

# Investigating the Transferability of Code Repair for Low-Resource Programming Languages

Kyle Wong<sup>1</sup>, Alfonso Amayuelas<sup>1</sup>, Liangming Pan<sup>2</sup>, William Yang Wang<sup>1</sup>

<sup>1</sup> University of California, Santa Barbara

<sup>2</sup> University of Arizona

{knw, amayuelas}@ucsb.edu

liangmingpan@arizona.edu, william@cs.ucsb.edu

## Abstract

Large language models (LLMs) have shown remarkable performance on code generation tasks. A recent use case is iterative code repair, where an LLM fixes an incorrect program by rationalizing about errors and generating new code. Recent works augment the code repair process by integrating modern techniques such as chain-of-thought reasoning or distillation, but only study their benefits on high-resource languages like Python, and ignore low-resource languages like Perl. To address this gap of knowledge, we investigate the benefits of distilling code repair for both high and low resource languages to determine if the techniques that are effective in a high resource setting are also applicable in a low resource setting. Our evaluation shows that distilling the ability to repair code has language dependent benefits. To explain this behavior, we perform a further analysis and find that contrary to preexisting beliefs, the correlation between reasoning ability and code correction ability is weak. We hypothesize this weak correlation is magnified in low-resource settings where base models lack deep knowledge of a programming language, leading to wavering benefits of code repair.

## 1 Introduction

While large language models (LLMs) like GPT-4 (OpenAI, 2024) display remarkable coding abilities on popular benchmarks like HumanEval (Chen et al., 2021), the results of smaller models like CodeLlama-7b (Rozière et al., 2024) lag behind. Thus, frameworks that improve code generation and can apply to smaller models have become increasingly useful (Ding et al., 2024; Shinn et al., 2023). Code repair is one such framework, inspired by human editing: erroneous feedback is provided through executing tests, while programmers rationalize about those errors to fix the code. A standard code repair pipeline is depicted in Figure 1.

Source code: [github.com/KyleWong288/Distill\\_LRPL](https://github.com/KyleWong288/Distill_LRPL)

Although code repair is seemingly effective, recent works conclude that it is bottlenecked by a model’s underlying ability to rationalize about errors (Olausson et al., 2024), leading to lesser improvements on weaker models. To further improve repairs for smaller LLMs, prior works transfer the ability to reason about incorrect code, either from humans (Chen et al., 2024a) or other LLMs (Olausson et al., 2024; Ren et al., 2024). However, the effects of knowledge transfer on code repair is primarily studied on high-resource programming languages (HRPLs), and its generalizability to low-resource programming languages (LRPLs) remains under-explored.

Distilling code repair is particularly useful in a low-resource setting because it improves code accuracy without requiring more pretraining data, which is the main limitation for LRPLs. For example, a modern code LLM DeepSeek-Coder (Guo et al., 2024) is trained on a dataset scraped from Github repositories, containing high-resource languages like Python and Java at rates of 15.12% and 18.63%, while other languages like Perl and Golang are at much lesser rates of 0.1% and 0.32%. Since code repair is effective at increasing pass rates without extra pretraining data, understanding when it works best is critical for improving low-resource code generation.

Our work investigates if the code repair techniques that work on HRPLs are equally applicable to LRPLs. If transferring reasoning abilities can sufficiently enhance LRPL code repair, then it could mitigate the problem of low representation in pretraining data. Thus, we aim to understand the efficacy of distilling code repair for LRPLs and offer a guideline on whether to prioritize reasoning or code knowledge. To do so, we compare two approaches to knowledge transfer: the first is transferring only reasoning (rationale-only), while the second is transferring both reasoning and code completions (rationale-plus-code).

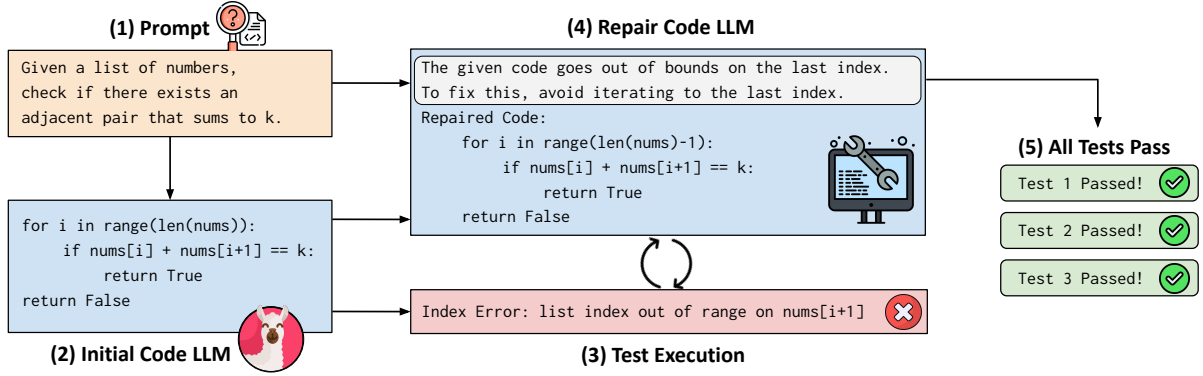


Figure 1: A standard code repair framework. In (1) and (2), a code LLM is given a question and generates a solution. In (3), test cases are executed and an error message is extracted. In (4), a repair LLM is given the question, incorrect solution, and error message, and generates a repair. A repair contains a rationale explaining why the old code was incorrect and how to fix it, followed by new code. If the new code is still incorrect, we iteratively generate new repairs using the code from previous repairs. In (5), we stop when all tests pass or after a fixed number of iterations.

We hypothesize that even if repair models are given a high quality rationale, they still often fail to fix incorrect code because they lack knowledge on the syntax and semantics of a programming language. This weakness is magnified in low-resource settings, so it is more beneficial to transfer code completions for LRPLs, but less necessary for HRPLs. To verify that the benefits of distilling code repair are language dependent, we perform a comprehensive suite of experiments and present our main research questions and findings below.

- **How effective is distillation for LRPLs?** Distilled repair models lead to higher pass rates. We see a relative increase in the average pass@1 of CodeLlama-7b-Instruct by 99.5% for Perl, 112.8% for Golang, and 144.5% for Swift after four rounds of repair on HumanEval.
- **How effective is distillation for LRPLs compared to HRPLs?** Rationale-plus-code distillation outperforms rationale-only transfer on LRPLs, but provides negligible improvements on HRPLs. On HumanEval, we see a relative increase in the pass@1 by 21.9% for Perl, 11.0% for Golang, and 16.3% for Swift, but see smaller increases on HRPLs.
- **Why does transferring both reasoning and code only outperform on LRPLs?** Even if repair models have a good rationale, they still struggle to make accurate code changes. This weakness is magnified in LRPLs, where models have a weaker understanding of the language. On LRPLs, rationale-only repair mod-

els generate a correct rationale 91.1% of the time, but only generate correct code 10.0% of the time, exposing a weak correlation between rationale correctness and code correctness.

## 2 Related Work

### 2.1 Repairing Code with LLMs

Iterative repair through feedback has been a well-studied area, surveyed in (Pan et al., 2023) and (Fernandes et al., 2023). The scenarios where self correction works best is also surveyed in (Kamoi et al., 2024). For code repair in specific, frameworks like Self-Repair (Olausson et al., 2024), CYCLE (Ding et al., 2024), CodeChain (Le et al., 2024), ILF (Chen et al., 2024a), Self-Edit (Zhang et al., 2023), Self-Debugging (Chen et al., 2023), and Reflexion (Shinn et al., 2023) have shown promising increases in pass rates.

### 2.2 Distillation for Code Repair

Distillation is the process of transferring knowledge from high capacity models to lower capacity models. Previous works show distillation can transfer the ability to reason and generate code (Sun et al., 2024; Wei et al., 2023b; Xu et al., 2023; Luo et al., 2023; Li et al., 2022a), but transferring the ability to repair code remains less explored. Recent methods like PERsD (Chen et al., 2024b) and LLM2LLM (Lee et al., 2024) use distillation to augment fine-tuning datasets. Self-Repair (Olausson et al., 2024) also experiments with transferring rationales from GPT-4 to CodeLlama-13b-Instruct in-context to improve the repair process. However, none of these approaches investigates the efficacy of distilling repair for low-resource code.

### 2.3 Low-Resource Programming Languages

Code repair experiments are usually evaluated on high-resource languages like Python, but our work investigates its efficacy on different languages. For evaluation, many works (Athiwaratkun et al., 2023; Orlanski et al., 2023; Zheng et al., 2023b) have created datasets to benchmark code generation in a multilingual setting. Since finding human written low-resource code is difficult, other approaches use capable LLMs to synthetically create low-resource code. Works like MultiPL-T (Cassano et al., 2024) and MultiPL-E (Cassano et al., 2022) translate popular pre-training datasets and monolingual benchmarks into a wide variety of different programming languages. Other works also study the transferability of coding ability between different languages (Baltaji et al., 2024; Gong et al., 2022).

## 3 Methodology

We provide an overview of a standard code repair framework, followed by an explanation of our distillation from a teacher model to a student model.

### 3.1 Code Repair Framework

We adopt a standard code repair pipeline as the basis of our method. While there exist more complex frameworks like those referenced in Section 2.1, our goal is to study the benefits of knowledge distillation, as opposed to inventing a new framework. Thus, we use a basic code repair pipeline as demonstrated in Figure 1, where the main components are the initial code generation, test execution, and iterative repair. We provide a formal explanation for each component.

First, we define  $M_{init}$  as the model generating initial answers. For a question  $q$ , we obtain  $n \geq 10$  initial samples, because it allows us to compute pass@10, along with lower variance pass@1 and pass@5 estimates. We define  $c_{t,i}$  as the  $i$ -th code sample generated on repair round  $t$ , where  $t = 0$  denotes the initial generation. Obtaining the initial code generations is formalized in expression 1.

$$M_{init}(q) \rightarrow \{c_{0,i}\}_{i=1}^n \quad (1)$$

Next, we define  $E$  as the code executor. Given a set of code samples, we execute the test cases associated with  $q$  on each sample. This produces a set of error messages, where  $e_{t,i}$  is the error message resulting from  $c_{t,i}$ . Obtaining the error messages is formalized in expression 2.

$$E(q, \{c_{t,i}\}_{i=1}^n) \rightarrow \{e_{t,i}\}_{i=1}^n \quad (2)$$

Finally, we define  $M_{repair}$  as the model generating repairs.  $M_{repair}$  has the same underlying model architecture as  $M_{init}$ . A repair is composed of a chain-of-thought (Wei et al., 2023a) rationale  $r_{t,i}$ , and the associated code completion  $c_{t,i}$ . Our work compares two different scenarios: transferring only reasoning (rationale-only transfer) vs transferring both reasoning and code completions (rationale-plus-code distillation).

For rationale-only transfer, we augment reasoning through in-context learning (Brown et al., 2020) by generating the rationale  $r_{t,i}$  from a separate larger model  $M_{teacher}$ , and appending this rationale to the prompt given to repair models. This process is formalized in expressions 3 and 4.

$$M_{teacher}(q, c_{t,i}, e_{t,i}) \rightarrow r_{t+1,i} \quad (3)$$

$$M_{repair}(q, c_{t,i}, e_{t,i}, r_{t+1,i}) \rightarrow c_{t+1,i} \quad (4)$$

For rationale-plus-code distillation,  $M_{repair}$  is responsible for independently generating both the rationale and code completion, and obtaining a repair is formalized in expression 5.

$$M_{repair}(q, c_{t,i}, e_{t,i}) \rightarrow (r_{t+1,i}, c_{t+1,i}) \quad (5)$$

### 3.2 Dataset Construction

To perform rationale-plus-code distillation from a teacher model to a student model, we construct a fine-tuning dataset. Our teacher model is GPT-3.5-Turbo (Ouyang et al., 2022; OpenAI, 2022), while our student models are CodeLlama-7b-Instruct (Rozière et al., 2024), CodeLlama-7b (Rozière et al., 2024), and Mistral-7b (Jiang et al., 2023). The fine-tuning datasets are constructed from MBXP (Athiwaratkun et al., 2023), which consists of multiple language specific benchmarks, each containing around 960 questions with corresponding test cases. An artificial train-test split is created by taking 800 random examples as potential training data and reserving the rest for testing. We process potential training examples into a finalized dataset, visualized in Figure 2. Our dataset is formally composed of five-tuples in the form  $(I, Q, A, E, R)$ , which we further explain.

**Instruction and Question.** Each five-tuple begins with a constant instruction  $I$ , informing the model to perform code repair. Next is a question  $Q$ , containing a problem description and function declaration. We collect  $Q$  by directly using the prompts provided in MBXP.

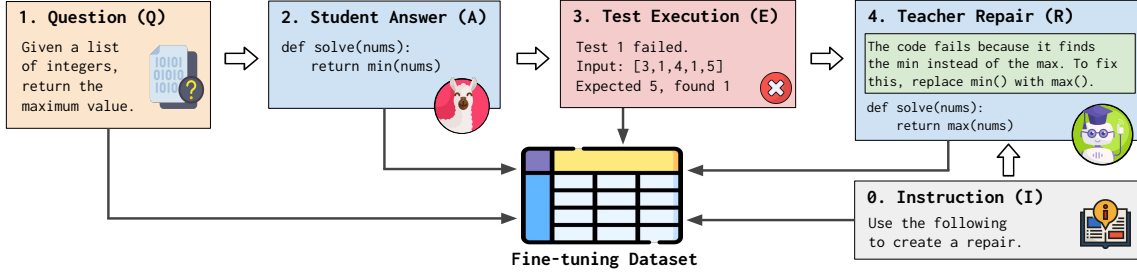


Figure 2: Our dataset construction pipeline. Examples in the fine-tuning dataset contain an instruction, the original question, the student’s incorrect answer, the execution feedback, and the teacher’s correct repair.

**Answer and Error.** The student’s incorrect answer is represented with  $A$ , which is collected by prompting a student model with  $Q$ . To ensure  $A$  is incorrect, we allow the student to continually generate independent samples, which are then immediately tested. Once a sample fails the given test cases, we select that sample as  $A$ . Then, we collect the associated error message  $E$  from the execution feedback.

**Repair.** Lastly, we finish with  $R$ , the teacher model’s repair. We collect  $R$  by providing a teacher model with  $(I, Q, A, E)$  and prompting it to generate two main components. The first component is a rationale that explains why the error occurred and a plan to fix it. The second component is updated code based on  $A$ , denoted with  $A'$ . To ensure  $A'$  is correct, we allow the teacher to continually generate independent repairs, which are then immediately tested. Once  $A'$  passes the given test cases, we select the associated repair as  $R$ . Our prompt format can be examined in Appendix D

**Quantity of Examples.** Although the original train split starts with 800 examples, our construction pipeline results in fine-tuning datasets with around 400 examples. Referencing Figure 2, this is because we may fail to obtain a usable  $A$  in step (2) or a usable  $R$  in step (4). In step (2), student models may consistently generate correct code. We allow a maximum of 10 samples before discarding the current example. Conversely, in step (4), teacher models may consistently generate incorrect code. We allow a maximum of 20 samples before discarding the current example. When prompting the teacher model, we use few-shot prompting (Brown et al., 2020) with three examples as an attempt to generate better repairs. The exact dataset sizes are listed in Appendix A.

## 4 Experiment

Our goal is to understand and compare the transferability of code repair for HRPLs and LRPLs, so we conduct a comprehensive experiment with three high-resource and low-resource languages. We identify Python, Javascript, and Java as high-resource, and identify Perl, Golang, and Swift as low-resource. These languages are picked based on having the highest three and lowest three pass rates observed in the original MBXP evaluations (Athiwaratkun et al., 2023), as well as cross referencing the percentage of each language in DeepSeek-Coder’s pretraining dataset (Guo et al., 2024), since it loosely reflects the distribution of programming languages found on Github. For each language, we perform our dataset construction and fine-tune a student model. Then, we generate an initial round of output and perform four rounds of code repair.

### 4.1 Experimental Setup

**Models.** To show our observations generalize to non-instruction-tuned and different model families, we run our experiments on CodeLlama-7b-Instruct (Rozière et al., 2024), CodeLlama-7b (Rozière et al., 2024), and Mistral-7b (Jiang et al., 2023). These models are used for the initial generation, and then a fine-tuned version of the same architecture is used as the distilled repair model.

**Benchmarks.** We evaluate on our MBXP (Athiwaratkun et al., 2023) test split from Section 3.2, containing around 160 programming problems for each language. Additionally, we evaluate on MultiLingual HumanEval (Athiwaratkun et al., 2023), a variation of HumanEval (Chen et al., 2021) transcompiled to different languages, which also contains around 160 programming problems for each language. Our evaluation on MultiLingual HumanEval (HumanEval for brevity) shows that fine-tuned repair models generalize to other datasets.



**Metrics.** We evaluate all generations using pass@k (Chen et al., 2021), a standard performance metric for code generation tasks. Since pass@k is prone to high variance, we use the unbiased estimator for pass@k, which estimates the probability that at least one out of k samples is correct. Given  $n \geq k$  code samples where  $c$  are correct, we compute pass@k using Equation 6.

$$\text{pass@k} := \mathbb{E}_{\text{Problems}} \left[ 1 - \frac{\binom{n-c}{k}}{\binom{n}{k}} \right] \quad (6)$$

**Training and Inference Details.** We perform a 90/10 train-dev split on the dataset resulting from Section 3.2, and train via LoRA fine-tuning (Hu et al., 2022). During the initial generation, we sample 10 answers for each question and compute pass@k using  $n=10$ , allowing us to measure certain baselines. However, we only perform code repair on the first 5 samples for later repair rounds and compute pass@1 using  $n=5$ , because we only care about the repair pass@1. To encourage diversity between samples, we use nucleus sampling with a threshold of 0.95 and sampling temperature of 0.2. Further training and inference hyperparameters are listed in Appendix B. For baselines that use a non-fine-tuned model for repair, we use one-shot prompting, whose format is shown in Appendix D.

## 4.2 Baselines

We compare the pass@1 of rationale-plus-code distillation to five different baselines. These baselines demonstrate how different repair approaches perform on HRPLs vs LRPLs, allowing us to analyze trade-offs and scenarios where rationale-plus-code distillation works best.

**Non-repair Independent Sampling.** We compare the efficiency of iterative repair rounds with independent sampling. This provides insights if distilled models can achieve higher pass rates with equal or fewer inference calls. Our experiment conducts 1 initial generation and 4 repair rounds for a total of 5 inference calls, so we compare the final pass@1 with the pass@5 and pass@10 of the initial generations.

**Basic Iterative Repair.** To measure the benefits of distillation, we evaluate how well the base model performs on iterative repair without any extra modifications. We use the same code repair workflow, but replace the distilled repair model with its non-fine-tuned counterpart.

**Rationale-only transfer** To measure the necessity of transferring code completions, we adopt an adjacent idea from Self-Repair (Olausson et al., 2024), where only the rationalization about an error is transferred in-context. First, a teacher model is prompted to generate the rationale portion of a repair. Then, a base repair model is prompted to generate the code portion of a repair, with the teacher’s rationale appended in-context. We use the same teacher and student models as Section 4.1, and our prompt to extract the teacher’s rationale is in Appendix C.

**Teacher Repair.** For demonstrating the limitations of distillation, we use the same code repair framework, but replace the student model with the teacher model used during dataset construction. This acts as a rough upper bound for the student model, and illustrates potential room for improvement.

## 4.3 Results

Our experiments provide empirical results demonstrating the pass@1 improvements of distilling both rationales and code completions, along with a wavering benefit of code repair between HRPLs and LRPLs. We report our results on CodeLlama-7b-Instruct in Figure 3, and similar results on CodeLlama-7b and Mistral-7b can be found in Appendix E and F.

**Impact of Distillation.** *Distilling code repair consistently outperforms repair with a base model.* Figure 3 shows both distillation plot lines steadily trending higher than the base repair plot lines. One explanation for this is that higher quality rationales may causally influence code correctness. Thus, weaker base models may not benefit as much from a framework like code repair, which requires strong reasoning to diagnose mistakes.

**Distilled Repair vs Independent Sampling.** *Distilled repair achieves higher pass rates than independently sampling.* Across all languages, four rounds of distilled code repair outperforms the initial pass@5. Furthermore, rationale-plus-code distilled models considerably outperform the initial pass@10 on LRPLs shown in Figure 3. Thus, when limited to a small amount of inference calls, distilling code repair can be a more efficient alternative than independent sampling for increasing pass rates.

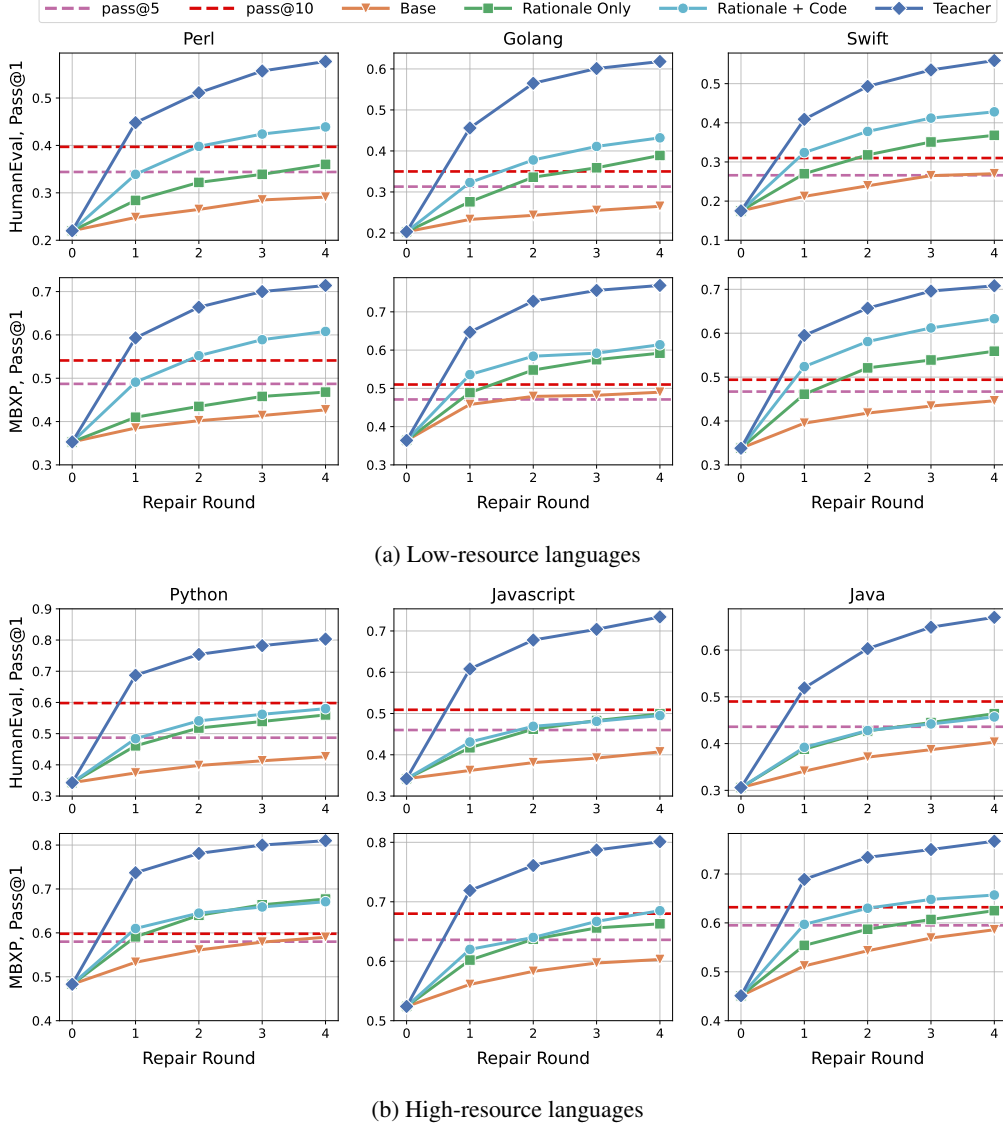


Figure 3: Mean pass@1 versus repair round for CodeLlama-7b-Instruct. Round 0 denotes the initial generation. Rationale-plus-code distillation outperforms rationale-only transfer on low-resource languages, but performs similarly on high-resource languages.

**Benefits on HRPLs vs LRPLs.** *Rationale-plus-code distillation consistently outperforms rationale-only transfer on LRPLs, but fails to do the same on HRPLs.* Although it might seem intuitive that transferring rationale-only would have lesser benefits than transferring rationale-plus-code, Table 1 quantifies how this is not always the case. Our results indicate distilling code repair provides a wavering benefit depending on the programming language, and spurs us to investigate why.

## 5 Analysis

After comparing the results of rationale-plus-code distillation with rationale-only transfer, we observe the benefits of distilling code repair de-

pends on whether the language is high-resource or low-resource. Thus, we seek an explanation on why rationale-plus-code distillation achieves higher pass rates on LRPLs, but not HRPLs. Previous research (Olausson et al., 2024; Ren et al., 2024) hypothesizes that code repair is bottlenecked by the model’s underlying ability to create a high quality rationale, which our experimental results support. However, there remains a lacking explanation of why repair models still generate incorrect code, even when given a high quality rationale.

We hypothesize there exists a second bottleneck: even if repair models are given good rationales, they fail to fix incorrect code because they lack the knowledge to convert suggestions from the ratio-

HumanEval Pass@1			
Language	Initial Round	Rationale Only Distillation	Rationale + Code Distillation
Perl	0.220	0.360	<b>0.439</b> ↑21.9%
Golang	0.203	0.389	<b>0.432</b> ↑11.3%
Swift	0.175	0.368	<b>0.428</b> ↑16.3%
Python	0.343	0.560	<b>0.580</b> ↑3.57%
Javascript	0.342	<b>0.499</b>	0.495 ↓0.80%
Java	0.306	<b>0.464</b>	0.457 ↓1.50%
MBXP Pass@1			
Perl	0.353	0.468	<b>0.608</b> ↑29.9%
Golang	0.364	0.592	<b>0.614</b> ↑3.71%
Swift	0.338	0.559	<b>0.633</b> ↑13.2%
Python	0.483	<b>0.677</b>	0.671 ↓0.88%
Javascript	0.524	0.663	<b>0.685</b> ↑3.31%
Java	0.451	0.625	<b>0.657</b> ↑5.12%

Table 1: Mean pass@1 of initial generations vs the two variants of distillation. Relative percent increases between the two variants of distillation is also provided. Rationale-plus-code distillation consistently outperforms on LRPLs, but performs around the same or slightly worse on HRPLs.

nale into accurate code modifications. This effect is magnified in a low-resource setting because base models are less knowledgeable about a language’s syntax and semantics, explaining why transferring code completions leads to more potent benefits for LRPLs. To support our hypothesis, we analyze the quality of rationales and a repair model’s knowledge of a language.

## 5.1 Correlation between Rationale and Code

To support our hypothesis that a bottleneck exists in a model’s ability to convert reasoning into code, we analyze the relationship between rationale quality and code correctness in Table 2. We quantitatively show that repair models frequently generate correct rationales, yet still generate incorrect code, exposing a weak correlation between the reasoning process and code editing process.

We use GPT-4 as an LLM judge to determine whether a rationale is sufficient or insufficient. Although human evaluation would be preferred, crowdsourcing participants well-versed in programming languages like Perl and Swift and capable of solving the coding problems found in HumanEval is challenging. We picked GPT-4 because the distilled rationales were generated with GPT-3.5-Turbo, and we aimed to use a more advanced and reliable model for better assessments (Zheng et al., 2023a).

To obtain judgments, we present a question, incorrect code, error message, and rationale to GPT-4, and instruct it to produce a verdict. A rationale is labeled good if it contains accurate information and includes sufficient detail to repair the given code, and bad otherwise. Our judgment prompt can be found in Appendix I. We obtain a verdict for all HumanEval rationales extracted between the initial generation and the first repair round.

Table 2 demonstrates that the rate of a good rationale leading to passing code is consistently low. For LRPLs, rationale-only models have a rate of 9.5%, while rationale-plus-code models have a rate of 15.4%. However, both students have equal rates of 14.1% on HRPLs. Thus, these results show how transferring code completions mitigate the weak correlation between rationale and code more effectively on LRPLs than for HRPLs. Furthermore, these results provide insight on the relationship between reasoning and code generation.

Many code repair frameworks follow some variation of the same steps: obtain an error, rationalize, generate code. Although prior works (Chen et al., 2024a; Olausson et al., 2024) focus solely on improving the weak link between error and rationale, our results highlight there is another weak link between rationale and code correctness. Since models have poor ability in converting a correct rationale into correct code modifications, augmenting the reasoning process alone leads to limited benefits, which has implications on popular techniques designed to improve reasoning, such as scaling inference-time compute. We display various examples where repair models provide good rationales, but the resulting code has clear syntactic or semantic errors in Appendix L.

## 5.2 Knowledge of LRPLs

Lastly, we analyze why transferring code completions only achieves consistent improvements on LRPLs. To support our hypothesis that a base model’s weak responsiveness is magnified in a low-resource setting, Table 3 shows how rationale-plus-code distilled models exhibit a deeper understanding of LRPLs. We use the frequency of syntax errors as a proxy for knowledge, since generating code with syntax errors is a blatant sign that a model lacks comprehension of a language.

To measure this, we first extract the set of syntax errors from a particular code repair run. Syntax errors are those occurring before execution and caught during compilation or interpretation time.

	RATIONALE + CODE			RATIONALE ONLY		
	Code Fails	Code Passes	Total	Code Fails	Code Passes	Total
LRPLS						
<b>Bad Rationale</b>	12.4%	1.0%	13.4%	8.4%	0.5%	8.9%
<b>Good Rationale</b>	71.2%	<b>15.4%</b>	86.6%	81.6%	<b>9.5%</b>	91.1%
<b>Total</b>	83.6%	16.4%		90.0%	10.0%	
HRPLS						
<b>Bad Rationale</b>	19.7%	2.3%	22.0%	9.3%	0.7%	10.0%
<b>Good Rationale</b>	63.9%	<b>14.1%</b>	78.0%	75.9%	<b>14.1%</b>	90.0%
<b>Total</b>	83.6%	16.4%		85.2%	14.8%	

Table 2: Empirical relationship between rationale quality and code correctness. Even when a good rationale is provided, the rate of producing passing code is significantly less than the rate of producing failing code. This exposes a weak correlation between rationale quality and code correctness.

HumanEval Average Syntax Errors				
Language	Initial Errors	Base Repair	Rationale Only	Rationale + Code
Perl	14.5	15.4 $\uparrow 0.9$	17.8 $\uparrow 3.3$	<b>9.20</b> $\downarrow 5.3$
Golang	44.7	70.4 $\uparrow 25.7$	48.7 $\uparrow 4.0$	<b>26.6</b> $\downarrow 18.1$
Swift	81.0	58.0 $\downarrow 23.0$	50.4 $\downarrow 30.6$	<b>37.4</b> $\downarrow 43.6$
Python	12.1	15.6 $\uparrow 3.5$	18.2 $\uparrow 6.1$	<b>14.2</b> $\uparrow 2.1$
Javascript	9.10	9.80 $\uparrow 0.7$	27.6 $\uparrow 18.5$	<b>9.00</b> $\downarrow 0.1$
Java	39.6	41.2 $\uparrow 1.6$	<b>37.0</b> $\downarrow 2.6$	41.2 $\uparrow 1.6$
MBXP Average Syntax Errors				
Perl	12.1	9.50 $\downarrow 2.6$	13.7 $\uparrow 1.6$	<b>2.70</b> $\downarrow 9.4$
Golang	33.2	29.2 $\downarrow 4.0$	26.8 $\downarrow 6.4$	<b>14.6</b> $\downarrow 18.6$
Swift	60.4	36.0 $\downarrow 24.0$	27.8 $\downarrow 32.6$	<b>11.0</b> $\downarrow 49.4$
Python	1.80	5.20 $\uparrow 3.4$	5.10 $\uparrow 3.3$	<b>3.60</b> $\uparrow 1.8$
Javascript	4.60	4.20 $\downarrow 0.4$	11.8 $\uparrow 7.6$	<b>3.60</b> $\downarrow 1.0$
Java	29.2	26.4 $\downarrow 2.8$	21.4 $\downarrow 5.0$	<b>20.4</b> $\downarrow 8.8$

Table 3: Mean number of syntax errors after code repair, along with the differences between the initial and final amount of errors. Rationale-plus-code distillation has a notably higher decline in syntax errors on LRPLs, but performs closer to baselines on HRPLs.

We can conveniently filter out syntax errors by parsing the execution feedback. Next, we compute the average amount of errors within the final repair round. Note that non-syntax errors can transform into syntax errors when repair models attempt to update code, leading to occasional increases. The average number of syntax errors for CodeLlama-7b-Instruct can be seen in Table 3, and similar results on CodeLlama-7b and Mistral-7b can be seen in Appendix J and K.

For LRPLs, the decrease in syntax errors when transferring code completions is higher than the other baselines. Averaging over the 3 LRPLs, rationale-plus-code models finish with 16.9 errors, while rationale-only models finish with 30.9 errors. Thus, for the case of LRPLs, boosting rationale

quality alone is not enough for encouraging a base model to generate a correct repair, and distilling code completions improves a model’s knowledge on a programming language.

For HRPLs, transferring code completions leads to less potent improvements. The decrease in syntax errors is smaller, and the final amount of errors among all baselines are relatively close. Averaging over the 3 HRPLs, rationale-plus-code models finish with 15.3 errors, while rationale-only models finish with 20.2 errors. Furthermore, even the base repair model performs comparably, finishing with an average of 17.1 errors. Thus, for the case of HRPLs, transferring code completions is less necessary because the base model already has sufficient knowledge on a programming language.

## 6 Conclusion

We distilled the ability to repair code and demonstrated that transferring only rationales is sufficient for high-resource languages, but further transferring code completions is necessary for low-resource languages. We also exposed that the correlation between rationale quality and code correctness is lower than previously perceived, especially in low-resource settings. This provides insight into the effectiveness of techniques that augment reasoning popular throughout research and industry, such as scaling inference-time compute. Rationale-plus-code distillation mitigates this weakness by improving a model’s understanding of a programming language, resulting in better responsiveness to feedback. Further research in distillation is important because it allows smaller models to gain fluency without costly human labeling, creating efficient and high-performing LLMs suitable for consumer-



grade devices. Such advancements would democratize the benefits of closed source research, making better code generation accessible for a wider range of languages, applications, and users.

## Limitations

One limitation is the lack of more challenging multilingual datasets. Other popular coding benchmarks like APPS (Hendrycks et al., 2021) and CodeContests (Li et al., 2022b) provide harder problems, which may demand stronger reasoning, but are only available in high-resource languages. Studying the benefits of distilling rationales on more reasoning heavy questions in low-resource languages would be an insightful future evaluation.

A natural limitation is the lack of instruction tuning datasets for LRPLs, as our fine-tuning datasets only contain around 400 examples. Although it would be ideal to have more examples, the amount of data available for these low-resource languages is naturally low. Our work counters concerns about the generalizability of our findings by evaluating on multiple models, languages, and benchmarks.

Another limitation in our evaluation are the stochastic processes within training and inference. To the best of our ability, we mitigate variance in our evaluation by seeding our training and inference, and by using the unbiased estimator of pass@k. However, since we use nucleus sampling for decoding, we observe there can be slight variations in our results.

Lastly, an underlying limitation is our hardware for training and inference. We use Nvidia Titan RTX GPUs with 24GB memory, so the size of student models that we can fine-tune is limited, which is why we choose 7b models for our experiments. Furthermore, since our evaluation has many dimensions (6 languages, 3 models, 5 baselines, 2 benchmarks, 160 questions each benchmark), we are limited in the amount of sampling we can do for each question. Although it may be interesting to obtain higher pass@k rates like k=10 or k=100, these are not time efficient to measure and do not contribute that much to our arguments. Thus, we choose to only show pass@1 for repair rounds.

## Ethics Statement

Since computing resources and research funding is extremely valuable, querying costly models like GPT-4 should be conducted responsibly. Estimating costs before running experiments and mak-

ing necessary adjustments is a responsible and resource-conscious approach to using such APIs.

Furthermore, there exists the possibility that users apply code repair for harmful applications. People with malicious intentions could use our research to improve code generation in certain domains that produce dangerous code, such as attacks on privacy and security. We encourage that code repair be used for socially responsible technology.

## References

- Ben Athiwaratkun, Sanjay Krishna Gouda, Zijian Wang, Xiaopeng Li, Yuchen Tian, Ming Tan, Wasi Uddin Ahmad, Shiqi Wang, Qing Sun, Mingyue Shang, Sujan Kumar Gonugondla, Hantian Ding, Varun Kumar, Nathan Fulton, Arash Farahani, Siddhartha Jain, Robert Giaquinto, Haifeng Qian, Murali Krishna Ramanathan, Ramesh Nallapati, Baishakhi Ray, Parminder Bhatia, Sudipta Sengupta, Dan Roth, and Bing Xiang. 2023. [Multi-lingual evaluation of code generation models](#). *Preprint*, arXiv:2210.14868.
- Razan Baltaji, Saurabh Pujar, Louis Mandel, Martin Hirzel, Luca Buratti, and Lav Varshney. 2024. [Learning transfers over several programming languages](#). *Preprint*, arXiv:2310.16937.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Federico Cassano, John Gouwar, Francesca Lucchetti, Claire Schlesinger, Anders Freeman, Carolyn Jane Anderson, Molly Q Feldman, Michael Greenberg, Abhinav Jangda, and Arjun Guha. 2024. [Knowledge transfer from high-resource to low-resource programming languages for code llms](#). *Preprint*, arXiv:2308.09895.
- Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, Arjun Guha, Michael Greenberg, and Abhinav Jangda. 2022. [Multipl-e: A scalable and extensible approach to benchmarking neural code generation](#). *Preprint*, arXiv:2208.08227.
- Angelica Chen, J  r  my Scheurer, Tomasz Korbak, Jon Ander Campos, Jun Shern Chan, Samuel R. Bow-

- man, Kyunghyun Cho, and Ethan Perez. 2024a. [Improving code generation by training with natural language feedback](#). *Preprint*, arXiv:2303.16749.
- Hailin Chen, Amrita Saha, Steven Hoi, and Shafiq Joty. 2024b. [Personalised distillation: Empowering open-sourced llms with adaptive learning for code generation](#). *Preprint*, arXiv:2310.18628.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. 2023. [Teaching large language models to self-debug](#). *Preprint*, arXiv:2304.05128.
- Yangruibo Ding, Marcus J. Min, Gail Kaiser, and Baishakhi Ray. 2024. [Cycle: Learning to self-refine the code generation](#). *Preprint*, arXiv:2403.18746.
- Patrick Fernandes, Aman Madaan, Emmy Liu, António Farinhas, Pedro Henrique Martins, Amanda Bertsch, José G. C. de Souza, Shuyan Zhou, Tongshuang Wu, Graham Neubig, and André F. T. Martins. 2023. [Bridging the gap: A survey on integrating \(human\) feedback for natural language generation](#). *Preprint*, arXiv:2305.00955.
- Zi Gong, Yinpeng Guo, Pingyi Zhou, Cuiyun Gao, Yasheng Wang, and Zenglin Xu. 2022. [Multicoder: Multi-programming-lingual pre-training for low-resource code completion](#). *Preprint*, arXiv:2212.09666.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. 2024. [Deepseek-coder: When the large language model meets programming – the rise of code intelligence](#). *Preprint*, arXiv:2401.14196.
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. 2021. [Measuring coding challenge competence with apps](#). *Preprint*, arXiv:2105.09938.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Ryo Kamoi, Yusen Zhang, Nan Zhang, Jiawei Han, and Rui Zhang. 2024. [When can llms actually correct their own mistakes? a critical survey of self-correction of llms](#). *Preprint*, arXiv:2406.01297.
- Hung Le, Hailin Chen, Amrita Saha, Akash Gokul, Doyen Sahoo, and Shafiq Joty. 2024. [Codechain: Towards modular code generation through chain of self-revisions with representative sub-modules](#). *Preprint*, arXiv:2310.08992.
- Nicholas Lee, Thanakul Wattanawong, Sehoon Kim, Kartikeya Mangalam, Sheng Shen, Gopala Anumanchipalli, Michael W. Mahoney, Kurt Keutzer, and Amir Gholami. 2024. [Llm2llm: Boosting llms with novel iterative data enhancement](#). *Preprint*, arXiv:2403.15042.
- Shiyang Li, Jianshu Chen, Yelong Shen, Zhiyu Chen, Xinlu Zhang, Zekun Li, Hong Wang, Jing Qian, Baolin Peng, Yi Mao, Wenhui Chen, and Xifeng Yan. 2022a. [Explanations from large language models make small reasoners better](#). *Preprint*, arXiv:2210.06726.
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d’Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. 2022b. [Competition-level code generation with alphacode](#). *Science*, 378(6624).
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023. [Wizardcoder: Empowering code large language models with evolve-instruct](#). *Preprint*, arXiv:2306.08568.
- Theo X. Olausson, Jeevana Priya Inala, Chenglong Wang, Jianfeng Gao, and Armando Solar-Lezama. 2024. [Is self-repair a silver bullet for code generation?](#) *Preprint*, arXiv:2306.09896.
- OpenAI. 2022. [Introducing chatgpt](#). Accessed: 2024-06-14.

- OpenAI. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Gabriel Orlanski, Kefan Xiao, Xavier Garcia, Jeffrey Hui, Joshua Howland, Jonathan Malmaud, Jacob Austin, Rishabh Singh, and Michele Catasta. 2023. [Measuring the impact of programming language distribution](#). *Preprint*, arXiv:2302.01973.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *Preprint*, arXiv:2203.02155.
- Liangming Pan, Michael Saxon, Wenda Xu, Deepak Nathani, Xinyi Wang, and William Yang Wang. 2023. [Automatically correcting large language models: Surveying the landscape of diverse self-correction strategies](#). *Preprint*, arXiv:2308.03188.
- Houxing Ren, Mingjie Zhan, Zhongyuan Wu, Aojun Zhou, Juntao Pan, and Hongsheng Li. 2024. [Reflectioncoder: Learning from reflection sequence for enhanced one-off code generation](#). *Preprint*, arXiv:2405.17057.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2024. [Code llama: Open foundation models for code](#). *Preprint*, arXiv:2308.12950.
- Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflexion: Language agents with verbal reinforcement learning](#). *Preprint*, arXiv:2303.11366.
- Zhihong Sun, Chen Lyu, Bolun Li, Yao Wan, Hongyu Zhang, Ge Li, and Zhi Jin. 2024. [Enhancing code generation performance of smaller models by distilling the reasoning ability of llms](#). *Preprint*, arXiv:2403.13271.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023a. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Yuxiang Wei, Zhe Wang, Jiawei Liu, Yifeng Ding, and Lingming Zhang. 2023b. [Magicoder: Source code is all you need](#). *Preprint*, arXiv:2312.02120.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhao Feng, Chongyang Tao, and Daxin Jiang. 2023. [Wizardlm: Empowering large language models to follow complex instructions](#). *Preprint*, arXiv:2304.12244.
- Kechi Zhang, Zhuo Li, Jia Li, Ge Li, and Zhi Jin. 2023. [Self-edit: Fault-aware code editor for code generation](#). *Preprint*, arXiv:2305.04087.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023a. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Preprint*, arXiv:2306.05685.
- Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Zihan Wang, Lei Shen, Andi Wang, Yang Li, Teng Su, Zhilin Yang, and Jie Tang. 2023b. [Codegeex: A pre-trained model for code generation with multilingual evaluations on humaneval-x](#). *Preprint*, arXiv:2303.17568.

## A Fine-tuning Dataset Sizes

Fine-tuning Dataset Sizes					
Language	Initial	Post-Student	Post-Teacher	Train	Dev
CodeLlama-7b-Instruct					
Perl	800	649	489	440	49
Golang	800	601	455	409	46
Swift	800	635	470	423	47
Python	800	559	446	401	45
Javascript	800	509	394	354	40
Java	800	667	510	459	51
CodeLlama-7b					
Perl	800	680	489	440	49
Golang	800	614	456	410	46
Swift	800	651	465	418	47
Python	800	596	470	423	47
Javascript	800	586	470	423	47
Java	800	642	499	449	50
Mistral-7b					
Perl	800	689	533	479	54
Golang	800	745	539	459	54
Swift	800	625	468	421	47
Python	800	602	487	438	49
Javascript	800	535	413	371	42
Java	800	573	439	395	44

Table 4: The final fine-tuning dataset sizes for each model, starting from the original MBXP train split of 800 questions. Intermediate sizes at each step of our dataset construction are also provided.

## B Training and Inference Hyperparameters

We provide our training and inference hyperparameters used throughout experiments. All training and inference are conducted on Nvidia Titan RTX (24GB) GPUs.

For training, we use LoRA fine-tuning with a rank of 128, lora alpha of 128, lora dropout of 0.1, maximum sequence length of 2048, batch size of 4, gradient accumulation steps of 2, weight decay of 0.01, cosine learning rate scheduler with warm up steps of 10, and checkpoint every 50 steps. For models in the CodeLlama family, we train for 8 epochs with a learning rate of  $2e-5$ , and for Mistral-7b, we train for 5 epochs with a learning rate of  $5e-6$ . To obtain our final distilled repair model, we pick the checkpoint with the lowest validation loss.

For inference, we use nucleus sampling with a threshold of 0.95, sampling temperature of 0.2, and limit the maximum new tokens to 800. When generating the initial 10 samples, we use random seeds from 1 to 10. For later repair rounds, all generations use a random seed of 17.

## C In-Context Rationale Prompt

```
### Instruction: You are given an incorrect python function and an error message.
Explain how to modify the code to fix the error, but do not write any actual code.
### Incorrect Code:
<Q,A>
### Error: <E>
### Response:
```

Figure 4: The prompt given to GPT-3.5-Turbo to generate the rationale portion of a repair. This is only used for the in-context learning baseline. <Q,A> is replaced with the question and previous answer, while <E> is replaced with the corresponding error.



## D Repair Prompt

We provide the general format of our repair prompt. When conducting code repair with a rationale-plus-code distilled model, we use zero-shot prompting. When conducting code repair with baselines, we use one-shot prompting. When creating a fine-tuning dataset with the teacher model, we use three-shot prompting. The few-shot examples change with each programming language.

```
### Instruction: You are given an incorrect python function and an error message. Explain how to fix
the error, and then write an updated python function with the correct code.

### Incorrect Code:
def check_String(str):
    """
    Write a python function to check whether a string has at least one letter and one number.
    >>> check_String('thishasboth29')
    True
    >>> check_String('python')
    False
    >>> check_String('string')
    False
    """
    if (str.isalnum()):
        return True
    else:
        return False

### Error:
Traceback (most recent call last):
  File "temp.py", line 27, in <module>
    check(check_String)
  File "temp.py", line 23, in check
    assert candidate('python') == False
AssertionError

### Response: The incorrect function is returning True for an input of "python", but it should return
False, since there is no number. This is because the incorrect code does not check for both a letter
and number. To fix this, the correct function can use two booleans to find both a letter and a number.
```python
def check_String(str):
    """
    Write a python function to check whether a string has at least one letter and one number.
    >>> check_String('thishasboth29')
    True
    >>> check_String('python')
    False
    >>> check_String('string')
    False
    """
    has_letter = False
    has_number = False
    for char in str:
        if char.isalpha():
            has_letter = True
        elif char.isdigit():
            has_number = True
        if has_letter and has_number:
            return True
    ...
    return False
```

-----

### Instruction: You are given an incorrect python function and an error message. Explain how to fix
the error, and then write an updated python function with the correct code.

### Incorrect Code: <Q,A>

### Error: <E>

### Response:
```

Figure 5: The prompt for generating a repair. For brevity, we only show a one-shot example. <Q,A> is replaced with the question and previous answer, while <E> is replaced with the corresponding error.

## E Evaluation on CodeLlama-7b

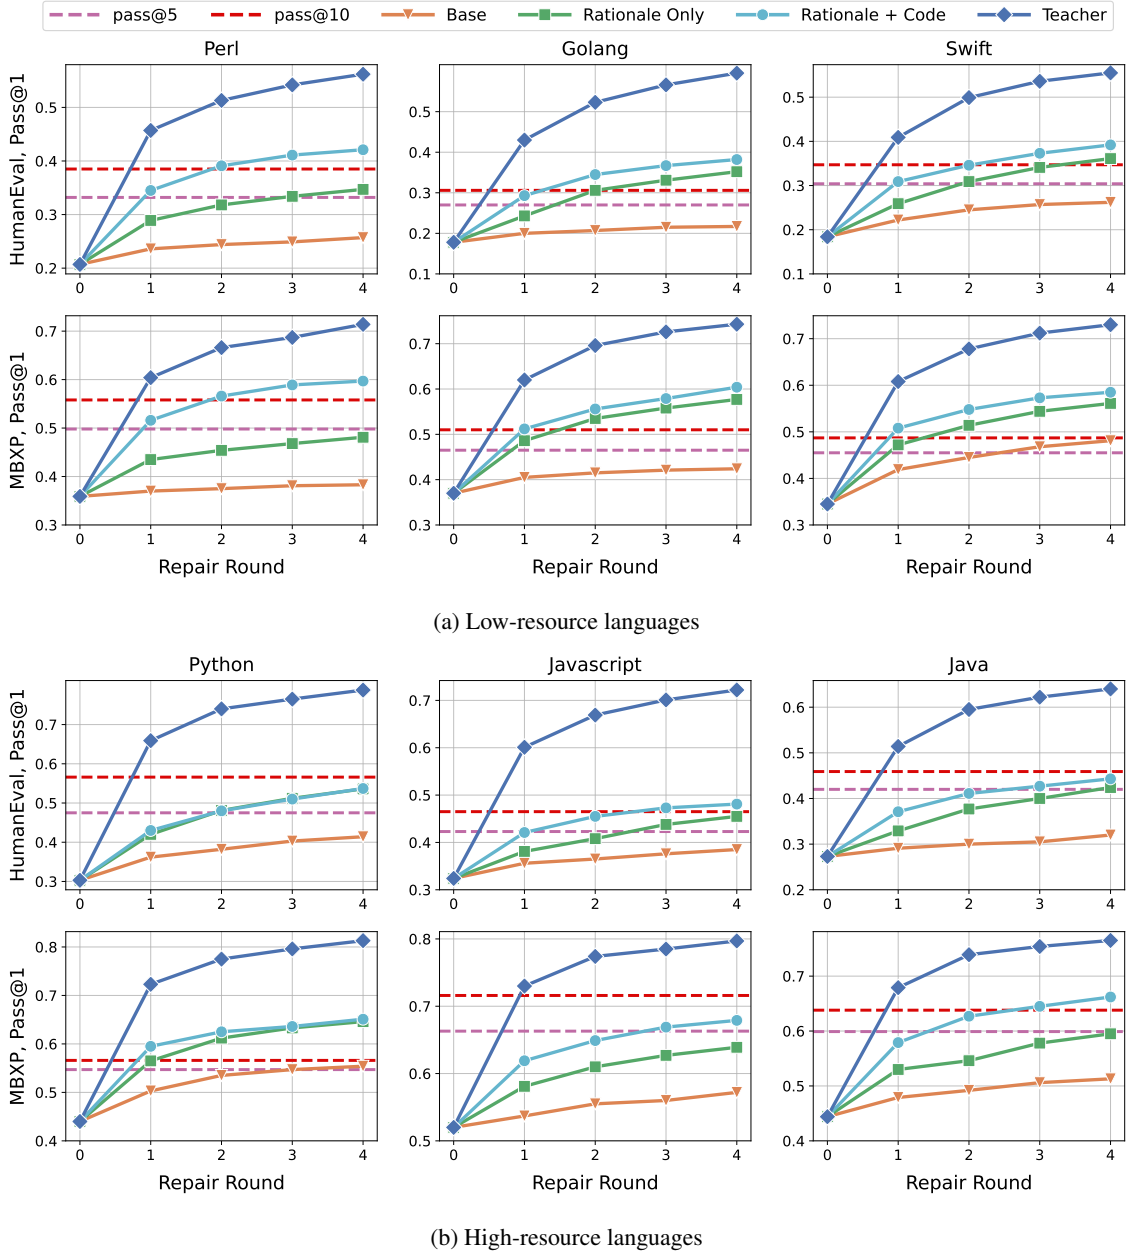


Figure 6: Mean pass@1 versus repair round for CodeLlama-7b. Round 0 denotes the initial generation. Rationale-plus-code distillation continues to outperform on all LRPLs. Compared to CodeLlama-7b-Instruct, transferring code completions sees better improvements in HRPLs. One possible explanation for this is that the base CodeLlama-7b has weaker responsiveness to our prompting, due to not being instruction-tuned.

## F Evaluation on Mistral-7b

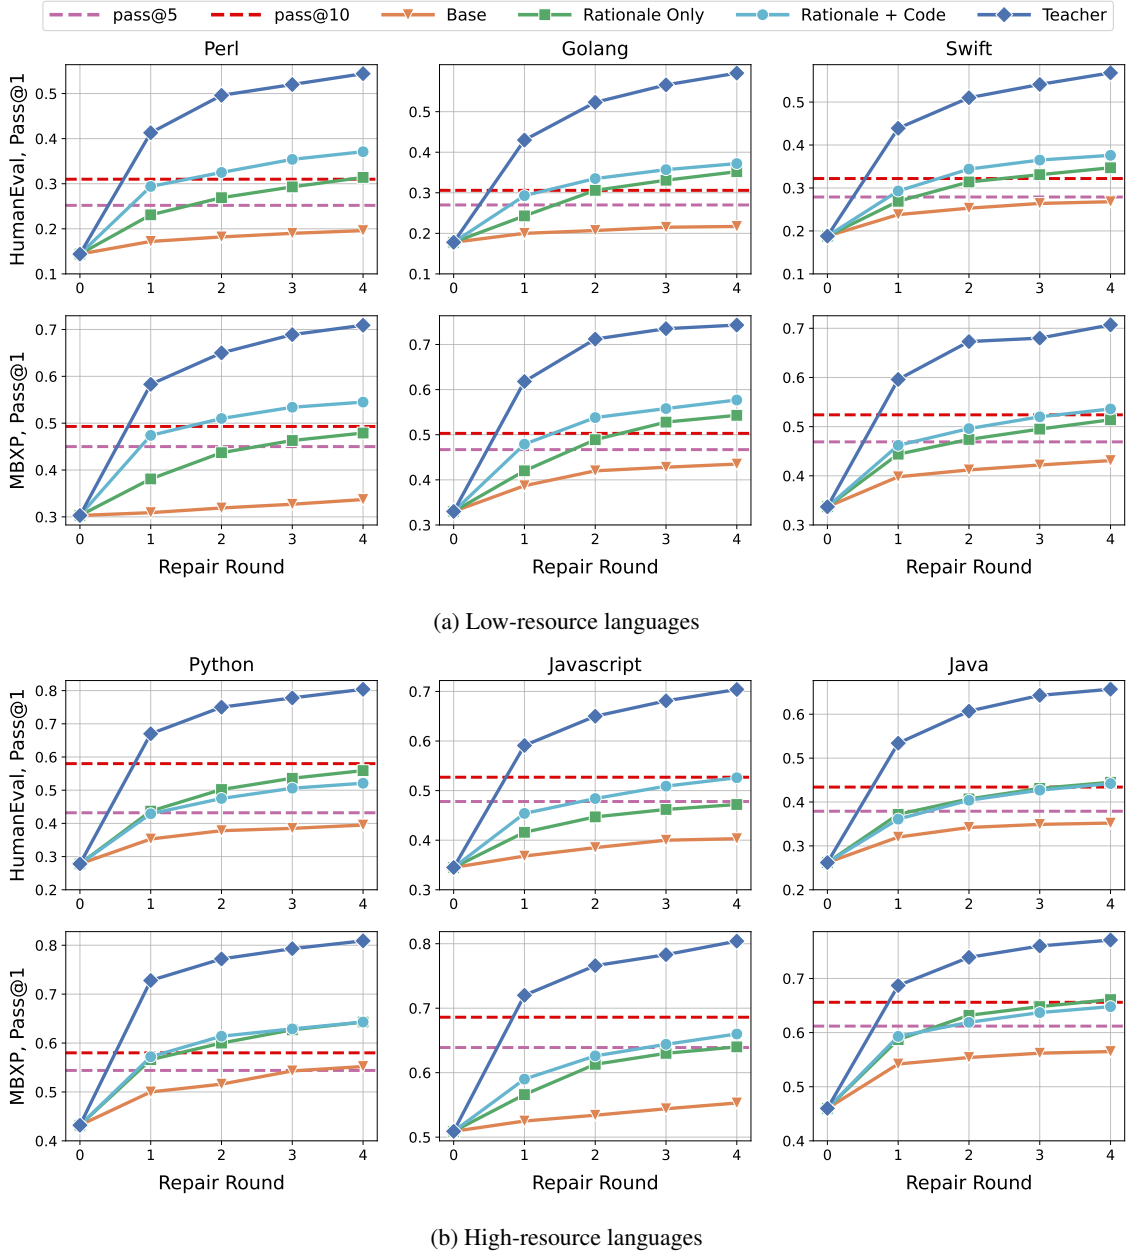


Figure 7: Mean pass@1 versus repair round for Mistral-7b. Round 0 denotes the initial generation. Rationale-plus-code distillation continues to outperforms on all LRPLs. Pass@1 improvements between LRPLs and HRPLs trend similarly to CodeLlama-7b-Instruct.

## G Pass@1 Comparison on CodeLlama-7b

| HumanEval Pass@1 |         |                |                     |
|------------------|---------|----------------|---------------------|
| Language         | Initial | Rationale Only | Rationale + Code    |
| Perl             | 0.207   | 0.347          | <b>0.421</b> ↑21.3% |
| Golang           | 0.178   | 0.352          | <b>0.372</b> ↑5.68% |
| Swift            | 0.184   | 0.361          | <b>0.392</b> ↑8.58% |
| Python           | 0.303   | 0.536          | <b>0.537</b> ↑1.86% |
| Javascript       | 0.324   | 0.455          | <b>0.481</b> ↑5.71% |
| Java             | 0.273   | 0.424          | <b>0.443</b> ↑4.48% |
| MBXP Pass@1      |         |                |                     |
| Perl             | 0.359   | 0.481          | <b>0.597</b> ↑24.1% |
| Golang           | 0.370   | 0.597          | <b>0.604</b> ↑1.17% |
| Swift            | 0.345   | 0.561          | <b>0.585</b> ↑4.27% |
| Python           | 0.440   | 0.646          | <b>0.651</b> ↑0.77% |
| Javascript       | 0.520   | 0.639          | <b>0.679</b> ↑6.25% |
| Java             | 0.444   | 0.595          | <b>0.662</b> ↑11.2% |

Table 5: Mean pass@1 of initial generations vs the two variants of distillation. Rationale-plus-code distillation sees consistent improvements on both LRPLs and HRPLs, likely because CodeLlama-7b has weaker instruction following, so fine-tuning improves its responsiveness.

## H Pass@1 Comparison on Mistral-7b

| HumanEval Pass@1 |         |                |                     |
|------------------|---------|----------------|---------------------|
| Language         | Initial | Rationale Only | Rationale + Code    |
| Perl             | 0.144   | 0.314          | <b>0.371</b> ↑18.1% |
| Golang           | 0.140   | 0.310          | <b>0.321</b> ↑3.55% |
| Swift            | 0.188   | 0.357          | <b>0.366</b> ↑2.52% |
| Python           | 0.278   | <b>0.559</b>   | 0.520 ↓6.97%        |
| Javascript       | 0.345   | 0.472          | <b>0.526</b> ↑11.4% |
| Java             | 0.262   | <b>0.445</b>   | 0.442 ↓0.67%        |
| MBXP Pass@1      |         |                |                     |
| Perl             | 0.303   | 0.479          | <b>0.545</b> ↑13.7% |
| Golang           | 0.330   | 0.543          | <b>0.576</b> ↑6.08% |
| Swift            | 0.337   | 0.514          | <b>0.536</b> ↑4.28% |
| Python           | 0.432   | <b>0.643</b>   | <b>0.643</b> ↓0.00% |
| Javascript       | 0.509   | 0.640          | <b>0.660</b> ↑3.12% |
| Java             | 0.460   | <b>0.661</b>   | 0.648 ↓2.73%        |

Table 6: Mean pass@1 of initial generations vs the two variants of distillation. A similar pattern as CodeLlama-7b-Instruct is observed, where rationale-plus-code distillation outperforms on LRPLs, but struggles to make consistent improvements on HRPLs.

## I GPT-4 Judgement Prompt

```

### Instruction: You are given an incorrect python function, an error message, and a
rationale to fix the error. Classify if the rationale is 'Good' or 'Bad'. If the rationale
provides enough detail to fix the code, output 'Good'. Otherwise, output 'Bad'.
### Incorrect Code: <Q,A>
### Error: <E>
### Rationale: <R>
### Response:

```

Figure 8: The prompt given to GPT-4 to judge rationale sufficiency. <Q,A> is replaced with the question and previous answer, <E> is replaced with the corresponding error, and <R> is replaced with the repair model’s rationale.



## J Syntax Errors for CodeLlama-7b

| HumanEval Average Syntax Errors |                |             |                |                   |                   |
|---------------------------------|----------------|-------------|----------------|-------------------|-------------------|
| Language                        | Initial Errors | Base Repair | Rationale Only | Rationale + Code  | Teacher Repair    |
| Perl                            | 21.2           | 21.0 ↓0.2   | 20.6 ↓0.6      | 12.8 ↓8.4         | <b>7.2</b> ↓14.0  |
| Golang                          | 39.1           | 72.8 ↑33.7  | 36.6 ↓2.5      | 30.8 ↓8.3         | <b>22.2</b> ↓16.9 |
| Swift                           | 78.1           | 57.2 ↓20.9  | 47.0 ↓31.1     | 48.4 ↓29.7        | <b>40.4</b> ↓37.7 |
| Python                          | 17.1           | 22.4 ↑5.3   | 23.5 ↑6.4      | 12.5 ↓4.6         | <b>7.8</b> ↓9.3   |
| Javascript                      | 10.6           | 10.0 ↓0.6   | 28.6 ↑18.0     | 13.2 ↑2.6         | <b>5.4</b> ↓5.2   |
| Java                            | 44.7           | 55.0 ↑10.3  | 44.6 ↓0.1      | 42.4 ↓2.3         | <b>20.4</b> ↓24.3 |
| MBXP Average Syntax Errors      |                |             |                |                   |                   |
| Perl                            | 16.4           | 15.0 ↓1.4   | 15.7 ↓0.7      | 6.2 ↓10.2         | <b>3.5</b> ↓12.9  |
| Golang                          | 30.7           | 48.2 ↑17.5  | 18.2 ↓12.5     | 15.2 ↓15.2        | <b>12.6</b> ↓18.1 |
| Swift                           | 62.4           | 37.6 ↓24.8  | 22.6 ↓39.8     | <b>19.0</b> ↓43.4 | 19.4 ↓43.0        |
| Python                          | 2.3            | 7.4 ↑5.1    | 5.7 ↑3.4       | 2.5 ↑0.2          | <b>1.3</b> ↓1.0   |
| Javascript                      | 7.1            | 5.0 ↓2.1    | 14.6 ↑7.5      | 7.0 ↓0.1          | <b>2.0</b> ↓5.1   |
| Java                            | 31.4           | 33.2 ↑1.8   | 24.0 ↓7.4      | 20.6 ↓10.8        | <b>9.6</b> ↓21.8  |

Table 7: Average number of syntax errors after code repair for CodeLlama-7b. We also include a column containing results from the GPT-3.5-Turbo teacher.

## K Syntax Errors for Mistral-7b

| HumanEval Average Syntax Errors |                |                 |                |                  |                   |
|---------------------------------|----------------|-----------------|----------------|------------------|-------------------|
| Language                        | Initial Errors | Base Repair     | Rationale Only | Rationale + Code | Teacher Repair    |
| Perl                            | 26.4           | 30.4 ↑4.0       | 24.0 ↓2.4      | 11.0 ↓15.4       | <b>9.4</b> ↓17.0  |
| Golang                          | 55.9           | 72.4 ↑16.5      | 48.2 ↓7.7      | 31.0 ↓24.9       | <b>25.2</b> ↓30.7 |
| Swift                           | 62.0           | 60.0 ↓2.0       | 54.0 ↓8.0      | 55.4 ↓6.6        | <b>39.8</b> ↓22.2 |
| Python                          | 14.5           | 17.2 ↑2.7       | 12.4 ↓2.1      | 13.5 ↓1.0        | <b>8.0</b> ↓6.5   |
| Javascript                      | 6.7            | <b>7.4</b> ↑0.7 | 16.6 ↑9.9      | 7.8 ↑1.1         | 7.8 ↑1.1          |
| Java                            | 41.4           | 42.2 ↑0.8       | 36.2 ↓5.2      | 31.4 ↓10.0       | <b>19.2</b> ↓22.2 |
| MBXP Average Syntax Errors      |                |                 |                |                  |                   |
| Perl                            | 26.3           | 25.2 ↓1.1       | 24.0 ↓2.3      | 6.0 ↓20.3        | <b>4.2</b> ↓22.1  |
| Golang                          | 43.6           | 40.0 ↓3.6       | 27.2 ↓16.4     | 13.8 ↓29.8       | <b>13.2</b> ↓30.4 |
| Swift                           | 49.3           | 36.6 ↓12.7      | 32.6 ↓16.7     | 25.6 ↓23.7       | <b>21.6</b> ↓27.7 |
| Python                          | 0.9            | 3.6 ↑2.7        | 4.2 ↑3.3       | 3.2 ↑2.3         | <b>2.2</b> ↑1.3   |
| Javascript                      | 7.8            | 7.6 ↓0.2        | 10.4 ↑2.6      | 7.4 ↓0.4         | <b>3.6</b> ↓4.2   |
| Java                            | 29.9           | 15.0 ↓14.9      | 11.2 ↓18.7     | 14.0 ↓15.9       | <b>7.2</b> ↓22.7  |

Table 8: Average number of syntax errors after code repair for Mistral-7b. We also include a column containing results from the GPT-3.5-Turbo teacher.

## L Distillation vs Baselines Examples

Our work hypothesizes that base models struggle to convert suggestions from the rationale into specific code corrections. We provide examples of this for each language. In particular, we highlight the differences between models distilled only rationales vs models distilled both rationales and code completions. In each example, the rationale-only model is provided a stellar rationale but generates incorrect code. Meanwhile, the rationale-plus-code model produces a weaker rationale but generates correct code. These examples support that the correlation between rationale quality and code correctness is weaker than previously perceived, and that distilling both rationales and code teaches models to respond better to feedback.

```
### Instruction: You are given an incorrect perl function and an error message. Explain how to fix the error, and then write an updated perl function with the correct code.

### Incorrect Code:
# You are an expert Perl programmer, and here is your task.
# # remove_vowels is a function that takes string and returns string without vowels.
# >>> remove_vowels('')
# ''
# >>> remove_vowels("abcdef\nghijklm")
# 'bcdf\nghjklm'
# >>> remove_vowels('abcdef')
# 'bcdf'
# >>> remove_vowels('aaaaa')
# ''
# >>> remove_vowels('aaBAA')
# 'B'
# >>> remove_vowels('zbcd')
# 'zbcd'
#
sub remove_vowels
{
    my ($text) = @_;
    $text =~ s/[aeiou]//g;
    return $text;
}


### Error: Exception -- test case 4 did not pass.
at temp.pl line 57.

### Response:
```

Repair Prompt


```
The incorrect function is using the wrong regex pattern to remove vowels. The correct function should use the pattern "[aeiou]".

sub remove_vowels
{
    my ($text) = @_;
    $text =~ s/[aeiou]//g;
    return $text;
}
```

Base 

```
To fix the error, we need to modify the regular expression pattern used in the substitution to be case-insensitive. Currently, it only matches lowercase vowels, so we need to make it match both lowercase and uppercase vowels. This can be achieved by adding the 'i' flag at the end of the pattern.

sub remove_vowels
{
    my ($text) = @_;
    $text =~ s/[aeiou]/i//g;
    return $text;
}
```

Rationale Only 

```
The issue with the incorrect code is that it is not handling the case where the input string contains uppercase vowels. To fix this, we need to modify the regular expression to include uppercase vowels as well.

sub remove_vowels
{
    my ($text) = @_;
    $text =~ s/[aeiouAEIOU]//g;
    return $text;
}
```


Rationale + Code 

Figure 9: Perl example from HumanEval question 51. The initial code is wrong because it does not remove uppercase vowels. From the base model, we see a weak rationale that fails to diagnose the uppercase issue. From the rationale-only model, we see a stellar rationale that proposes using the "i" regex flag for case insensitivity. However, the generated code incorrectly modifies the regex to `"/[aeiou]/i//g"` instead of `"/[aeiou]//gi"`, displaying a lack of knowledge. From the rationale-plus-code model, we see a good rationale that suggests adding uppercase letters to the regex, followed by correct code modifications.

### Instruction: You are given an incorrect golang function and an error message. Explain how to fix the error, and then write an updated golang function with the correct code.

### Incorrect Code:

```
// You're an expert Golang programmer
// Return the largest prime factor of n.
// Assume n > 1 and is not a prime.
// >>> largest_prime_factor(13195)
// 29
// >>> largest_prime_factor(2048)
// 2
//
func largest_prime_factor (n int) int {

    for n % 2 == 0 {
        n = n / 2
    }
    for n % 3 == 0 {
        n = n / 3
    }
    for n % 5 == 0 {
        n = n / 5
    }
    for n % 7 == 0 {
        n = n / 7
    }
    return n
}
```


### Error: panic: Exception --- test case 0 failed to pass

### Response:

Repair Prompt

The incorrect function is wrong because it does not check for the prime factors of n. To fix this, the correct function should check for the prime factors of n, and return the largest one.


```
func largest_prime_factor (n int) int {
    for i := 2; i < n; i++ {
        if n % i == 0 {
            return i
        }
    }
    return n
}
```

Base 

The incorrect code is not correctly finding the prime factors of "n". Instead of hardcoding prime factors, you need to generalize the loop to check for all possible prime factors.

One way to achieve this is to iterate through all numbers starting from 2 up to the square root of "n" and check if each number is a factor of "n" and a prime number. This way, we can ensure that we are finding the largest prime factor of the given number "n" accurately.

```
func largest_prime_factor (n int) int {
    if n == 1 {
        return 1
    }
    if n % 2 == 0 {
        return 2
    }
    for i := 3; i <= int(n ** 0.5); i += 2 {
        for n % i == 0 {
            n /= i
        }
    }
    return n
}
```

Rationale Only 

Test case 0 failed because the function is not correctly finding the largest prime factor of n. The function is currently returning n instead of the actual largest prime factor. To fix this, we need to update the function to iterate over all prime factors of n and return the largest one. Here is the corrected code:

```
func largest_prime_factor(n int) int {
    if n <= 1 {
        return n
    }
    for i := 2; i*i <= n; i++ {
        if n%i == 0 {
            n = n / i
            i--
        }
    }
    return n
}
```


Rationale + Code 

Figure 10: Golang example from HumanEval question 59. The initial code is wrong because it hardcodes prime factors. From the base model, we see a weak rationale that is unable to diagnose the hardcoding issue. From the rationale-only model, we see a stellar rationale that suggests iterating over all numbers from 2 to  $\sqrt{n}$ . However, the generated code results in a compilation error due to using "n \*\* 0.5" to obtain the square root (which does not work in golang), displaying a lack of knowledge. From the rationale-plus-code model, we see a good rationale that provides a plan of iterating over prime factors, followed by correct code modifications.

### Instruction: You are given an incorrect swift function and an error message. Explain how to fix the error, and then write an updated swift function with the correct code.

### Incorrect Code:

```
/**
 * You are an expert Swift programmer, and here is your task.
 * Return list of prime factors of given integer in the order
 * from smallest to largest.
 * Each of the factors should be listed number of times
 * corresponding to how many times it appears in factorization.
 * Input number should be equal to the product of all factors
 * >>> factorize(8)
 * [2, 2, 2]
 * >>> factorize(25)
 * [5, 5]
 * >>> factorize(70)
 * [2, 5, 7]
 *
 */
func factorize(n : Int) -> [Int] {

    var factors = [Int]()
    var i = 2
    while i * i <= n {
        while n % i == 0 {
            factors.append(i)
            n /= i
        }
        i += 1
    }
    if n > 1 {
        factors.append(n)
    }
    return factors
}
```


### Error: temp.swift:21:15: error: left side of mutating operator isn't mutable: 'n' is a 'let' constant  
 n /= i  
 ~ ^

### Response:

Repair Prompt


The incorrect function is trying to modify the input number, which is a constant. To fix this, the correct function should return a new array of prime factors instead of modifying the input number.

```
func factorize(n : Int) -> [Int] {
    var factors = [Int]()
    var i = 2
    while i * i <= n {
        while n % i == 0 {
            factors.append(i)
            n /= i
        }
        i += 1
    }
    if n > 1 {
        factors.append(n)
    }
    return factors
}
```

Base 

To fix the error in the given Swift function, you need to change the parameter declaration of 'n' from 'let' to 'var'. This change will make 'n' mutable so that you can reassign it when performing the division operation inside the while loop. By making this change, you will be able to update the value of 'n' during the factorization process.

```
func factorize(n : Int) -> [Int] {
    var factors = [Int]()
    var i = 2
    while i * i <= n {
        while n % i == 0 {
            factors.append(i)
            n /= i
        }
        i += 1
    }
    if n > 1 {
        factors.append(n)
    }
    return factors
}
```

Rationale Only 

The error is due to trying to modify a constant 'n' in the function. To fix this, we need to create a mutable copy of 'n' before modifying it. Here is the corrected function:

```
func factorize(n: Int) -> [Int] {
    var n = n
    var factors = [Int]()
    var i = 2
    while i * i <= n {
        while n % i == 0 {
            factors.append(i)
            n /= i
        }
        i += 1
    }
    if n > 1 {
        factors.append(n)
    }
    return factors
}
```


Rationale + Code 

Figure 11: Swift example from HumanEval question 25. The initial code is wrong because it tries to modify the immutable input variable `n`. For swift, input parameters are immutable by default. From the base model, we see a weak rationale that does not specify how to fix the error. From the rationale-only model, we see a stellar rationale that explicitly mentions using the "var" keyword. However, the generated code remains the same because the base model is unable to make the necessary code modifications, displaying a lack of responsiveness. From the rationale-plus-code model, we see a good rationale that proposes making a mutable copy, followed by correct code modifications through the additional line "`var n = n`".



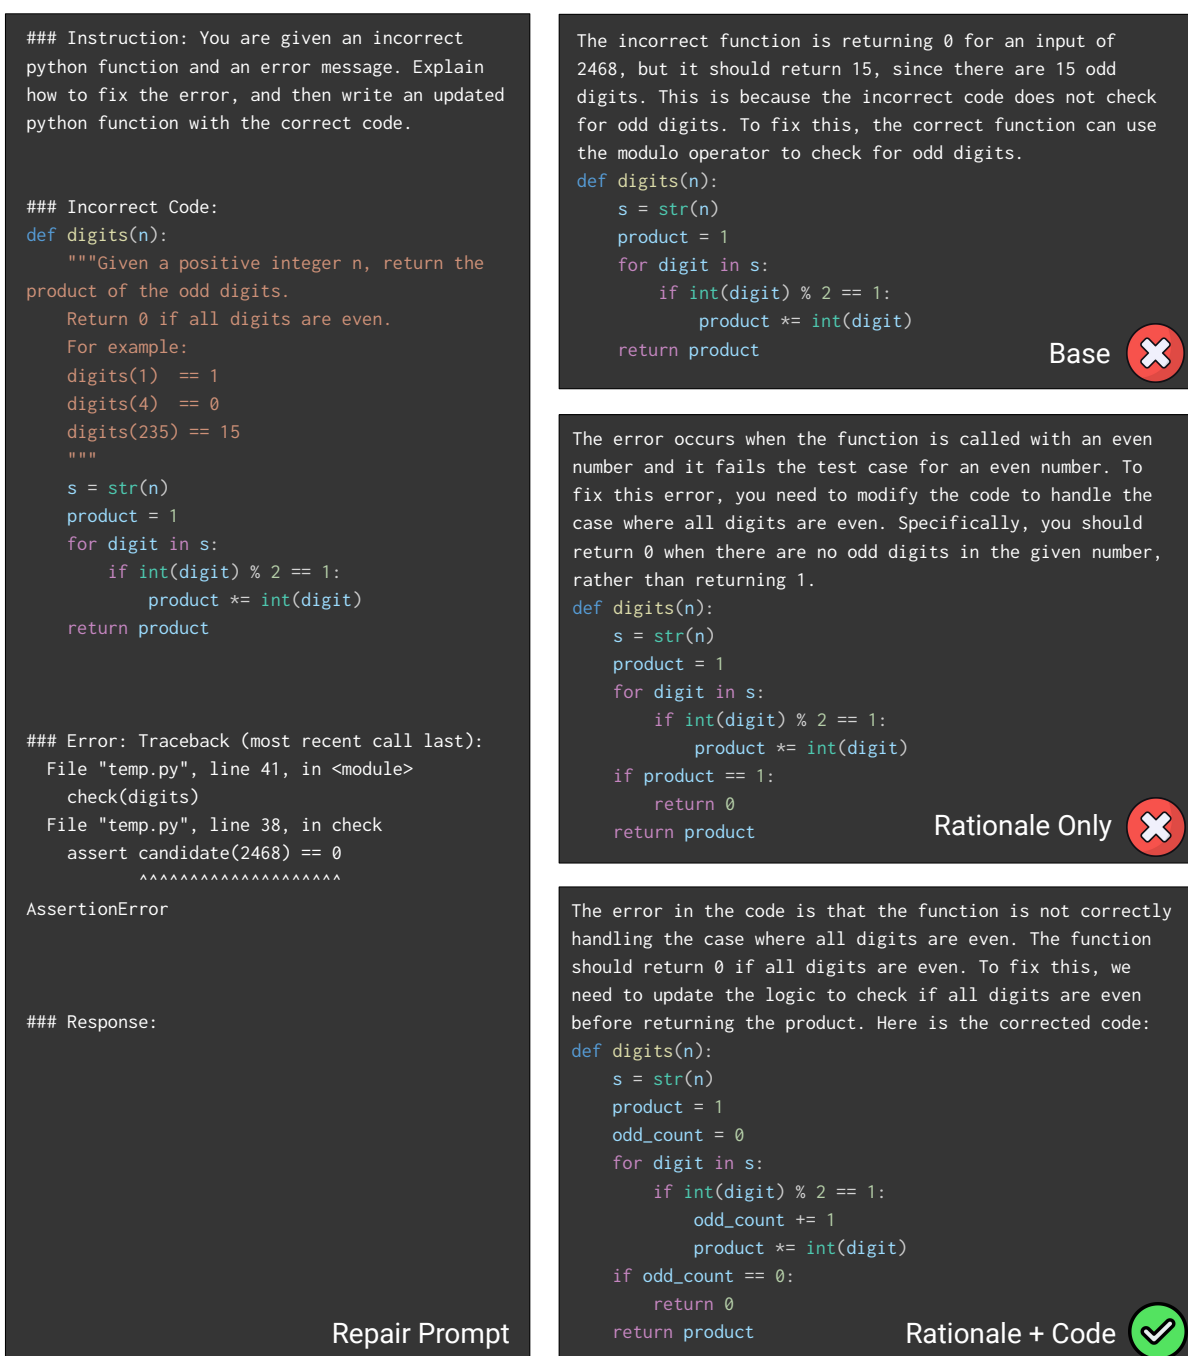


Figure 12: Python example from HumanEval question 131. The initial code is wrong because it returns 1 if there are no odd digits, instead of 0. From the base model, we see a weak rationale that states wrong information. From the rationale-only model, we see a stellar rationale that addresses the edge case and suggests returning 0 if there are no odd digits. However, the generated code implements this incorrectly, displaying weak responsiveness. From the rationale-plus-code model, we see a good rationale that suggests checking if all digits are even, followed by correct code modifications via the additional variable "odd\_count".

### Instruction: You are given an incorrect javascript function and an error message. Explain how to fix the error, and then write an updated javascript function with the correct code.

### Incorrect Code:

```
/**
 * Evaluate whether the given number n
 can be written as the sum of exactly 4
 positive even numbers
 * Example
 * is_equal_to_sum_even(4) == False
 * is_equal_to_sum_even(6) == False
 * is_equal_to_sum_even(8) == True
 */
function isEqualToSumEven(n) {
  for (let i = 2; i <= n; i += 2) {
    if (n - i === 2) {
      return true;
    }
  }
  return false;
}
```

### Error:


```
temp.js:29
    throw 'Error at 1th assert
statement. Value = ' +
JSON.stringify(x0)
    ^
Error at 1th assert statement. Value =
true
(Use 'node --trace-uncaught ...' to
show where the exception was thrown)
```

### Response:

Repair Prompt


The incorrect code is using a for loop to iterate through the numbers from 2 to n, and then checking if the number is equal to 2. However, the function description wants to check if the number can be written as the sum of exactly 4 positive even numbers. To fix this, the correct code should use a different approach to check if the number can be written as the sum of 4 positive even numbers.

```
function isEqualToSumEven(n) {
  let arr = [2, 4, 6, 8];
  let sum = 0;
  for (let i = 0; i < arr.length; i++) {
    if (n - arr[i] === 0) {
      return true;
    }
    sum += arr[i];
  }
  return false;
}
```

Base 

The error message is indicating that the code provided is not returning the correct output for the test case where 'n' is 8. Currently, the code checks if 'n - i' is equal to 2, which doesn't cover all possible combinations. You should update the logic in the function to iterate over possible combinations of 4 positive even numbers and check if their sum equals 'n'. You can do this by considering all possible even numbers less than or equal to 'n/4' and checking if there exists a combination of 4 even numbers that sum up to 'n'. By making these adjustments, you can ensure that the function correctly checks if the given number 'n' can be represented as the sum of exactly 4 positive even numbers.

```
function isEqualToSumEven(n) {
  let count = 0;
  for (let i = 2; i <= n; i += 2) {
    if (count === 4) return true;
    for (let j = 2; j <= n; j += 2) {
      if (count === 4) return true;
      for (let k = 2; k <= n; k += 2) {
        if (count === 4) return true;
        for (let l = 2; l <= n; l += 2) {
          if (count === 4) return true;
          if (i + j + k + l === n) {
            count++;
          }
        }
      }
    }
  }
  return false;
}
```

Rationale Only 

The error indicates that the function is failing the first test case. The issue lies in the logic of the function, which is not correctly checking if the given number can be written as the sum of exactly 4 positive even numbers. To fix this, we need to adjust the condition in the loop to ensure that the sum of 4 even numbers is equal to the given number. Here is the corrected code for the function:

```
function isEqualToSumEven(n) {
  for (let i = 2; i <= n; i += 2) {
    for (let j = i + 2; j <= n; j += 2) {
      for (let k = j + 2; k <= n; k += 2) {
        if (k + 2 === n) {
          return true;
        }
      }
    }
  }
  return false;
}
```


Rationale + Code 

Figure 13: Javascript example from HumanEval question 138. The initial code is wrong because it checks if  $n$  is a sum of 2 even numbers, instead of 4. From the base model, we see a decent rationale that lacks a highly detailed plan. From the rationale-only model, we see a stellar rationale suggesting a highly detailed plan. However, the generated code implements this incorrectly by misusing the "count" variable, displaying weak responsiveness. From the rationale-plus-code model, we see a decent rationale that lacks a highly detailed plan, but correct code modifications anyways, displaying strong responsiveness.

### Instruction: You are given an incorrect java function and an error message. Explain how to fix the error, and then write an updated java function with the correct code.

### Incorrect Code:

```

/**
 * This function takes a list l and returns a list l' such
 * that
 * l' is identical to l in the odd indicies, while its
 * values at the even indicies are equal
 * to the values of the even indicies of l, but sorted.
 * >>> sort_even([1, 2, 3])
 * [1, 2, 3]
 * >>> sort_even([5, 6, 3, 4])
 * [3, 6, 5, 4]
 */
public static List<Integer> sortEven(List<Integer> l) {
    List<Integer> l2 = new ArrayList<Integer>();
    for (int i = 0; i < l.size(); i++) {
        if (i % 2 == 0) {
            l2.add(l.get(i));
        }
    }
    Collections.sort(l2);
    for (int i = 0; i < l.size(); i++) {
        if (i % 2 != 0) {
            l2.add(l.get(i));
        }
    }
    return l2;
}

```

### Error:

```

Exception in thread "main" java.lang.Exception: Exception
-- test case 0 did not pass. x0 = [1, 3, 2]
   at Main.main(Main.java:24)

```


### Response:

The incorrect function is trying to sort the even indices of the list, but it is not correctly identifying the even indices. To fix this, the correct function should use the modulo operator to check if the index is even.

```

public static List<Integer> sortEven(List<Integer> l) {
    List<Integer> l2 = new ArrayList<Integer>();
    for (int i = 0; i < l.size(); i++) {
        if (i % 2 == 0) {
            l2.add(l.get(i));
        }
    }
    Collections.sort(l2);
    for (int i = 0; i < l.size(); i++) {
        if (i % 2 != 0) {
            l2.add(l.get(i));
        }
    }
    return l2;
}

```


Base 

The error message indicates that the test case [1, 3, 2] did not pass. To fix this issue in the code, the function sortEven should correctly handle the even and odd indices when constructing the new list l2. Make sure that the values at even indices of l are added to l2 in sorted order, while the values at odd indices are retained from the original list. This adjustment will ensure that the function behaves as intended for sorting even indices.

```

public static List<Integer> sortEven(List<Integer> l) {
    List<Integer> l2 = new ArrayList<Integer>();
    for (int i = 0; i < l.size(); i++) {
        if (i % 2 == 0) {
            l2.add(l.get(i));
        } else {
            l2.add(l.get(i));
        }
    }
    Collections.sort(l2);
    return l2;
}

```


Rationale Only 

The error indicates that the test case is failing because the output is not correct. The issue in the incorrect code is that the sorting is happening incorrectly. To fix this, we need to sort the even indices of the list separately and then add them back to the original list in the correct order.

```

public static List<Integer> sortEven(List<Integer> l) {
    List<Integer> l2 = new ArrayList<Integer>();
    for (int i = 0; i < l.size(); i++) {
        if (i % 2 == 0) {
            l2.add(l.get(i));
        }
    }
    Collections.sort(l2);
    for (int i = 0; i < l.size(); i++) {
        if (i % 2 == 0) {
            l.set(i, l2.get(i / 2));
        }
    }
    return l;
}

```

Rationale + Code 

Repair Prompt

Figure 14: Java example from HumanEval question 37. The initial code is wrong because it first appends even indices and then appends odd indices, instead of interleaving them. From the base model, we see a weak rationale and no code modifications. From the rationale-only model, we see a stellar rationale suggesting to interleave odd/even indices. However, the generated code incorrectly implements the sorting of even indices by sorting the entire list at the end, displaying weak responsiveness. From the rationale-plus-code model, we see a good rationale suggesting to interleave odd/even indices, followed by correct code modifications.