EXECUTABLE FUNCTIONAL ABSTRACTIONS: INFERRING GENERATIVE PROGRAMS FOR ADVANCED MATH PROBLEMS

Anonymous authorsPaper under double-blind review

000

001

002

004

006

008 009 010

011

013

014

016

017

018

019

021

025

026

027

028

029

031

034

040

041

042

043

044

046

047

048

049

051

052

ABSTRACT

Scientists often infer abstract procedures from specific instances of problems and use the abstractions to generate new, related instances. For example, programs encoding the formal rules and properties of a system have been useful in fields ranging from reinforcement learning (procedural environments) to physics (simulation engines). These programs can be seen as functions which execute to different outputs based on their parameterizations (e.g., gridworld configuration or initial physical conditions). We introduce the term EFA (Executable Functional Abstraction) to denote such programs for math problems. EFA-like constructs have been shown to be useful for mathematical reasoning as problem generators for stress-testing models. However, prior work has been limited to automatically constructing abstractions for grade-school math (whose simple rules are easy to encode in programs), while generating EFAs for advanced math has thus far required human engineering. We explore the automatic construction of EFAs for advanced mathematics problems by developing EFAGen, which operationalizes the task of automatically inferring an EFA for a given seed problem and solution as a program synthesis task. We first formalize the properties of any valid EFA as executable unit tests. Using execution feedback from the unit tests, we search over candidate programs sampled from a large language model (LLM) to find EFA programs that are faithful to the generalized problem and solution class underlying the seed problem. We then apply the tests as a reward signal, training LLMs to become better writers of EFAs. We show that EFAs inferred by EFAGen are faithful to the seed problems, produce learnable problem variations, and that EFAGen can infer EFAs across diverse sources of competition-level math problems. Finally, we show uses of model-written EFAs, such as finding problem variations that are harder or easier for a learner to solve, as well as data generation.1

1 Introduction

In many fields, experts abstract specific instances into general procedures that can generate a wide range of related cases. For example, physicists distill observations of falling objects into equations of motion capable of predicting trajectories under varying initial conditions (Smith, 2024). This ability is not limited to certain domain experts: in fact, the ability to infer underlying compositional structures from surface forms is a core component of human language and intelligence (Chomsky, 1957; Montague et al., 1970; Partee, 2008; Lake et al., 2017). The outcome of this process of abstraction is often a data-generating program whose execution is controlled by parameters, such as a gridworld generator that produces different world layouts given different configuration files. In fields such as reinforcement learning, notable instances of data generating programs such as Holodeck (Yang et al., 2024) and BabyAI (Chevalier-Boisvert et al., 2018) have become important parts of the research ecosystem for their capability to endlessly generate well-formed randomized task instances.

We introduce Executable Functional Abstraction (EFA), a programmatic abstraction that encapsulates the logic of a math problem in a parameterized form and enables the automated sampling of problems variants. Although similar abstractions have been used in other domains, automatic

¹We will open-source the code and data upon acceptance.




Figure 1: **Left**: The generative process underlying computational math problems, where the different instances share the same underlying problem-solving logic (function) but differ in parameter values. We introduce **executable functional abstractions (EFAs)** to model this latent structure. **Right**: we study the task of inferring EFAs; i.e., recovering the underlying problem-solving function and parameters from math problems expressed in natural language.

construction of EFAs for generating fresh, diverse math problems remains largely unexplored. The property enabling the construction of EFAs for mathematics is that many math problems are a surface form of a more abstract deep structure. For example, consider the problem in Fig. 1 (left), which asks for positive integers m and n with a greatest common divisor (GCD) of 6 and a least common multiple (LCM) of 126, seeking the minimum value of m+n, which we denote as LcmGcdMinSum(gcd=6, lcm=126). This specific problem is a special case of a more general problem LcmGcdMinSum(gcd=g, lcm=l) where $l,g \in \mathbb{N}$ can be any natural numbers. Inferring an EFA requires transforming the LcmGcdMinSum(gcd=6, lcm=126) problem about a specific pair of numbers into a program that generates valid LcmGcdMinSum problems with varying parameters while implementing a general solution procedure that solves any specific instances of the general problem, such as LcmGcdMinSum(gcd=7, lcm=42). In this paper, we explore the automatic creation of EFAs for higher-level math problems. Our central research question is:

How can we automatically transform static math problems into their corresponding executable functional abstractions (EFAs)?

The task of automatically transforming static math problems into an EFA is nontrivial. Recent work has made progress with grade-school level math problems (Zhang et al., 2024; Mirzadeh et al., 2025) by taking advantage of the simple computational graphs of their solutions. Higher-level problems with more complex computational graphs have thus far required human involvement to lift problems into functional forms (Shah et al., 2024; Srivastava et al., 2024). An automated approach for mathematical problems more complex than grade-school arithmetic has not been developed. Such automatic construction of EFAs requires simultaneously solving multiple subproblems: identifying which numerical values should be parameterized, discovering the constraints between these parameters to maintain problem validity, abstracting the solution procedure to handle all valid parameterizations, and ensuring mathematical correctness across the entire parameter space. For example, in Fig. 1, m and n are not parameters of the problem despite already being abstract variables, as they are dependent on the values of the gcd and lcm given. Nor can the gcd or lcm values be allowed to vary arbitrarily. Some parameterizations of the gcd and lcm may yield trivial problems (if the gcd is 1 and the lcm is a prime), while other parameterizations are simply invalid (such as gcd > lcm or gcd and lcm being relatively prime).

We operationalize the task of inferring EFAs as a program synthesis task using large language models (LLMs). Our method, EFAGen, conditions an LLM on a static seed math problem and its step-by-step solution to generate candidate programs implementing an EFA for the seed math problem. To generate a correct EFA, the program synthesizer must identify which numerical values in the static problem should be treated as parameters, determine appropriate sampling distributions for these parameters, and encode the constraints between them to ensure problem validity (Fig. 1). We formalize mathematical properties a well-formed EFA must possess as unit tests that can automatically detect violations of these properties. We can then adopt an overgenerate-then-filter approach (Li




Figure 2: **EFAGen generalizes prior work on constructing arithmetic word problems to automatically constructing more complex, higher-level math problems.** Given a math problem and solution, EFAGen infers an underlying abstraction whose construction and general solution may involve arbitrary computations beyond fixed sequences of arithmetic operations. For example, the abstraction underlying the eigenvalue problem on the right is that of a tridiagonal 3×3 matrix. The general solution requires a symbolic computation composed with a numerial root-finding procedure. Details of inferred EFA code in Fig. 6.

et al., 2022), first generating a large number of candidate programs implementing EFAs for a seed problem, and then rejecting EFAs that fail our tests. Finally, we conduct a series of experiments probing properties of the EFAs constructed by EFAGen, demonstrating the utility of model-written EFAs and testing whether LLMs can be trained to be more successful writers of EFAs.

We first show that EFAs have properties signaling their coherence. EFAs are faithful to the seed problem they were derived from: the verifiable problems sampled from an EFA help a model solve the seed problem the EFA was constructed from. Similarly, the verifiable problems produced by an EFA are learnable: when sampling a train and test set from the same EFA, a model is able to improve on the test set when given step-by-step solutions of the training problems.

Because EFAs allow us to sample a large number of verified problems, we can also use them to create more instances of a problem that a model struggles with, or to refresh a static dataset by first constructing an EFA from a problem that the model already can solve, and then sampling fresh variants using the EFA that the model struggles with, thereby stress-testing models on similar data. We show that EFAGen can be applied to multiple sources of competition-level mathematics problems to automatically construct EFAs. This applicability to multiple kinds of problems allows us to use EFAs as a data augmentation for mathematical problem solving on MATH-500 (Hendrycks et al., 2021) and FnEval (Srivastava et al., 2024), where we show EFA-based augmentation yields consistent improvements. Finally, we show that models can improve at inducing EFAs from math problems by using the execution feedback from automatic tests in EFAGen as rewards in a simple reinforced self-training scheme (Zelikman et al., 2022; Singh et al., 2023; Dong et al., 2023).

Our contributions in this paper as illustrated in Fig. 2 are as follows:

- We formalize the notion of Executable Functional Abstractions (EFAs) in Sec. 2.2, and develop EFAGen (Sec. 2.3, Fig. 3), an approach that automatically infers EFAs from advanced math problems, providing a scalable approach to generate verifiable problem variants with automatic tests for validity and correctness.
- We show that these tests can be used as a reward signal for training LLMs to improve at the task of inferring EFAs from static problems (Sec. 3.1).




Figure 3: Left: Representation of an executable functional abstraction (EFA) as a Python class. Right: Overview of EFAGen, a method for automatically inferring EFAs from a math problem. In EFAGen, we (a) over-generate multiple EFA candidates with an LLM and (b) filter out invalid candidates that fail automated tests. The EFA can generate new problem variants by sampling parameters and executing the solver. Full code is in Appendix G.

• We show that EFAGen generates faithful (Sec. 3.2) and learnable (Sec. 3.3) EFAs and can automatically infer EFAs from diverse sources of math data (Appendix C), and that EFAs can be used as a data augmentation (Sec. 3.4).

2 EXECUTABLE FUNCTIONAL ABSTRACTIONS (EFAs)

Our goal is to automatically convert math problems with static numerical values into **parameterized abstractions** that can generate variants of the original problems. We refer to these parameterized abstractions as **Executable Functional Abstractions** (**EFAs**). EFAs enable the systematic generation of new problem instances by varying numerical parameters while preserving the underlying problemsolving logic. We operationalize the task of inferring an EFA for a static math problem as a program synthesis task where the goal is to write a class implementing the EFA. We use LLMs to generate many candidate EFA implementations for a static problem and use a suite of automatic unit tests to filter the candidates by rejecting mathematically unsound ones. Below, we describe the desired properties of EFAs (Sec. 2.1), how an EFA is represented as a Python class (Sec. 2.2), and how we infer EFAs from static math problems using LLMs (Sec. 2.3).

2.1 Desired Properties of Abstractions

An effective abstraction of a math problem must support variation, preserve validity, and enable automated problem-solving. We identify three core properties of an EFA:

- **Structured parameter space**: The abstraction should define a set of parameters that characterize the problem and specify valid relationships among them. This includes identifying which parameters are independent, how dependent parameters are derived, and what constraints must be satisfied to ensure valid problem instances. Such structure enables systematic variation, ensuring that changes to parameters yield meaningful variants with potentially different solutions.
- **Procedural generation of instances**: The abstraction should support random sampling of a set of valid parameters (e.g., EFA.sample() in Sec. 2.2) and converting the abstract problem into natural language form (e.g., EFA.render() in Sec. 2.2), to help users generate valid problem instances by sampling parameter values within defined constraints. These constraints are problem-specific and crucial for generating diverse but coherent examples.
- Executable solution logic: The abstraction should include a method (e.g., EFA.solve() in Sec. 2.2) that computes the correct answer for any valid parameter configuration. This solution logic is typically derived from the chain-of-thought (Wei et al., 2022) used for the static version of the problem and can be implemented as an executable program.

2.2 EFA AS A PYTHON CLASS

As shown in Fig. 3(a), each EFA is implemented as a Python class that contains the logic of a math problem in a parameterized form. The class defines a list of parameters along with three key methods:

- **EFA.sample()** → **parameters**: Samples a valid set of parameters representing problem variants, respecting all constraints specified in the abstraction.
- EFA.render(parameters) → question: Provides a natural language problem statement, given a specific (sampled) parameter set. This ensures that each generated instance is presented in a format suitable for human or model consumption. In most cases, this involves reusing the problem statement of the seed instruction and mutating the numerical values to be consistent with the given parameters.
- **EFA. solve (parameters)** \rightarrow **answer**: Computes the correct answer for a given parameter configuration. The solution is expressed as a numerical expression derived through deterministic computations over the parameters. The solver does not need to access the natural language problem statement, as the solution is only dependent on the parameterization of the problem, which is a structured object.

These methods operationalize the abstraction and enable automated generation, rendering, and evaluation of math problems.

2.3 EFAGEN: INFERRING EFAS FROM MATH PROBLEMS

We introduce EFAGen, a framework for automatically constructing EFAs from static math problems. Given a problem statement and its solution procedure (typically expressed as chain-of-thought reasoning), EFAGen uses a large language model (LLM) to generate a candidate EFA implementation that captures the logic and structure of the original problem. This process relies on supervision that is readily available in many math datasets.

Since generating correct and robust code is challenging for LLMs, EFAGen adopts an overgenerate-and-filter approach inspired by AlphaCode (Li et al., 2022). As described in Fig. 3 (a), for each problem, we sample N (e.g., 50) EFA candidates and apply a suite of automated tests to discard invalid abstractions. EFAGen uses the following tests to validate candidate EFAs, as illustrated in Fig. 3 (b):

- is_extractable (response): Verifies that the class contains all required methods.
- is_executable (EFA): Confirms that the class can be instantiated and executed without errors, and methods like EFA.sample() and EFA.solve() can be called without errors.
- has_dof (EFA): Ensures that sampled parameters differ, rejecting EFAs with zero degrees of freedom that cannot produce new problems.
- is_single_valued (EFA): Confirms that identical parameters yield equivalent solutions, rejecting impermissible implementations including multivalued functions or incoherent abstractions.
- matches_original (EFA, orig_params, orig_sol): Validates that the abstraction, when instantiated with the original parameters, produces the original problem and solution. This serves as a cycle-consistency or soundness check.

Any program that fails these tests cannot logically be a valid implementation of an EFA. EFAGen enables generation of EFAs at scale, as shown in Fig. 3 (c), as large numbers of candidate EFAs can be generated and filtered automatically. Over time, these tests can also be used to fine-tune LLMs toward better abstraction generation, such as with reinforced self-training (Singh et al., 2023; Dong et al., 2023) or reinforcement learning with verifiable rewards (Lambert et al., 2024).

3 EXPERIMENTS & RESULTS

Below, we show experiments on self-improving at inferring EFAs (Sec. 3.1), faithfulness (Sec. 3.2) and learnability (Sec. 3.3) of EFAs, the complementarity of EFAs with existing data generation methods (Sec. 3.4, Sec. 3.5). In the Appendix, we analyze the quality of EFA-generated data (Appendix B), scaling experiments (Appendix E), applying EFAs to find hard variants of problems (Appendix D), ablations (Appendix F), and EFAGen inference on olympiad-level problems (Appendix C).

Datasets. Throughout this section, we use the following datasets in our experiments:

• MATH (Hendrycks et al., 2021). Competition math dataset with a test set of 5k math problems described in text comprising different categories and five levels of difficulty. We show in Sec. 3.1 that LLMs struggle with task of EFA generation and we improve their performance by training on the EFA generation task using the MATH train set consisting of 7.5k problems.




Figure 4: **LLMs can use our tests to self-improve at inferring EFAs.** We plot the percentage of constructed EFAs passing all tests across iterations of self-training, grouped by MATH problem difficulty (left) and by problem category (right). Harder difficulty levels and problem categories are harder to infer EFAs for and improve more during training.

- FnEval (Srivastava et al., 2024). A functional version of the MATH benchmark designed to evaluate generalization. It consists of multiple "snapshots", each containing variations of problems from the MATH dataset. These variations preserve the abstract reasoning structure of the original problems. We use two snapshots to test if our method can capture the underlying abstractions of a problem and generalize to unseen, related instances.
- MATH-Hard is a subset of MATH test problems of the highest difficulty (level 5) across all categories (1387 problems).

Metrics. To evaluate the performance of LLMs we use the following metrics:

- EFA Success Rate. We measure the ability of LLMs to generate valid, high-quality EFAs (defined in Sec. 2.1) as the frequency (%) of EFAs generated that past all the diagnostic tests (c.f. Sec. 2.3).
- Pass@k Rate (%). Following Chen et al. (2021), we measure the ability of LLMs to solve math problems by sampling 25 generations with temperature sampling and estimating the unbiased pass@k rate, i.e., the likelihood that out of k generated solutions any one yields the correct answer.

3.1 SELF-IMPROVEMENT: LMs IMPROVE AT EFA INFERENCE WITH EXECUTION FEEDBACK

Inferring valid EFAs across diverse math problems is challenging, especially as the difficulty and complexity of topics increases. For instance, as shown in Fig. 4, Llama3.1-8B-Instruct (Llama Team, 2024) struggles to generate valid EFAs for Level 5 problems and for topics such as Precalculus in the MATH dataset, where it is only able to infer valid EFAs for $\approx 35\%$ of Precalculus questions. In Sec. 2, we introduce a number of unit tests (i.e., verifiable rewards) that indicate whether a generated EFA is valid. Here, we show that we can train models to improve on inferring valid EFAs by self-training according to these tests. Specifically, we use a rejection-finetuning approach (Zelikman et al., 2022; Singh et al., 2023; Dong et al., 2023), in which we sample EFA candidates from a model and filter according to our rewards to construct a training dataset of correct examples. We begin with the MATH training set (7,500 problems) and sample 10 candidate EFAs per problem. Candidates failing any of the reward checks are discarded. The remaining valid examples form a dataset for supervised fine-tuning. This process – sampling, filtering, and retraining – is repeated over 5 iterations (see Appendix H.2 for details).

We report the EFA success rates across iterations in Fig. 4, where we group by difficulty levels (left) and by annotated problem category (right). Success rates steadily improve over training iterations, especially for harder problems. At iteration 0 (before training), we observe that harder problems (e.g., Level 5) are also harder to infer EFAs for, with EFA success rates $\approx 17\%$ lower for Level 5 than Level 1 problems. Similarly, certain categories like 'Intermediate Algebra', 'Counting' and 'Probability' are harder to infer EFAs for. These domains generally see the most significant increases from training. Between iteration 1 and iteration 5, the Intermediate Algebra's EFA success rate showed the most significant increase, rising from 52.93% to 81.38%, and Geometry improved from 65.71% to 85.71%.

Table 1: **EFAs faithfully capture the solutions of the problems they were derived from (left), and problem variants constructed by EFAs share learnable structure (right).** Left: Giving solutions to problems variants from an EFA as in-context examples nearly doubles the solve rate of an LLM on the seed problem the EFA was derived from. Right: Giving solutions to problem variants from an EFA as in-context examples helps an LLM solve a holdout set of variants from the same EFA. See Sec. 3.2 and Sec. 3.3 for details.

Faithfulness (S	ec. 3.2): EFA helps o	n the original problem	Learnability (Sec. 3.3): EFA helps on its variants			
Initial Pass@1	al Pass@1 +Data from EFA Sample Size			Initial Pass@1 +Data from EFA		
15.66	38.73 (+23.07%)	307	14.58	31.23 (+16.65%)	1,000	

Additionally, the pass rate for Level 5 problems increased from 65.95% to 78.73%. These changes indicate substantial improvements in the model's ability to infer EFA across these dataset slices. The final model trained for 5 iterations becomes the basis for our EFAGen method.

3.2 FAITHFULNESS: EFAS CAPTURE THE REASONING REQUIRED TO SOLVE SEED PROBLEMS

To evaluate the faithfulness of EFAs, we ask: can the generated variant problems improve a model's solve rate on the original seed problem? We select all of problems from MATH-Hard for which Llama3.1-8B-Instruct's pass@5 rate <50% and for which EFAGen can successfully infer an EFA using the gold solution.² For each problem, we sample additional problem variants (we ensure their parameters differ from the seed problem) until Llama3.1-8B-Instruct solves one correctly. We then check if Llama3.1-8B-Instruct can solve the original problem, given the variant and its solution as an in-context example. Results in Table 1 (left) show a 23.07% absolute improvement in pass@1 rate, i.e., EFA-generated variants can teach model the problem-solving needed for the seed problem.

3.3 Learnability: Generated Problem Performance Increases with Experience

An effective problem abstraction should enable a model to solve both the original seed problem and its variants. Thus, we test whether training on EFA-generated problem variants helps a model solve additional variants that are drawn from the same EFA but are different from the seed problem. We sample 1k EFAs inferred from the MATH-Hard test set and generate one new variant per EFA, forming a held-out set. For each EFA, we also identify one variant that Llama3.1-8B-Instruct solves correctly. We then test Llama3.1-8B-Instruct's performance on the held-out set, with and without access to that solved variant as an in-context example. As shown in Table 1, access to one correctly-solved variant improves the model's pass rate on other variants by 16.65% on average; demonstrating that reasoning learned from one variant reliably transfers to others within the same abstraction.

3.4 AUGMENTATION: EFAS ARE EFFECTIVE AT EXPANDING STATIC MATH DATASETS

While high-quality math datasets exist, these are often expensive to construct. EFAGen offers a scalable solution by generating diverse, faithful problem variants through EFAs, thereby augmenting existing datasets. To demonstrate this, we fine-tune Llama3.1-8B-Base using EFA-generated data derived from the MATH training set. Concretely, we annotate 7,500 training problems with step-by-step reasoning from a teacher model (Llama3.1-8B-Instruct). We ensure that the reasoning is correct by filtering out the reasoning that yields incorrect answers. Then, for each of the 7,500 problems, we construct an EFA and sample one problem variant. We compare two training settings. In the first setting, we use only the teacher-labeled seed data. In the second, we augment the seed data by adding EFA-generated examples in a 1:1 ratio. We perform experiments with both 33% (2,500) and 100% (7,500) of the seed data and evaluate performance on three benchmarks: MATH-500 split (Lightman et al., 2023) and the November and December splits of FnEval, each containing perturbed versions of MATH problems. See Appendix H.4 for hyperameter details.

Table 2 shows that EFA-based augmentation leads to consistent improvements across all evaluation metrics: Pass@1, Pass@10 rate, and Majority@25 (Wang et al., 2022), e.g., in the 33% seed setting,

²Based on the intuition that testing for faithfulness requires an EFA (i.e., requires a problem that can be solved in principle) but improving requires a problem that is not solved 100% of the time.

Table 2: **EFAs are effective at data augmentation.** Comparison with and without synthetic data augmentation using problems drawn from generated EFAs. The table shows performance across MATH-500 and FnEval benchmarks (November and December snapshots). When augmenting, we use a 1:1 ratio of examples drawn from training data vs. from an EFA, and report results using 33% of the MATH train set and 100% of the train set.

	MATH-500			FnEval (November Split)			FnEval (December Split)		
Training Data	Pass @ 1	Pass @ 10	Maj @ 25	Pass @ 1	Pass @ 10	Maj @ 25	Pass @ 1	Pass @ 10	Maj @ 25
MATH (33%) +EFA (1:1)	22.4 24.3 (+1.9%)	56.4 58.3 (+1.9%)	36.8 38.8 (+2.0%)	24.5 26.7 (+2.2%)	55.3 59.2 (+3.9%)	39.6 41.8 (+2.2%)	24.4 26.6 (+2.2%)	55.4 57.3 (+1.9%)	39.3 41.2 (+1.9%)
MATH (100%) +EFA (1:1)	24.3 26.1 (+1.8%)	57.8 60.6 (+2.8%)	37.0 40.4 (+3.4%)	26.8 29.3 (+2.5%)	58.6 60.1 (+1.5%)	43.1 44.3 (+1.2%)	26.5 28.8 (+2.3%)	57.6 59.6 (+2.0%)	41.5 43.7 (+2.2%)

Pass@1 improves by +1.9 on MATH-500 and by +2.2 on both FnEval splits. In the 100% seed setting, the gain still holds, underscoring the value of EFAs in enhancing data quality and model performance.

3.5 EFAGEN COMPLEMENTS EXISTING SYNTHETIC DATA GENERATION APPROACHES

EFAs are designed to complement, not replace, existing synthetic data generation approaches. To demonstrate this complementary relationship, we conduct experiments with high-quality synthetic data from NuminaMath (Li et al., 2024), which aggregates synthetic data from various sources, showing that EFAGen can infer EFAs for synthetic data and use these EFAs to augment synthetic datasets at different scales.

We sample 1k, 2.5k, and 5k problems with step-by-step solutions from the synthetic_math and synthetic_amc sources in NuminaMath. For each sample, we apply EFAGen to infer EFAs, generate one problem variant from each EFA, and use rejection sampling to create training data from the EFAs. We train three models at each scale: one trained only on the NuminaMath synthetic data (*NuminaMath Synthetic*), one trained only on data derived from EFAs (*EFA Generated*), and one trained on the NuminaMath synthetic data augmented with our EFA-derived data (*NuminaMath Synthetic + EFA Generated*).

Results on MATH-500 are shown in Table 3. At each scale, the model trained on synthetic data augmented with EFA-generated data performs best across most metrics. Notably, the EFA-generated data typically outperforms the original synthetic NuminaMath data, suggesting that the EFA inference process produces high-quality problem variants that enhance model learning. These results demonstrate that EFAGen provides a scalable approach for augmenting existing synthetic datasets, effectively complementing current synthetic data generation methods.

4 RELATED WORK

Symbolic Approaches to Math Reasoning. A distinct line of prior work has focused on assessing the true mathematical reasoning capabilities of LLMs, specifically by measuring the "reasoning gap" or the drop in math reasoning performance after perturbing questions in existing datasets (Shi et al., 2023; Zhou et al., 2025; Huang et al., 2025; Ye et al., 2025). One prominent approach is to generate different or difficult math questions conditioned on an existing question but test skills by employing frontier models (Zhang et al., 2024; Patel et al., 2025) or human annotators (Srivastava et al., 2024; Shah et al., 2024; Huang et al., 2025). For instance, Srivastava et al. (2024) propose FnEval dataset by manually functionalizing select problems from the MATH dataset (Hendrycks et al., 2021) that can be subsequently used to sample multiple distinct math problems testing similar skills (albeit with different numerical variables). Similarly, Mirzadeh et al. (2025) release the GSM-Symbolic dataset that augments the existing GSM8K dataset (Cobbe et al., 2021) with templates containing placeholders for several numeric and textual variables and can be used to sample distinct math word problems for a robust evaluation of LLM's reasoning abilities. In contrast, to this line of work requiring expensive annotations from humans or frontier models (thereby, hindering scalability) and tailored to specific, predefined math datasets (c.f. Fig. 2); we propose EFAGen that automatically functionalizes any math problem using relatively small language models making it widely-applicable and scalable, i.e., able to sample a potentially infinite number of related math problems from any

Table 3: **EFAGen complements existing synthetic data generation approaches.** Performance comparison across different data scales (1k, 2.5k, 5k) when training models on: NuminaMath synthetic data alone, EFA-generated data alone, and both combined. The combined approach typically performs best, with EFA-generated data generally outperforming the original synthetic data. The (+%) values show absolute improvements over the NuminaMath Synthetic baseline within each scale. 1st-place is **bold**, 2nd is *italicized*.

		MATH-500 Performance						
Scale	Data Mix	Pass@1	Pass@5	Pass@10	MV Acc			
	NuminaMath Synthetic	20.8	45.6	56.4	38.6			
1k	EFA Generated	24.0	48.5	58.7	38.6			
	NuminaMath Synthetic + EFA Generated	24.4	48.5	58.2	40.6			
	,	(+3.7%)	(+2.9%)	(+1.8%)	(+2.0%)			
	NuminaMath Synthetic	23.0	47.6	58.5	38.8			
2.5k	EFA Generated	23.1	47.0	57.2	35.8			
	NuminaMath Synthetic + EFA Generated	24.9	50.5	61.1	41.6			
	•	(+1.9%)	(+2.9%)	(+2.6%)	(+2.8%)			
	NuminaMath Synthetic	20.9	46.3	57.0	39.8			
5k	EFA Generated	23.6	48.6	59.2	39.8			
	NuminaMath Synthetic + EFA Generated	26.7	51.9	62.1	44.0			
	•	(+5.8%)	(+5.6%)	(+5.0%)	(+4.2%)			

distribution or dataset. Moreover, the prior work only focuses on the evaluation of LLMs, whereas we extend the concept of abstraction for downstream applications via training, as shown in Sec. 3.4.

Data and Environment Generation. Past work has generally approached improving models on reasoning tasks like math by generating large amounts of broad-coverage training data. This trend builds on work in generating instruction-tuning data (Wang et al., 2023), where model-generated instructions have been used to teach models to follow prompts. Luo et al. (2023) introduced generation method based on Evol-Instruct (Xu et al., 2023), which augmented a seed dataset of math problems by generating easier and harder problems. Related lines of work have sought to expand datasets by augmenting existing math datasets (Yu et al., 2024), adding multiple reasoning strategies (Yue et al., 2024), covering challenging competition problems (Li et al., 2024), or curating responses (Liu et al., 2024). The data generated in these settings differs from our data in a number of respects: first, it is generally broad-coverage, focusing on large-scale diverse data, as opposed to targeted, instance-specific data. This direction was also explored by Khan et al. (2025), who define data generation agents that can generate specific data based on a particular model's weaknesses, covering math and several other domains. Finally, past work that has augmented a seed dataset (e.g., Yu et al. (2024); Yue et al. (2024)) has done so by modifying problems in the surface form, whereas our method first infers a latent structure and then creates problems by sampling from the structure. In contrast, EFAGen focuses on generating similar examples of existing data by inferring an underlying structure from an example; we show that this has applications to data generation for augmentation but also for stress-testing or measuring the performance gap of models on similar problems.

5 Conclusion

We introduce Executable Functional Abstractions (EFA), a representation of the abstracted logic of a math problem in a parameterized form, enabling the automated sampling of variant problems. We then propose EFAGen, a framework that infers EFAs via program synthesis using large language models (LLMs) that we train using rewards from EFA execution. Our approach over-generate EFA candidates with an LLM and filters them using a suite of property tests that verify their validity. We show that EFAGen successfully infers EFAs for diverse math problems and incorporating execution feedback as a reward in a simple self-training scheme further improves its performance. Models trained on EFA-generated math problems not only perform better on the generated variants but also improve accuracy on the original seed problems. Finally, we show that EFAs provide a scalable solution for augmenting diverse problem variants across various math datasets.

ETHICS STATEMENT

In this work, we propose an inference-time method, EFAGen that can be used sample additional math problems for training or testing. Consequently, the LLMs utilized by EFAGen may still exhibit stereotypes, biases, and other negative traits inherent in their pre-training data (Weidinger et al., 2021), over which we have no control. Therefore, the outputs produced by EFAGen carry the same potential for misuse as those from other test-time methods. Further research is necessary to assess and mitigate these biases in LLMs. Additionally, care must be taken when executing LLM-generated code which can be erroneous and cause unrecoverable changes to the system files.

REPRODUCIBILITY STATEMENT

We will open source our code and data to aid replication of our findings. We also provide implementation details of EFAGen in Sec. 2 and prompts in Appendix H. The math datasets we use are all publicly available.

REFERENCES

- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. Babyai: A platform to study the sample efficiency of grounded language learning. *arXiv preprint arXiv:1810.08272*, 2018.
- Noam Chomsky. Syntactic Structures. Mouton, The Hague, 1957.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, KaShun Shum, and Tong Zhang. RAFT: Reward ranked finetuning for generative foundation model alignment. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=m7p507zblY.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *NeurIPS*, 2021.
- Kaixuan Huang, Jiacheng Guo, Zihao Li, Xiang Ji, Jiawei Ge, Wenzhe Li, Yingqing Guo, Tianle Cai, Hui Yuan, Runzhe Wang, et al. Math-perturb: Benchmarking llms' math reasoning abilities against hard perturbations. *arXiv preprint arXiv:2502.06453*, 2025.
- Zaid Khan, Elias Stengel-Eskin, Jaemin Cho, and Mohit Bansal. Dataenvgym: Data generation agents in teacher environments with student feedback. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40:e253, 2017.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tulu 3: Pushing Frontiers in Open Language Model Post-Training, December 2024. URL http://arxiv.org/abs/2411.15124. arXiv:2411.15124 [cs].

Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif Rasul, Longhui Yu, Albert Q Jiang, Ziju Shen, et al. Numinamath: The largest public dataset in ai4maths with 860k pairs of competition math problems and solutions. *Hugging Face repository*, 13:9, 2024.

Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. Competition-level code generation with alphacode. *Science*, 378(6624):1092–1097, 2022.

- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.
- Zihan Liu, Yang Chen, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Acemath: Advancing frontier math reasoning with post-training and reward modeling. *arXiv preprint arXiv:2412.15084*, 2024.
- Llama Team. The Llama 3 Herd of Models, 2024.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*, 2023.
- Seyed Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad Farajtabar. GSM-symbolic: Understanding the limitations of mathematical reasoning in large language models. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Richard Montague et al. Universal grammar. 1974, pp. 222–46, 1970.
- Barbara H Partee. Compositionality in formal semantics: Selected papers. John Wiley & Sons, 2008.
- Arkil Patel, Siva Reddy, and Dzmitry Bahdanau. How to get your llm to generate challenging problems for evaluation. *arXiv preprint arXiv:2502.14678*, 2025.
- Vedant Shah, Dingli Yu, Kaifeng Lyu, Simon Park, Nan Rosemary Ke, Michael Curtis Mozer, Yoshua Bengio, Sanjeev Arora, and Anirudh Goyal. AI-assisted generation of difficult math questions. In *The 4th Workshop on Mathematical Reasoning and AI at NeurIPS'24*, 2024.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael Schärli, and Denny Zhou. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pp. 31210–31227. PMLR, 2023.
- Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, et al. Beyond human data: Scaling self-training for problem-solving with language models. *arXiv preprint arXiv:2312.06585*, 2023.
- George Smith. Newton's Philosophiae Naturalis Principia Mathematica. In Edward N. Zalta and Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2024 edition, 2024.
- Saurabh Srivastava, Anto PV, Shashank Menon, Ajay Sukumar, Alan Philipose, Stevin Prince, and Sooraj Thomas. Functional benchmarks for robust evaluation of reasoning performance, and the reasoning gap. *arXiv preprint arXiv:2402.19450*, 2024.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-Consistency Improves Chain of Thought Reasoning in Language Models. 2022. URL http://arxiv.org/abs/2203.11171.

Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13484–13508, 2023.

- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359*, 2021.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. URL https://www.aclweb.org/anthology/2020.emnlp-demos.6.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. Wizardlm: Empowering large language models to follow complex instructions. *arXiv* preprint arXiv:2304.12244, 2023.
- Yue Yang, Fan-Yun Sun, Luca Weihs, Eli VanderBilt, Alvaro Herrasti, Winson Han, Jiajun Wu, Nick Haber, Ranjay Krishna, Lingjie Liu, et al. Holodeck: Language guided generation of 3d embodied ai environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16227–16237, 2024.
- Tian Ye, Zicheng Xu, Yuanzhi Li, and Zeyuan Allen-Zhu. Physics of language models: Part 2.1, grade-school math and the hidden reasoning process. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=Tn5B6Udq3E.
- Longhui Yu, Weisen Jiang, Han Shi, YU Jincheng, Zhengying Liu, Yu Zhang, James Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhu Chen. Mammoth: Building math generalist models through hybrid instruction tuning. In *The Twelfth International Conference on Learning Representations*, 2024.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488, 2022.
- Zhehao Zhang, Jiaao Chen, and Diyi Yang. Darg: Dynamic evaluation of large language models via adaptive reasoning graph. *arXiv* preprint arXiv:2406.17271, 2024.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, Zhangchi Feng, and Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand, 2024. Association for Computational Linguistics. URL http://arxiv.org/abs/2403.13372.
- Yang Zhou, Hongyi Liu, Zhuoming Chen, Yuandong Tian, and Beidi Chen. Gsm-infinite: How do your llms behave over infinitely increasing context length and reasoning complexity? *arXiv* preprint arXiv:2502.05252, 2025.

Table 4: **Low-quality EFAs are naturally filtered out during rejection sampling.** We compare the training data yield rates (percentage of responses that receive non-zero rewards) between good and bad EFAs. Bad EFAs are identified using LLM-based heuristics that flag trivial problems, extraneous variables, or hard-coded values. The low yield rates of bad EFAs mean they contribute minimally to training data.

	Good EFAs	Bad EFAs	Good to Bad Data Ratio
Training Data Yield Rate (1 Answer Attempt)	27.0%	5.04%	5.36 to 1
Training Data Yield Rate (5 Answer Attempts)	39.9%	8.85%	4.51 to 1

A APPENDIX

The section Adversarial Search (Fig. 7) outlines how EFAs can generate challenging problem variants to probe model weaknesses. The Scaling section (Appendix E) investigates the effect of the number of sampled variants per EFA, showing how performance trends with increased augmentation. The Ablation section (Appendix F) analyzes the impact of applying unit tests during EFA generation on downstream data quality. Qualitative Examples (Appendix G) presents representative EFAs spanning several MATH domains, including algebra, number theory, and probability, illustrating the range and structure captured by the method. The Experimental Details section describes all data generation, augmentation, and model training settings—EFA generation (box H.1), rejection finetuning and variant sampling protocols (Appendix H.2), math inference configuration, and details for math-specific training (Appendix H.4).

B QUALITY ANALYSIS: LOW-QUALITY EFAS ARE NATURALLY FILTERED OUT

A potential concern with EFAGen is that the automated EFA generation process may produce low-quality abstractions that could negatively impact training. To address this, we analyze how rejection sampling naturally filters out problematic EFAs during the training data generation process.

We identify "bad" EFAs using an LLM with heuristics that flag abstractions exhibiting common failure modes: trivial problems, extraneous variables, or hard-coded values. We then compare the training data yield rates (the percentage of responses that receive non-zero rewards during rejection sampling) between good and bad EFAs.

As shown in Table 4, low-quality EFAs have significantly lower yield rates compared to good EFAs. With a single answer attempt, bad EFAs contribute training data only 5.04% of the time, compared to 27.0% for good EFAs – a ratio of over 5 to 1 in favor of good data. Even when allowing up to 5 answer attempts, the ratio remains favorable at 4.51 to 1. This demonstrates that as long as rejection sampling or reinforcement learning is used, noisy EFAs naturally filter themselves out, ensuring that good data significantly outnumbers bad data in the final training set.

To further validate the quality of EFA-generated data, we conduct a direct comparison between training exclusively on problem variants generated by EFAs versus training exclusively on real problems from the MATH training set. As shown in Table 5, despite potential noise in rejection-sampled EFA data, models trained on synthetic data achieve nearly identical performance to those trained on real data (22.6% vs 22.4% Pass@1 on MATH-500). This shows that EFA-generated data is as effective as existing math data for model training.

C GENERALITY: EFAGEN CAN WORK ACROSS DIVERSE MATH DOMAINS

Importantly, EFAGen generalizes beyond the distribution of questions in the MATH dataset. As detailed in Fig. 5, our approach successfully infers EFAs across various math sources from the NuminaMath dataset (Li et al., 2024) – ranging from grade-school problems (GSM8K) to national/international competitions (e.g., AMC, AIME, IMO). This demonstrates the broad applicability of EFAs for structuring and scaling math data across diverse domains. We generally see that easier

Table 5: **EFA-generated data performs comparably to real data.** Direct comparison of training exclusively on problem variants generated by EFAs versus training exclusively on real problems from the MATH training set. Despite potential noise in rejection-sampled EFA data, models trained on synthetic data achieve nearly identical performance to those trained on real data.

	MATH-500			Fn	Eval (Novem	ber)	FnEval (December)		
Training Data	Pass @ 1	Pass @ 10	Maj @ 25	Pass @ 1	Pass @ 10	Maj @ 25	Pass @ 1	Pass @ 10	Maj @ 25
Real Data Only	22.4	56.4	36.8	24.4	55.4	39.3	24.5	55.3	39.6
Synthetic Data Only	22.6	58.0	37.8	24.9	56.6	38.3	25.5	57.2	40.0




Figure 5: **EFAGen** can infer **EFAs** for diverse sources of math problems. Here, we show the results of applying EFAGen to infer EFAs for the NuminaMath (Li et al., 2024) dataset, which contains a mix of math problems from a diversity of sources ranging from grade school mathematics (GSM8K) to national/international olympiads (olympiads). EFAGen achieves a nonzero success rate across all sources of problems.

math domains like GSM8K are easier to infer EFAs for than harder domains like AIME or Olympiad problems; nevertheless, EFAGen can infer some successful EFAs even on the hardest domain.

To further demonstrate the scalability of EFAGen, we evaluate its performance on a larger set of 10,000 competition-level math problems from NuminaMath. As shown in Table 6, we are able to successfully infer EFAs at rates of 38.4%, 50.9%, and 40.6% for the Olympiads, Synthetic AMC, and AMC-AIME sources in NuminaMath, respectively. The 95% confidence intervals for each source are significantly above 0% (the lowest is 33.7%), demonstrating that EFAGen can reliably infer EFAs for the hardest problems in large math training datasets.

D ADVERSARIAL SEARCH: EFAGEN CAN FIND HARD PROBLEM VARIANTS

EFAs can also be used for evaluation or as a source of targeted training data by finding hard instances that models struggle with.

To demonstrate this, we randomly sample problems from the MATH training that are correctly solved by a strong model (GPT-4o); we sample N=20 of both Level 1 (easiest) and Level 5 (hardest) problems. For each problem, we construct an EFA using EFAGen and then sample 50 variants from the EFA. We attempt to solve each variant with GPT-4o, and measure for what fraction of problems we are able to find variants among the 50 samples that GPT-4o cannot solve. This is an estimate of the probability that we can use an EFA to sample problems that cannot be solved by the model, even when the seed problem is solvable. The results are shown in Fig. 7 where we see that there is a non-zero probability of finding hard variants to a given problem, even for easy problems (i.e., Level 1 in MATH) and with a strong model like GPT-4o.

EFA for Arithmetic Word Problem

Original Problem When Sophie watches her nephew, she gets out a variety of toys for him. The bag of building blocks has 31 blocks in it. The bin of stuffed animals has 8 stuffed animals has 9 multicolored rings on it. Sophie recently bought a tube of bouncy balls, bringing her total number of toys for her nephew up to 62. How many bouncy balls came in the tube?

```
class Problem(BaseModel):
blocks: int
animals: int
rings: int
total: int

@classmethod
def original(cls) → Self:
    return cls(blocks=31, animals=8, rings=9, total=62)

@classmethod
def sample(cls) → Self:
    blocks = random.randint(20, 50)
    animals = random.randint(5, 15)
    rings = random.randint(5, 15)
    balls = random.randint(5, 15)
    balls = random.randint(10, 40)
    total = blocks + animals + rings + balls
    return cls(blocks=blocks, animals=animals, rings=rings, total=total)

def render(self) → str:
    return (
    f'When Sophie watches her nephew, she gets out a variety of toys for him."
    f'The bag of building blocks has {self.blocks} blocks in it."
    f'The bin of stuffed animals has {self.animals} stuffed animals inside."
    f'The tower of stacking rings has {self.animals} stuffed animals inside."
    f'The tower of stacking rings has {self.nings}
multicolored rings on it."
    "Sophie recently bought a tube of bouncy balls, bringing her total number of toys"
    ffor her nephew up to {self.total}. How many bouncy balls came in the tube?"
    balls = self.total - known
    return str(balls)
```

EFA for Tridiagonal Matrix Eigenvalue Problem




Figure 6: EFAs inferred for problems shown in Fig. 2. On the left is an EFA for a grade-school level math word problem. On the right is an EFA for the tridiagonal matrix eigenvalue problem. EFAs are able to represent both types of problems, despite the wide gap in problem complexity. The sample method constructs mathematical objects with required properties, while the solve method implements a generalized solution for any object constructible by the sample method. See Sec. 2.2 for a more detailed explanation.

Table 6: **EFAGen can infer EFAs for large-scale competition-level mathematics.** Across 10,000 competition-level problems in NuminaMath, we successfully infer EFAs at substantial rates across different sources. The 95% confidence intervals are significantly above 0% (lowest is 33.7%), demonstrating that EFAGen can reliably infer EFAs for the hardest problems available in large math training datasets.

Source	Functionalization Rate (%)	95% CI (%)	Num Problems
Olympiads	38.4	[37.2%, 39.5%]	6,950
Synthetic AMC	50.9	[49.0%, 52.7%]	2,881
AMC-AIME	40.6	[33.7%, 48.4%]	169




Figure 7: **EFAs can find harder variants of problems.** We infer an EFA for a sample of Level 1 (easiest) and Level 5 (hardest) seed problems GPT-40 solves correctly, and generate k variants of each problem. We plot the percentage of seed problems for which a variant that GPT-40 solved incorrectly was found.

E SCALING: EFAGEN SCALES EFFECTIVELY UP TO 16 EXAMPLES PER EFA

To understand the scaling behavior of EFA-based data augmentation, we investigate how performance varies with the number of problem variants generated per EFA. We sample 100 unique EFAs from the MATH training set and vary the number of problem variants generated by each EFA from 1 to 64. For each scaling setting, we train Llama3.1-8B-Base on the generated data and evaluate on MATH-500.

As shown in Table 7, we observe smooth scaling improvements as we increase the number of variants from 1 to 16 examples per EFA, with performance gains plateauing beyond 16 examples. Specifically, Pass@1 improves from 14.1% with 1 example per EFA to 23.8% with 16 examples, while Pass@10 increases from 48.5% to 57.6% over the same range. However, scaling begins to saturate at 32 and 64 examples per EFA, suggesting that sampling too many problem variants from each EFA uniformly may hurt diversity and lead to diminishing returns. The optimal scaling point appears to be around 16 examples per EFA, where three of the four metrics achieve their peak performance.

F ABLATION: UNIT TESTS IMPROVE EFA-BASED DATA AUGMENTATION QUALITY

Despite some errors in EFA generation, we find that the current EFAs are effectively improving performance. When we lower the quality by removing our unit tests, the performance gains from augmentation also decrease. As shown in Table 8, applying unit tests consistently improves performance across all benchmarks and metrics. The unit tests provide an average improvement of 2.2 percentage points on MATH-500 Pass@1, 1.7 percentage points on FnEval November Pass@1, and 2.9 percentage points on FnEval December Pass@1.

Table 7: **EFAGen scales effectively up to 16 examples per EFA.** We train Llama3.1-8B-Base on varying numbers of problem variants generated from each EFA and evaluate on MATH-500. Performance improves smoothly from 1 to 16 examples per EFA, with diminishing returns beyond that point. Bold numbers indicate the best performance for each metric.

Training Data per EFA	Pass@1	Pass@5	Pass@10	Majority Vote Accuracy
1	14.1	37.2	48.5	29.6
2	19.1	42.8	53.3	34.0
4	21.9	45.1	54.7	35.4
8	22.9	46.9	57.4	35.6
16	23.8	47.6	57.6	37.4
32	24.3	46.6	56.4	37.2
64	23.9	45.6	55.2	36.2

Table 8: **Unit tests improve EFA-based data augmentation quality.** We compare the performance of EFA-based data augmentation with and without the unit tests that filter out low-quality EFAs. The unit tests consistently improve performance across all benchmarks, demonstrating their effectiveness in maintaining data quality.

MATH-500			FnEval (November)			FnEval (December)			
Unit Tests	Pass @ 1	Pass @ 10	Maj @ 25	Pass @ 1	Pass @ 10	Maj @ 25	Pass @ 1	Pass @ 10	Maj @ 25
False	20.4	55.2	35.4	24.1	54.7	35.7	22.0	55.2	37.1
True	22.6	58.0	37.8	24.9	56.6	38.3	25.5	57.2	40.0
Δ	(+2.2%)	(+2.8%)	(+2.4%)	(+0.8%)	(+1.9%)	(+2.6%)	(+3.5%)	(+2.0%)	(+2.9%)

In general, we believe there is a tradeoff between the level of noise in generated data and the cost of data generation, and EFAs occupy a generally useful point on the tradeoff curve. We can change the tradeoff and reduce noise by increasing the cost of filtering and data generation. These same issues occur with synthetic data generation approaches. The value of our approach is that data generation can be replaced with program execution rather than a call to a frontier LLM.

G QUALITATIVE EXAMPLES

In this section, we display qualitative examples of EFAs across the MATH training set which were validated by our tests.

Box G.1 | EFA (Algebra)

Original Problem

Solve the equation:

$$\sqrt{5x - 4} + \frac{15}{\sqrt{5x - 4}} = 8$$

Original Solution

Let $u = \sqrt{5x - 4}$. We can rewrite the equation in terms of u:

$$u + \frac{15}{u} = 8$$

Multiply through by u to remove the fraction:

$$u^2 + 15 = 8u$$

Reorganize into a standard quadratic form:

$$u^2 - 8u + 15 = 0$$

Factorize the quadratic:

918 919 (u-5)(u-3) = 0920 Thus, u = 5 or u = 3. For u = 5: 921 922 $\sqrt{5x-4} = 5 \Rightarrow 5x-4 = 25 \Rightarrow 5x = 29 \Rightarrow x = \frac{29}{5}$ 923 924 For u=3: 925 $\sqrt{5x-4} = 3 \Rightarrow 5x-4 = 9 \Rightarrow 5x = 13 \Rightarrow x = \frac{13}{5}$ 926 927 Therefore, the solutions are $x = \frac{29}{5}, \frac{13}{5}$. Conclusion: The possible values of x are $\frac{29}{5}$ and $\frac{13}{5}$, which satisfy the original equation when we substitute back, confirming the correctness of 928 929 our solution. 930 931 Inferred EFA 932 class Problem(BaseModel): 933 coefficient: int 934 constant: int 935 term: int 936 937 @classmethod def original(cls) ->Self: 938 return cls(coefficient=5, constant=4, term=15) 939 940 @classmethod 941 def sample(cls) ->Self: coefficient = random.randint(3, 10) 942 constant = random.randint(1, 5) 943 term = random.randint(10, 20) 944 return cls(coefficient=coefficient, constant=constant, term= 945 term) 946 def render(self) ->str: 947 return f"""Solve the equation: 948 \\[\\sqrt{{{self.coefficient}x-{self.constant}}} + \\frac{{{self. 949 \\]""" 951 952 def solve(self) ->str: u = sympy.symbols('u')953 equation = sympy.Eq(u + self.term / u, 8)954 u_values = sympy.solve(equation, u) 955 $x_solutions = []$ 956 for u_val in u_values: $x_{equation} = sympy.Eq(u_val ** 2, self.coefficient * sympy.$ 957 958 symbols('x') - self.constant) 959 $x_solution = sympy.solve(x_equation, sympy.symbols('x'))$ 960 x_solutions.extend(x_solution) return ', '.join(961 f'\\frac{{{sol.as_numer_denom()[0]}}}{{{sol.as_numer_denom 962 ()[1]}}}' 963 for sol in x_solutions) 964 Variant generated by EFA 965 Solve the equation: 966 $\sqrt{7x - 5} + \frac{14}{\sqrt{7x - 5}} = 8$ 967 968

Solution:

969

970 971

$$\frac{23-8\sqrt{2}}{7}, \frac{8\sqrt{2}+23}{7}$$

Box G.2 | EFA (Number Theory)

Original Problem

How many positive divisors does 8! have?

Original Solution

First, calculate $8!: 8! = 1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8 = 40320$. Next, find the prime factorization of $40320: 40320 = 2^7 \times 3^2 \times 5^1 \times 7^1$. Now, apply the formula for counting the divisors: If $n = p^a \times q^b \times r^c \times \ldots$, then the number of divisors t(n) is given by:

$$t(n) = (a+1)(b+1)(c+1)\dots$$

Here a=7,b=2,c=1,d=1 for the primes 2,3,5, and 7 respectively. Applying the formula:

$$t(40320) = (7+1)(2+1)(1+1)(1+1) = 8 \times 3 \times 2 \times 2 = 96$$

Conclusion: The result is consistent with the factorial and prime factorization, providing a logically correct count of divisors.

Inferred EFA

```
class Problem(BaseModel):
   n: int.
   @classmethod
   def original(cls) ->Self:
      return cls(n=8)
   @classmethod
   def sample(cls) ->Self:
      n = random.randint(4, 10)
      return cls(n=n)
   def render(self) ->str:
      return f'How many positive divisors does {self.n}! have?'
   def solve(self) ->str:
      factorial_value = math.factorial(self.n)
      factors = sympy.factorint(factorial_value)
      divisor count = 1
      for exponent in factors.values():
         divisor_count *= exponent + 1
      return str(divisor_count)
```

Variant generated by EFA

How many positive divisors does 9! have? *Solution*:

Box G.3 | EFA (Probability)

Original Problem

Two 8-sided dice are tossed. What is the probability that the sum of the numbers shown on the dice is a prime number? Express your answer as a common fraction.

Original Solution

Let d_1 and d_2 be the outcomes of the two 8-sided dice, where $d_1, d_2 \in \{1, 2, \dots, 8\}$. The total number of possible outcomes in the sample space is:

$$|\Omega| = 8 \times 8 = 64$$

We want to find the number of outcomes where the sum $S=d_1+d_2$ is a prime number. The smallest possible sum is 1+1=2 and the largest is 8+8=16. The prime numbers in this range are 2,3,5,7,11, and 13.

```
1026
           We enumerate the pairs (d_1, d_2) corresponding to each prime sum:
1027
1028
         Sum = 2: \{(1,1)\}
                                                                                      (1 way)
1029
         Sum = 3: \{(1,2),(2,1)\}
                                                                                     (2 ways)
1030
         Sum = 5: \{(1,4), (4,1), (2,3), (3,2)\}
                                                                                     (4 ways)
1031
         Sum = 7: \{(1,6), (6,1), (2,5), (5,2), (3,4), (4,3)\}
                                                                                     (6 ways)
1032
        Sum = 11: \{(3,8), (8,3), (4,7), (7,4), (5,6), (6,5)\}
                                                                                     (6 ways)
1033
1034
        Sum = 13: \{(5,8),(8,5),(6,7),(7,6)\}
                                                                                     (4 ways)
1035
           The total number of favorable outcomes is the sum of the ways for each prime:
1036
                                  N_{\rm favorable} = 1 + 2 + 4 + 6 + 6 + 4 = 23
1037
           The probability is the ratio of the number of favorable outcomes to the total number of
1038
           outcomes:
1039
                                   P(\text{Sum is prime}) = \frac{N_{\text{favorable}}}{|\Omega|} =
1040
1041
           The final answer is:
1042
1043
1044
           Inferred EFA
1045
1046
           class Problem (BaseModel):
1047
              sides: int
1048
              @classmethod
1049
              def original(cls) ->Self:
1050
                  return cls(sides=8)
1051
1052
              @classmethod
              def sample(cls) ->Self:
1053
                  sides = random.randint(4, 20)
1054
                  return cls(sides=sides)
1055
1056
              def solve(self) ->str:
1057
                  total_outcomes = self.sides * self.sides
1058
                  def is_prime(n: int) ->bool:
1059
                      if n <= 1:
                         return False
1061
                      for i in range(2, int(math.sqrt(n)) + 1):
1062
                         if n % i == 0:
                             return False
1063
                      return True
1064
                  primal_sum_occurrences = 0
1065
                  for die1 in range(1, self.sides + 1):
1066
                      for die2 in range(1, self.sides + 1):
1067
                         sum\_of\_dice = die1 + die2
                         if is_prime(sum_of_dice):
1068
                             primal_sum_occurrences += 1
1069
                  probability = primal_sum_occurrences / total_outcomes
1070
                  fraction = sympy.Rational(primal_sum_occurrences,
1071
                      total_outcomes)
1072
                  return f' \setminus frac\{\{\{fraction.numerator\}\}\}\{\{\{fraction.numerator\}\}\}\}
                      denominator}}'
1073
1074
              def render(self) ->str:
1075
                  return (
1076
                      f'Two {self.sides}-sided dice are tossed. What is the
1077
                          probability that the sum of the numbers shown on the
                          dice is a prime number? Express your answer as a common
1078
                            fraction.'
1079
```

1084

1085

1086

1087 1088 1089

1090 1091

1092 1093

1094

1095

1096

1098 1099 1100

1101

1102

1103

1104 1105

1106

1107

1108 1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

11221123

1124

1125

1126

1127 1128

1129

1130

1131

1132 1133 **Variant generated by EFAGen** Two 19-sided dice are tossed. What is the probability that the sum of the numbers shown on the dice is a prime number? Express your answer as a common fraction.

Solution:

 $\frac{105}{361}$

H EXPERIMENTAL DETAILS

H.1 GENERATING EFAS

When generating EFAs, we use the prompt in box H.1. To sample multiple candidates for EFAs, we use beam search with a temperature of 0.7 and a max generation length of 4096. We extract the resulting EFAs from the LLMs response by looking for a markdown code block and extracting all markdown code blocks that have the necessary class structure.

Box H.1 | Prompt for Inferring EFAs

- # Instructions for Math Problem Functionalization
- Your task is to convert a mathematical problem and its solution into a reusable Python class that can generate similar problems. Follow these steps:
- Create a Python class that inherits from BaseModel with parameters that can vary in the problem. These parameters should capture the core numerical or mathematical values that could be changed while maintaining the same problem structure.
- 2. Implement the following required methods:
 - 'original()': A class method that returns the original problem'
 s parameters
 - 'sample()': A class method that generates valid random
 parameters for a similar problem
 - 'render()': An instance method that produces the problem statement as a formatted string
 - 'solve()': An instance method that computes and returns the solution
- 3. For the 'sample()' method:
 - Generate random parameters that maintain the problem's mathematical validity
 - Include appropriate constraints and relationships between parameters
 - Use reasonable ranges for the random values
- 4. For the 'render()' method:
 - Format the problem statement using f-strings
 - Include proper mathematical notation using LaTeX syntax where appropriate
 - Maintain the same structure as the original problem
- 5. For the 'solve()' method:
 - Implement the solution logic using the instance parameters
 - Return the final answer in the expected format (string, typically)
 - Include any necessary helper functions within the method

```
1134
1135
          6. Consider edge cases and validity:
            - Ensure generated problems are mathematically sound
1136
            - Handle special cases appropriately
1137
            - Maintain reasonable complexity in generated problems
1138
1139
         7. Do not import any libraries! The following libraries have been
1140
             imported. Use fully qualified names for all imports:
            - pydantic.BaseModel is imported as 'BaseModel'
1141
            - random is imported as 'random'
1142
            - math is imported as 'math'
1143
            - numpy is imported as 'np'
1144
            - sympy is imported as 'sympy'
            - typing. Self is imported as 'Self'
1145
1146
         Example usage:
1147
          '''python
1148
         problem = MyMathProblem.original() # Get original problem
1149
         variant = MyMathProblem.sample() # Generate new variant
         question = variant.render() # Get problem statement
1150
         answer = variant.solve() # Compute solution
1151
1152
1153
         The goal is to create a class that can both reproduce the original
1154
             problem and generate mathematically valid variations of the same
              problem type.
1155
1156
          # Example 1
1157
          ## Problem Statement
1158
         Evaluate \frac{1^5+i^{-25}+i^{45}}{.}
1159
          ## Solution
1160
         We have i^5 = i^4 \cdot i = 1 \cdot (i) = i. We also have i
1161
              \{-25\} = 1/i^{25} = 1/(i^{24})  = 1/[1\cdot(i)] = 1/i =
1162
              \frac1{i}\cdot\frac{i}{i} = i/(-1) = -i and i^{45} =
1163
          (i^{44})\ cdot i = 1\ and . So, adding these three
1164
             results gives i^5 + i^{-25} + i^{45} = i+-i+i = \boxed{i}.\
             nFinal Answer: The final answer is $\\boxed{ i }$.
1165
1166
          ## Functionalization
1167
          '''python
1168
         class Problem(BaseModel):
1169
            exponent1: int.
1170
            exponent2: int
            exponent3: int
1171
1172
            @classmet.hod
1173
            def original(cls) -> Self:
1174
                return cls(exponent1=5, exponent2=-25, exponent3=45)
1175
            @classmethod
1176
            def sample(cls) -> Self:
1177
                exponent1 = random.randint(-100, 100)
1178
                exponent2 = random.randint(-100, 100)
1179
                exponent3 = random.randint(-100, 100)
1180
                return cls(exponent1=exponent1, exponent2=exponent2,
                   exponent3=exponent3)
1181
1182
            def render(self) -> str:
1183
                return f"Evaluate $i^{{{self.exponent1}}} + i^{{{self.}}
1184
                    exponent2}}} + i^{{self.exponent3}}}."
1185
            def solve(self) -> str:
1186
                # Compute the values of i^n mod 4 cycle
1187
```

```
1188
1189
                def compute_i_power(exp: int) -> complex:
                    cycle = [1, 1j, -1, -1j] # 1, i, -1, -i
1190
                    return cycle[exp % 4]
1191
1192
                # Compute each term
1193
                term1 = compute_i_power(self.exponent1)
1194
                term2 = compute_i_power(self.exponent2)
                term3 = compute_i_power(self.exponent3)
1195
1196
                # Calculate the sum
1197
                result = term1 + term2 + term3
1198
1199
                # Express as LaTeX
                result_latex = (
1200
                    f"{result:.0f}" if result.imag == 0 else str(result).
1201
                        replace("j", "i")
1202
1203
                return f"{result_latex}"
1204
1205
          # Example 2
1206
          ## Problem Statement
1207
          Altitudes \operatorname{AX}\ and \operatorname{BY}\ of acute triangle
1208
              ABC intersect at H. If \alpha BAC = 43^\circ circ and \alpha GAC
              ABC = 67^\circ circ, then what is \alpha = 67^\circ circ?
1209
          ## Solution
1210
          First, we build a diagram:
1211
1212
1213
          size(150); defaultpen(linewidth(0.8));
          pair B = (0,0), C = (3,0), A = (1.2,2), P = foot(A,B,C), Q = foot(B,
1214
              A,C),H = intersectionpoint(B--Q,A--P);
1215
          draw(A--B--C--cycle);
1216
          draw(A--P^{-B}-Q);
1217
          pair Z;
1218
          Z = foot(C,A,B);
          draw(C--Z);
1219
          label("$A$",A,N); label("$B$",B,W); label("$C$",C,E); label("$X$",P,
1220
              S); label("\$Y$",Q,E); label("\$H$",H+(0,-0.17),SW);
1221
          label("$Z$",Z,NW);
1222
          draw(rightanglemark(B,Z,H,3.5));
1223
          draw(rightanglemark(C,P,H,3.5));
1224
          draw(rightanglemark(H,Q,C,3.5));
1225
1226
          Since altitudes \operatorname{AX} and \operatorname{BY} intersect at
1227
              $H$, point $H$ is the orthocenter of $\triangle ABC$. Therefore,
1228
               the line through $C$ and $H$ is perpendicular to
          side \odots overline{AB}$, as shown. Therefore, we have \adots me HCA = \
1229
              angle ZCA = 90^\circ circ - 43^\circ circ = \boxed{47^\circ circ}$.
1230
1231
          ## Functionalization
1232
          '''python
1233
          class Problem(BaseModel):
             angle_BAC: int # angle BAC in degrees
1234
             angle_ABC: int # angle ABC in degrees
1235
1236
             @classmethod
1237
             def original(cls) -> Self:
1238
                return cls(angle_BAC=43, angle_ABC=67)
1239
             @classmethod
1240
             def sample(cls) -> Self:
1241
```

```
1242
1243
                # Generate random acute angles that form a valid triangle
                # Sum of angles must be less than 180
1244
                angle1 = random.randint(30, 75) # Keep angles acute
1245
                angle2 = random.randint(30, 75)
1246
                # Ensure the third angle is also acute
1247
                if angle1 + angle2 >= 150:
1248
                   angle1 = min(angle1, 60)
                   angle2 = min(angle2, 60)
1249
                return cls(angle_BAC=angle1, angle_ABC=angle2)
1250
1251
            def solve(self) -> str:
1252
                # The angle HCA is complementary to angle BAC
1253
                # This is because H is the orthocenter and CH is
                    perpendicular to AB
1254
                angle_HCA = 90 - self.angle_BAC
1255
                return f"{angle_HCA}"
1256
1257
            def render(self) -> str:
                return (
1258
                   f"Altitudes \langle AX \rangle \ and \langle BY \rangle \ of
1259
                       acute triangle $ABC$ "
1260
                   f"intersect at $H$. If $\\angle BAC = {self.angle_BAC}^\\
1261
                       circ$ and "
1262
                   f"\$\\ ABC = \{self.angle\_ABC\}^\\ then \ what \ is \ \$\\ 
1263
                       angle HCA$?"
1264
          . . .
1265
1266
          # Example 3
1267
          ## Problem Statement
         On a true-false test of 100 items, every question that is a
1268
             multiple of 4 is true, and all others are false. If a student
1269
             marks every item that is a multiple of 3 false and all others
1270
             true, how many of the 100 items will be correctly answered?
1271
          ## Solution
1272
         The student will answer a question correctly if
1273
         Case 1: both the student and the answer key say it is true. This
1274
             happens when the answer is NOT a multiple of 3 but IS a multiple
1275
              of 4.
1276
1277
         Case 2. both the student and the answer key say it is false. This
             happens when the answer IS a multiple of 3 but is NOT a multiple
1278
              of 4.
1279
1280
         Since the LCM of 3 and 4 is 12, the divisibility of numbers (in our
1281
              case, correctness of answers) will repeat in cycles of 12. In
1282
              the first 12 integers, $4$ and $8$ satisfy Case 1
         and $3,6,$ and $9$ satisfy Case 2, so for every group of 12, the
1283
             student will get 5 right answers. Since there are 8 full groups
1284
             of 12 in 100, the student will answer at least $8
1285
          \coloredge \cdot 5 = 40$ questions correctly. However, remember that we must
1286
             also consider the leftover numbers 97, 98, 99, 100 and out of
1287
             these, $99$ and $100$ satisfy one of the cases. So
         our final number of correct answers is $40 + 2 = \dots (42)$.
1288
1289
          ## Functionalization
1290
          '''python
1291
         class Problem(BaseModel):
1292
            total questions: int # Total number of questions
            multiple1: int # First multiple (4 in original problem)
1293
            multiple2: int # Second multiple (3 in original problem)
1294
1295
```

```
1296
1297
             @classmet.hod
             def original(cls) -> Self:
1298
                return cls(total_questions=100, multiple1=4, multiple2=3)
1299
1300
             @classmethod
1301
             def sample(cls) -> Self:
1302
                # Generate reasonable random parameters
                total = random.randint(50, 200) # Reasonable test length
1303
                # Choose coprimes or numbers with small LCM for interesting
1304
                    results
1305
                mult1 = random.randint(2, 6)
1306
                mult2 = random.randint(2, 6)
                while mult1 == mult2: # Ensure different numbers
1307
                   mult2 = random.randint(2, 6)
1308
                return cls(total_questions=total, multiple1=mult1, multiple2=
1309
                    mult2)
1310
1311
             def solve(self) -> str:
1312
                def lcm(a: int, b: int) -> int:
                   def gcd(x: int, y: int) -> int:
1313
                      while y:
1314
                         x, y = y, x % y
1315
                      return x
1316
1317
                   return abs(a * b) // gcd(a, b)
1318
                # Find cycle length (LCM)
1319
                cycle_length = lcm(self.multiple1, self.multiple2)
1320
1321
                # Count correct answers in one cycle
1322
                correct_per_cycle = 0
                for i in range(1, cycle_length + 1):
1323
                   answer_key_true = i % self.multiple1 == 0
1324
                   student_true = i % self.multiple2 != 0
1325
                   if answer_key_true == student_true:
1326
                      correct_per_cycle += 1
1327
                # Calculate complete cycles and remainder
1328
                complete_cycles = self.total_questions // cycle_length
1329
                remainder = self.total_questions % cycle_length
1330
1331
                # Calculate total correct answers
1332
                total_correct = complete_cycles * correct_per_cycle
1333
                # Add correct answers from remainder
1334
                for i in range(1, remainder + 1):
1335
                   answer_key_true = i % self.multiple1 == 0
1336
                   student_true = i % self.multiple2 != 0
1337
                   if answer_key_true == student_true:
                      total_correct += 1
1338
1339
                return str(total_correct)
1340
1341
             def render(self) -> str:
               return (
1342
                   f"On a true-false test of {self.total_questions} items,
1343
                   f"every question that is a multiple of {self.multiple1} is
1344
                        true, "
1345
                   f"and all others are false. If a student marks every item
1346
                       that is "
                   f"a multiple of {self.multiple2} false and all others true,
1347
                        how "
1348
1349
```

```
1350
                    f"many of the {self.total_questions} items will be
1351
                        correctly answered?"
1352
1353
1354
1355
          # Your Turn
1356
          Functionalize the following problem:
1357
          ## Problem Statement
1358
          [% problem_statement %]
1359
1360
          ## Solution
          [ 응
             solution %]
             Functionalization
1363
1364
```

H.2 EFAGEN TRAINING DETAILS

When doing rejection finetuning, we sample 20 candidate EFAs programs from the LLM for each seed problem during the rejection sampling phase. We sample 20 variants from each EFA in order to run the has_dof (EFA) and is_single_valued (EFA) tests. When finetuning on the EFAs that pass all tests, we use the the same prompt box H.1 as the instruction and the extracted code of the EFA as the response. We use Transformers (Wolf et al., 2020) and Llama-Factory (Zheng et al., 2024) libraries for training. We format all data in the Alpaca format (Taori et al., 2023) as instruction-response pairs. We use the Adam optimizer with a batch size of 16 and a cosine learning rate scheduler with a warmup ratio of 0.1 and train for 3 epochs in the FP16 datatype. We apply LoRA to all linear layers with a rank of 16 and an alpha of 32, no bias, and a dropout of 0.05. We truncate all training examples to a maximum length of 4096 tokens with a batch size of 32.

H.3 MATH INFERENCE SETTINGS

When doing 0-shot inference with Llama3.1-8B-Instruct, we use the official Llama3.1 prompt in box H.2. When doing few-shot inference with Llama3.1-8B-Instruct, we use a modified version of the official prompt, shown in box H.3. When sampling multiple responses, we use beam search with a temperature of 0.7 and a max generation length of 2048. When sampling a single response, we use beam search with a temperature of 0.0 and a max generation length of 2048. In all cases, we check for equality of answers using the math-verify library.

```
1386
          Box H.2 | Llama3.1 0-shot MATH Prompt
1387
1388
          Solve the following math problem efficiently and clearly:
1389
1390
          - For simple problems (2 steps or fewer):
1391
         Provide a concise solution with minimal explanation.
1392
          - For complex problems (3 steps or more):
1393
         Use this step-by-step format:
1394
1395
          ## Step 1: [Concise description]
          [Brief explanation and calculations]
          ## Step 2: [Concise description]
1398
          [Brief explanation and calculations]
1399
1400
1401
          Regardless of the approach, always conclude with:
1402
1403
```

```
Therefore, the final answer is: $\boxed{answer}$. I hope it is
    correct.

Where [answer] is just the final number or expression that solves
    the problem.

Problem: {{ instruction }}
```

Box H.3 | Llama3.1 N-shot MATH Prompt Solve the following math problem efficiently and clearly: - For simple problems (2 steps or fewer): Provide a concise solution with minimal explanation. - For complex problems (3 steps or more): Use this step-by-step format: \#\# Step 1: [Concise description] [Brief explanation and calculations] \#\# Step 2: [Concise description] [Brief explanation and calculations] . . . Regardless of the approach, always conclude with: Therefore, the final answer is: \$\boxed{answer}\$. I hope it is correct. Where [answer] is just the final number or expression that solves the problem. Here are some examples: {% for few_shot_example in few_shot_examples %} Problem: {{ few_shot_example.instruction }} {{ few_shot_example.response }} {% endfor %}

H.4 MATH TRAINING DETAILS

Problem: {{ instruction }}

We use the same hyperparameters and chat data format as in Appendix H.2, except we cutoff training data over 2048 tokens. However, we use a simpler prompt template, shown in box H.4 to format the teacher responses. When annotating with a Llama3.1-8B-Instruct teacher, we sample 5 responses per math problem with a temperature of 0.7. We check for equality of answers using the math-verify library.

```
Box H.4 | Minimal instruction-tuning prompt used for augmentation experiments

Question: {{ question }}
Step-by-step Answer
```