

000 EXECUTABLE FUNCTIONAL ABSTRACTIONS: INFERRING 001 GENERATIVE PROGRAMS FOR ADVANCED MATH 002 PROBLEMS 003 004

005
 006 **Anonymous authors**
 007 Paper under double-blind review
 008
 009
 010
 011

ABSTRACT

013 Abstract Interpretation provides a framework for approximating the behavior of discrete systems by establishing a correspondence between concrete execution traces
 014 and abstract properties. We apply this framework to mathematics to address the inverse problem: automatically synthesizing a general program (the *abstraction*)
 015 from a single concrete example, which executes to produce specific, valid problem instances (the *concretization*). Prior approaches to capturing this structure rely
 016 on hand-crafted templates, a labor-intensive process that restricts the technique to arithmetic word problems or small datasets. We introduce EFAGen, a method that
 017 operationalizes this inference as a program synthesis task, generating Executable
 018 Functional Abstractions (EFAs) that encode the parameters, constraints, and solution
 019 procedure of the seed problem. Because formal verification of synthesized code is intractable, we filter candidates using executable unit tests that enforce
 020 necessary properties. We demonstrate that these inferred abstractions enable data
 021 augmentation that complements existing strong data mixes for math reasoning and
 022 facilitate adversarial search to discover problem variants that models fail to solve.
 023
 024
 025
 026
 027

1 INTRODUCTION

030 Abstract Interpretation (Cousot & Cousot, 1977) provides a rigorous framework for approximating the behavior of discrete systems. It establishes a correspondence between a *concrete domain* of
 031 specific execution traces and an *abstract domain* of general properties. In the context of mathematics,
 032 we can view a specific problem instance (e.g., “Find the GCD of 6 and 126”) as a point in the concrete
 033 domain. The underlying general logic (variables, constraints, and solution procedure) serves as its
 034 representation in the abstract domain. We term this programmatic representation an Executable
 035 Functional Abstraction (EFA). Possessing the EFA for a problem is powerful. It allows one to analyze
 036 the general class of the problem rather than a single instance and enables the generation of infinite
 037 valid variants through *concretization* functions. These variants have the potential to be useful as a
 038 source of training data or to construct challenging benchmarks for evaluation.
 039

040 However, a reliable “abstraction function”, a mechanism to automatically lift a concrete problem
 041 into a valid EFA, does not exist for complex mathematics. Current approaches to obtaining these
 042 abstractions, such as GSM-Symbolic (Mirzadeh et al., 2025) and FnEval (Srivastava et al., 2024),
 043 rely heavily on manual engineering. Humans must painstakingly identify variables, define domains,
 044 and write code for every problem template. This manual reliance restricts the abstractions to simple
 045 grade-school arithmetic or small, curated datasets is not scalable. Constructing an abstraction function
 046 for complex mathematics poses two fundamental challenges. First, synthesis is difficult: identifying
 047 the correct parameters, discovering non-trivial constraints, and generalizing the solution logic must
 048 all succeed simultaneously. Getting any component wrong yields an invalid abstraction. Second,
 049 verification is intractable: formally proving correctness of these synthesized programs is beyond
 050 current capabilities.

051 Our key insight is to reformulate this open-ended inference problem as a tractable search problem de-
 052 fined by executable code and operational verification. We introduce EFAGen, which operationalizes
 053 the Abstract Interpretation relationship as a program synthesis task. Our inference pipeline acts as an
abstraction function α . It takes a concrete problem instance x as input and synthesizes an EFA that

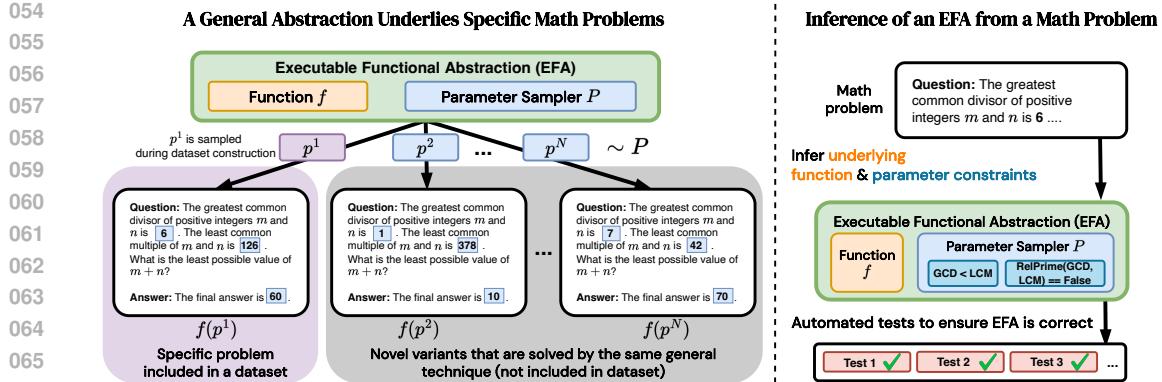


Figure 1: **Left:** We view math problems through the lens of Abstract Interpretation: specific problem instances with concrete values lie in the *concrete domain*, while **executable functional abstractions (EFAs)** represent the *abstract domain* of parameterized logic and constraints. The *concretization function* γ (via `sample()`) generates valid concrete instances from an EFA. **Right:** We study the task of automating the *abstraction function* α that lifts a concrete problem instance into its corresponding EFA, automatically inferring parameters, constraints, and general solution procedures from natural language problems. We approach this as a program synthesis task, and show the validity of the inferred EFAs as well as their utility in downstream tasks.

transforms specific numerical values into typed parameters, encodes constraints between them, and implements a general solution procedure valid for any parameterization satisfying these constraints. Each synthesized EFA implements a *concretization function* γ via its `sample()` method, which instantiates the abstract schema into concrete problem instances. Rather than formal verification, we implement *operational soundness checks* as executable unit tests that verify necessary conditions for validity. These checks ensure that $x \in \gamma(\alpha(x))$ —the abstraction can reproduce the original instance—but also that sampled variants are non-trivial (distinct from the seed), solvable (match expected answers), and valid (satisfy domain constraints). We generate multiple candidate programs using an overgenerate-and-filter approach (Li et al., 2022), treating each EFA as a hypothesis and selecting those that pass all operational checks. This search procedure enables us to discover abstractions that are operationally sound with respect to the seed problem.

We confirm the internal validity of the inferred EFAs by measuring the faithfulness of the generated variants to the seed problem and their utility in training models. We then demonstrate the applications of EFAs to two downstream tasks. Specifically, we show that EFAs can be used for adversarial search to discover harder problem variants and for data augmentation. In the latter, we demonstrate that EFA-generated data is high quality and complementary to existing data augmentation methods. Our experiments show that EFA-based augmentation combined with NuminaMath (?) yields better performance than using NuminaMath alone. This suggests that the inferred abstractions capture structural patterns distinct from those in standard corpora.

We make the following contributions.

- We formalize EFAs and develop EFAGen, an automated approach to infer executable abstractions from competition-level math problems by treating abstraction as a synthesis and verification task.
- We demonstrate that the execution feedback from our validity tests acts as a reward signal. This enables LLMs to self-improve at the abstraction task via reinforcement learning.
- We empirically show that inferred EFAs provide a complementary data source to strong baselines like NuminaMath. Data augmentation with EFAs improves performance on MATH-500 and FnEval, and EFAs can be used to search for easier harder or easier variants of problems.

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

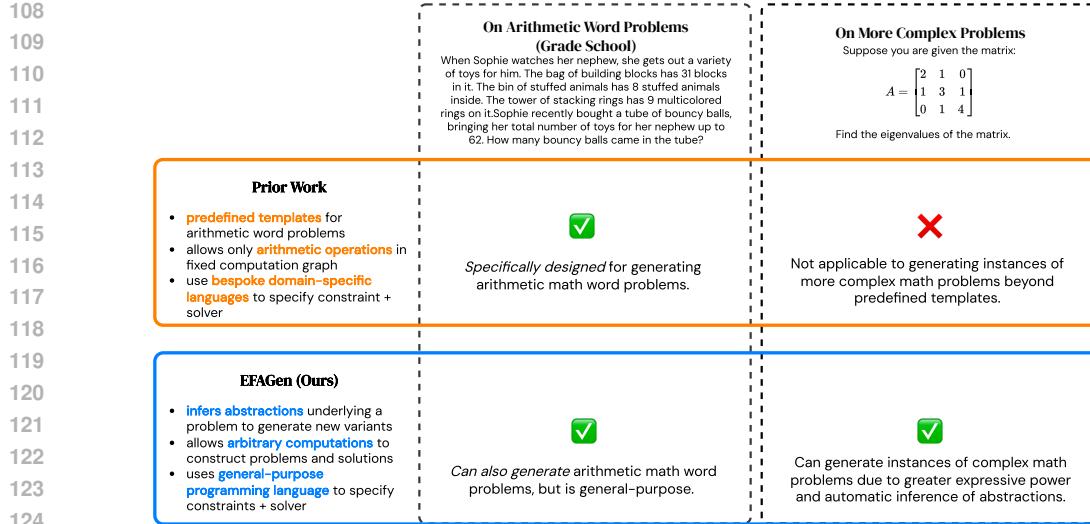


Figure 2: **EFAGen generalizes prior work on constructing arithmetic word problems to automatically constructing more complex, higher-level math problems.** Given a math problem and solution, EFAGen infers an underlying abstraction whose construction and general solution may involve arbitrary computations 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 numerical root-finding procedure. Details of inferred EFA code in Fig. 7.

of new problem instances by varying numerical parameters while preserving the underlying problem-solving 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:

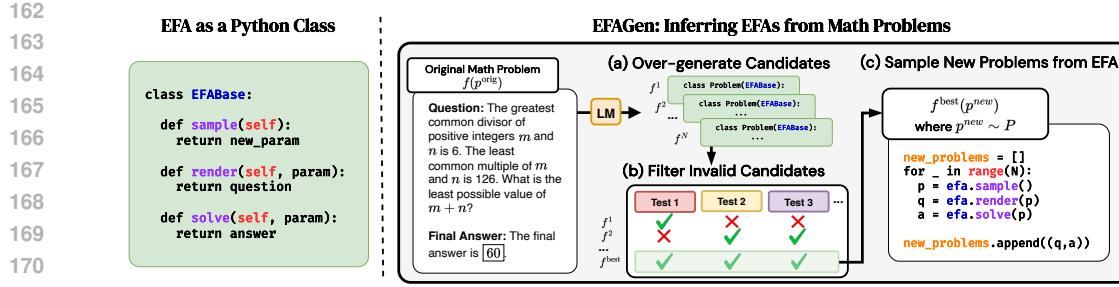


Figure 3: **Left: Representation of an executable functional abstraction (EFA) as a Python class.** **Right: Overview of EFAGen, a method for automatically inferring EFAs from a math problem.** In EFAGen, we (a) over-generate multiple EFA candidates with an LLM and (b) filter out invalid candidates that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

- **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) → 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

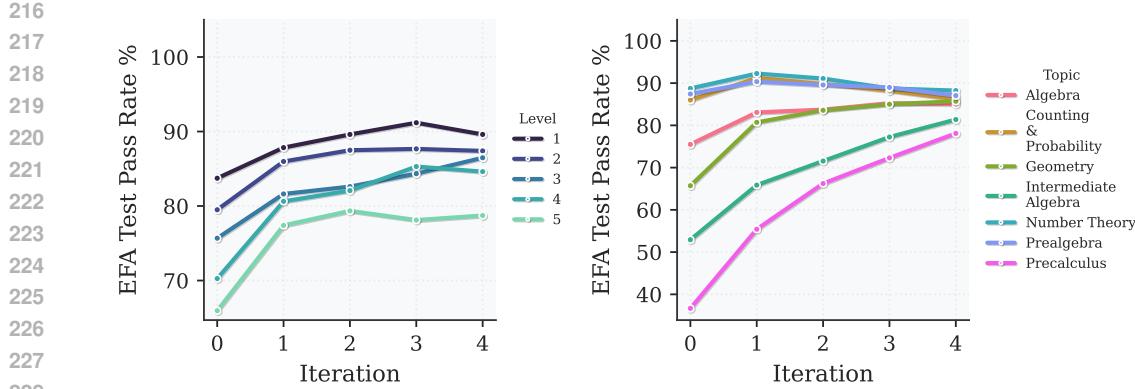


Figure 4: **LLMs can use our tests to self-improve at inferring EFAs.** We plot the percentage of constructed EFAs passing all tests across iterations of self-training, grouped by MATH problem difficulty (left) and by problem category (right). Harder difficulty levels and problem categories are harder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

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 C), applying EFAs to find hard variants of problems (Sec. 3.7), ablations (Appendix D), and EFAGen inference on olympiad-level problems (Sec. 3.6).

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.
- **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 pass 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

270
 271 **Table 1: EFAs faithfully capture the solutions of the problems they were derived from (left), and**
 272 **problem variants constructed by EFAs share learnable structure (right).** Left: Giving solutions
 273 to problems variants from an EFA as in-context examples nearly doubles the solve rate of an LLM
 274 on the seed problem the EFA was derived from. Right: Giving solutions to problem variants from an
 275 EFA as in-context examples helps an LLM solve a holdout set of variants from the same EFA. See
 276 Sec. 3.2 and Sec. 3.3 for details.

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

280
 281
 282 **Table 2: EFAs are effective at data augmentation.** Comparison with and without synthetic data
 283 augmentation using problems drawn from generated EFAs. The table shows performance across
 284 MATH-500 and FnEval benchmarks (November and December snapshots). When augmenting, we
 285 use a 1:1 ratio of examples drawn from training data vs. from an EFA, and report results using 33%
 286 of the MATH train set and 100% of the train set.

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

294
 295 MATH dataset, where it is only able to infer valid EFAs for $\approx 35\%$ of Precalculus questions. In Sec. 2,
 296 we introduce a number of unit tests (i.e., verifiable rewards) that indicate whether a generated EFA is
 297 valid. Here, we show that we can train models to improve on inferring valid EFAs by self-training
 298 according to these tests. Specifically, we use a rejection-finetuning approach (Zelikman et al., 2022;
 299 Singh et al., 2023; Dong et al., 2023), in which we sample EFA candidates from a model and filter
 300 according to our rewards to construct a training dataset of correct examples. We begin with the
 301 MATH training set (7,500 problems) and sample 10 candidate EFAs per problem. Candidates failing
 302 any of the reward checks are discarded. The remaining valid examples form a dataset for supervised
 303 fine-tuning. This process – sampling, filtering, and retraining – is repeated over 5 iterations (see
 304 Appendix F.2 for details).

305 We report the EFA success rates across iterations in Fig. 4, where we group by difficulty levels (left)
 306 and by annotated problem category (right). Success rates steadily improve over training iterations,
 307 especially for harder problems. At iteration 0 (before training), we observe that harder problems (e.g.,
 308 Level 5) are also harder to infer EFAs for, with EFA success rates $\approx 17\%$ lower for Level 5 than Level
 309 1 problems. Similarly, certain categories like ‘Intermediate Algebra’, ‘Counting’ and ‘Probability’
 310 are harder to infer EFAs for. These domains generally see the most significant increases from training.
 311 Between iteration 1 and iteration 5, the Intermediate Algebra’s EFA success rate showed the most
 312 significant increase, rising from 52.93% to 81.38%, and Geometry improved from 65.71% to 85.71%.
 313 Additionally, the pass rate for Level 5 problems increased from 65.95% to 78.73%. These changes
 314 indicate substantial improvements in the model’s ability to infer EFA across these dataset slices. The
 315 final model trained for 5 iterations becomes the basis for our EFAGen method.

3.2 EFAS FAITHFULLY CAPTURE THE REASONING PATTERNS OF SEED PROBLEMS

316
 317 We expect that valid EFAs should be able to capture the reasoning patterns of the seed problem, a
 318 property we call *faithfulness*. We measure faithfulness by checking if seeing solutions to problem
 319 variants generated from an EFA can improve a model’s solve rate on the original seed problem. If an
 320 EFA is faithful, then seeing solutions to problem variants generated from it should improve a model’s
 321 solve rate on the original seed problem. We select all of problems from MATH-Hard for which
 322 Llama3.1-8B-Instruct’s pass@5 rate $< 50\%$ and for which EFAGen can successfully infer an EFA

324
 325 **Table 3: EFAGen complements existing synthetic data generation approaches.** Performance
 326 comparison across different data scales (1k, 2.5k, 5k) when training models on: NuminaMath
 327 synthetic data alone, EFA-generated data alone, and both combined. The combined approach
 328 typically performs best, with EFA-generated data generally outperforming the original synthetic data.
 329 The *(+%)* values show absolute improvements over the NuminaMath Synthetic baseline within each
 330 scale. 1st-place is **bold**, 2nd is *italicized*.

		MATH-500 Performance			
Scale	Data Mix	Pass@1	Pass@5	Pass@10	MV Acc
1k	NuminaMath Synthetic	20.8	45.6	56.4	38.6
	EFA Generated	24.0	48.5	58.7	38.6
	NuminaMath Synthetic + EFA Generated	24.4	48.5	58.2	40.6
2.5k	NuminaMath Synthetic	23.0	47.6	58.5	38.8
	EFA Generated	23.1	47.0	57.2	35.8
	NuminaMath Synthetic + EFA Generated	24.9	50.5	61.1	41.6
5k	NuminaMath Synthetic	20.9	46.3	57.0	39.8
	EFA Generated	23.6	48.6	59.2	39.8
	NuminaMath Synthetic + EFA Generated	26.7	51.9	62.1	44.0

348 using the gold solution.¹ For each problem, we sample additional problem variants (we ensure their
 349 parameters differ from the seed problem) until Llama3.1-8B-Instruct solves one correctly. We then
 350 check if Llama3.1-8B-Instruct can solve the original problem, given the variant and its solution as an
 351 in-context example. Results in Table 1 (left) show a 23.07% absolute improvement in pass@1 rate,
 352 i.e., EFA-generated variants **demonstrate faithfulness to the reasoning pattern required for the seed**
 353 **problem.**

3.3 EFAs ENCODE LEARNABLE, SHARED STRUCTURE

354
 355 We expect that valid EFAs should generate problem variants that share common structure. While
 356 this is hard to define formally, we can informally measure this by checking the *learnability* of the
 357 EFAs. An EFA is learnable if seeing solutions to problem variants generated from it can improve
 358 a model’s solve rate on other variants generated from the same EFA. We sample 1k EFAs inferred
 359 from the MATH-Hard test set and generate one new variant per EFA, forming a held-out set. For
 360 each EFA, we also identify one variant that Llama3.1-8B-Instruct solves correctly. We then test
 361 Llama3.1-8B-Instruct’s performance on the held-out set, with and without access to that solved variant
 362 as an in-context example. As shown in Table 1, access to one correctly-solved variant improves the
 363 model’s pass rate on other variants by 16.65% on average. **This provides evidence that the EFAs**
 364 **encode learnable, shared structure.**

3.4 AUGMENTATION: EFAS ARE EFFECTIVE AT EXPANDING STATIC MATH DATASETS

365
 366 While high-quality math datasets exist, these are often expensive to construct. EFAGen offers a
 367 scalable solution by generating diverse, faithful problem variants through EFAs, thereby augmenting
 368 existing datasets. To demonstrate this, we fine-tune Llama3.1-8B-Base using EFA-generated data
 369 derived from the MATH training set. Concretely, we annotate 7,500 training problems with step-by-
 370 step reasoning from a teacher model (Llama3.1-8B-Instruct). We ensure that the reasoning is correct
 371 by filtering out the reasoning that yields incorrect answers. Then, for each of the 7,500 problems, we
 372 construct an EFA and sample one problem variant. We compare two training settings. In the first
 373 setting, we use only the teacher-labeled seed data. In the second, we augment the seed data by adding
 374

375 ¹Based on the intuition that testing for faithfulness requires an EFA (i.e., requires a problem that can be
 376 solved in principle) but improving requires a problem that is not solved 100% of the time.

378
 379 **Table 4: EFAGen can infer EFAs for large-scale competition-level mathematics.** Across 10,000
 380 competition-level problems in NuminaMath, we successfully infer EFAs at substantial rates across
 381 different sources. The 95% confidence intervals are significantly above 0% (lowest is 33.7%),
 382 demonstrating that EFAGen can reliably infer EFAs for the hardest problems available in large math
 383 training datasets.

Source	Success 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

384 EFA-generated examples in a 1:1 ratio. We perform experiments with both 33% (2,500) and 100%
 385 (7,500) of the seed data and evaluate performance on three benchmarks: MATH-500 split (Lightman
 386 et al., 2023) and the November and December splits of FnEval, each containing perturbed versions of
 387 MATH problems. See Appendix F.4 for hyperparameter details.

388 Table 2 shows that EFA-based augmentation leads to consistent improvements across all evaluation
 389 metrics: Pass@1, Pass@10 rate, and Majority@25 (Wang et al., 2022), e.g., in the 33% seed setting,
 390 Pass@1 improves by +1.9 on MATH-500 and by +2.2 on both FnEval splits. In the 100% seed setting,
 391 the gain still holds, underscoring the value of EFAs in enhancing data quality and model performance.

397 3.5 EFAGen Complements Existing Synthetic Data Generation Approaches

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

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

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

418 3.6 Generality: EFAGen Can Work Across Diverse Math Domains

419 Importantly, EFAGen generalizes beyond the distribution of questions in the MATH dataset. As
 420 detailed in Fig. 5, our approach successfully infers EFAs across various math sources from the
 421 NuminaMath dataset (Li et al., 2024) – ranging from grade-school problems (GSM8K) to nation-
 422 al/international competitions (e.g., AMC, AIME, IMO). This demonstrates the broad applicability
 423 of EFAs for structuring and scaling math data across diverse domains. We generally see that easier
 424 math domains like GSM8K are easier to infer EFAs for than harder domains like AIME or Olympiad
 425 problems; nevertheless, EFAGen can infer some successful EFAs even on the hardest domain.

426 To further demonstrate the scalability of EFAGen, we evaluate its performance on a larger set of
 427 10,000 competition-level math problems from NuminaMath. As shown in Table 4, we are able to
 428 successfully infer EFAs at rates of 38.4%, 50.9%, and 40.6% for the Olympiads, Synthetic AMC, and
 429 AMC-AIME sources in NuminaMath, respectively. The 95% confidence intervals for each source are
 430 significantly above 0% (the lowest is 33.7%), demonstrating that EFAGen can reliably infer EFAs
 431 for the hardest problems in large math training datasets.

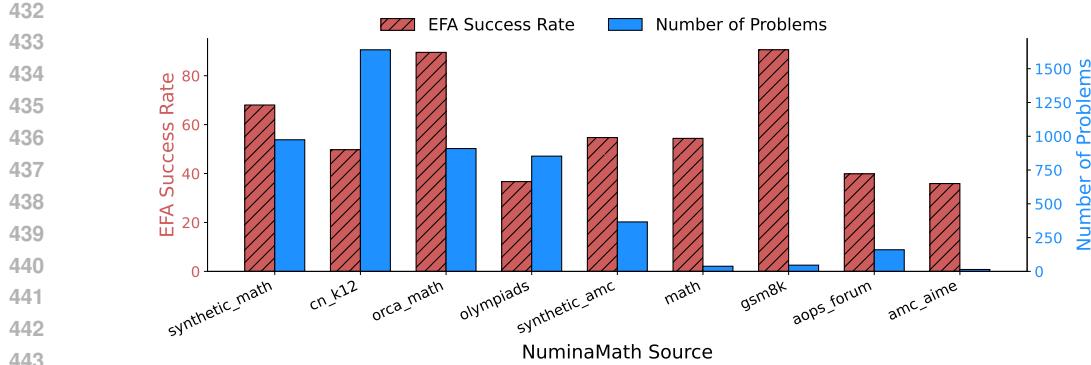


Figure 5: **EFAGen can infer EFAs for diverse sources of math problems.** Here, we show the results of applying EFAGen to infer EFAs for the NuminaMath (Li et al., 2024) dataset, which contains a mix of math problems from a diversity of sources ranging from grade school mathematics (GSM8K) to national/international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

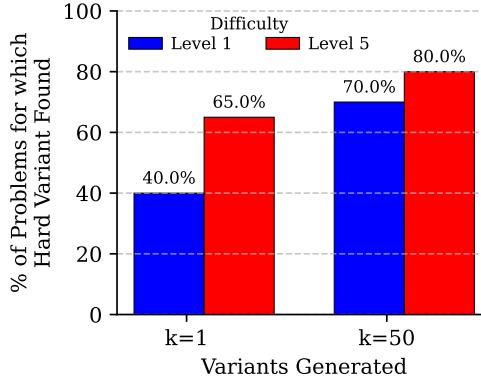


Figure 6: **EFAs can find harder variants of problems.** We infer an EFA for a sample of Level 1 (easiest) and Level 5 (hardest) seed problems GPT-4o solves correctly, and generate k variants of each problem. We plot the percentage of seed problems for which a variant that GPT-4o solved incorrectly was found.

3.7 ADVERSARIAL SEARCH: EFA_{GEN} 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. 6 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.

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.,

486 2023; Zhou et al., 2025; Huang et al., 2025; Ye et al., 2025). One prominent approach is to generate
 487 different or difficult math questions conditioned on an existing question but test skills by employing
 488 frontier models (Zhang et al., 2024; Patel et al., 2025) or human annotators (Srivastava et al., 2024;
 489 Shah et al., 2024; Huang et al., 2025). For instance, Srivastava et al. (2024) propose FnEval dataset
 490 by manually functionalizing select problems from the MATH dataset (Hendrycks et al., 2021) that
 491 can be subsequently used to sample multiple distinct math problems testing similar skills (albeit
 492 with different numerical variables). Similarly, Mirzadeh et al. (2025) release the GSM-Symbolic
 493 dataset that augments the existing GSM8K dataset (Cobbe et al., 2021) with templates containing
 494 placeholders for several numeric and textual variables and can be used to sample distinct math word
 495 problems for a robust evaluation of LLM’s reasoning abilities. In contrast, to this line of work
 496 requiring expensive annotations from humans or frontier models (thereby, hindering scalability) and
 497 tailored to specific, predefined math datasets (c.f. Fig. 2); we propose EFAGen that automatically
 498 functionalizes *any* math problem using relatively small language models making it *widely-applicable*
 499 and *scalable*, i.e., able to sample a potentially infinite number of related math problems from any
 500 distribution or dataset. Moreover, the prior work only focuses on the evaluation of LLMs, whereas
 501 we extend the concept of abstraction for downstream applications via training, as shown in Sec. 3.4.

502 **Data and Environment Generation.** Past work has generally approached improving models on
 503 reasoning tasks like math by generating large amounts of broad-coverage training data. This trend
 504 builds on work in generating instruction-tuning data (Wang et al., 2023), where model-generated
 505 instructions have been used to teach models to follow prompts. Luo et al. (2023) introduced generation
 506 method based on Evol-Instruct (Xu et al., 2023), which augmented a seed dataset of math problems
 507 by generating easier and harder problems. Related lines of work have sought to expand datasets
 508 by augmenting existing math datasets (Yu et al., 2024), adding multiple reasoning strategies (Yue
 509 et al., 2024), covering challenging competition problems (Li et al., 2024), or curating responses (Liu
 510 et al., 2024). The data generated in these settings differs from our data in a number of respects:
 511 first, it is generally broad-coverage, focusing on large-scale diverse data, as opposed to targeted,
 512 instance-specific data. This direction was also explored by Khan et al. (2025), who define data
 513 generation agents that can generate specific data based on a particular model’s weaknesses, covering
 514 math and several other domains. Finally, past work that has augmented a seed dataset (e.g., Yu et al.
 515 (2024); Yue et al. (2024)) has done so by modifying problems in the surface form, whereas our
 516 method first infers a latent structure and then creates problems by sampling from the structure. In
 517 contrast, EFAGen focuses on generating similar examples of existing data by inferring an underlying
 518 structure from an example; we show that this has applications to data generation for augmentation
 519 but also for stress-testing or measuring the performance gap of models on similar problems.

5 CONCLUSION

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

531 ETHICS STATEMENT

532 In this work, we propose an inference-time method, EFAGen that can be used sample additional
 533 math problems for training or testing. Consequently, the LLMs utilized by EFAGen may still exhibit
 534 stereotypes, biases, and other negative traits inherent in their pre-training data (Weidinger et al.,
 535 2021), over which we have no control. Therefore, the outputs produced by EFAGen carry the same
 536 potential for misuse as those from other test-time methods. Further research is necessary to assess
 537 and mitigate these biases in LLMs. Additionally, care must be taken when executing LLM-generated
 538 code which can be erroneous and cause unrecoverable changes to the system files.

540 REPRODUCIBILITY STATEMENT
541542 We will open source our code and data to aid replication of our findings. We also provide implemen-
543 tation details of EFAGen in Sec. 2 and prompts in Appendix F. The math datasets we use are all
544 publicly available.
545546 REFERENCES
547548 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared
549 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large
550 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.551 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
552 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
553 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,
554 2021.555 Patrick Cousot and Radhia Cousot. Abstract interpretation: a unified lattice model for static analysis
556 of programs by construction or approximation of fixpoints. In *Proceedings of the 4th ACM
557 SIGACT-SIGPLAN symposium on Principles of Programming Languages, POPL '77*, pp. 238–252,
558 New York, NY, USA, 1977. ACM. doi: 10.1145/512950.512973.
559560 Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao,
561 Jipeng Zhang, KaShun Shum, and Tong Zhang. RAFT: Reward ranked finetuning for generative
562 foundation model alignment. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.
563 URL <https://openreview.net/forum?id=m7p507zb1Y>.564 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
565 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *NeurIPS*,
566 2021.567 Kaixuan Huang, Jiacheng Guo, Zihao Li, Xiang Ji, Jiawei Ge, Wenzhe Li, Yingqing Guo, Tianle
568 Cai, Hui Yuan, Runzhe Wang, et al. Math-perturb: Benchmarking llms' math reasoning abilities
569 against hard perturbations. *arXiv preprint arXiv:2502.06453*, 2025.
570571 Zaid Khan, Elias Stengel-Eskin, Jaemin Cho, and Mohit Bansal. Dataenvgym: Data generation
572 agents in teacher environments with student feedback. In *The Thirteenth International Conference
573 on Learning Representations*, 2025.574 Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman,
575 Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria
576 Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca
577 Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tulu 3: Pushing
578 Frontiers in Open Language Model Post-Training, December 2024. URL <http://arxiv.org/abs/2411.15124>. arXiv:2411.15124 [cs].
579580 Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif
581 Rasul, Longhui Yu, Albert Q Jiang, Ziju Shen, et al. Numinamath: The largest public dataset in
582 ai4maths with 860k pairs of competition math problems and solutions. *Hugging Face repository*,
583 13:9, 2024.585 Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittweis, Rémi Leblond, Tom
586 Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. Competition-level code generation
587 with alphacode. *Science*, 378(6624):1092–1097, 2022.588 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
589 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth
590 International Conference on Learning Representations*, 2023.592 Zihan Liu, Yang Chen, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Acemath: Advancing
593 frontier math reasoning with post-training and reward modeling. *arXiv preprint arXiv:2412.15084*,
2024.

594 Llama Team. The Llama 3 Herd of Models, 2024.
 595

596 Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng,
 597 Qingwei Lin, Shifeng Chen, and Dongmei Zhang. Wizardmath: Empowering mathematical
 598 reasoning for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*,
 599 2023.

600 Seyed Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and
 601 Mehrdad Farajtabar. GSM-symbolic: Understanding the limitations of mathematical reasoning in
 602 large language models. In *The Thirteenth International Conference on Learning Representations*,
 603 2025.

604 Arkil Patel, Siva Reddy, and Dzmitry Bahdanau. How to get your llm to generate challenging
 605 problems for evaluation. *arXiv preprint arXiv:2502.14678*, 2025.

606 Vedant Shah, Dingli Yu, Kaifeng Lyu, Simon Park, Nan Rosemary Ke, Michael Curtis Mozer, Yoshua
 607 Bengio, Sanjeev Arora, and Anirudh Goyal. AI-assisted generation of difficult math questions. In
 608 *The 4th Workshop on Mathematical Reasoning and AI at NeurIPS'24*, 2024.

609 Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael
 610 Schärli, and Denny Zhou. Large language models can be easily distracted by irrelevant context. In
 611 *International Conference on Machine Learning*, pp. 31210–31227. PMLR, 2023.

612 Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J
 613 Liu, James Harrison, Jaehoon Lee, Kelvin Xu, et al. Beyond human data: Scaling self-training for
 614 problem-solving with language models. *arXiv preprint arXiv:2312.06585*, 2023.

615 Saurabh Srivastava, Anto PV, Shashank Menon, Ajay Sukumar, Alan Philipose, Stevin Prince, and
 616 Sooraj Thomas. Functional benchmarks for robust evaluation of reasoning performance, and the
 617 reasoning gap. *arXiv preprint arXiv:2402.19450*, 2024.

618 Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy
 619 Liang, and Tatsunori B Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.

620 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowd-
 621 hery, and Denny Zhou. Self-Consistency Improves Chain of Thought Reasoning in Language
 622 Models. 2022. URL <http://arxiv.org/abs/2203.11171>.

623 Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and
 624 Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. In
 625 *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume*
 626 *1: Long Papers)*, pp. 13484–13508, 2023.

627 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 628 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
 629 *neural information processing systems*, 35:24824–24837, 2022.

630 Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra
 631 Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. Ethical and social risks of harm from
 632 language models. *arXiv preprint arXiv:2112.04359*, 2021.

633 Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi,
 634 Pierrick Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von
 635 Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama
 636 Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language
 637 processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Pro-
 638 cessing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational
 639 Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.

640 Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Dixin
 641 Jiang. Wizardlm: Empowering large language models to follow complex instructions. *arXiv*
 642 *preprint arXiv:2304.12244*, 2023.

648 Tian Ye, Zicheng Xu, Yuanzhi Li, and Zeyuan Allen-Zhu. Physics of language models: Part 2.1,
 649 grade-school math and the hidden reasoning process. In *The Thirteenth International Conference*
 650 *on Learning Representations*, 2025. URL <https://openreview.net/forum?id=Tn5B6Udq3E>.

652 Longhui Yu, Weisen Jiang, Han Shi, YU Jincheng, Zhengying Liu, Yu Zhang, James Kwok, Zhenguo
 653 Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for
 654 large language models. In *The Twelfth International Conference on Learning Representations*,
 655 2024.

656 Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen.
 657 Mammoth: Building math generalist models through hybrid instruction tuning. In *The Twelfth*
 658 *International Conference on Learning Representations*, 2024.

659 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with
 660 reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488, 2022.

661 Zhehao Zhang, Jiaao Chen, and Difyi Yang. Darg: Dynamic evaluation of large language models via
 662 adaptive reasoning graph. *arXiv preprint arXiv:2406.17271*, 2024.

663 Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyuan Luo, Zhangchi Feng, and
 664 Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings*
 665 *of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System*
 666 *Demonstrations)*, Bangkok, Thailand, 2024. Association for Computational Linguistics.
 667 URL <http://arxiv.org/abs/2403.13372>.

668 Yang Zhou, Hongyi Liu, Zhuoming Chen, Yuandong Tian, and Beidi Chen. Gsm-infinite: How
 669 do your llms behave over infinitely increasing context length and reasoning complexity? *arXiv*
 670 *preprint arXiv:2502.05252*, 2025.

671 A APPENDIX

672 The section Adversarial Search (Fig. 6) outlines how EFAs can generate challenging problem variants
 673 to probe model weaknesses. The Scaling section (Appendix C) investigates the effect of the number
 674 of sampled variants per EFA, showing how performance trends with increased augmentation. The
 675 Ablation section (Appendix D) analyzes the impact of applying unit tests during EFA generation
 676 on downstream data quality. Qualitative Examples (Appendix E) presents representative EFAs
 677 spanning several MATH domains, including algebra, number theory, and probability, illustrating
 678 the range and structure captured by the method. The Experimental Details section describes all
 679 data generation, augmentation, and model training settings—EFA generation (box F.1), rejection
 680 finetuning and variant sampling protocols (Appendix F.2), math inference configuration, and details
 681 for math-specific training (Appendix F.4).

682 B QUALITY ANALYSIS: LOW-QUALITY EFAS ARE NATURALLY FILTERED 683 OUT

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

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

691 As shown in Table 5, low-quality EFAs have significantly lower yield rates compared to good EFAs.
 692 With a single answer attempt, bad EFAs contribute training data only 5.04% of the time, compared
 693 to 27.0% for good EFAs – a ratio of over 5 to 1 in favor of good data. Even when allowing up to 5
 694 answer attempts, the ratio remains favorable at 4.51 to 1. This demonstrates that as long as rejection

702
703
704
705
706
707
708
709

EFA for Arithmetic Word Problem

```

710
711     Original Problem
712
713     When Sophie watches her nephew, she gets out a
714     variety of toys for him. The bag of building blocks
715     has 31 blocks in it. The bin of stuffed animals has 8
716     stuffed animals inside. The tower of stacking rings
717     has 9 multicolored rings on it. Sophie recently
718     bought a tube of bouncy balls, bringing her total
719     number of toys for her nephew up to 62. How many
720     bouncy balls came in the tube?
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

```

```

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(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.rings}
            multicolored rings on it."
            f"\"Sophie recently bought a tube of bouncy balls, bringing
            her total number of toys "
            f"for her nephew up to {self.total}. How many bouncy
            balls came in the tube?"
        )

    def solve(self) → str:
        known = self.blocks + self.animals + self.rings
        balls = self.total - known
        return str(balls)

```

EFA for Tridiagonal Matrix Eigenvalue Problem

```

Original Problem
Suppose you are given the matrix:

$$A = \begin{bmatrix} 2 & 1 & 0 \\ 1 & 3 & 1 \\ 0 & 1 & 4 \end{bmatrix}$$

Find the eigenvalues of the matrix.

class Problem(BaseModel):
    d1: int
    d2: int
    d3: int
    e1: int
    e2: int

    @classmethod
    def original(cls) → Self:
        return cls(d1=2, d2=3, d3=4, e1=1, e2=1)

    @classmethod
    def sample(cls) → Self:
        while True:
            d1 = random.randint(1, 6)
            d2 = random.randint(1, 6)
            d3 = random.randint(1, 6)
            e1 = random.randint(-3, 3)
            e2 = random.randint(-3, 3)
            if e1 == 0 or e2 == 0:
                continue
            if len(set([d1, d2, d3])) < 3:
                continue
            break
        return cls(d1=d1, d2=d2, d3=d3, e1=e1, e2=e2)

    def render(self) → str:
        return (
            f"Suppose you are given the matrix:\n\n"
            f"\u2212\n"
            f"  A = \begin{bmatrix} a & e_1 & 0 \\ e_1 & b & e_2 \\ 0 & e_2 & c \end{bmatrix}\n"
            f"  EFAGen identifies a tridiagonal matrix where a, b, c, e1, and e2 are real numbers and e1, e2 \u2260 0 as a suitable abstraction to create variants.\n"
        )

    def solve(self) → str:
        lam = sympy.Symbol('lambda')
        A = sympy.Matrix([
            [self.d1, self.e1, 0],
            [self.e1, self.d2, self.e2],
            [0, self.e2, self.d3]
        ])
        char_poly = A.charpoly(lam)
        roots = sympy.solve(char_poly.as_expr(), lam)
        def pretty_latex(x):
            if hasattr(x, 'is_number') and x.is_number:
                return sympy.latex(sympy.N(x, 6))
            else:
                return sympy.latex(x)
        roots_str = '\u2212'.join(pretty_latex(r) for r in roots)
        return f"The eigenvalues are: $\\boxed{{{roots_str}}}$"

```

Figure 7: 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.

756
 757 **Table 5: Low-quality EFAs are naturally filtered out during rejection sampling.** We compare the
 758 training data yield rates (percentage of responses that receive non-zero rewards) between good and
 759 bad EFAs. Bad EFAs are identified using LLM-based heuristics that flag trivial problems, extraneous
 760 variables, or hard-coded values. The low yield rates of bad EFAs mean they contribute minimally to
 761 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

762
 763
 764
 765
 766 **Table 6: EFA-generated data performs comparably to real data.** Direct comparison of training
 767 exclusively on problem variants generated by EFAs versus training exclusively on real problems from
 768 the MATH training set. Despite potential noise in rejection-sampled EFA data, models trained on
 769 synthetic data achieve nearly identical performance to those trained on real data.

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

775
 776 sampling or reinforcement learning is used, noisy EFAs naturally filter themselves out, ensuring that
 777 good data significantly outnumbers bad data in the final training set.

778
 779 To further validate the quality of EFA-generated data, we conduct a direct comparison between
 780 training exclusively on problem variants generated by EFAs versus training exclusively on real
 781 problems from the MATH training set. As shown in Table 6, despite potential noise in rejection-
 782 sampled EFA data, models trained on synthetic data achieve nearly identical performance to those
 783 trained on real data (22.6% vs 22.4% Pass@1 on MATH-500). This shows that EFA-generated data
 784 is as effective as existing math data for model training.

785 C SCALING: EFAGEN SCALES EFFECTIVELY UP TO 16 EXAMPLES PER EFA

786
 787 To understand the scaling behavior of EFA-based data augmentation, we investigate how performance
 788 varies with the number of problem variants generated per EFA. We sample 100 unique EFAs from the
 789 MATH training set and vary the number of problem variants generated by each EFA from 1 to 64. For
 790 each scaling setting, we train Llama3.1-8B-Base on the generated data and evaluate on MATH-500.

791
 792 As shown in Table 7, we observe smooth scaling improvements as we increase the number of variants
 793 from 1 to 16 examples per EFA, with performance gains plateauing beyond 16 examples. Specifically,
 794 Pass@1 improves from 14.1% with 1 example per EFA to 23.8% with 16 examples, while Pass@10
 795 increases from 48.5% to 57.6% over the same range. However, scaling begins to saturate at 32 and
 796 64 examples per EFA, suggesting that sampling too many problem variants from each EFA uniformly
 797 may hurt diversity and lead to diminishing returns. The optimal scaling point appears to be around 16
 798 examples per EFA, where three of the four metrics achieve their peak performance.

799 D ABLATION: UNIT TESTS IMPROVE EFA-BASED DATA AUGMENTATION 800 QUALITY

801
 802 Despite some errors in EFA generation, we find that the current EFAs are effectively improving
 803 performance. When we lower the quality by removing our unit tests, the performance gains from
 804 augmentation also decrease. As shown in Table 8, applying unit tests consistently improves perfor-
 805 mance across all benchmarks and metrics. The unit tests provide an average improvement of 2.2
 806 percentage points on MATH-500 Pass@1, 1.7 percentage points on FnEval November Pass@1, and
 807 2.9 percentage points on FnEval December Pass@1.

808
 809 In general, we believe there is a tradeoff between the level of noise in generated data and the cost of
 810 data generation, and EFAs occupy a generally useful point on the tradeoff curve. We can change the

810

811

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.

815

816

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

823

824

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.

825

826

827

828

Unit Tests	MATH-500			FnEval (November)			FnEval (December)		
	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%)

833

834

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.

835

836

E QUALITATIVE EXAMPLES

840

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

843

844

Box E.1| EFA (Algebra)

845

846

Original Problem

847

Solve the equation:

848

849

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

850

Original Solution

851

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

852

853

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

854

Multiply through by u to remove the fraction:

855

856

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

857

858

Reorganize into a standard quadratic form:

859

860

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

861

862

Factorize the quadratic:

863

$$(u - 5)(u - 3) = 0$$

864
865
866
867
868
869Thus, $u = 5$ or $u = 3$. For $u = 5$:

$$\sqrt{5x - 4} = 5 \Rightarrow 5x - 4 = 25 \Rightarrow 5x = 29 \Rightarrow x = \frac{29}{5}$$

For $u = 3$:

$$\sqrt{5x - 4} = 3 \Rightarrow 5x - 4 = 9 \Rightarrow 5x = 13 \Rightarrow x = \frac{13}{5}$$

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 our solution.870
871
872
873
874
875**Inferred EFA**876
877
878
879
880
881
882
883
884

```
class Problem(BaseModel):
    coefficient: int
    constant: int
    term: int

    @classmethod
    def original(cls) ->Self:
        return cls(coefficient=5, constant=4, term=15)

    @classmethod
    def sample(cls) ->Self:
        coefficient = random.randint(3, 10)
        constant = random.randint(1, 5)
        term = random.randint(10, 20)
        return cls(coefficient=coefficient, constant=constant, term=term)

    def render(self) ->str:
        return f"""
Solve the equation:
\sqrt{{self.coefficient}x-{self.constant}} + \frac{{self.term}}{{\sqrt{{self.coefficient}x-{self.constant}}}} = 8.
"""

    def solve(self) ->str:
        u = sympy.symbols('u')
        equation = sympy.Eq(u + self.term / u, 8)
        u_values = sympy.solve(equation, u)
        x_solutions = []
        for u_val in u_values:
            x_equation = sympy.Eq(u_val ** 2, self.coefficient * sympy.

                symbols('x') - self.constant)
            x_solution = sympy.solve(x_equation, sympy.symbols('x'))
            x_solutions.extend(x_solution)
        return ', '.join(
            f'\frac{{sol.as_numer_denom()[0]}}{{sol.as_numer_denom()
                [1]}}'
            for sol in x_solutions)
```

900
901
902
903
904
905
906
907
908
909**Variant generated by EFA**

Solve the equation:

910
911
912
913
914
915
916
917

$$\sqrt{7x - 5} + \frac{14}{\sqrt{7x - 5}} = 8$$

Solution:

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

918 Box E.2| EFA (Number Theory)
 919

920 **Original Problem**

921 How many positive divisors does $8!$ have?

922 **Original Solution**

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

$$926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad 935 \quad 936 \quad 937 \quad 938 \quad 939 \quad 940 \quad 941 \quad 942 \quad 943 \quad 944 \quad 945 \quad 946 \quad 947 \quad 948 \quad 949 \quad 950 \quad 951 \quad 952 \quad 953 \quad 954 \quad 955 \quad 956 \quad 957 \quad 958 \quad 959 \quad 960 \quad 961 \quad 962 \quad 963 \quad 964 \quad 965 \quad 966 \quad 967 \quad 968 \quad 969 \quad 970 \quad 971 \quad t(n) = (a+1)(b+1)(c+1)\dots$$

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

$$931 \quad 932 \quad 933 \quad 934 \quad 935 \quad 936 \quad 937 \quad 938 \quad 939 \quad 940 \quad 941 \quad 942 \quad 943 \quad 944 \quad 945 \quad 946 \quad 947 \quad 948 \quad 949 \quad 950 \quad 951 \quad 952 \quad 953 \quad 954 \quad 955 \quad 956 \quad 957 \quad 958 \quad 959 \quad 960 \quad 961 \quad 962 \quad 963 \quad 964 \quad 965 \quad 966 \quad 967 \quad 968 \quad 969 \quad 970 \quad 971 \quad t(40320) = (7+1)(2+1)(1+1)(1+1) = 8 \times 3 \times 2 \times 2 = 96$$

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

934 **Inferred EFA**

```
935
936 class Problem(BaseModel):
937     n: int
938
939     @classmethod
940     def original(cls) ->Self:
941         return cls(n=8)
942
943     @classmethod
944     def sample(cls) ->Self:
945         n = random.randint(4, 10)
946         return cls(n=n)
947
948     def render(self) ->str:
949         return f'How many positive divisors does {self.n}! have?'
950
951     def solve(self) ->str:
952         factorial_value = math.factorial(self.n)
953         factors = sympy.factorint(factorial_value)
954         divisor_count = 1
955         for exponent in factors.values():
956             divisor_count *= exponent + 1
957         return str(divisor_count)
```

958 **Variant generated by EFA**

959 How many positive divisors does $9!$ have?

960 *Solution:*

961 160

962 Box E.3| EFA (Probability)

963 **Original Problem**

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

966 **Original Solution**

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

$$969 \quad 970 \quad 971 \quad |\Omega| = 8 \times 8 = 64$$

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

972

We enumerate the pairs (d_1, d_2) corresponding to each prime sum:

973

Sum = 2: $\{(1, 1)\}$ (1 way)

974

Sum = 3: $\{(1, 2), (2, 1)\}$ (2 ways)

975

Sum = 5: $\{(1, 4), (4, 1), (2, 3), (3, 2)\}$ (4 ways)

976

Sum = 7: $\{(1, 6), (6, 1), (2, 5), (5, 2), (3, 4), (4, 3)\}$ (6 ways)

977

Sum = 11: $\{(3, 8), (8, 3), (4, 7), (7, 4), (5, 6), (6, 5)\}$ (6 ways)

978

Sum = 13: $\{(5, 8), (8, 5), (6, 7), (7, 6)\}$ (4 ways)

979

The total number of favorable outcomes is the sum of the ways for each prime:

980

$$N_{\text{favorable}} = 1 + 2 + 4 + 6 + 6 + 4 = 23$$

981

The probability is the ratio of the number of favorable outcomes to the total number of outcomes:

982

$$P(\text{Sum is prime}) = \frac{N_{\text{favorable}}}{|\Omega|} = \frac{23}{64}$$

983

The final answer is:

984

$$\frac{23}{64}$$

985

Inferred EFA

986

```

class Problem(BaseModel):
    sides: int

    @classmethod
    def original(cls) ->Self:
        return cls(sides=8)

    @classmethod
    def sample(cls) ->Self:
        sides = random.randint(4, 20)
        return cls(sides=sides)

    def solve(self) ->str:
        total_outcomes = self.sides * self.sides

        def is_prime(n: int) ->bool:
            if n <= 1:
                return False
            for i in range(2, int(math.sqrt(n)) + 1):
                if n % i == 0:
                    return False
            return True
        primal_sum_occurrences = 0
        for die1 in range(1, self.sides + 1):
            for die2 in range(1, self.sides + 1):
                sum_of_dice = die1 + die2
                if is_prime(sum_of_dice):
                    primal_sum_occurrences += 1
        probability = primal_sum_occurrences / total_outcomes
        fraction = sympy.Rational(primal_sum_occurrences,
                                   total_outcomes)
        return f'\frac{{{{fraction.numerator}}}}{{{{fraction.denominator}}}}'

    def render(self) ->str:
        return (
            f'Two {self.sides}-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.'
)

```

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

)

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:

$$\frac{105}{361}$$

F EXPERIMENTAL DETAILS

F.1 GENERATING EFAS

When generating EFAs, we use the prompt in box F.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 F.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:

1. 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
```

```

1080
1081     6. Consider edge cases and validity:
1082         - Ensure generated problems are mathematically sound
1083         - Handle special cases appropriately
1084         - Maintain reasonable complexity in generated problems
1085
1086     7. Do not import any libraries! The following libraries have been
1087         imported. Use fully qualified names for all imports:
1088         - pydantic.BaseModel is imported as 'BaseModel'
1089         - random is imported as 'random'
1090         - math is imported as 'math'
1091         - numpy is imported as 'np'
1092         - sympy is imported as 'sympy'
1093         - typing.Self is imported as 'Self'
1094
1095     Example usage:
1096     '''python
1097         problem = MyMathProblem.original() # Get original problem
1098         variant = MyMathProblem.sample() # Generate new variant
1099         question = variant.render() # Get problem statement
1100         answer = variant.solve() # Compute solution
1101     '''
1102
1103     The goal is to create a class that can both reproduce the original
1104     problem and generate mathematically valid variations of the same
1105     problem type.
1106
1107     # Example 1
1108     ## Problem Statement
1109     Evaluate  $i^5 + i^{-25} + i^{45}$ .
1110
1111     ## Solution
1112     We have  $i^5 = i^4 \cdot i = 1 \cdot i = i$ . We also have  $i^{-25} = 1/i^{25} = 1/(i^{24} \cdot i) = 1/[1 \cdot i] = 1/i = \frac{1}{i}$ . So, adding these three results gives  $i^5 + i^{-25} + i^{45} = i + i + i = 3i$ . The final answer is  $\boxed{i}$ .
1113
1114     ## Functionalization
1115     '''python
1116         class Problem(BaseModel):
1117             exponent1: int
1118             exponent2: int
1119             exponent3: int
1120
1121             @classmethod
1122             def original(cls) -> Self:
1123                 return cls(exponent1=5, exponent2=-25, exponent3=45)
1124
1125             @classmethod
1126             def sample(cls) -> Self:
1127                 exponent1 = random.randint(-100, 100)
1128                 exponent2 = random.randint(-100, 100)
1129                 exponent3 = random.randint(-100, 100)
1130                 return cls(exponent1=exponent1, exponent2=exponent2,
1131                           exponent3=exponent3)
1132
1133             def render(self) -> str:
1134                 return f"Evaluate  $i^{{self.exponent1}} + i^{{self.exponent2}} + i^{{self.exponent3}}$ ."
1135
1136             def solve(self) -> str:
1137                 # Compute the values of  $i^n \pmod 4$  cycle

```



```

1188
1189     # Generate random acute angles that form a valid triangle
1190     # Sum of angles must be less than 180
1191     angle1 = random.randint(30, 75) # Keep angles acute
1192     angle2 = random.randint(30, 75)
1193     # Ensure the third angle is also acute
1194     if angle1 + angle2 >= 150:
1195         angle1 = min(angle1, 60)
1196         angle2 = min(angle2, 60)
1197         return cls(angle_BAC=angle1, angle_ABC=angle2)
1198
1199     def solve(self) -> str:
1200         # The angle HCA is complementary to angle BAC
1201         # This is because H is the orthocenter and CH is
1202             perpendicular to AB
1203         angle_HCA = 90 - self.angle_BAC
1204         return f"{{angle_HCA}}"
1205
1206     def render(self) -> str:
1207         return (
1208             f"Altitudes $\\overline{{AX}}$ and $\\overline{{BY}}$ of
1209                 acute triangle $ABC$"
1210             f"intersect at $H$. If $\\angle BAC = {self.angle_BAC}^\\circ$ and "
1211             f"$\\angle ABC = {self.angle_ABC}^\\circ$, then what is $\\angle HCA$?"
1212         )
1213
1214     # Example 3
1215     ## Problem Statement
1216     On a true-false test of 100 items, every question that is a
1217         multiple of 4 is true, and all others are false. If a student
1218         marks every item that is a multiple of 3 false and all others
1219             true, how many of the 100 items will be correctly answered?
1220     ## Solution
1221     The student will answer a question correctly if
1222
1223     Case 1: both the student and the answer key say it is true. This
1224         happens when the answer is NOT a multiple of 3 but IS a multiple
1225             of 4.
1226
1227     Case 2. both the student and the answer key say it is false. This
1228         happens when the answer IS a multiple of 3 but is NOT a multiple
1229             of 4.
1230
1231     Since the LCM of 3 and 4 is 12, the divisibility of numbers (in our
1232         case, correctness of answers) will repeat in cycles of 12. In
1233             the first 12 integers, $4$ and $8$ satisfy Case 1
1234         and $3, 6, 9$ and $99$ satisfy Case 2, so for every group of 12, the
1235             student will get 5 right answers. Since there are 8 full groups
1236                 of 12 in 100, the student will answer at least $8
1237             \\cdot 5 = 40$ questions correctly. However, remember that we must
1238                 also consider the leftover numbers 97, 98, 99, 100 and out of
1239                     these, $99$ and $100$ satisfy one of the cases. So
1240                         our final number of correct answers is $40 + 2 = \\boxed{42}$.
1241
1242     ## Functionalization
1243     '''python
1244     class Problem(BaseModel):
1245         total_questions: int # Total number of questions
1246         multiple1: int # First multiple (4 in original problem)
1247         multiple2: int # Second multiple (3 in original problem)

```

```

1242
1243     @classmethod
1244     def original(cls) -> Self:
1245         return cls(total_questions=100, multiple1=4, multiple2=3)
1246
1247     @classmethod
1248     def sample(cls) -> Self:
1249         # Generate reasonable random parameters
1250         total = random.randint(50, 200) # Reasonable test length
1251         # Choose coprimes or numbers with small LCM for interesting
1252         # results
1253         mult1 = random.randint(2, 6)
1254         mult2 = random.randint(2, 6)
1255         while mult1 == mult2: # Ensure different numbers
1256             mult2 = random.randint(2, 6)
1257         return cls(total_questions=total, multiple1=mult1, multiple2=
1258             mult2)
1259
1260     def solve(self) -> str:
1261         def lcm(a: int, b: int) -> int:
1262             def gcd(x: int, y: int) -> int:
1263                 while y:
1264                     x, y = y, x % y
1265                 return x
1266
1267             return abs(a * b) // gcd(a, b)
1268
1269             # Find cycle length (LCM)
1270             cycle_length = lcm(self.multiple1, self.multiple2)
1271
1272             # Count correct answers in one cycle
1273             correct_per_cycle = 0
1274             for i in range(1, cycle_length + 1):
1275                 answer_key_true = i % self.multiple1 == 0
1276                 student_true = i % self.multiple2 != 0
1277                 if answer_key_true == student_true:
1278                     correct_per_cycle += 1
1279
1280             # Calculate complete cycles and remainder
1281             complete_cycles = self.total_questions // cycle_length
1282             remainder = self.total_questions % cycle_length
1283
1284             # Calculate total correct answers
1285             total_correct = complete_cycles * correct_per_cycle
1286
1287             # Add correct answers from remainder
1288             for i in range(1, remainder + 1):
1289                 answer_key_true = i % self.multiple1 == 0
1290                 student_true = i % self.multiple2 != 0
1291                 if answer_key_true == student_true:
1292                     total_correct += 1
1293
1294             return str(total_correct)
1295
1296     def render(self) -> str:
1297         return (
1298             f"On a true-false test of {self.total_questions} items, "
1299             f"every question that is a multiple of {self.multiple1} is "
1300             f"true, "
1301             f"and all others are false. If a student marks every item "
1302             f"that is "
1303             f"a multiple of {self.multiple2} false and all others true, "
1304             f"how "
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691
2692
2693
2694
2695
2696
2697
2698
2699
2700
2701
2702
2703
2704
2705
2706
2707
2708
2709
2710
2711
2712
2713
2714
2715
2716
2717
2718
2719
2720
2721
2722
2723
2724
2725
2726
2727
2728
2729
2730
2731
2732
2733
2734
2735
2736
2737
2738
2739
2740
2741
2742
2743
2744
2745
2746
2747
2748
2749
2750
2751
2752
2753
2754
2755
2756
2757
2758
2759
2760
2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788
2789
2790
2791
2792
2793
2794
2795
2796
2797
2798
2799
2800
2801
2802
2803
2804
2805
2806
2807
2808
2809
2810
2811
2812
2813
2814
2815
2816
2817
2818
2819
2820
2821
2822
2823
2824
2825
2826
2827
2828
2829
2830
2831
2832
2833
2834
2835
2836
2837
2838
2839
2840
2841
2842
2843
2844
2845
2846
2847
2848
2849
2850
2851
2852
2853
2854
2855
2856
2857
2858
2859
2860
2861
2862
2863
2864
2865
2866
2867
2868
2869
2870
2871
2872
2873
2874
2875
2876
2877
2878
2879
2880
2881
2882
2883
2884
2885
2886
2887
2888
2889
2890
2891
2892
2893
2894
2895
2896
2897
2898
2899
2900
2901
2902
2903
2904
2905
2906
2907
2908
2909
2910
2911
2912
2913
2914
2915
2916
2917
2918
2919
2920
2921
2922
2923
2924
2925
2926
2927
2928
2929
2930
2931
2932
2933
2934
2935
2936
2937
2938
2939
2940
2941
2942
2943
2944
2945
2946
2947
2948
2949
2950
2951
2952
2953
2954
2955
2956
2957
2958
2959
2960
2961
2962
2963
2964
2965
2966
2967
2968
2969
2970
2971
2972
2973
2974
2975
2976
2977
2978
2979
2980
2981
2982
2983
2984
2985
2986
2987
2988
2989
2990
2991
2992
2993
2994
2995
2996
2997
2998
2999
2999

```

```

1296
1297         f"many of the {self.total_questions} items will be
1298         correctly answered?"
1299     )
1300
1301     # Your Turn
1302     Functionalize the following problem:
1303
1304     ## Problem Statement
1305     [% problem_statement %]
1306
1307     ## Solution
1308     [% solution %]
1309
1310     ## Functionalization
1311
1312
1313

```

F.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 same prompt box F.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.

F.3 MATH INFERENCE SETTINGS

When doing 0-shot inference with Llama3.1-8B-Instruct, we use the official Llama3.1 prompt in box F.2. When doing few-shot inference with Llama3.1-8B-Instruct, we use a modified version of the official prompt, shown in box F.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.

Box F.2| Llama3.1 0-shot MATH Prompt

```

1332
1333
1334     Solve the following math problem efficiently and clearly:
1335
1336     - For simple problems (2 steps or fewer):
1337         Provide a concise solution with minimal explanation.
1338
1339     - For complex problems (3 steps or more):
1340         Use this step-by-step format:
1341
1342         ## Step 1: [Concise description]
1343         [Brief explanation and calculations]
1344
1345         ## Step 2: [Concise description]
1346         [Brief explanation and calculations]
1347
1348     ...
1349
1350     Regardless of the approach, always conclude with:

```

```

1350
1351     Therefore, the final answer is: $\boxed{answer}$. I hope it is
1352     correct.
1353
1354     Where [answer] is just the final number or expression that solves
1355     the problem.
1356
1357     Problem: {{ instruction }}
1358

```

Box F.3| Llama3.1 N-shot MATH Prompt

```

1360
1361     Solve the following math problem efficiently and clearly:
1362
1363     - For simple problems (2 steps or fewer):
1364         Provide a concise solution with minimal explanation.
1365
1366     - For complex problems (3 steps or more):
1367         Use this step-by-step format:
1368
1369         \#\# Step 1: [Concise description]
1370         [Brief explanation and calculations]
1371
1372         \#\# Step 2: [Concise description]
1373         [Brief explanation and calculations]
1374
1375         ...
1376
1377     Regardless of the approach, always conclude with:
1378
1379     Therefore, the final answer is: $\boxed{answer}$. I hope it is
1380     correct.
1381
1382     Where [answer] is just the final number or expression that solves
1383     the problem.
1384
1385     Here are some examples:
1386
1387     {% for few_shot_example in few_shot_examples %}
1388     Problem: {{ few_shot_example.instruction }}
1389     {{ few_shot_example.response }}
1390     {% endfor %}
1391
1392     Problem: {{ instruction }}
1393

```

F.4 MATH TRAINING DETAILS

We use the same hyperparameters and chat data format as in Appendix F.2, except we cutoff training data over 2048 tokens. However, we use a simpler prompt template, shown in box F.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 F.4| Minimal instruction-tuning prompt used for augmentation experiments

```

1397
1398     Question: {{ question }}
1399     Step-by-step Answer
1400
1401
1402
1403

```