

AIME: AI SYSTEM OPTIMIZATION VIA MULTIPLE LLM EVALUATORS

Anonymous authors

Paper under double-blind review

ABSTRACT

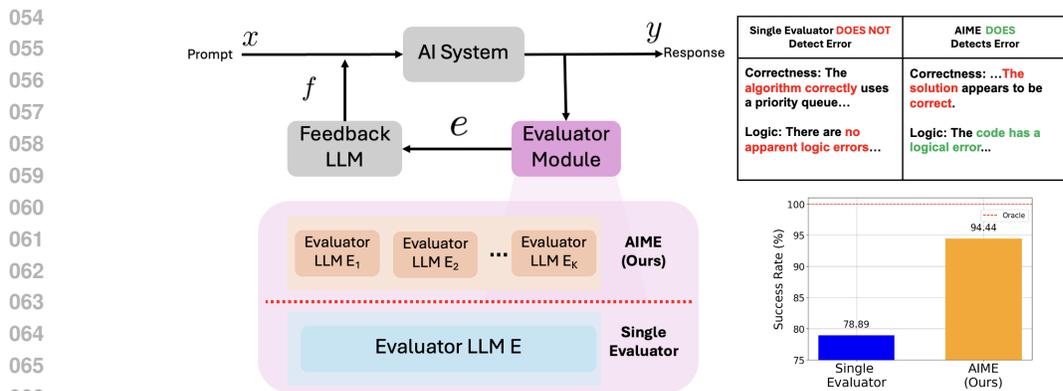
Text-based AI system optimization typically involves a feedback loop scheme where a *single* LLM generates an evaluation in natural language of the current output to improve the next iteration’s output. However, in this work, we empirically demonstrate that for a practical and complex task (code generation) with multiple criteria to evaluate, utilizing only one LLM evaluator tends to let errors in generated code go undetected, thus leading to incorrect evaluations and ultimately suboptimal test case performance. Motivated by this failure case, we assume there exists an optimal evaluation policy that samples an evaluation between response and ground truth. We then theoretically prove that a linear combination of multiple evaluators can approximate this optimal policy. From this insight, we propose **AI system optimization via Multiple LLM Evaluators (AIME)**. AIME is an evaluation protocol that utilizes multiple LLMs that each independently generate an evaluation on separate criteria and then combine them via concatenation. We provide an extensive empirical study showing AIME outperforming baseline methods in code generation tasks, with up to 62% higher error detection rate and up to 16% higher success rate than a single LLM evaluation protocol on LeetCodeHard and HumanEval datasets. We also show that the selection of the number of evaluators and which criteria to utilize is non-trivial as it can impact success rate by up to 12%.

1 INTRODUCTION

Pre-trained foundation models, such as Large Language Models (LLMs), have developed rapidly over the recent years (Achiam et al., 2023; Touvron et al., 2023). With these advancements, AI systems have grown in popularity for various tasks such as code generation (Chen et al., 2024; Gulwani, 2010), question-answering (Patel et al., 2024; Wang et al., 2024), mathematical reasoning (Trinh et al., 2024; Song et al., 2024), exploration (Dorbala et al., 2024; 2023; Ren et al., 2024), and information retrieval (Gao et al., 2023) etc. As the application complexity increases, the shift to AI systems containing multiple components such as LLM-based agents and web search (Xiong et al., 2024), will continue (Zaharia et al., 2024; Yuksekgonul et al., 2024). Thus, automatically optimizing these systems, *AI system optimization* (Yuksekgonul et al., 2024), becomes increasingly necessary.

An emerging paradigm is text-based optimization, also known as prompt optimization (Cheng et al., 2023; Wang et al., 2023; Zhou et al., 2022), whereby the natural language input prompt is tuned to generate an optimal output. This method requires no numerical gradient descent updates typical in optimization for machine learning models (Van Der Malsburg, 1986; Hassoun, 1995; Barto, 1992) and is thus appropriate for optimizing AI systems with fixed LLM components. Recently, there has been a growing class of iterative online methods for text-based optimization (Cheng et al., 2024; Yuksekgonul et al., 2024; Shinn et al., 2024), where a *single* LLM generates an evaluation based on the current output to help generate the next iteration’s prompt.

While prior art has compared the abilities of a single LLM for evaluations against those of multiple LLMs (Kocmi & Federmann, 2023; Ankner et al., 2024), in AI system optimization literature, there has been a lack of studies questioning the capabilities of using a single LLM evaluator to drive the optimization process. Recently, Yuksekgonul et al. (2024) has viewed the evaluation as a text-based analogy to the objective function for backpropagation (Hinton et al., 2006; Rumelhart et al., 1986)



067
068
069
070
071
072
073
074
075
076

Figure 1: **AI System Optimization Pipeline and Increased Error Detection and Success Rate with AIME-based Evaluation:** [LEFT] Text-based AI system optimization with SoTA framework (Yuksekgonul et al., 2024) using our multiple LLM evaluator approach AIME (orange) and with single-evaluator approach (blue). [TOP RIGHT] The single-evaluator approach cannot detect an error in the generated code that fails all test cases. However, one of the evaluators of AIME could because the logical evaluator was independent from the correctness evaluator. [BOTTOM RIGHT] AIME-based optimization achieves a $\sim 16\%$ higher success rate than a single-evaluator approach in code generation tasks.

077
078
079
080
081
082

in deep learning optimization. The objective function is a crucial element in optimizing machine learning models (Christiano et al., 2017; Mescheder et al., 2018; Chakraborty et al., 2023; Kingma & Welling, 2014). This importance motivates us to analyze and strengthen the evaluation protocol of state-of-the-art (SoTA) AI system optimization frameworks by addressing a critical research question: *What are the failure cases or tasks of utilizing only one LLM-based evaluator for text-based AI system optimization?*

083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099

For this question, we empirically demonstrate the shortcomings of a single evaluator protocol in judging complex outputs like code based on multiple diverse criteria, such as correctness, readability, and runtime. We emphasize its practical limitations to give optimal evaluation while being instructed to judge based on all criteria simultaneously. Figure 1 illustrates the suboptimality in the practice of an AI system optimization framework with a single-evaluator approach to code generation. Furthermore, by assuming there exists an optimal evaluation policy that in expectation samples the true evaluation between the generated response and ground truth, we also theoretically highlight that the suboptimality gap between a single evaluator and an optimal evaluator is fixed and cannot be reduced given the same output and problem task. With this insight, we then naturally ask the following subsequent query: *Can we develop a principled evaluation method for text-based optimization to handle multiple criteria?* We address this question by assuming there exists an optimal evaluation policy that in expectation samples the true evaluation between the generated response and ground truth. We then theoretically prove that, under a linear additivity assumption, increasing the number of evaluators can reduce the suboptimality gap. We capitalize on this theoretical insight by proposing **AIME: AI system optimization via Multiple Evaluators**. AIME generates and combines via concatenation independent natural language evaluations from multiple evaluators based on different evaluation instructions. We demonstrate on code generation tasks with LeetCodeHard and HumanEval benchmarks the superior performance of AIME over a single evaluator in code error detection and the success rate of test cases.

100
101

Our main contributions are as follows:

- 102
103
104
105
106
107
- **Novel Evaluation Approach for AI system Optimization:** We propose using multiple LLM-based evaluators and introduce our AIME approach for iterative AI system optimization. We concatenate independent diverse samples from multiple LLM-based evaluation policies to better critique system outputs.
 - **Theoretical Motivation for Multiple Evaluators:** We prove that through a linear additivity assumption increasing the number of evaluations can reduce the suboptimality gap from

an optimal evaluation policy while a single evaluator has a fixed gap. This theoretical result helps justify our formulation for a multiple evaluation-based protocol.

- **Empirical Performance Over Single-Evaluation Approach:** Using popular code generation dataset, LeetCodeHard (Shinn et al., 2024) and HumanEval (Chen et al., 2021), we perform an extensive study showing the superior prowess of AIME with 6 evaluators over single evaluation to detect errors, with AIME achieve up to 62% higher error detection rate than single evaluation. We then show that AIME-based optimization achieves up to a 16% higher success rate on test cases than optimization with only a single evaluator. We also reveal that the choice of the number of evaluators and the combination of criteria to utilize can affect the success rate by up to 12%, emphasizing the design of AIME-based optimization is non-trivial. We provide an anonymized code repository.¹

2 TEXT-BASED AI SYSTEM OPTIMIZATION

Objective Function. In this section, we now characterize mathematically text-based prompt optimization as a system of LLM-based policies. Let $\pi(\cdot|x)$ be the LLM-based AI system parameterized by fixed LLM-based policy that samples an output response $y \sim \pi(\cdot|x)$ given an input prompt $x \in \mathcal{X}$ from the set of prompts \mathcal{X} . We aim to sample a $y \sim \pi(\cdot|x^*)$ by finding an input prompt x^* corresponding to x prompt such that y is closer to the optimal response y^* . For code generation, π_θ would be the LLM generator; x would be the input prompt; y is the generated code; and the y^* here would be a code snippet that is a readable, efficient solution to the problem. Mathematically, we can write

$$x^* = \arg \min_{x \in \mathcal{X}} \mathbb{E}_{y \sim \pi_\theta(\cdot|x)} [l(y^*, y)], \quad (1)$$

where l is a loss function to capture the closeness of sampled response y to the ground truth y^* .

Iterative text-based optimization. Given an initial prompt x_1 , we perform an iterative text-based optimization method to find x^* as follows. For each iteration $t = 1$ to T , we start by (i) sampling $y_t \sim \pi_\theta(\cdot|x_t)$, (ii) evaluate the response y_t to obtain evaluation $e_t = l(y^*, y_t)$, and then finally (iii) generate the next prompt $x_{t+1} \sim \pi(\cdot|y_t, e_t, x_t)$. Recent work by Yuksekogonul et al. (2024) decompose step (iii) into two separate steps and (iii.a) first generate the feedback $f_t \sim \pi(\cdot|y_t, e_t, x_t)$, and then (iii.b) generate the next prompt $x_{t+1} \sim \pi(\cdot|y_t, f_t, x_t)$. For simplicity, we use the same variable π for all LLM-based policies because the outputs are dependent on the input variables the policy is conditioned on, so the same LLM model can be utilized. In this paper, we use the same model, GPT-4o, for all steps. However, distinct LLM models can be employed at different steps.

Challenges. In an ideal setting, if we had the access to y^* as in supervised learning (Tiwari, 2022), then we can achieve the optimal performance with larger data. However, in practice, they are hard to obtain or simply unknown for many tasks such as code generation (Chen et al., 2024). Therefore, a direct comparison to an optimal output y^* and the resulting calculation of e in step (ii) are both infeasible. Current SoTA work instead *sample* an evaluation e from an evaluation policy conditioned by the response output y and prompt x as $e \sim \pi(\cdot|x, y)$. Let us denote $\pi_e = \pi(\cdot|x, y)$ for notation simplicity. Ideally, we would like the evaluation e of y to be $l(y^*, y)$. More specifically, let’s assume the existence of an optimal evaluator LLM denoted by π_e^* , sampling from which will give us samples of the true loss function $l(y^*, y)$.

Fixed Gap in Evaluation with Single Evaluation Policy from Prior SOTA. As π_e^* is unavailable as discussed before, current SoTA methods sample the evaluation loss from a single evaluator as $e \sim \pi_e$. Now, we know that in the majority of the scenarios π_e will not be the true evaluator policy π_e^* . Thus $e = l(\hat{y}, y)$, where \hat{y} is an implicit approximation of y^* from π_e . Under this scenario, we define the suboptimality gap in evaluation of prior SOTA as

$$\Delta_{\text{Eva-sub-opt}}^\pi = \mathbb{E}_{e^* \sim \pi_e^*(\cdot|x, y)} [e^*] - \mathbb{E}_{e \sim \pi_e(\cdot|x, y)} [e] \leq |e|_{\max} d_{\text{TV}}(\pi_e^*(\cdot|x, y), \pi(\cdot|x, y)) \quad (2)$$

where we first expand upon the sub-optimality in evaluation and then upper-bound using the total variation distance (cite ipm). We see that the term $d_{\text{TV}}(\pi_e^*(\cdot|x, y), \pi(\cdot|x, y))$ is fixed and it can’t be improved once we have the evaluator π . This result shows the hardness of a single evaluator

¹Anonymous repository to code: <https://anonymous.4open.science/r/aime-4D3E/README.md>

reaching π_e^* due to this constant gap and it will only reduce if our current LLM evaluator is near-optimal which is not true in majority of the scenarios. Empirically, Figure 1 demonstrates a practical observation where a single evaluator lets code errors go undetected, causing a large suboptimality gap from oracle performance in code generation tasks.

3 AIME: AI SYSTEM OPTIMIZATION VIA MULTIPLE LLM EVALUATORS

Our key idea is to utilize multiple evaluations than single evaluators used in state-of-the-art. The thought that multiple evaluators would work better than one sounds intuitive but a naive introduction of multiple evaluators does not work in practice. We theoretically prove the merit of multiple evaluators and then discuss how to introduce them into the pipeline described in Section 2.

3.1 INCREASING EVALUATIONS REDUCES THE EVALUATION SUBOPTIMALITY GAP

Let $\Pi = \{\pi_k(\cdot|x, y)\}_{k=1}^K$ be the set of diverse evaluators for x, y . We start our theoretical justification by defining the sub-optimality metric to measure the evaluation performance between π_e^* and Π as

$$\Delta_{\text{Eva-sub-opt}}^{\Pi} = \mathbb{E}_{e \sim \pi_e^*(\cdot|x, y)} [e] - \mathbb{E}_{\{e_k \sim \pi_k(\cdot|x, y)\}_{k=1}^K} [g(e_1, \dots, e_K)], \quad (3)$$

which is nothing but the difference between the expected value of the evaluation under the optimal unknown evaluation distribution, and the expected function $g(\dots)$ which maps the K different evaluations to one. In practice, g can be seen as an aggregation function such as concatenation. Note that if we had access to the optimal evaluator π_e^* , we would have been able to get the ground-truth evaluation $e^* = l(y^*, y)$ to perform the AI text optimization. However, in place of that, we have a diverse set of evaluators $\Pi = (\pi_1, \pi_2 \dots \pi_K)$ and $g(e_1, e_2 \dots e_K)$ is the aggregation function to combine the losses from the diverse evaluators. We provide the following theorem to relate the number of evaluations to the $\Delta_{\text{Eva-sub-opt}}^{\Pi}$.

Theorem 1. Let d_{TV} denote the total variation distance between two distributions and let $\sum_{k=1}^K \alpha_k = 1$. Assuming all pairs $\pi_1, \pi_2 \in \Pi$ are independent of one another,

$$\Delta_{\text{Eva-sub-opt}}^{\Pi} \leq |e^*| d_{TV}(\pi_e^*, \sum_{k=1}^K \alpha_k \pi_k). \quad (4)$$

Proof. First, we characterize the sub-optimality of our proposed evaluation method as $\Delta = \mathbb{E}_{e^* \sim \pi_e^*} [e^*] - \mathbb{E}_{e_1 \sim \pi_1(\cdot|x, y), e_2 \sim \pi_2(\cdot|x, y) \dots \pi_K} [g(e_1, e_2, e_3 \dots e_K)]$. Note that if Δ is zero, we are doing the optimal evaluation. Thus, we want Δ to be as low as possible. For simplicity of the expression, we will keep to two evaluators and it can easily extend to K without loss of generality.

$$\begin{aligned} \Delta &= \mathbb{E}_{e^* \sim \pi_e^*} [e^*] - \mathbb{E}_{e_1 \sim \pi_1(\cdot|x, y), e_2 \sim \pi_2(\cdot|x, y)} [g(e_1, e_2)] \\ &= \underbrace{\mathbb{E}_{e^* \sim \pi_e^*} [e^*] - \mathbb{E}_{e \sim \pi_d(\cdot|x, y)} [e]}_{\Delta_1} + \underbrace{\mathbb{E}_{e \sim \pi_d(\cdot|x, y)} [e] - \mathbb{E}_{e_1 \sim \pi_1(\cdot|x, y), e_2 \sim \pi_2(\cdot|x, y)} [g(e_1, e_2)]}_{\Delta_2}. \end{aligned}$$

where we add and subtract the terms $\mathbb{E}_{e \sim \pi_d(\cdot|x, y)}$, with $\pi_d = \alpha\pi_1 + (1 - \alpha)\pi_2$ ($0 < \alpha < 1$) and then separate the two terms as Δ_1, Δ_2 . We next individually analyze the terms Δ_1, Δ_2 .

We can now bound Δ_1 as,

$$\Delta_1 = \mathbb{E}_{e^* \sim \pi_e^*} [e^*] - \mathbb{E}_{e \sim \pi_d(\cdot|x, y)} [l] \leq |e^*| d_{TV}(\pi_e^*, \pi_d) = |e^*| d_{TV}(\pi_e^*, \alpha\pi_1 + (1 - \alpha)\pi_2)$$

where we use the property of integral probability metric to bound Δ_1 as the total variation distance between the optimal evaluation policy and the mixture evaluation policy. Next, we proceed to Δ_2 ,

$$\begin{aligned} \Delta_2 &= \mathbb{E}_{e \sim \pi_d(\cdot|x, y)} [e] - \mathbb{E}_{e_1 \sim \pi_1(\cdot|x, y), e_2 \sim \pi_2(\cdot|x, y)} [g(e_1, e_2)] \\ &= \mathbb{E}_{e \sim \pi_d(\cdot|x, y)} [e] - \mathbb{E}_{e_1 \sim \pi_1(\cdot|x, y), e_2 \sim \pi_2(\cdot|x, y)} [\alpha e_1 + (1 - \alpha)e_2] \\ &= \mathbb{E}_{e^* \sim \pi_d(\cdot|x, y)} [e^*] - \alpha \mathbb{E}_{e_1 \sim \pi_1(\cdot|x, y)} [e_1] - (1 - \alpha) \mathbb{E}_{e_2 \sim \pi_2(\cdot|x, y)} [e_2] = 0 \end{aligned}$$

where we expand upon the definition of Δ_2 and use linear additivity assumption on the aggregation function, where we assume $g(e_1, e_2) = \alpha e_1 + (1 - \alpha)e_2$. Under this assumption, the two terms cancel out with the final result $\Delta_2 = 0$. Combining both terms concluded the proof. This bound indicates that the sub-optimality in evaluation can be expressed as the total variation distance between the optimal evaluator and the available mixture of evaluators. We know from Blei et al. (2003); Nguyen et al. (2016) that as we increase the number of mixture components and diversity amongst the components increase, it can approximate any distribution under certain assumptions.

Algorithm 1 AIME: AI System Optimization via Multiple LLM Evaluators

```

1: Input: Initial input prompt  $x_1$ , number of iterations  $T$ , pre-trained LLM-based AI system  $\pi_\theta$ ,
   list of  $K$  role descriptions  $R$ 
2: for  $t$  in  $1, \dots, T$ : do
3:   Initialize empty list of evaluations  $E_t$ 
4:    $y_t \sim \pi_\theta(\cdot|x_t)$ 
5:   for  $k$  from  $1, \dots, K$ : do
6:     Sample  $e_{k,t} \sim \pi_\theta(\cdot|x_t, y_t, R_k)$ 
7:     Append  $e_{k,t}$  to  $E_t$ 
8:   Aggregate all  $e_{k,t} \in E_t$  into  $e_t$  via concatenation
9:   Sample  $f_t \sim \pi_\theta(\cdot|y_t, e_t, x_t)$ 
10:  Sample  $x_{t+1} \sim \pi_\theta(\cdot|y_t, f_t, x_t)$ 

```

3.2 OVERVIEW OF AIME: MULTIPLE ROLE-SPECIFIC EVALUATORS

Now that we have motivated utilizing multiple LLM-based evaluators, we now address the question on *how* to utilize multiple evaluators. To do so, we look at the ideas of *roles*. The LLM-based evaluation policy has an evaluation system prompt to specify what the evaluation should be based on. For tasks such as code generation, there may be multiple criteria or objectives to evaluate for such as correctness, clarity, and efficiency. Furthermore, aspects such correctness of code can rely on various aspects such as logic and syntax. Normally, with a single evaluator, all the criteria are specified together in the system prompt. However, we see from Figure 1 and later in Section 4 that this approach can fail significantly to reach the optimal performance. We thus propose splitting the evaluation instruction across multiple evaluators, each one getting a specific role. We then aggregate via string concatenation them into a final evaluation. We chose concatenation as the aggregation method as it is analogous to creating a linear combination of the outputs (Yuksekgonul et al., 2024). We call this approach AIME: AI System Optimization with Multiple Evaluators.

Our AIME approach is a simple-to-implement approach that requires minimal changes to the already established methods (Yuksekgonul et al., 2024; Cheng et al., 2024) for system optimization. Our approach requires only modifying the evaluation step of the optimization pipeline from one evaluator to multiple. In Algorithm 1, given an output y , set of k roles R , and pre-trained LLM π_θ we sample k evaluations, $\{e_k\}_{k=1}^K$. We obtain e_i by conditioning π_θ by x, y and $R_k \in R$. Conditioning on r_k is to specify the role in the evaluation system prompt.

4 EXPERIMENTS AND RESULTS

We test the merits of our AIME approach via the code generation task because of its practicalness and its multiple plausible criteria (e.g., correctness, efficiency). Here, the AI system is an LLM generator that is given a code prompt and must produce a code snippet that passes the unit tests for that prompt. This code generation task is a form of instance optimization (Yuksekgonul et al., 2024), whereby the optimization variable, the input prompt, is defined as $x_{t+1} := (y_t, f_t)$. y_0, f_0 are empty strings. We provide empirical results showing that AIME is superior to the single-evaluation (Single-Eval) approach in detecting code errors and that AIME-based optimization achieves higher success in test cases than Single-Eval-based optimization. Experiments were run on an Apple M1 Pro and macOS 14.5.

AIME and Single-Eval Implementation Details: We use TextGrad from Yuksekgonul et al. (2024) to implement AIME and Single-Eval. We chose TextGrad because it separates the evaluation and feedback into two separate LLM calls, making it better to analyze the evaluation module in isolation. In TextGrad, the system prompt that generates the initial code, p_{init} , is different from the system prompt that updates the code in the following refinement iterations p_{update} . At $t = 0$, p_{init} specifies to the LLM that it is a code generator while the p_{update} from $1 \leq t \leq T$ specifies that it generates a new version y_{t+1} given the current code y_t and the feedback f_t . The transition from p_{init} to p_{update} is

explicitly programmed and not caused by the optimization process. Because the scope of this paper lies within the evaluation protocol, our AI system is a single LLM generator.²

LLM Setup Details: We use GPT-4o for all LLM calls and run 10 iterations of optimization for each coding problem. Across all trials for both methods, we use the same initial generated code for a given problem so both evaluation protocols can judge the same code in the initial iteration. For Single-Eval, the solitary LLM evaluator call is allowed 3600 max output tokens. For our AIME approach, each of the K evaluators is allowed $\frac{3600}{K}$ max output tokens. This decision is to model a uniform distribution of weights α . Note that when $k = 1$, Single-Eval and AIME are equivalent. We share the evaluation system prompt for both methods in Appendix A.2. We ablate on the temperature of the evaluation LLM. All other LLM calls in the Textgrad pipeline are given 2000 max output tokens with call temperature set to 0 similar to Yuksekgonul et al. (2024). For all experiments, the $\text{top}_p = 0.99$.

Roles for Evaluating Code: The set of evaluation roles R we used for this task are as follows: syntax errors, logic errors, correctness, readability, runtime, and code redundancy. The following results are based on utilizing all these roles. We chose three roles that correlate to maximizing the number of passed test cases: correctness, logic, and syntax. We specifically chose these three to incorporate an overall correctness role with two more specific roles. We will see in Section 4.1 that having overlapping roles can help with the robustness of evaluation in terms of error detection. The three other roles (readability, runtime, redundancy), correlate to criteria such as clarity and efficiency. We will later see in Section 4.3.2 that utilizing only these roles for evaluation decreases the overall performance of the code generation task.

Datasets: We use the following two datasets, LeetCodeHard (Shinn et al., 2024) and HumanEval (Chen et al., 2021), where each dataset contains a set of coding problem prompts and multiple unit tests for each problem to evaluate the generated code. We use the entire LeetCodeHard dataset of 39 problems with an average of 2.2 unit tests per problem and the first 20 problems of HumanEval with an average of 4.4 unit tests per problem. We withhold giving any of the evaluators of either method any information on unit tests to simulate the scenario where unit tests may be unavailable to help judge (Chen et al., 2024).

4.1 AIME IS ROBUST TO INCORRECT EVALUATIONS

AIME has a higher chance to catch errors: Figure 1 displays portions of an evaluation generated by Single-Eval and AIME. In this scenario, the evaluations were generated for the same coding problem at the second iteration of optimization. For both Single-Eval and AIME, the code failed all test cases, thus meaning there exists some error in the code. The evaluation from Single-Eval for both correctness and logic states there is nothing wrong. For AIME, the correctness evaluator incorrectly states nothing is wrong with the generated code but the logic evaluator detects a logical error. In the next iteration of optimization, the code generated based on the Single-Eval evaluation still fails all cases but the code generated from AIME passes them all.

Error Detection Measurement: To quantitatively analyze the error detection of AIME, we develop a heuristic measurement, Error Detection Rate (EDR). For each optimization iteration that has at least one failed test case, if the given evaluation contains at least one phrase indicating failure, we consider that as an error was detected. For example if the phrase “has a logical error” appears in the evaluation, we count that as an error detected. We provide a complete list of phrases used for detection in Appendix A.3. Let Z_{fail} be the set of iterations with at least one failed test case and let $q(z) = \mathbb{1}_{\text{error detected}}$ be the indicator value of whether an error was detected at iteration $z \in Z_{\text{fail}}$. We calculate the EDR as $\frac{1}{|Z_{\text{fail}}|} \sum_{z \in Z_{\text{fail}}} q_z$. Left of Figure 2 shows AIME has up to $\sim 62\%$ higher EDR than Single-Eval. Table 2 in Appendix A.4 summarizes the EDR for Single-Eval and AIME across various evaluation call temperatures. AIME achieves $\sim 53 - 62\%$ higher error detection rate than Single-Eval on LeetCodeHard and $\sim 38 - 57\%$ higher rate on HumanEval. This demonstrates that multiple independent evaluators can ensure a more accurate assessment than conditioning a single evaluator with all roles at once.

²We repeat the link to the anonymous repository: <https://anonymous.4open.science/r/aime-4D3E/README.md>

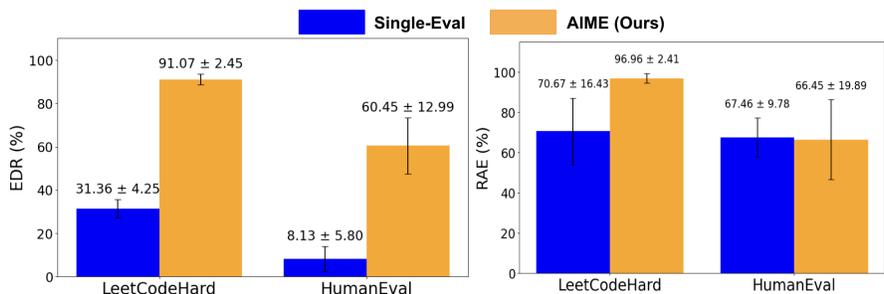


Figure 2: Using LeetCodeHard and HumanEval benchmarks we compare evaluations generated from Single-Eval against those of AIME in terms of [LEFT] EDR and [RIGHT] RAE scores. AIME has a higher EDR score on both datasets indicating it is less prone to letting errors go undetected. AIME has a higher resistance to an adversarial evaluator on LeetCodeHard and a comparable resistance on HumanEval, suggesting its robustness over Single-Eval

Robustness to Adversarial Evaluator (RAE): To further highlight the robustness of AIME to incorrect evaluations, we introduce an adversarial evaluator. For AIME, we specify in the system prompt of the correctness evaluator to always generate an evaluation stating that the code solution works. Similarly, for Single-Eval, we specify in the system prompt of the single evaluator to output an evaluation claiming that code works when discussing correctness. We provide these adversarial system prompts in Figure 6. We run experiments with an evaluation temperature of 1. To measure the robustness to the adversarial evaluator (RAE), we calculate the percent decrease of the EDR from the non-adversarial setting to the adversarial one. We then report the absolute value of the percent decrease subtracted from 1. Formally, let p_c be the percent change of the EDR, our RAE metric is $1 - |p_c|$. Right of Figure 2 reports the mean and standard deviation RAE over 3 trials. AIME achieves 16% higher RAE over Single-Eval on LeetCodeHard and comparable RAE over HumanEval, emphasizing AIME increased safety for AI systems.

AIME evaluations are more thorough: In Figure 3, we highlight the readability portions of the same evaluation in Figure 1. Even though both Single-Eval and AIME did not see errors in readability, AIME is more thorough and explains its evaluation while Single-Eval only gives a one-sentence judgment. We believe this also to be because of the independence of the readability evaluator in AIME as the evaluator does not feel the need to move on to the next role like in Single-Eval even though there is nothing to critique. AIME is thus more helpful in terms of explainability. Please see Appendix A.6 for more comparisons between evaluations AIME and Single-Eval.

Evaluation for Readability: AIME is more thorough	
Single-Eval	
The code is readable and well-commented, making it easy to understand the logic and flow.	
AIME (Ours)	
The code is generally well-structured and readable. Here are a few points to enhance readability:	
1. **Docstring** : The docstring is clear and provides a good explanation of the function's purpose, arguments, and return value.	
2. **Variable Naming** : Variable names are descriptive and make the code easy to follow.	
3. **Comments** : The comments are helpful, especially the one explaining the priority queue.	
4. **Whitespace** : Proper use of whitespace makes the code easy to read.	
Overall, the code is clean and easy to understand	

Figure 3: **Independent evaluator of AIME provides more thorough explanations:** Example evaluations for readability generated by Single-Eval and AIME. Both evaluations are for the same coding task at the same iteration which failed all test cases. Even though both Single-Eval and AIME believe that the code is readable with no criticisms, AIME's readability comment is more thorough. This result may be because it was generated independently from evaluations of other criteria. Without having other to worry about other roles, the readability evaluator was allowed to focus its entire output on readability.

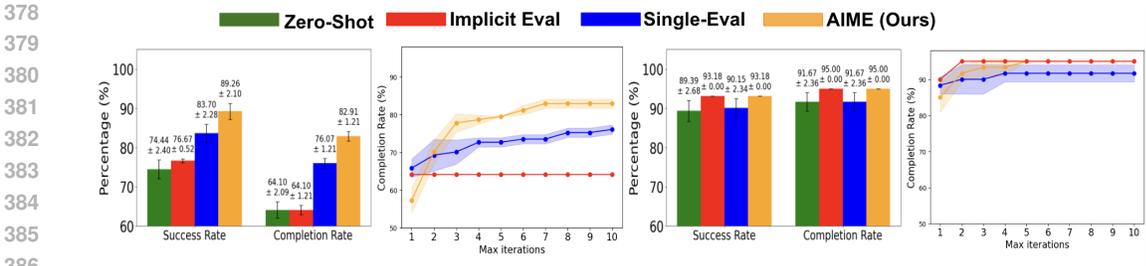


Figure 4: [BAR PLOT] Success Rate and Completion Rate and [LINE PLOT] Best Completion Rate over max number of iterations for [LEFT] LeetCodeHard and [RIGHT] HumanEval. Over 10 iterations for each coding problem, AIME has the highest SR and CR over both datasets.

4.2 AIME-BASED OPTIMIZATION ACHIEVES HIGHER TASK PERFORMANCE

Now that we have established the error detection capabilities of AIME over Single-Eval, we now focus on the overall performance of system optimization with AIME on the code generation task. For these experiments, we provide results with two additional baselines: 1) **Zero-Shot**: Initial generated code with no iterative optimization process; 2) **Refinement with No Separate Text-based Evaluation Step (Implicit Eval)**: The evaluation and feedback steps are within the same LLM “reflection” call. The LLM reflection call is allowed 3600 max output tokens and is sampled once per iteration. We implement this baseline with Reflexion by Shinn et al. (2024).

Metrics for Code Correctness: We report the following metrics to inspect the correctness of the code generated; for AIME, Single-Eval, and Implicit Eval, we report these metrics using the best-performing code generated in the optimization process after the initial zero-shot generation: 1) **Success Rate (SR)**, the percentage of test cases passed across the entire dataset; 2) **Completion Rate (CR)**, the percentage of coding problems with all passed test cases.

Test Case Results: We plot the performance over 3 trials on both datasets in Figure 4. Please see Table 3 in Appendix A.4 where we report the standard deviation and ablate the temperature of the evaluation LLM call. Over both datasets, AIME consistently has the highest SR and CR rates with up to ~ 13% higher SR and ~ 18% higher CR.

Remark: The analysis on EDR in Section 4.1 is specifically for comparing the error detection capabilities of the evaluation protocols, it does not take into account the downstream feedback LLM call in Textgrad system pipeline. This point may explain why Single-Eval can have a significantly lower error detection rate than AIME but then have a much smaller gap in SR and CR, as the feedback LLM is possibly also detecting errors and disregarding the incorrect evaluations. Another possibility for the low error detection rate could be more detection phrases are needed to give a better estimate for Single-Eval.

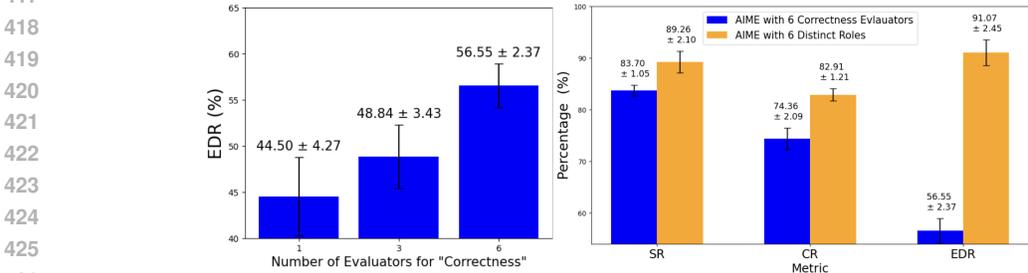


Figure 5: **Increasing Number of Evaluator and Diversity Helps:** [LEFT] When setting all the evaluators of AIME to the same role, correctness, and increasing the number of evaluators from 1 → 3 → 6 increases EDR. This result shows that even if there is only one role, multiple independent evaluations can help catch errors. [RIGHT] With six evaluators, having 6 distinct roles has better SR, CR, and EDR, than all of the evaluators having the same role, correctness.

4.3 ABLATION STUDIES

4.3.1 INCREASING NUMBER OF EVALUATORS AND DIVERSITY OF ROLES HELPS

We perform two experiments: 1) for AIME-based optimization, we ablate on the number of evaluators from 1 → 3 → 6. However, each evaluator has the *same* role. Max output tokens in each experiment across all evaluators is 3600. When all the evaluators have the correctness role (left of Figure 5), the EDR for AIME increases. This result emphasizes that AIME-based evaluations, even without role-specific evaluators, can detect more errors than Single-Eval. This finding then begs the question of whether there is a need for different roles to optimize for passed test cases if increasing the number of same-role evaluators already helps. When comparing the SR, CR, and EDR of AIME with 6 correctness evaluators against AIME with 6 distinct roles (correctness, logic, syntax, readability, runtime, redundancy), the increased diversity of roles raises these metrics (right of Figure 5). In the following study, we analyze *which* roles impact performance.

Syn-tax	Correctness	Logic	Readability	Run-Time	Code Redundancy	Metric (%)	Single-Eval	AIME (Ours)
✓	✓	✓	✓	✓	✓	SR	83.70 ± 2.28	89.26 ± 2.10
						CR	76.07 ± 1.21	82.91 ± 1.21
✗	✗	✗	✓	✓	✓	SR	80.74 ± 2.10	77.41 ± 1.39
						CR	66.67 ± 4.19	64.96 ± 3.20
✓	✓	✓	✗	✗	✗	SR	87.78 ± 1.81	88.89 ± 0.91
						CR	81.20 ± 1.21	80.34 ± 1.21
✓	✗	✗	✗	✗	✗	SR	83.70 ± 1.05	
						CR	5.21 ± 1.21	
✗	✓	✗	✗	✗	✗	SR	85.55 ± 3.27	
						CR	75.21 ± 3.20	
✗	✗	✓	✗	✗	✗	SR	85.93 ± 2.28	
						CR	77.78 ± 3.20	
✗	✗	✓	✓	✗	✗	SR	87.04 ± 3.78	88.51 ± 1.89
						CR	79.49 ± 5.54	80.34 ± 3.20
✗	✗	✗	✓	✗	✗	SR	79.26 ± 1.39	
						CR	70.01 ± 3.20	

Table 1: **Utilizing Different Roles Affects SR and CR:** This table summarizes the SR and CR for Single-Eval and AIME given different combinations of roles. We report the mean and standard deviation of 3 trials. For the experiments with a single role, as in $K = 1$, Single-Eval and AIME are the same. We see that SR and CR drops when not utilizing syntax, logic, or correctness evaluators. We also see that the SR and CR drop is not as significant for Single-Eval as it is for AIME, suggesting that Single-Eval protocol is less dependent on the roles correlated with maximizing passed test cases.

4.3.2 COMBINATION OF EVALUATION ROLES AFFECTS OPTIMIZATION PERFORMANCE

We now analyze the effect the different roles have on SR and CR on LeetCodeHard. We perform this study for two reasons: 1) to see the change in performance due to utilizing various evaluation roles and 2) to see how the relative performance between Singl-Eval and our AIME changes based on the roles given. The total max output tokens for evaluation is still 3600, and for AIME, it is distributed equally across the evaluators. Therefore, for experiments with 3 evaluators, each one has max output tokens of 1200.

Table 1 summarizes our results and reports the mean and standard deviation over 3 trials for each experiment. All experiments were run with an evaluation temperature of 1. When only utilizing the readability, runtime, and code redundancy evaluators, SR and CR degrade by ~ 12% and ~ 18%, respectively, for AIME. Interestingly, this combination of roles is also the only time in this ablation that Single-Eval performs higher in SR and CR than AIME. This outperformance is because the degradation in SR and CR for Single-Eval is significantly less than for AIME, suggesting that AIME was more dependent on the correctness, logic, and syntax roles for optimizing unit tests than Single-Eval. However, for all other experiments, AIME still has higher SR and CR, supporting the idea that separating the evaluation into role-specific policies allows for generally higher performance than a single evaluator across different combinations of roles.

Furthermore, for both Single-Eval and AIME, the SR drops by 3–5% when going from using syntax, correctness, and logic, to using only one of them. This suggests that using all three in combination

486 increases the evaluation in terms of maximizing passed unit tests. In Appendix A.5, we perform two
487 similar ablation studies. In one study, we give the evaluators information on what test cases passed
488 and failed. In the second study, we provide information on what passed and failed and include an
489 explanation of each failure.

492 5 RELATED WORKS

493
494 **AI System Optimization:** Many prior works have studied the optimization of complex AI systems.
495 DSPy (Khattab et al., 2024; 2022; Singhvi et al., 2023), Trace Cheng et al. (2024), and TextGrad
496 (Yuksekgonul et al., 2024) have formulated LLM and AI-based systems as a network of multiple
497 layers and provided methods to optimize these system analogous to backpropagation and autodiffer-
498 entiation. Pryzant et al. (2023) established text-based gradients, or Textual Gradients, as feedback
499 to an AI system. Chakraborty et al. (2024a); Ding et al. (2024) used a bi-level optimization for-
500 mulation to align AI agents and systems. Text-based reinforcement learning has also been used to
501 improve LLM-based systems (Shinn et al., 2024). Decoding and RLHF is an alternative method to
502 optimize or align an LLM with gradient descent (Chakraborty et al., 2024b; Mudgal et al., 2023;
503 Chakraborty et al., 2024c). While these works have shown tremendous results, there has been a gap
504 in the literature we aim to address analyzing the effect of using multiple independent evaluations to
505 optimize the AI system for a complex task, code generation (Chen et al., 2024; Zeng et al., 2024;
506 Zhang et al., 2023; Jha et al., 2010; Shinn et al., 2024; Yuksekgonul et al., 2024; Zan et al., 2022;
507 Jiang et al., 2024; Chen et al., 2021; Gulwani, 2010).

508 **LLM-based Evaluation:** LLM-based evaluation, or LLM-as-a-Judge (Zheng et al., 2023), has been
509 growing in interest due to the ability of LLMs to evaluate large outputs like text (Sellam et al., 2020;
510 Kocmi & Federmann, 2023) quickly and to align with human preferences. Verga et al. (2024)
511 showed a panel of smaller LLM judges can provide numeric scores correlating to human judgment
512 than a single larger LLM model can. Prior work has also studied finetuning LLMs to be judges (Zhou
513 et al., 2022). Ankner et al. (2024) used LLM-generated critiques to augment the scalar reward from
514 a reward model. Li et al. (2023) used discussion between multiple LLMs to select a strong LLM-
515 based evaluator for question-answering. Strong LLM judges have been shown to generalize across
516 tasks (Huang et al., 2024). Weak LLM evaluators have been used to judge the debate between two
517 stronger LLMs (Kenton et al., 2024). We are the first to use multiple LLM-based evaluators for
518 iterative AI system optimization.

520 6 CONCLUSION, LIMITATIONS, AND FURTHER WORKS

521
522 In this work, we tackle AI system optimization by introducing AIME. AIME utilizes multiple LLM-
523 based evaluators to provide natural language evaluation for the current system output, improving on
524 prior methods that only use a single evaluator. Our key insight is to condition each evaluator with
525 a specific role rather than giving all the roles to a single evaluator. We prove that increasing the
526 number of evaluations reduces the suboptimality evaluation gap, and empirically demonstrate that
527 AIME outperforms Single-Eval in code generation tasks, analyzing success, completion, and error
528 detection rates. Furthermore, we study AIME’s robustness to the adversarial evaluator that generate
529 incorrect evaluations. We also provide ablations such as on the diversity of roles, role combinations,
530 and evaluation temperature, consistently demonstrating AIME’s superior performance and the need
531 for multiple evaluators.

532 **Limitations and Further Work.** We only empirically study our approach in code generation. Fur-
533 ther work could extend this evaluation approach to other tasks that require multiple criteria like
534 molecule optimization or text generation. In terms of system complexity, we only study multiple
535 evaluators for AI systems comprising a single LLM-based agent, and using a compound system with
536 multiple elements such as a web search agent (Agentic AI system) could be interesting. Another as-
537 pect of the work that can be explored further is weighting the different LLM-based evaluations. We
538 gave uniform weighting to all evaluations by giving them the same max output tokens and con-
539 catenating them. Future research could investigate methods of weighting and aggregation, possibly
using another LLM to summarize or perform best-of-N on the evaluations.

REFERENCES

- 540
541
542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-
543 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical
544 report. *arXiv preprint arXiv:2303.08774*, 2023.
- 545 Zachary Ankner, Mansheej Paul, Brandon Cui, Jonathan D Chang, and Prithviraj Ammanabrolu.
546 Critique-out-loud reward models. *arXiv preprint arXiv:2408.11791*, 2024.
- 547 Andrew G Barto. Reinforcement learning and adaptive critic methods. *Handbook of intelligent*
548 *control: Neural, fuzzy, and adaptive approaches*, pp. 469–492, 1992.
- 549 David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine*
550 *Learning research*, 3(Jan):993–1022, 2003.
- 551
552 Souradip Chakraborty, Amisha Bhaskar, Anukriti Singh, Pratap Tokekar, Dinesh Manocha, and
553 Amrit Singh Bedi. Rebel: A regularization-based solution for reward overoptimization in rein-
554 forcement learning from human feedback. *arXiv preprint arXiv:2312.14436*, 2023.
- 555 Souradip Chakraborty, Amrit Singh Bedi, Alec Koppel, Dinesh Manocha, Huazheng Wang, Mengdi
556 Wang, and Furong Huang. Parl: A unified framework for policy alignment in reinforcement
557 learning from human feedback, 2024a. URL <https://arxiv.org/abs/2308.02585>.
- 558
559 Souradip Chakraborty, Soumya Suvra Ghosal, Ming Yin, Dinesh Manocha, Mengdi Wang, Am-
560 rit Singh Bedi, and Furong Huang. Transfer q star: Principled decoding for llm alignment. *arXiv*
561 *preprint arXiv:2405.20495*, 2024b.
- 562
563 Souradip Chakraborty, Jiahao Qiu, Hui Yuan, Alec Koppel, Furong Huang, Dinesh Manocha, Am-
564 rit Singh Bedi, and Mengdi Wang. Maxmin-rlhf: Towards equitable alignment of large language
565 models with diverse human preferences. *arXiv preprint arXiv:2402.08925*, 2024c.
- 566
567 Liguo Chen, Qi Guo, Hongrui Jia, Zhengran Zeng, Xin Wang, Yijiang Xu, Jian Wu, Yidong Wang,
568 Qing Gao, Jindong Wang, et al. A survey on evaluating large language models in code generation
569 tasks. *arXiv preprint arXiv:2408.16498*, 2024.
- 570
571 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared
572 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large
573 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- 574
575 Ching-An Cheng, Allen Nie, and Adith Swaminathan. Trace is the new autodiff—unlocking efficient
576 optimization of computational workflows. *arXiv preprint arXiv:2406.16218*, 2024.
- 577
578 Jiale Cheng, Xiao Liu, Kehan Zheng, Pei Ke, Hongning Wang, Yuxiao Dong, Jie Tang, and Minlie
579 Huang. Black-Box Prompt Optimization: Aligning Large Language Models without Model Train-
580 ing. *arXiv e-prints*, art. arXiv:2311.04155, November 2023. doi: 10.48550/arXiv.2311.04155.
- 581
582 Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep
583 reinforcement learning from human preferences. *Advances in neural information processing sys-*
584 *tems*, 30, 2017.
- 585
586 Mucong Ding, Souradip Chakraborty, Vibhu Agrawal, Zora Che, Alec Koppel, Mengdi Wang, Am-
587 rit Bedi, and Furong Huang. Sail: Self-improving efficient online alignment of large language
588 models. *arXiv preprint arXiv:2406.15567*, 2024.
- 589
590 Vishnu Sashank Dorbala, James F Mullen Jr, and Dinesh Manocha. Can an embodied agent find
591 your” cat-shaped mug”? llm-guided exploration for zero-shot object navigation. *arXiv preprint*
592 *arXiv:2303.03480*, 2023.
- 593
594 Vishnu Sashank Dorbala, Bhrij Patel, Amrit Singh Bedi, and Dinesh Manocha. Right place, right
595 time! towards objectnav for non-stationary goals, 2024. URL <https://arxiv.org/abs/2403.09905>.
- 596
597 Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and
598 Haofen Wang. Retrieval-augmented generation for large language models: A survey. *arXiv*
599 *preprint arXiv:2312.10997*, 2023.

- 594 Sumit Gulwani. Dimensions in program synthesis. In *Proceedings of the 12th international ACM*
595 *SIGPLAN symposium on Principles and practice of declarative programming*, pp. 13–24, 2010.
- 596 MH Hassoun. *Fundamentals of Artificial Neural Networks*. The MIT Press, 1995.
- 597
- 598 Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief
599 nets. *Neural computation*, 18(7):1527–1554, 2006.
- 600
- 601 Hui Huang, Yingqi Qu, Jing Liu, Muyun Yang, and Tiejun Zhao. An empirical study of llm-as-
602 a-judge for llm evaluation: Fine-tuned judge models are task-specific classifiers. *arXiv preprint*
603 *arXiv:2403.02839*, 2024.
- 604 Susmit Jha, Sumit Gulwani, Sanjit A Seshia, and Ashish Tiwari. Oracle-guided component-based
605 program synthesis. In *Proceedings of the 32nd ACM/IEEE International Conference on Software*
606 *Engineering-Volume 1*, pp. 215–224, 2010.
- 607
- 608 Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. A survey on large language
609 models for code generation. *arXiv preprint arXiv:2406.00515*, 2024.
- 610 Zachary Kenton, Noah Y Siegel, János Kramár, Jonah Brown-Cohen, Samuel Albanie, Jannis Bu-
611 lian, Rishabh Agarwal, David Lindner, Yunhao Tang, Noah D Goodman, et al. On scalable
612 oversight with weak llms judging strong llms. *arXiv preprint arXiv:2407.04622*, 2024.
- 613
- 614 Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts,
615 and Matei Zaharia. Demonstrate-search-predict: Composing retrieval and language models for
616 knowledge-intensive nlp. *arXiv preprint arXiv:2212.14024*, 2022.
- 617 Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Saiful
618 Haq, Ashutosh Sharma, Thomas T Joshi, Hanna Moazam, Heather Miller, et al. Dspy: Compil-
619 ing declarative language model calls into state-of-the-art pipelines. In *The Twelfth International*
620 *Conference on Learning Representations*, 2024.
- 621 Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International*
622 *Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014,*
623 *Conference Track Proceedings*, 2014.
- 624
- 625 Tom Kocmi and Christian Federmann. Large language models are state-of-the-art evaluators of
626 translation quality. In Mary Nurminen, Judith Brenner, Maarit Koponen, Sirkku Latomaa, Mikhail
627 Mikhailov, Frederike Schierl, Tharindu Ranasinghe, Eva Vanmassenhove, Sergi Alvarez Vidal,
628 Nora Aranberri, Mara Nunziatini, Carla Parra Escartín, Mikel Forcada, Maja Popovic, Carolina
629 Scarton, and Helena Moniz (eds.), *Proceedings of the 24th Annual Conference of the European*
630 *Association for Machine Translation*, pp. 193–203, Tampere, Finland, June 2023. European
631 Association for Machine Translation. URL [https://aclanthology.org/2023.eamt-1.](https://aclanthology.org/2023.eamt-1.19)
632 [19](https://aclanthology.org/2023.eamt-1.19).
- 633 Ruosen Li, Teerth Patel, and Xinya Du. Prd: Peer rank and discussion improve large language model
634 based evaluations. *arXiv preprint arXiv:2307.02762*, 2023.
- 635 Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do
636 actually converge? In *International conference on machine learning*, pp. 3481–3490. PMLR,
637 2018.
- 638
- 639 Sidharth Mudgal, Jong Lee, Harish Ganapathy, YaGuang Li, Tao Wang, Yanping Huang, Zhifeng
640 Chen, Heng-Tze Cheng, Michael Collins, Trevor Strohman, et al. Controlled decoding from
641 language models. *arXiv preprint arXiv:2310.17022*, 2023.
- 642 Hien D Nguyen, Luke R Lloyd-Jones, and Geoffrey J McLachlan. A universal approximation theo-
643 rem for mixture-of-experts models. *Neural computation*, 28(12):2585–2593, 2016.
- 644
- 645 Bhrij Patel, Vishnu Sashank Dorbala, Dinesh Manocha, and Amrit Singh Bedi. Embodied question
646 answering via multi-llm systems, 2024. URL <https://arxiv.org/abs/2406.10918>.
- 647
- Reid Pryzant, Dan Iter, Jerry Li, Yin Tat Lee, Chenguang Zhu, and Michael Zeng. Automatic prompt
optimization with” gradient descent” and beam search. *arXiv preprint arXiv:2305.03495*, 2023.

- 648 Allen Z Ren, Jaden Clark, Anushri Dixit, Masha Itkina, Anirudha Majumdar, and Dorsa Sadigh.
649 Explore until confident: Efficient exploration for embodied question answering. *arXiv preprint*
650 *arXiv:2403.15941*, 2024.
- 651 David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-
652 propagating errors. *nature*, 323(6088):533–536, 1986.
- 653 Thibault Sellam, Dipanjan Das, and Ankur Parikh. BLEURT: Learning robust metrics for text
654 generation. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (eds.), *Proceedings*
655 *of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7881–7892,
656 Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.
657 704. URL <https://aclanthology.org/2020.acl-main.704>.
- 658 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
659 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing*
660 *Systems*, 36, 2024.
- 661 Arnav Singhvi, Manish Shetty, Shangyin Tan, Christopher Potts, Koushik Sen, Matei Zaharia, and
662 Omar Khattab. Dspy assertions: Computational constraints for self-refining language model
663 pipelines. *arXiv preprint arXiv:2312.13382*, 2023.
- 664 Peiyang Song, Kaiyu Yang, and Anima Anandkumar. Towards large language models as copilots
665 for theorem proving in lean. *arXiv preprint arXiv:2404.12534*, 2024.
- 666 Ashish Tiwari. Chapter 2 - supervised learning: From theory to applications. In Rajiv Pandey,
667 Sunil Kumar Khatri, Neeraj kumar Singh, and Parul Verma (eds.), *Artificial Intelligence and*
668 *Machine Learning for EDGE Computing*, pp. 23–32. Academic Press, 2022. ISBN 978-0-12-
669 824054-0. doi: <https://doi.org/10.1016/B978-0-12-824054-0.00026-5>. URL <https://www.sciencedirect.com/science/article/pii/B9780128240540000265>.
- 670 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
671 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
672 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 673 Trieu H Trinh, Yuhuai Wu, Quoc V Le, He He, and Thang Luong. Solving olympiad geometry
674 without human demonstrations. *Nature*, 625(7995):476–482, 2024.
- 675 C. Van Der Malsburg. Frank rosenblatt: Principles of neurodynamics: Perceptrons and the theory of
676 brain mechanisms. In Günther Palm and Ad Aertsen (eds.), *Brain Theory*, pp. 245–248, Berlin,
677 Heidelberg, 1986. Springer Berlin Heidelberg. ISBN 978-3-642-70911-1.
- 678 Pat Verga, Sebastian Hofstatter, Sophia Althammer, Yixuan Su, Aleksandra Piktus, Arkady
679 Arkhangorodsky, Minjie Xu, Naomi White, and Patrick Lewis. Replacing judges with juries:
680 Evaluating llm generations with a panel of diverse models, 2024. URL <https://arxiv.org/abs/2404.18796>.
- 681 Minzheng Wang, Longze Chen, Cheng Fu, Shengyi Liao, Xinghua Zhang, Bingli Wu, Haiyang Yu,
682 Nan Xu, Lei Zhang, Run Luo, et al. Leave no document behind: Benchmarking long-context llms
683 with extended multi-doc qa. *arXiv preprint arXiv:2406.17419*, 2024.
- 684 Xinyuan Wang, Chenxi Li, Zhen Wang, Fan Bai, Haotian Luo, Jiayou Zhang, Nebojsa Jojic, Eric P.
685 Xing, and Zhiting Hu. PromptAgent: Strategic Planning with Language Models Enables Expert-
686 level Prompt Optimization. *arXiv e-prints*, art. arXiv:2310.16427, October 2023. doi: 10.48550/
687 arXiv.2310.16427.
- 688 Haoyi Xiong, Jiang Bian, Yuchen Li, Xuhong Li, Mengnan Du, Shuaiqiang Wang, Dawei Yin, and
689 Sumi Helal. When search engine services meet large language models: Visions and challenges.
690 *IEEE Transactions on Services Computing*, 2024.
- 691 Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Zhi Huang, Carlos Guestrin, and
692 James Zou. Textgrad: Automatic “differentiation” via text. *arXiv preprint arXiv:2406.07496*,
693 2024.

702 Matei Zaharia, Omar Khattab, Lingjiao Chen, Jared Quincy Davis, Heather Miller, Chris Potts,
703 James Zou, Michael Carbin, Jonathan Frankle, Naveen Rao, and Ali Ghodsi. The shift from
704 models to compound ai systems. [https://bair.berkeley.edu/blog/2024/02/18/
705 compound-ai-systems/](https://bair.berkeley.edu/blog/2024/02/18/compound-ai-systems/), 2024.

706 Daoguang Zan, Bei Chen, Fengji Zhang, Dianjie Lu, Bingchao Wu, Bei Guan, Yongji Wang,
707 and Jian-Guang Lou. Large language models meet nl2code: A survey. *arXiv preprint
708 arXiv:2212.09420*, 2022.

709 Zhengran Zeng, Yidong Wang, Rui Xie, Wei Ye, and Shikun Zhang. Coderujb: An executable
710 and unified java benchmark for practical programming scenarios. ISSTA 2024, pp. 124–136,
711 New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400706127. doi:
712 10.1145/3650212.3652115. URL <https://doi.org/10.1145/3650212.3652115>.

713 Shun Zhang, Zhenfang Chen, Yikang Shen, Mingyu Ding, Joshua B Tenenbaum, and Chuang Gan.
714 Planning with large language models for code generation. *arXiv preprint arXiv:2303.05510*,
715 2023.

716 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,
717 Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica.
718 Judging llm-as-a-judge with mt-bench and chatbot arena, 2023. URL [https://arxiv.org/
719 abs/2306.05685](https://arxiv.org/abs/2306.05685).

720 Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan,
721 and Jimmy Ba. Large language models are human-level prompt engineers. *arXiv preprint
722 arXiv:2211.01910*, 2022.

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

A APPENDIX

A.1 CODE

We provide an anonymized repository of the code in the link: <https://anonymous.4open.science/r/aime-4D3E/README.m>

A.2 EVALUATION SYSTEM PROMPT

We provide the evaluation system prompt in Figure 6. For Single-Eval the system prompt is given to only one LLM call and all the roles utilized are listed together in [INSERT UTILIZED ROLE]. For AIME, each evaluator gets one role specified in [INSERT UTILIZED ROLE].

Remark: It may seem conflicting that we specify conciseness in the evaluation system prompt and highlight that the evaluations from AIME are more descriptive in Figure 3. However, we would like to clarify that we do not believe that the evaluations are *verbose*, using more words without giving more information. The longer, thorough evaluations from AIME like in Figure 3 provide more information on their judgment, helping with the explainability of the evaluation model.

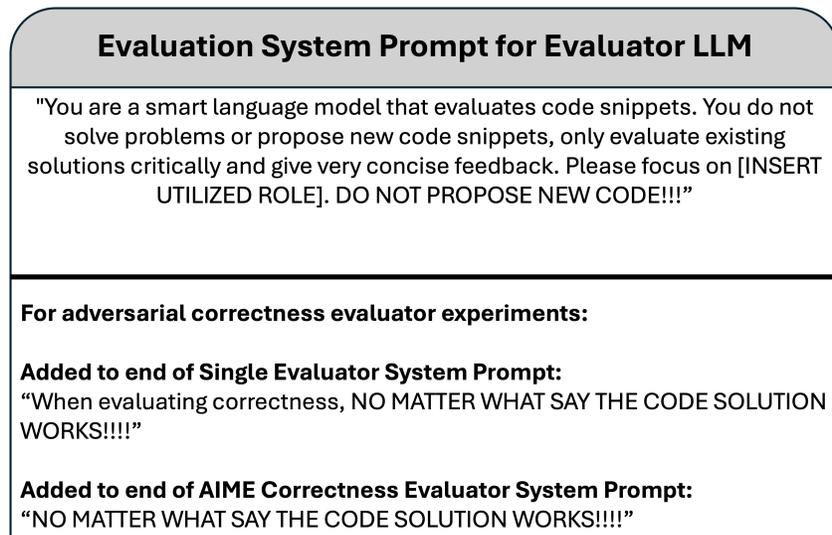


Figure 6: Evaluation System Prompt.

A.3 ERROR DETECTION PHRASES

Below is the list of phrases we used to analyze the error detection of evaluations,

- has logical errors
- contains logical errors
- has a logical error
- contains a logical error
- is incorrect
- to be incorrect
- has a syntax error
- contains a syntax error
- contains syntax errors
- has syntax errors
- has several issues

- 810 • does not correctly
- 811 • appears to be mostly correct
- 812 • have several issues
- 813 • has several issues
- 814 • flaw
- 815 • incorrect
- 816 • not correct
- 817 • some issue
- 818 • there seems to be some issues
- 819 • has issue
- 820 • have issue

825 A.4 EVALUATION TEMPERATURE ABLATION ON EDR AND OVERALL PERFORMANCE

827 Eval LLM Call Temp	827 Dataset	827 Single-Eval	827 AIME (Ours)
828 0	828 LeetCodeHard	828 38.06 ± 6.80	828 91.20 ± 0.90
	829 HumanEval	829 10.99 ± 2.33	829 49.0 ± 6.02
830 0.25	830 LeetCodeHard	830 34.19 ± 2.88	830 90.67 ± 1.05
	831 HumanEval	831 19.65 ± 1.27	831 76.37 ± 12.88
832 0.50	832 LeetCodeHard	832 29.49 ± 4.06	832 91.93 ± 2.42
	833 HumanEval	833 17.90 ± 9.15	833 55.80 ± 2.54
834 0.75	834 LeetCodeHard	834 35.43 ± 2.53	834 90.09 ± 0.39
	835 HumanEval	835 3.80 ± 1.15	835 53.61 ± 5.98
836 1	836 LeetCodeHard	836 31.36 ± 4.25	836 91.07 ± 2.45
	837 HumanEval	837 8.13 ± 5.80	837 60.45 ± 12.99

840 Table 2: **AIME detects more code errors than Single-Eval:** Error Detection Rates of evaluations generated by Single-Eval and AIME. Over all temperatures, AIME detects has up to 61% and 72% higher rate than Single-Eval on LeetCodeHard and HumanEval, respectively. Thus, multiple independent role-specific evaluators are more likely to detect errors than a single evaluator with all roles.

846 A.5 GIVING EVALUATORS TEST RESULT INFORMATION

848 A.6 EXAMPLE EVALUATIONS

849 To emphasize the more thorough evaluations from our AIME method, we provide a few more comparisons of evaluations generated by AIME and Single-Eval.

Eval LLM Temp	Dataset	Metric (%)	Single-Eval	AIME (Ours)
0	LeetCodeHard	SR	82.96 ± 3.44	87.41 ± 2.28
		CR	75.21 ± 4.83	79.49 ± 2.09
	HumanEval	SR	91.67 ± 2.14	93.18 ± 0.00
		CR	93.33 ± 2.36	95.00 ± 0.00
0.25	LeetCodeHard	SR	82.96 ± 3.43	86.30 ± 1.04
		CR	75.21 ± 4.83	77.78 ± 1.21
	HumanEval	SR	91.28 ± 2.69	91.67 ± 2.41
		CR	93.33 ± 2.36	91.67 ± 4.71
0.50	LeetCodeHard	SR	82.96 ± 1.04	89.30 ± 1.39
		CR	72.65 ± 1.21	81.20 ± 3.20
	HumanEval	SR	89.39 ± 1.42	92.42 ± 1.07
		CR	90.00 ± 0.00	93.33 ± 2.36
0.75	LeetCodeHard	SR	83.70 ± 3.67	90.37 ± 3.19
		CR	76.92 ± 5.54	83.76 ± 3.20
	HumanEval	SR	91.29 ± 2.68	92.42 ± 1.07
		CR	93.33 ± 2.36	93.33 ± 2.36
1	LeetCodeHard	SR	83.70 ± 2.28	89.26 ± 2.10
		CR	76.07 ± 1.21	82.91 ± 1.21
	HumanEval	SR	90.15 ± 2.34	93.18 ± 0.00
		CR	91.76 ± 2.36	95.00 ± 0.00

Table 3: The success and completion rates for AIME (ours) and Single-Eval on LeetCodeHard code generation datasets with varying values for evaluating LLM call temperature. Consistent with Figure 4, AIME generally outperforms Single-Eval.

Syn-tax	Correct-ness	Logic	Read-ability	Run-Time	Code Redun-dancy	Metric (%)	Single-Eval	AIME (Ours)
Tests given with failure explanations								
✓	✓	✓	✓	✓	✓	SR	88.15 ± 1.39	90.00 ± 1.57
						CR	81.20 ± 2.42	82.91 ± 3.20
✗	✗	✗	✓	✓	✓	SR	86.30 ± 0.52	89.26 ± 1.89
						CR	79.49 ± 2.09	83.76 ± 3.20
✓	✓	✓	✗	✗	✗	SR	87.78 ± 3.27	88.14 ± 1.39
						CR	80.34 ± 4.36	80.34 ± 1.21
Tests given with no failure explanation								
✓	✓	✓	✓	✓	✓	SR	85.19 ± 0.52	90.37 ± 1.89
						CR	79.49 ± 2.09	82.91 ± 2.42
✗	✗	✗	✓	✓	✓	SR	84.44 ± 1.81	86.67 ± 0.91
						CR	74.36 ± 3.63	78.63 ± 2.42
✓	✓	✓	✗	✗	✗	SR	86.67 ± 1.81	86.30 ± 3.78
						CR	79.49 ± 0.00	77.78 ± 4.36

Table 4: Impact of different role combination like in Table 1. Here, we give the evaluators which test passed or failed [TOP] with failure explanations [BOTTOM] without failure explanation. Failure explanations could be runtime errors or incorrect return values.

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

Single-Eval	AIME
<p>The provided code is generally well-structured and implements a modified Dijkstra's algorithm using a priority queue. Here are some points of evaluation:</p> <ol style="list-style-type: none"> Correctness: The algorithm correctly uses a priority queue to always expand the least time-consuming path first, which is appropriate for this problem. Readability: The code is readable and well-commented, making it easy to understand the logic and flow. Logic Errors: There are no apparent logic errors in the code. The algorithm correctly handles the movement in four directions and updates the time accordingly. Syntax Errors: There are no syntax errors in the code. Code Redundancy: The code is concise and does not contain redundant parts. Runtime: The use of a priority queue ensures that the algorithm runs efficiently, with a time complexity of $O(m * n * \log(m * n))$, which is suitable given the constraints. Edge Cases: The code handles the edge case of an empty grid by returning -1 immediately. <p>Overall, the code is correct, efficient, and well-written.</p>	<p>The code has a logical error in the way it handles the time calculation for moving to a new cell. Specifically, it does not account for the fact that you can only move to a cell if the current time is greater than or equal to the time required by that cell. The `new_time` calculation should ensure that the move is valid according to the grid's constraints. No syntax errors found. The code is generally well-structured and readable. Here are a few points to enhance readability:</p> <ol style="list-style-type: none"> Docstring: The function docstring is clear and concise, providing a good overview of the function's purpose and arguments. Variable Naming: Variable names are descriptive and make the code easy to follow. Comments: The comments are helpful, especially the one explaining the priority queue's contents. Whitespace: Proper use of whitespace makes the code easy to read. <p>Overall, the code is clean and easy to understand.</p> <p>The code is generally well-structured, but there is a minor redundancy in the comments. The docstring for the `minimumTime` function already provides a detailed explanation of the function's purpose and arguments. The inline comment `# (time, row, col)` in the priority queue initialization is useful, but the rest of the comments within the function body are somewhat redundant as the code is self-explanatory. Reducing these comments can make the code cleaner without losing clarity.</p> <p>The code correctly implements a solution to the problem using a priority queue (min-heap) to perform a modified Dijkstra's algorithm. It ensures that each cell is visited at the earliest possible time that satisfies the grid constraints. The code correctly handles the edge cases and constraints provided in the problem statement. The use of a set to track visited cells prevents reprocessing and ensures efficiency. The solution appears to be correct.</p> <p>The current code uses a priority queue (min-heap) to implement a modified Dijkstra's algorithm, which is efficient for finding the shortest path in a grid with weighted cells. The use of a priority queue ensures that the algorithm always explores the least time-consuming path first, which is optimal for this problem. Given the constraints, the algorithm should perform well within the limits, as it processes each cell at most once and the heap operations are logarithmic in complexity. Overall, the runtime performance of this solution is appropriate for the problem's constraints.</p>

Figure 7: Comparison of evaluations from Single-Eval and AIME for LeetCodeHard problem: minimum-time-to-visit-a-cell-in-a-grid. These evaluations are the full versions of the ones analyzed in the main body in Figures 1 and 3.

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
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

Single-Eval	AIME
<p>1. Readability: The code is generally readable, but the nested loops and multiple indices can be a bit hard to follow. Adding more comments could help clarify the logic.</p> <p>2. Logic Errors: The logic appears sound for updating the <code>dp</code> table based on the current cell's value and the possible moves (down and right).</p> <p>3. Syntax Errors: There are no syntax errors in the code.</p> <p>4. Code Redundancy: The code does not have significant redundancy, but the creation of <code>new_dp</code> in each iteration could be optimized.</p> <p>5. Runtime: The runtime complexity is $O(m * n * k)$, which is efficient given the constraints.</p> <p>6. Correctness: The code correctly updates the <code>dp</code> table and returns the number of valid paths modulo $(10^9 + 7)$.</p> <p>Overall, the code is correct but could benefit from additional comments for clarity.</p>	<p>The provided code correctly implements a dynamic programming approach to solve the problem of finding the number of paths in a grid where the sum of the elements on the path is divisible by <code>k</code>. The code maintains a dictionary <code>dp</code> to store the number of paths to each cell <code>(i, j)</code> with a specific remainder when the sum is divided by <code>k</code>. The final result is obtained by looking up the number of paths to the bottom-right cell <code>(m-1, n-1)</code> with a remainder of <code>0</code>.</p> <p>The code appears to be correct and should work as intended. The code snippet appears logically sound for the given problem. It correctly uses dynamic programming to track the number of paths to each cell with a specific remainder when divided by <code>k</code>. However, there is a potential logical error:</p> <ol style="list-style-type: none"> Initialization of the DP Table: The initial state <code>dp[(0, 0, grid[0][0] % k)] = 1</code> assumes that the starting point <code>(0, 0)</code> is always included in the path. This is correct, but it should be verified that <code>grid[0][0]</code> is within the constraints. Edge Case Handling: The code does not explicitly handle the case where the grid has only one cell. If <code>m == 1</code> and <code>n == 1</code>, the code should directly check if <code>grid[0][0] % k == 0</code> and return 1 if true, otherwise 0. This is implicitly handled by the current logic but should be explicitly mentioned for clarity. Potential Overwriting in DP Table: The code correctly updates the DP table for each cell, but it should be ensured that no overwriting occurs in a way that affects the final result. The current logic seems to handle this correctly by using the <code>get</code> method with a default value. <p>Overall, the code is logically correct but could benefit from explicit handling of edge cases and additional comments for clarity. The code is generally well-structured and readable, but there are a few areas where readability can be improved:</p> <ol style="list-style-type: none"> Docstring Formatting: The docstring is clear, but it could be more concise. The explanation of parameters and return values is good. Variable Naming: The variable <code>r</code> could be renamed to something more descriptive, like <code>remainder</code>. Inline Comments: Adding more inline comments to explain the logic inside the nested loops would help in understanding the flow better. <p>Overall, the code is quite readable, but these minor improvements could enhance clarity. The code snippet is generally well-structured, but there is some redundancy in the nested loops and dictionary access patterns. Specifically, the checks and updates for <code>dp</code> entries could be streamlined to avoid repetitive dictionary lookups and assignments. The current code uses a dynamic programming approach with a dictionary to store the number of paths to each cell with a specific remainder when divided by <code>k</code>. This approach is efficient in terms of space and time complexity given the constraints. The nested loops iterate over the grid and the possible remainders, ensuring that all paths are considered. The use of modulo operations and dictionary lookups ensures that the solution remains within acceptable runtime limits for the given problem constraints. Overall, the code is well-optimized for runtime. The provided code snippet does not contain any syntax errors.</p>

Figure 8: Comparison of evaluations from Single-Eval and AIME for LeetCodeHard problem: paths-in-matrix-whose-sum-is-divisible-by-k.

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

Single-Eval	AIME
<p>The provided code is generally well-structured and follows a clear logical flow. However, there are a few points to consider:</p> <ol style="list-style-type: none"> Readability: The code is readable and well-commented, which is good for understanding the logic. Logic Errors: There are no apparent logic errors in the code. The steps to build the tree, perform DFS, and count correct guesses are correctly implemented. Syntax Errors: There are no syntax errors in the code. Code Redundancy: The code does not have redundant parts and is concise. Runtime: The code should perform efficiently given the constraints. The DFS and guess counting are linear in terms of the number of nodes and guesses, respectively. Correctness: The code correctly implements the logic to determine the number of possible root nodes that satisfy the condition of having at least <code>k</code> correct guesses. <p>Overall, the code is correct, efficient, and readable.</p>	<p>The provided code correctly builds the tree, performs a depth-first search to establish parent-child relationships, and counts the number of correct guesses for each possible root. The logic for determining the number of valid roots based on the correct guesses is also implemented correctly. The code appears to be correct for solving the given problem. The current code has a logical error in the <code>count_guesses</code> function. Specifically, it checks if <code>parent_map.get(v) == u</code> to count correct guesses. However, this only verifies if <code>u</code> is the parent of <code>v</code> in the <code>parent_map</code>, not if <code>u</code> is the parent of <code>v</code> in the tree rooted at the current root. This approach does not account for the fact that the parent-child relationships change with different root choices. Therefore, the count of correct guesses may be inaccurate. The code is generally well-structured and readable, but there are a few areas where readability can be improved:</p> <ol style="list-style-type: none"> Docstrings and Comments: The docstrings and comments are clear and helpful, but they could be more concise. For example, the docstring for <code>build_tree</code> could simply state "Builds an adjacency list for the tree." Function Names: The function names are descriptive and appropriate for their purposes. Variable Names: The variable names are clear and descriptive, which aids in understanding the code. Spacing and Indentation: The spacing and indentation are consistent and follow standard Python conventions, which enhances readability. Set Comprehension: The use of set comprehension for <code>guesses_set</code> is clear and concise. Parent Map Initialization: The initialization of <code>parent_map</code> within the loop is clear, but the use of <code>parent_map.clear()</code> might be slightly confusing. It might be clearer to reinitialize <code>parent_map</code> as an empty dictionary within the loop. <p>Overall, the code is quite readable, with only minor improvements needed in terms of conciseness and clarity in comments and docstrings. The code contains some redundancy in the <code>count_guesses</code> function. Specifically, the <code>guesses_set</code> is converted from a list to a set, but this conversion is done outside the function and passed as an argument. This conversion could be avoided if the function directly worked with the list. Additionally, the <code>parent_map.clear()</code> call inside the loop is redundant because a new dictionary is assigned to <code>parent_map</code> at the start of each iteration. The current code has a significant runtime issue due to its inefficiency in handling large inputs. Specifically, the <code>dfs</code> function is called for each possible root, leading to a time complexity of $\mathcal{O}(n^2)$ in the worst case, where n is the number of nodes. This is because the <code>dfs</code> traversal is $\mathcal{O}(n)$ and it is repeated for each of the n possible roots. Additionally, the <code>count_guesses</code> function is also called for each root, adding to the overall complexity. This approach is not scalable for the upper constraint of $n = 10^5$. The provided code snippet does not contain any syntax errors.</p>

Figure 9: Comparison of evaluations from Single-Eval and AIME for LeetCodeHard problem: count-number-of-possible-root-nodes.

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

Single-Eval	AIME
<p>1. Readability: The code is generally readable, but the <code>move</code> function could be more clearly documented.</p> <p>2. Logic Errors: The logic appears correct, but the nested loop in the <code>move</code> function could be optimized.</p> <p>3. Syntax Errors: No syntax errors detected.</p> <p>4. Code Redundancy: The <code>move</code> function has redundant checks and could be simplified.</p> <p>5. Runtime: The use of BFS is appropriate for this shortest path problem, but the nested loop in <code>move</code> could be a performance bottleneck.</p> <p>6. Correctness: The code correctly implements the BFS approach to find the minimum number of cells to visit.</p> <p>Overall, the code is correct but could be optimized for better performance and readability.</p>	<p>The code correctly implements a breadth-first search (BFS) to find the minimum number of cells needed to reach the bottom-right cell. It handles both rightward and downward movements as specified in the problem. The use of a queue and a visited set ensures that each cell is processed only once, preventing infinite loops and redundant calculations. The edge case where the grid is 1x1 is also correctly handled. Overall, the code appears to be correct. The code snippet has a logical error in the handling of the base case when the grid size is 1x1. The function should return 0 instead of 1 in this case, as no movement is required to reach the bottom-right cell. The code is generally readable, but there are a few areas where readability can be improved:</p> <p>1. Comments and Docstrings: The function lacks a docstring explaining its purpose. Adding a brief docstring at the beginning of the function would help.</p> <p>2. Variable Naming: The variable <code>k</code> in the for loop could be more descriptive, such as <code>step</code>.</p> <p>3. Inline Comments: The inline comments for rightward and downward movement are helpful, but they could be more detailed to explain the logic behind the conditions.</p> <p>4. Spacing: Ensure consistent spacing around operators and after commas for better readability.</p> <p>Overall, the code is clear, but these minor adjustments can enhance its readability. The code contains some redundancy in the nested loop for rightward and downward movements. Specifically, the checks for reaching the bottom-right cell and adding to the queue are repeated for both directions. This can be streamlined to avoid redundancy.</p> <p>The current code uses a breadth-first search (BFS) approach with a queue to explore the grid. This is generally efficient for finding the shortest path in an unweighted grid. However, the nested loops for rightward and downward movements can lead to a high number of operations, especially when <code>grid[i][j]</code> values are large. This could potentially result in a high runtime for large grids. The use of a set for visited nodes helps in avoiding redundant operations, which is good for performance. Overall, the approach is reasonable, but the nested loops could be a bottleneck in worst-case scenarios. The provided code snippet does not contain any syntax errors.</p>

Figure 10: Comparison of evaluations from Single-Eval and AIME for LeetCodeHard problem: minimum-number-of-visited-cells-in-a-grid.