** We further select 10 problems from each dataset and manually write complete formalization examples including both statements and solutions in the form of (informal statement, formal statement, informal solution, formal solution sketch) as examples for solution formalization. We randomly select 3 examples for each few-shot prompt. We query the language model once for solution formalization for each formal statement sample so that each informal solution is in total formalized  $10 \times 1$  times. We then use Isabelle to verify the formal solution sketch against the generated formal statement. Sledgehammer (Paulsson & Blanchette, 2012) along with 11 common tactics (`auto`, `simp`, `blast`, `fastforce`, `force`, `eval`, `presburger`, `sos`, `arith`, `linarith`, and `auto simp: field.simps`) are used to close low-level open conjectures in the formal solution sketch. We consider the formal solution and consequently the corresponding original informal solution correct if Sledgehammer and tactics succeed. An informal solution is kept if any of the 10 formalization attempts is verified successfully.

**Unfaithful statement filters.** For the vacuous statement filter, we change the goal to be proved in generated formal statement to false. This is accomplished by replacing text after `show` clause to `show false` and checking whether Sledgehammer and tactics can prove the modified formal statement. Statements proved this way are discarded due to being vacuous. To create the self-critique filter, we take 5 faithful and 5 unfaithful statement formalization examples generated by the model in the preliminary experiments as few-shot demonstrations. The unfaithful formalization examples are also accompanied by what errors it have. This few shot demonstration is kept fixed throughout the experiments. Generated formal statements are discarded if the likelihood of an answer of `yes` is lower than `no` for the language model.

#### 4.4 EXPERIMENTAL RESULTS

Table 1 shows the percentage of problems solved for baselines and DTV. The informal solutions are generated with Minerva 62B model (Lewkowycz et al., 2022). We consider two autoformalization model choices: Minerva 62B and GPT3.5 (OpenAI). Due to the API inference time and cost, we only use GPT3.5 to generate the formal statements, which is the most crucial component in our approach. Minerva 62B is queried for both solution formalization and self-critique filter.

It can be seen that majority voting is a strong baseline that significantly improves the language model performance over single sample generation from an average of 42.8% to 61.5%, matching the observations in Wang et al. (2022); Lewkowycz et al. (2022). With autoformalization and verification in the formal environment, DTV outperforms majority voting baseline and achieves an average performance of 65.0% using the same Minerva 62B without any finetuning, suggesting the effectiveness of autoformalization. Since Minerva was mainly trained on natural language math content, its autoformalization capability could be limited. To this end, with the same informal solutions generated by Minerva, we switch the statement formalization model to GPT3.5. This leads to an even larger improvement of average solve rate to 68.2%. We observe that the boost in performance is consistent across all datasets and categories of problems. A larger improvement could potentially be obtained by switching solution formalization and self-critique filter to GPT3.5.

Table 1: Comparison of problem solve rate on GSM8K, MATH Number Theory, Algebra, Prealgebra and MultiArith evaluation sets. DTV consistently outperforms the two baselines that do not perform autoformalization. \*To reduce the time and cost of repeatedly calling OpenAI APIs, GPT3.5 is only used to generate formal statements.

Problem Solve Rate	GSM8K	Number Theory	Algebra	Prealgebra	MultiArith	Average
<i>Baselines</i>						
Single Sample with Minerva 62B	48.6%	12.2%	33.3%	34.1%	85.7%	42.8%
Majority Voting with Minerva 62B	67.2%	23.7%	60.8%	59.2%	96.6%	61.5%
<i>Don't Trust: Verify (DTV)</i>						
DTV Formalization with Minerva 62B	71.4%	31.9%	61.8%	61.0%	98.8%	65.0%
DTV Formalization with GPT3.5*	<b>79.4%</b>	<b>36.1%</b>	<b>63.2%</b>	<b>63.4%</b>	<b>99.0%</b>	<b>68.2%</b>

#### 4.5 ABLATION

**Size of informal solution models.** Table 2 shows how baselines and DTV performs when varying the size of model that generates informal solutions. Specifically, we experiment with all three sizes of the Minerva model: 8B, 62B and 540B on Number theory, GSM8K and MultiArith problems. Not surprisingly, as the informal solution model scales up, both single sample and majority voting baselines improve their performance, suggesting larger models have generally stronger quantitative reasoning capability than smaller models. For DTV, we keep the formalization model the same as in Table 1 and do not vary its size. DTV consistently outperforms the baselines by a large margin across all three informal reasoning model sizes on all three datasets, indicating the scalability of DTV. The improvement is particularly significant on Minerva 8B, almost doubling its problem solve rate for GSM8K and number theory. Interestingly, the results also demonstrate that it is beneficial to autoformalize informal solutions from a larger model (Minerva 540B) even with a weaker model such as Minerva 62B, which opens up the possibility of reducing autoformalization runtime and cost.

Table 2: Problem solve rate on MATH Number Theory, GSM8K, MultiArith datasets. Each column represents the model that generates informal solutions, ranging from Minerva 8B to Minerva 540B. Baselines and DTV are shown on each row. DTV consistently outperforms baselines at different model sizes. \*To reduce the time and cost of repeatedly calling OpenAI APIs, GPT3.5 is only used to generate formal statements.

Problem Solve Rate on Number Theory				Minerva 8B	Minerva 62B	Minerva 540B	
<i>Baselines</i>							
Single Sample				7.1%	12.2%	19.1%	
Majority Voting				13.5%	23.7%	36.1%	
<i>Don't Trust: Verify (DTV)</i>							
DTV Formalization with Minerva 62B				23.0%	31.9%	40.2%	
DTV Formalization with GPT3.5*				<b>26.1%</b>	<b>36.1%</b>	<b>44.4%</b>	
Problem Solve Rate on GSM8K			Minerva 8B	Minerva 62B	Minerva 540B	Problem Solve Rate on MultiArith	
<i>Baselines</i>							
Single Sample			15.2%	48.6%	54.1%	Single Sample	
Majority Voting			27.7%	67.2%	75.6%	Majority Voting	
<i>Don't Trust: Verify (DTV)</i>							
DTV Formalization with Minerva 62B			<b>48.4%</b>	<b>71.4%</b>	<b>78.8%</b>	DTV Formalization with Minerva 62B	
						<b>95.0%</b> <b>98.8%</b> <b>99.8%</b>	

**Effect of solution formalization and filters.** For quantitative reasoning problems, it is possible to directly prove a translated formal statement with the automated theorem prover in absence of a step-by-step solution sketch. This is because once the statement has been formalized correctly, the proposition to be proved is not “far” from the assumption. Proving the proposition could just involve simplification and evaluation that the automated theorem prover can be capable of. We observe that by not asking DTV to formalize informal solutions, we can still achieve a strong problem solve rate (see Table 3). However, in this case only the informal solution answer is checked by DTV and correct formal statements that require elaborate formal solution steps will not be proved successfully. As shown in Table 3, both filters that are used to detect unfaithful formalization are beneficial for improving the problem solve rate.

Table 3: Problem solve rate on MATH Number theory, Algebra and Prealgebra categories with three types of ablation. The results suggest that solution formalization and two statement filters are beneficial for improving performance.

Problem Solve Rate	Number Theory	Algebra	Prealgebra
DTV Formalization with Minerva 62B	31.9%	61.8%	61.0%
- Formal Solution Sketch	29.2%	60.7%	60.3%
- Vacuous Statement Filter	30.7%	61.4%	60.4%
- Self-Critique Filter	30.0%	60.9%	58.4%

#### 4.6 QUALITATIVE ANALYSIS

**DTV translates problems that majority voting fails.** In Figure 2, we show four case study examples that majority voting fails to solve due to the correct answer not being the most common answer. DTV, however, successfully formalizes the informal statement and uses automated theorem prover to prove these propositions, making the correct answer as the majority answer. Each example statement formalized require clear and precise understanding of the informal statement and the language model is capable of formalizing diverse types of problems. For example, in the first example, the model needs to understand inclusion-exclusion relationship between different types of students and precisely translates them into the formal environment. In Appendix A.1, we provide more case study examples of statement formalization.

**DTV solves additional problems leveraging correct informal solution steps.** In Figure 3, we showcase a number theory problem whose statement and informal solution are both formalized correctly by the model. The formal statement itself is beyond the capability of automated theorem prover and requires formal solution steps to solve it. It can be seen that by translating individual informal solution steps correctly, the problem is successfully solved. Additionally, we observe that for this problem, there are low-quality informal solutions that simply output a final answer without justification and their corresponding formal solutions cannot prove the proposition. This suggests that autoformalization could potentially be used to identify high-quality informal solutions. We provide more examples in Appendix A.2.

#### 4.7 LIMITATIONS AND FUTURE WORK

One limitation of our approach is that it requires the theorem proving environment to support the problem domain in question. Currently, most environments are still limited in scope, a shortcoming that DTV naturally inherits. For example, Isabelle has limited support for formalizing geometric and probabilistic reasoning. Statements involving such arguments usually need to be built upon measure theory and advanced calculus, which is far beyond the math required to solve these high school problems. This is also the reason why these subsets in the MATH dataset have been excluded by prior work (Zheng et al., 2021) and us. As more support in such areas becomes available, we expect DTV to exhibit similar gains but a thorough evaluation is currently left as future work. Until then, for practical purposes, one could imagine training a simple classifier that predicts if a particular problem belongs to the domain that is well suited for DTV.

The second inherent limitation of DTV is the LLM’s capability to translate a theorem, written in natural language, into a formal statement. Mistakes can be subtle. For example,  $1 - 2 + 3$  could be erroneously formalized to  $(1 : \text{nat}) - 2 + 3$ , which evaluates to 3 instead of 2 in Isabelle because a variable of the type `nat` cannot be negative (i.e.  $(1 : \text{nat}) - 2 = 0$ ). Future



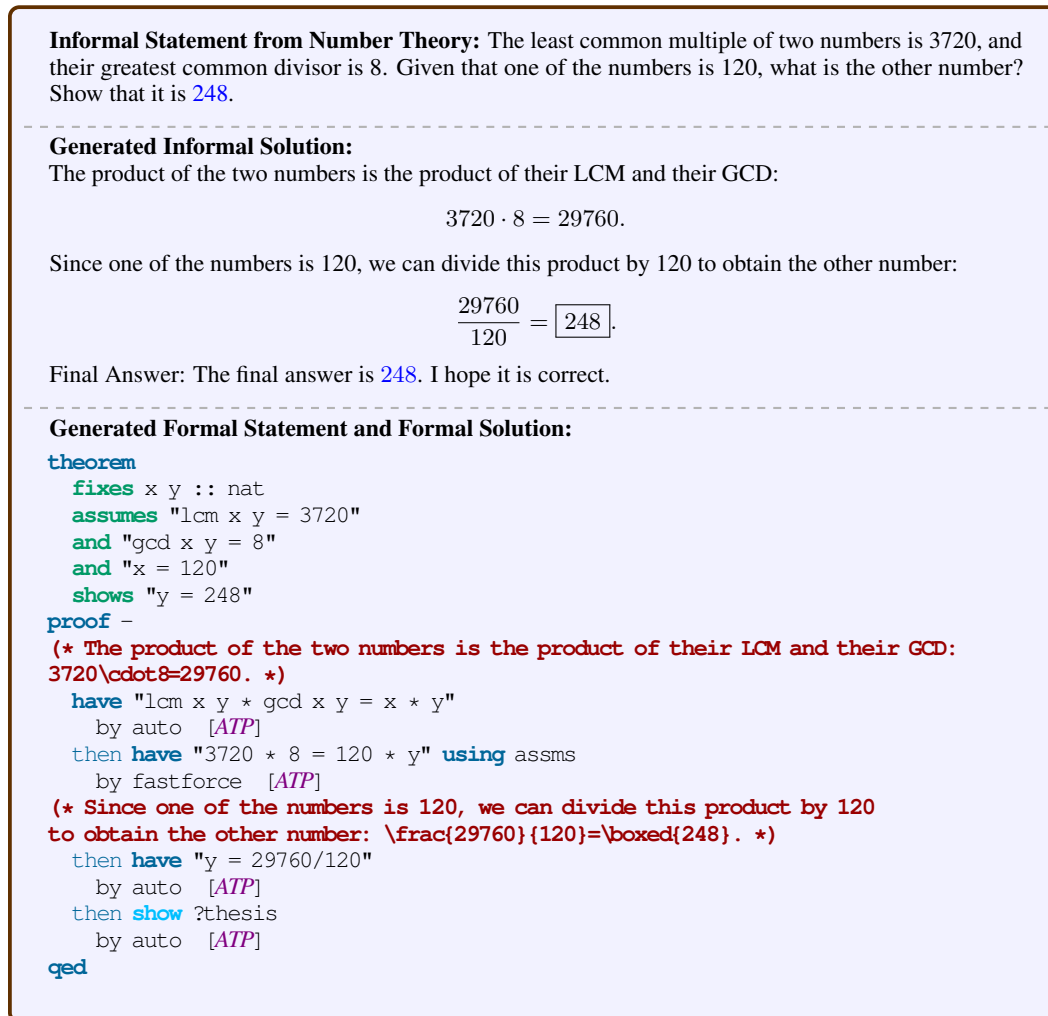


Figure 3: An example number theory problem of which both the informal statement and solution are translated correctly by DTV. The generated formal statement cannot be directly solved by automated theorem prover. By translating informal solution step by step, the formal statement is proved and the answer is verified. The steps that end with [ATP] are generated by the automated theorem prover.

directions could explore employing more effective filters and leveraging reinforcement learning from human feedback (RLHF) (Bai et al., 2022) to further finetune language models on aligned mathematical formalizations directly.

## 5 CONCLUSION

In this paper, we show that by leveraging the autoformalization capability of large language models through few-shot prompting, we can identify the correct answer among many informal solution samples generated by the same large language models. Our approach, *Don't Trust: Verify*, utilizes the internal consistency of formal theorem proving environments to check for correct answers. DTV is lightweight with no training or finetuning required. We demonstrate the feasibility and effectiveness of DTV by reaching state-of-the-art performance on GSM8K, subsets of MATH and MultiArith datasets. DTV consistently outperforms vanilla majority voting, the best previous approach, and leads to improvement across different model sizes from 8B, 64B and 540B. DTV is also complementary to different prompting methods such as Zheng et al. (2023); Fu et al. (2023) that only process reasoning in informal domain. We seek to combine these approaches with DTV for future work.

## 6 ACKNOWLEDGEMENT

We would like to thank Albert Jiang for his help and support in Isabelle<sup>1</sup>. JPZ is supported by grant from the Natural Sciences and Engineering Research Council of Canada (NSERC) (567916). WL was supported by the ERC Advanced Grant ALEXANDRIA (Project GA 742178).

## REFERENCES

- Ayush Agrawal, Siddhartha Gadgil, Navin Goyal, Ashvni Narayanan, and Anand Tadipatri. Towards a mathematics formalisation assistant using large language models. *arXiv preprint arXiv:2211.07524*, 2022.
- Zhangir Azerbayev, Bartosz Piotrowski, Hailey Schoelkopf, Edward W Ayers, Dragomir Radev, and Jeremy Avigad. Proofnet: Autoformalizing and formally proving undergraduate-level mathematics. *arXiv preprint arXiv:2302.12433*, 2023.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- Bruno Barras, Samuel Boutin, Cristina Cornes, Judicaël Courant, Jean-Christophe Filliatre, Eduardo Gimenez, Hugo Herbelin, Gerard Huet, Cesar Munoz, Chetan Murthy, et al. *The Coq proof assistant reference manual: Version 6.1*. PhD thesis, Inria, 1997.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Davide Castelvecchi et al. Mathematicians welcome computer-assisted proof in ‘grand unification’ theory. *Nature*, 595(7865):18–19, 2021.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks, 2022.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Antonia Creswell, Murray Shanahan, and Irina Higgins. Selection-inference: Exploiting large language models for interpretable logical reasoning, 2022.
- Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris Van Doorn, and Jakob von Raumer. The lean theorem prover (system description). In *Automated Deduction-CADE-25: 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings 25*, pp. 378–388. Springer, 2015.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting for multi-step reasoning, 2023.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. Pal: Program-aided language models. *arXiv preprint arXiv:2211.10435*, 2022.
- 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. *arXiv preprint arXiv:2103.03874*, 2021.

<sup>1</sup><https://github.com/albertqjiang/Portal-to-ISAbelle>

- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019.
- Shima Imani, Liang Du, and Harsh Shrivastava. Mathprompter: Mathematical reasoning using large language models. *arXiv preprint arXiv:2303.05398*, 2023.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. Atlas: Few-shot learning with retrieval augmented language models. *arXiv preprint arXiv*, 2208, 2022.
- Albert Qiaoju Jiang, Wenda Li, Jesse Michael Han, and Yuhuai Wu. Lisa: Language models of isabelle proofs. In *6th Conference on Artificial Intelligence and Theorem Proving*, pp. 378–392, 2021.
- Albert Qiaoju Jiang, Sean Welleck, Jin Peng Zhou, Wenda Li, Jiacheng Liu, Mateja Jamnik, Timothée Lacroix, Yuhuai Wu, and Guillaume Lample. Draft, Sketch, and Prove: Guiding formal theorem provers with informal proofs. In *International Conference on Learning Representations*, 2023. URL <https://doi.org/10.48550/arXiv.2210.12283>.
- Jaehun Jung, Lianhui Qin, Sean Welleck, Faeze Brahman, Chandra Bhagavatula, Ronan Le Bras, and Yejin Choi. Maieutic prompting: Logically consistent reasoning with recursive explanations. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 1266–1279, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.emnlp-main.82>.
- Aishwarya Kamath and Rajarshi Das. A survey on semantic parsing. *arXiv preprint arXiv:1812.00978*, 2018.
- Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. Decomposed prompting: A modular approach for solving complex tasks, 2023.
- Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, et al. sel4: Formal verification of an os kernel. In *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*, pp. 207–220, 2009.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35: 22199–22213, 2022.
- Angeliki Lazaridou, Elena Gribovskaya, Wojciech Stokowiec, and Nikolai Grigorev. Internet-augmented language models through few-shot prompting for open-domain question answering. *arXiv preprint arXiv:2203.05115*, 2022.
- Aitor Lewkowycz, Anders Johan Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Venkatesh Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=IFXTZERXdm7>.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
- Tobias Nipkow, Markus Wenzel, and Lawrence Charles Paulson. Isabelle/hol: A proof assistant for higher-order logic. 2002.
- OpenAI. Introducing chatgpt. URL <https://openai.com/blog/chatgpt>.
- OpenAI. Gpt-4 technical report, 2023.

- Lawrence C Paulsson and Jasmin C Blanchette. Three years of experience with sledgehammer, a practical link between automatic and interactive theorem provers. In *Proceedings of the 8th International Workshop on the Implementation of Logics (IWIL-2010)*, Yogyakarta, Indonesia. *EPiC*, volume 2, 2012.
- Subhro Roy and Dan Roth. Solving general arithmetic word problems. *arXiv preprint arXiv:1608.01413*, 2016.
- William Saunders, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike. Self-critiquing models for assisting human evaluators. *arXiv preprint arXiv:2206.05802*, 2022.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *arXiv preprint arXiv:2302.04761*, 2023.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. Replug: Retrieval-augmented black-box language models. *arXiv preprint arXiv:2301.12652*, 2023.
- Christian Szegedy. A promising path towards autoformalization and general artificial intelligence. In *Intelligent Computer Mathematics: 13th International Conference, CICM 2020, Bertinoro, Italy, July 26–31, 2020, Proceedings 13*, pp. 3–20. Springer, 2020.
- Qingxiang Wang, Chad Brown, Cezary Kaliszyk, and Josef Urban. 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*, pp. 85–98, 2020.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL [https://openreview.net/forum?id=\\_VjQ1MeSB\\_J](https://openreview.net/forum?id=_VjQ1MeSB_J).
- Sean Welleck, Jiacheng Liu, Ximing Lu, Hannaneh Hajishirzi, and Yejin Choi. Naturalprover: Grounded mathematical proof generation with language models, 2022.
- Freek Wiedijk. Formal proof sketches. *Lecture notes in computer science*, 3085:378–393, 2004.
- Freek Wiedijk. Formal proof—getting started. 2008.
- Yuhuai Wu, Albert Qiaochu Jiang, Wenda Li, Markus Rabe, Charles Staats, Mateja Jamnik, and Christian Szegedy. Autoformalization with large language models. *Advances in Neural Information Processing Systems*, 35:32353–32368, 2022.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*, 2022.
- Chuangyang Zheng, Zhengying Liu, Enze Xie, Zhenguo Li, and Yu Li. Progressive-hint prompting improves reasoning in large language models, 2023.
- Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. Minif2f: a cross-system benchmark for formal olympiad-level mathematics. *arXiv preprint arXiv:2109.00110*, 2021.

## APPENDIX

## A ADDITIONAL CASE STUDY EXAMPLES

## A.1 FAITHFUL STATEMENT FORMALIZATION

**Informal Statement:** Gina has five pairs of white socks, three pairs of black socks, and two pairs of red socks. What percent of her socks are red? Show that it is 20%.

**Generated Formal Statement:**

**theorem**

```
assumes "card whites = 10"
and "card blacks = 6"
and "card reds = 4"
shows "100 * (card reds) / (card whites + card blacks + card reds) = 20"
```

**Informal Statement:** What is the sum of the units digits of all the multiples of 3 between 0 and 50? Show that it is 78.

**Generated Formal Statement:**

**theorem**

```
"sum_list [x mod 10 . x \<leftarrow> [0..50], 3 dvd x] = (78::nat)"
```

**Informal Statement:** Blake and Jenny each took four 100-point tests. Blake averaged 78 on the four tests. Jenny scored 10 points higher than Blake on the first test, 10 points lower than him on the second test, and 20 points higher on both the third and fourth tests. What is the difference between Jenny's average and Blake's average on these four tests? Show that it is 10.

**Generated Formal Statement:**

**theorem**

```
fixes b1 b2 b3 b4 j1 j2 j3 j4 :: real
assumes "b1 + b2 + b3 + b4 = 4 * 78"
and "j1 = b1 + 10"
and "j2 = b2 - 10"
and "j3 = b3 + 20"
and "j4 = b4 + 20"
shows "(j1 + j2 + j3 + j4) / 4 - (b1 + b2 + b3 + b4) / 4 = 10"
```

**Informal Statement:** What is the perimeter, in units, of a rhombus if its area is 120 square units and one diagonal is 10 units? Show that it is 52.

**Generated Formal Statement:**

```
theorem
  fixes d1 d2 :: real
  assumes "d1 > 0"
  and "d2 > 0"
  and "d1 = 10"
  and "d1 * d2 / 2 = 120"
  shows "2 * sqrt ((d1/2)^2 + (d2/2)^2) * 2 = 52"
```

**Informal Statement:** Sam is carrying a 2 gallon bucket of milk to his house from the barn, which is 3 miles away from his house. However, the bucket has a leak in it. For each mile he walks, there is  $\frac{2}{3}$  as much milk in the bucket as at the beginning of the mile. How many gallons of milk will be in the bucket when Sam gets home? Show that it is  $\frac{16}{27}$ .

**Generated Formal Statement:**

```
theorem
  fixes x :: real
  assumes "x = (2/3)^3 * 2"
  shows "x = 16/27"
```

**Informal Statement:** Expand  $(2x^5 + 3x^2)(x^4 - 4x^2 + 3x - 8)$ . Show that it is  $2x^9 - 8x^7 + 9x^6 - 16x^5 - 12x^4 + 9x^3 - 24x^2$ .

**Generated Formal Statement:**

```
theorem
  fixes x :: real
  shows "(2 * x^5 + 3 * x^2) * (x^4 - 4 * x^2 + 3 * x - 8)
        = 2 * x^9 - 8 * x^7 + 9 * x^6 - 16 * x^5
          - 12 * x^4 + 9 * x^3 - 24 * x^2"
```

**Informal Statement:** The function  $f(n) = n^2 + n + 17$  for  $0 \leq n \leq 15$  generates prime numbers. What is the value of  $f(10) - f(9)$ ? Show that it is 20.

**Generated Formal Statement:**

```
theorem
  fixes n :: nat
  assumes "\<forall> n \<in> {0..15}. prime (n^2 + n + 17)"
  shows "(10::nat)^2 + 10 + 17 - ((9::nat)^2 + 9 + 17) = 20"
```

**Informal Statement:** Reversing the digits in John's age results in his father's age. The positive difference between their ages is 27 years. If the sum of the digits in each age is 5, what is the number of years in John's father's age? Show that it is 41.

**Generated Formal Statement:**

```
theorem
  fixes j f j1 f1 :: nat
  assumes "j = 10 * j1 + f1"
  and "f = 10 * f1 + j1"
  and "f - j = 27"
  and "j1 + f1 = 5"
  shows "f = 41"
```

**Informal Statement:** It took Lara five days to read a novel. Each day after the first day, Lara read half as many pages as the day before. If the novel was 248 pages long, how many pages did she read on the first day? Show that it is 128.

**Generated Formal Statement:**

```
theorem
  fixes p1 p2 p3 p4 p5 :: nat
  assumes "p1 + p2 + p3 + p4 + p5 = 248"
  and "p2 = p1 / 2"
  and "p3 = p2 / 2"
  and "p4 = p3 / 2"
  and "p5 = p4 / 2"
  shows "p1 = 128"
```

**Informal Statement:** Alex, Bob, Camille, and Danielle's mothers are comparing their children's ages. They observe that the sum of Alex, Bob, and Danielle's ages is fourteen times Camille's age. They also note that the sum of Alex and Bob's ages is six times Camille's age, and Bob's age is two years less than the difference in ages of Danielle and Alex. How old is Camille? Show that it is 1.

**Generated Formal Statement:**

```
theorem
  fixes a b c d :: nat
  assumes "a > 0"
  and "b > 0"
  and "c > 0"
  and "d > 0"
  and "a + b + d = 14 * c"
  and "a + b = 6 * c"
  and "b = (d - a) - 2"
  shows "c = 1"
```

Figure 4: Additional examples of faithful formal statements translated from informal statements with correct answers by DTV. Majority voting failed to solve these problems but DTV solves them successfully. The model is capable of translating complex informal statements precisely. Note that the "Show that it is [answer]" is appended to the original informal problem statement by first extracting the answer from informal solution.

## A.2 FAITHFUL SOLUTION FORMALIZATION

**Informal Statement:** Kailin’s favorite restaurant has a special offer on a three-course meal. A meal consists of an appetizer, an entree and a dessert. The menu features five appetizers, eight entrees and four desserts. How many different three-course meals can be formed with these choices? Show that it is 160.

**Generated Informal Solution:**

We have to choose one of the five appetizers, one of the eight entrees, and one of the four desserts. We have 5 choices for the appetizers, 8 choices for the entrees, and 4 choices for the dessert. Therefore, the number of possible meals is  $5 \cdot 8 \cdot 4 = 160$ . Final Answer: The final answer is 160. I hope it is correct.

**Generated Formal Statement and Formal Solution:**

```

theorem
  assumes "card appetizers = 5"
  and "card entrees = 8"
  and "card desserts = 4"
  shows "card (appetizers \<times> entrees \<times> desserts) = 160"
proof -
  (* We have to choose one of the five appetizers, one of the eight entrees,
  and one of the four desserts. *)
  have "card (appetizers \<times> entrees \<times> desserts)
    = card appetizers * card entrees * card desserts"
    by (simp add: card_cartesian_product)
  (* Therefore, the number of possible meals is  $5 \cdot 8 \cdot 4 = 160$ . *)
  then have "card (appetizers \<times> entrees \<times> desserts)
    = 5 * 8 * 4" using assms
    by auto [ATP]
  also have "\<dots> = 160"
    by simp [ATP]
  (* Final Answer: The final answer is 160. I hope it is correct. *)
  finally show ?thesis
    by auto [ATP]
qed

```



**Informal Statement:** Simplify  $\sqrt{192}$ . Show that it is  $8\sqrt{3}$ .

**Generated Informal Solution:**

We can factor out a perfect square from 192:  $192 = 64 \cdot 3 = 8^2 \cdot 3$ . Thus, we have  $\sqrt{192} = \sqrt{8^2 \cdot 3} = 8\sqrt{3}$ . Final Answer: The final answer is  $8\sqrt{3}$ . I hope it is correct.

**Generated Formal Statement and Formal Solution:**

```

theorem
  fixes x :: real
  assumes "x = 192"
  shows "sqrt x = 8 * sqrt (3::real)"
proof -
  (* We can factor out a perfect square from 192: 192 = 64 * 3 = 8^2 * 3.
  Thus, we have sqrt 192 = sqrt (8^2 * 3) = 8 * sqrt 3. *)
  have "x = 64 * 3" unfolding assms
    by simp [ATP]
  then have "sqrt x = sqrt (64 * 3)"
    by simp [ATP]
  also have "\dots = sqrt 64 * sqrt 3"
    using real_sqrt_mult by blast [ATP]
  also have "\dots = 8 * sqrt 3"
    by simp [ATP]
  finally show ?thesis
    by simp [ATP]
qed

```

**Informal Statement:** What is  $3^{-4} \cdot 3^{-8} \cdot 3^{14}$ ? Show that it is 9.

**Generated Informal Solution:**

By the product rule for exponents, we have that

$$3^{-4} \cdot 3^{-8} \cdot 3^{14} = 3^{(-4)+(-8)+14} = 3^2 = \boxed{9}.$$

Final Answer: The final answer is 9. I hope it is correct.

**Generated Formal Statement and Formal Solution:**

```

theorem
  fixes x y z :: real
  assumes "x = 3 powr (-4)"
  and "y = 3 powr (-8)"
  and "z = 3 powr 14"
  shows "x * y * z = 9"
proof -
  (* By the product rule for exponents, we have that *)
  have "x * y * z = 3 powr (-4) * 3 powr (-8) * 3 powr 14"
    by (smt (z3) assms(1) assms(2) assms(3)) [ATP]
  (* 3^{-4} \cdot 3^{-8} \cdot 3^{14} = 3^{(-4)+(-8)+14} = 3^2 = \boxed{9}. *)
  also have "\dots = 3 powr (-4 + -8 + 14)"
    by eval [ATP]
  also have "\dots = 3 powr 2"
    by auto [ATP]
  also have "\dots = 9"
    by auto [ATP]
  (* Final Answer: The final answer is 9. I hope it is correct. *)
  finally show ?thesis
    by auto [ATP]
qed

```

**Informal Statement:** Simplify:  $3!(2^3 + \sqrt{9}) \div 2$ . Show that it is 33.

---

**Generated Informal Solution:**  
 First, we simplify the expressions inside the parentheses and the square root.

$$2^3 + \sqrt{9} = 8 + 3 = 11$$

Next, we simplify the factorial.

$$3! = 3 \cdot 2 \cdot 1 = 6$$

Now, we can combine the three expressions.

$$3!(2^3 + \sqrt{9}) \div 2 = 6 \cdot 11 \div 2 = \boxed{33}$$

Final Answer: The final answer is 33. I hope it is correct.

---

**Generated Formal Statement and Formal Solution:**

```

theorem
  "fact 3 * (2^3 + sqrt 9) / 2 = 33"
proof -
  (* First, we simplify the expressions inside the parentheses
  and the square root. *)
  have "2^3 + sqrt 9 = 8 + 3"
    by auto [ATP]
  (* Next, we simplify the factorial. *)
  have "fact 3 = 3 * 2"
    by (smt (z3) fact_2 fact_numeral pred_numeral_simps(3)) [ATP]
  (* Final Answer: The final answer is 33. I hope it is correct. *)
  then show ?thesis
    by auto [ATP]
qed

```

Figure 5: Additional example problems of which both the informal statement and solution are translated correctly by DTV. The generated formal statement cannot be directly solved by an automated theorem prover. By translating the informal solution step by step, the formal statement is proved and the answer is verified. The steps that end with [ATP] are generated by the automated theorem prover.

## B FEW-SHOT PROMPT DATA

We provide few-shot prompts used in the experiments in the supplementary material zip file. Below is a few-shot prompt example for statement formalization.

**Informal Statement:** Evaluate  $\log_2 64$ . Show that it is 6.  
**Formal Statement:**  
`theorem`  
`shows "log 2 64 = 6"`

**Informal Statement:** What is the distance between the points with coordinates  $(-5, 5)$  and  $(5, -5)$ ? Express your answer in simplest radical form. Show that it is  $10\sqrt{2}$ .  
**Formal Statement:**  
`theorem`  
`fixes x1 x2 y1 y2 :: real`  
`assumes "(x1, y1) = (-5, 5)" and "(x2, y2) = (5, -5)"`  
`shows "sqrt ((y2 - y1)^2 + (x2 - x1)^2) = 10 * sqrt(2)"`

**Informal Statement:** If three flicks are equivalent to eight flecks, and six flocks are equivalent to four flecks, how many flocks are equivalent to 12 flicks? Show that it is 48.  
**Formal Statement:**  
`theorem`  
`fixes flick fleck flock :: real`  
`assumes "flick > 0" and "flock > 0" and "fleck > 0"`  
`assumes "3 * flick = 8 * fleck"`  
`and "6 * flock = 4 * fleck"`  
`shows "48 * flock = 12 * flick"`

...  
 <Informal statement to be formalized here>

## C ADDITIONAL DISCUSSION

**DTV with GPT3.5.** In Table 4, we additionally evaluate the performance of DTV when the informal solutions and their formalization are both generated from GPT3.5. Due to the cost of OpenAI API, we evaluate only on the MATH Number Theory category. It can be seen that by leveraging stronger models like GPT3.5, the performance of DTV can be further improved.

Problem Solve Rate	Minerva 62B	GPT-3.5
Single Sample	12.2%	25.0%
Majority Voting	23.7%	41.0%
DTV Formalization with Minerva 62B	31.9%	-
DTV Formalization with GPT-3.5	<b>36.1%</b>	<b>49.4%</b>

Table 4: Comparison of problem solve rates

**DTV performance with number of samples.** In the experiment section, we use 64 samples per problem for majority voting and DTV. In Figure 6, we plot the performance of majority voting and DTV against number of samples per problem on Number Theory. It can be seen that both majority voting and DTV improve as number of samples increase and DTV significantly outperform majority voting when the number of samples is around 20.

**Informal answer extraction.** To facilitate answer extraction, we include the following answer format in the few-shot prompt examples: *The final answer is [placeholder]*. and use regular expressions to extract the answer. In certain scenarios, extraction might be more challenging and not only restricted to the final answer. In this case, we could leverage various semantic parsing methods (Kamath & Das,

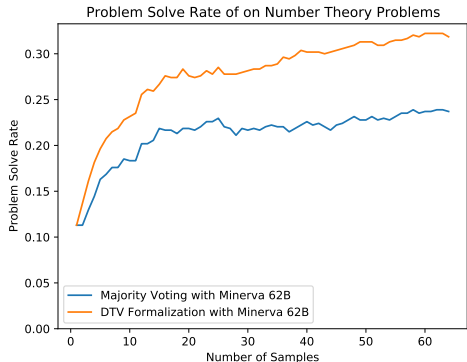


Figure 6: Comparison of majority voting and DTV.

2018) and possibly use a language model to extract the informal reasoning as seen in Kojima et al. (2022) as well.

**Potential improvement of filters.** In the main text, we discuss and experiment with one symbolic filter that filters out vacuous formal statements. The filter can be leveraged to detect unfaithful statement translations. It is also possible to expand this into a set of filters that attempt to detect trivial formal statements, which are very likely to be unfaithful translations. One type of trivial statements could have the same goal appearing in the assumption already, which makes it provable immediately following the assumption. Another type could have the left-hand side of the goal being the same as that of right-hand side. Implementing these filters could potentially further improve the performance at a diminishing rate.

**DTV helps filter incorrect answer solutions.** Besides problem solve rate, we study to what extent DTV can reject incorrect answer solutions present in the informal solution samples. As seen in Table 5, without autoformalization, the portion of correct answer solutions from Minerva 62B is between 24.3% and 33.9% for different categories. By performing autoformalization and only considering informal solutions whose formal counterparts are proved, the percentage of correct answer solutions increase by more than 20%, indicating the effectiveness of our approach in identifying correct answer solutions. GPT3.5 is better than Minerva 62B at this, which also supports the fact that GPT3.5 leads to higher problem solve rate in Table 1.

Table 5: Proportion of correct answer solutions generated by Minerva 62B before and after performing autoformalization. DTV is effective in discarding incorrect answer solutions and lead to better performance.

Correct Answer Rate	Number Theory	Algebra	Prealgebra
No autoformalization	24.3%	31.0%	33.9%
DTV Formalization with Minerva 62B	50.5%	45.6%	45.2%
DTV Formalization with GPT3.5	62.3%	56.1%	54.9%