

# ECCO: Can We Improve Model-Generated Code Efficiency Without Sacrificing Functional Correctness?

Anonymous ARR submission

## Abstract

Although large language models (LLMs) have been largely successful in generating functionally correct programs, conditioning models to produce efficient solutions while ensuring correctness remains a challenge. Further, unreliability in benchmarking code efficiency is a hurdle across varying hardware specifications for popular interpreted languages such as Python. In this paper, we present ECCO, a reproducible benchmark for evaluating program efficiency via two paradigms: natural language (NL) based code generation and history-based code editing. On ECCO, we adapt and thoroughly investigate the three most promising existing LLM-based approaches: in-context learning, iterative refinement with execution or NL feedback, and fine-tuning conditioned on execution and editing history. While most methods degrade functional correctness and moderately increase program efficiency, we find that adding execution information often helps maintain functional correctness, and NL feedback enhances more on efficiency. We release our benchmark to support future work on LLM-based generation of efficient code.

## 1 Introduction

The ability to write efficient code is a cornerstone of software development (Li et al., 2022). While large language models (LLMs) have shown remarkable progress in generating functionally correct code (Roziere et al., 2023; Guo et al., 2024), the ability to generate solutions that are both correct and efficient remains elusive (Shypula et al., 2021, 2024).

Current methods for optimizing program efficiency improve performance measured by execution time. However, this apparent success often comes at the cost of severely decreasing the functional correctness (Shypula et al., 2024). An example of this issue is illustrated in Figure 1: When optimizing the program on the left, models sometimes perform *spurious optimizations* that, although they

reduce the runtime, make the program no longer functionally correct so that it fails all test cases. On the other hand, a correct optimization (bottom right) — that improves efficiency while maintaining functional correctness — is often harder to achieve for current LMs. This *spurious optimization* is certainly undesirable in practice, and can even increase debugging time for software developers (Li et al., 2022; Cummins et al., 2023a). To achieve the goal of real and robust program optimization, we ask: *Can LMs improve program efficiency without sacrificing functional correctness?*

In this work, we curate an efficiency-oriented programming benchmark **ECCO**, short for Ensuring Correctness in Code Optimizations, which enables program evaluation in three aspects: execution correctness, runtime efficiency, and memory efficiency. ECCO supports two optimization paradigms: (i) history-based code editing: based on a previous version of the program, test if an LM can further optimize the code while maintaining its correctness, and (ii) NL-based code generation: test the efficiency of a program generated by an LM given a programming problem described in NL. We collect over 50k Python solution pairs, where each pair consists of a slower and faster version written by the same user for a given problem. These pairs span 1.3k competitive programming problems from IBM CodeNet (Puri et al., 2021), with an average of 3.1 public and 20.0 private test cases to support reliable execution-based evaluations of correctness and efficiency.

Further, to perform reliable and reproducible executions, we introduce an evaluation setup using a cloud-hosted code execution engine, JUDGE0 (Došilović and Mekterović, 2020), which produces stable execution output on correctness, runtime, and memory usage, thanks to its agnostic nature to local hardware specifications. It supports up to 66 programming languages (PLs), allowing future work to extend to other languages.

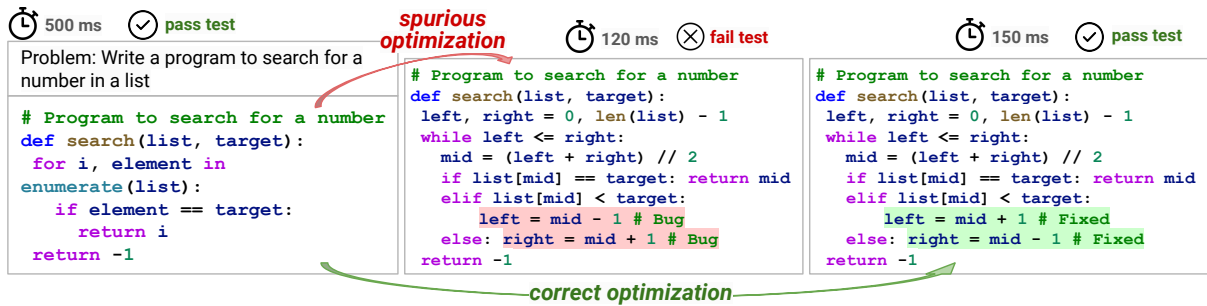


Figure 1: *Correctness-preserving* versus *spurious optimization* when optimizing a linear search algorithm with binary search on a sorted list. Spurious optimization can reduce runtime, but add errors that cause the program to be incorrect. In contrast, a true optimization reduces runtime and remains correct, as we emphasized in ECCO.

To explore various correctness-preserving program optimization methods, we evaluate three classes of methods on ECCO— in-context learning, iterative refinement, and fine-tuning, across a suite of open-source language models pre-trained on code — which previous works have used to attempt to improve efficiency, while however overlooking their effects on correctness. We find that execution information and fine-tuning help maintain functional correctness, and NL-involved prompting often yields higher efficiency improvements. However, we broadly reconfirm findings that no existing methods can improve time/space efficiency without sacrificing functional correctness. We hope ECCO can serve as a solid testbed for program optimization, and call for more efforts in advancing correctness-preserving program optimizations.

## 2 Related Work

**Benchmarks for Code Efficiency** Some works have proposed benchmarks for optimizing program assembly code (Bunel et al., 2016; Shypula et al., 2021; Shi and Zhang, 2020; Cummins et al., 2023b). More recently, Shypula et al. (2024) target C++ program speedups, and Huang et al. (2024) evaluate the efficiency of Python solutions for LeetCode coding interview problems (Niu et al., 2024). Although most efforts on LLM-based code generation focus on evaluating functional correctness (e.g., Chen et al. 2021), some works evaluate code efficiency (Moudgalya et al., 2023; Sikka et al., 2020; Jeon et al., 2023; Baik et al., 2024) by classifying the time complexity of programs. However, these works are limited in their single-reference evaluation paradigm, assembly language support, or by the limited problem space on LeetCode (Coignon et al., 2024). Our work supports reliable evaluation across arbitrary coding problems and the widely-used Python language.

**Evaluating Program Efficiency** It is challenging to robustly evaluate program efficiency, due to varying hardware platforms and setups. Previous works have evaluated the efficiency of code by executing code in a local software environment (Singhal et al., 2024; Huang et al., 2024) or using Docker containerized environments on local hardware (Khan et al., 2023), but this can result in varying runtime and memory usage across hardware, thus causing unreliable and irreproducible efficiency evaluations. An alternative approach is to use an architecture simulator (Shypula et al., 2024) which ensures the execution of each program is exactly simulated at the hardware level on the x86 architecture, but is limited to compiled languages such as C++. For more popular interpreted languages such as Python and Java (JVM), some use LeetCode’s execution engine (Niu et al., 2024; Coignon et al., 2024), but with a restricted space of testable problems. In our work, we propose an evaluation setup using an accessible cloud computing instance that ensures consistent virtual hardware and reliable benchmarking.

**Program Optimization Approaches** To start, some works explore in-context learning to optimize program efficiency (Shypula et al., 2024; Huang et al., 2024). Beyond vanilla prompting, iterative prompting methods (Madaan et al., 2024; Shinn et al., 2024; Ridnik et al., 2024) have been explored to improve specific aspects of generation, by incorporating feedback or reflection from an LM or external modules. Meanwhile, Shypula et al. (2024) propose fine-tuning with self-play synthetic data, with retrieved in-context examples via a nearest neighbor search. However, none of these methods have been rigorously studied in the correctness-preserving optimization setup. We fill in this gap and provide systematic studies of all methods.

### 3 The ECCO Benchmark

In this section, we first introduce our evaluation platform (§3.1), then describe the construction process of our ECCO benchmark (§3.2), and lastly, introduce our two task formulations with corresponding evaluation metrics (§3.3).

#### 3.1 Evaluation Platform

To reliably evaluate program efficiency in both runtime and memory usage, we need to first establish a robust and reproducible evaluation platform. However, evaluating program efficiency is challenging, as resource usage statistics vary greatly across hardware and setups (Singhal et al., 2024; Huang et al., 2024).

To ensure stable execution for interpreted languages such as Python, we propose to use a reproducible cloud compute instance that ensures the same virtual hardware, as illustrated in Figure 2. Specifically, we use an EC2 instance (detailed in §B), and execute the code within a code execution engine JUDGE0 (Došilović and Mekterović, 2020). Note that our recipe can easily extend to over 60 programming languages that the JUDGE0 engine supports. This is set up as a sandboxed Docker container within the instance, which thus ensures an isolated setup for secure and reproducible code execution. Our platform is similar to evaluating generated code on LeetCode execution console (Niu et al., 2024; Coignon et al., 2024), but applies to arbitrary coding problems and is not limited to questions available on LeetCode.<sup>1</sup>




Figure 2: Our evaluation platform using JUDGE0.

#### 3.2 Benchmark Construction

Our goal is to collect programming problems, each with an NL description, and multiple functionally-correct solutions at varied efficiency levels.

**Problem Selection** We collect programming problems from the IBM CodeNet dataset (Puri

<sup>1</sup>We make the Amazon Machine Image (AMI) for the setup available to enable reproducible benchmarks. We leave evaluations in other languages for future work.

et al., 2021), which contains competitive programming problems with NL descriptions, user program submissions, and other metadata, scraped from the AIZU and AtCoder online judging systems. CodeNet problems mostly require algorithmic techniques such as data structure optimization.

Specifically, we first convert CodeNet to  $\sim 187k$  (slow, fast) Python code pairs following Shypula et al. (2024), where each pair of programs has two solutions for the same coding problem. We converted all Python 2 solutions to Python 3 using lib2to3.<sup>2</sup> Lastly, we filter out spurious program pairs in which the ‘fast’ code was in fact slower when evaluated on our setup.<sup>3</sup>

Next, we split the pairs and group programs by their associated problem ID. We filter out all problems with less than two solutions to ensure that each NL problem description has multiple associated solutions, to enable program optimization based on code editing history. We then remove the programs that are repetitive or cannot successfully execute due to syntax errors or test case failures.

Our curated dataset was partitioned into three subsets: train, validation, and test, with each split consisting of a distinct set of problems. In the end, the process yields 1,387 unique problems in ECCO in total.

**Test Case Collection** To evaluate functional correctness, we require test cases. We collect (i) the original test cases for each problem from CodeNet, and (ii) the additional tests from the AlphaCode project (Li et al., 2022). Each test case contains the program inputs as well as expected outputs when executing canonical program solutions on these inputs. With these two sets of test cases, we simulate a realistic coding setting where one can refer to (i) as the public test cases for debugging or other accuracy-improving purposes,  $T_{public}$ , and (ii) as private test cases to conduct final execution-based evaluations on the programs,  $T_{private}$ .

#### 3.3 Task Formulation and Evaluation

We propose two formulations for the program optimization task, namely *NL-instructed generation* and *history-based program editing*. In this section, we introduce the data we use for each formulation, and our evaluations of program correctness, runtime, and memory usage.

<sup>2</sup><https://docs.python.org/3/library/2to3.html>

<sup>3</sup>Speed statistics reported in CodeNet may be inconsistent.

### 3.3.1 History-Based Program Editing

Our first paradigm follows previous work on program optimization (Shypula et al., 2024), where we facilitate a history-based editing paradigm. Concretely, we give a previous, presumably slow, version of the solution program,  $p_{in}$ . We then prompt LMs to edit the code to generate a more efficient version  $p_{out}$ , denoted as  $CodeLM(p_{in}) \rightarrow p_{out}$ , where  $p_{out}$  is expected to run faster than  $p_{in}$ .

#### Evaluating Speedup and Memory Reduction

Using the (slow, fast) program pairs remaining after post-processing in §3.2, we evaluate the relative speedup and memory reduction of the model-generated program against the input program on private test cases  $T_{private}$ .

We adopt the *speedup* metric introduced by Shypula et al. (2024), which is formulated as:

$$Speedup = \frac{\text{Runtime of } p_{in}}{\text{Runtime of } p_{out}} \quad (1)$$

Similarly, to evaluate improvement in memory usage, we introduce a *memory reduction* metric as:

$$Memory\ Reduction = \frac{\text{Memory of } p_{in}}{\text{Memory of } p_{out}} \quad (2)$$

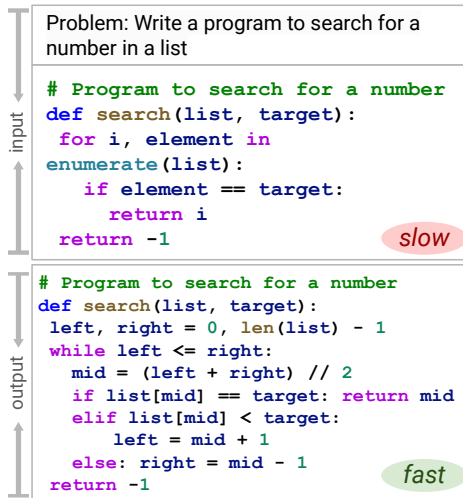


Figure 3: Illustration of history-based editing.

### 3.3.2 NL-Instructed Generation

In addition, we support the most common NL-to-code generation setup: given the NL description  $d$  of a problem, we ask the LM to generate the program solution  $p$ , as  $CodeLM(d) \rightarrow p$ . Our goal is for the code LM to generate an efficient and correct solution  $p$ . We execute  $p$  on the public test cases to evaluate its performance<sup>4</sup>.

<sup>4</sup>We evaluate the NL-instructed generation paradigm on public test cases due to resource limitations

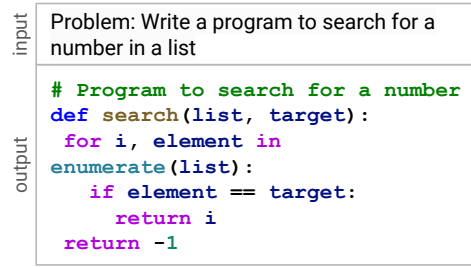


Figure 4: Illustration of NL-based generation.

**Solution Program Spectrum** To evaluate relative runtime and memory efficiency, we measure where a model-generated program lies on the spectrum of all user-submitted programs to that problem. We use the JUDGE0 evaluation platform (§3.1) to measure the runtime and memory usage.

**Evaluating Percentile over the Spectrum** We introduce *runtime and memory percentile* to measure the efficiency of the model-generated program over the solution spectrum for a given problem as:

$$Runtime\ \% = \frac{\# \text{ Slower user programs}}{\text{Total \# of user programs}} \quad (3)$$

$$Memory\ \% = \frac{\# \text{ Programs w/ more memory}}{\text{Total \# of user programs}} \quad (4)$$

### 3.3.3 Evaluating Functional Correctness

To measure if program correctness is preserved, a key metric is the functional correctness of model-generated programs. We adopt the pass@1 metric introduced by Chen et al. (2021), which samples one program from the model and measures whether the generated program passes all test cases.

### 3.4 ECCO Feature Analysis

After filtering the problem description dataset, we split the dataset into train, test, and validation sets for experiments. We ensure that no problem IDs overlap across these splits, to avoid data contamination. As shown by the detailed statistics of ECCO in Table 1, ECCO contains 1.3k problems and over 50k program pairs for code optimization evaluation.

Split	# Problems	# Pairs	# Avg. Test Cases	
			Public	Private
Train	1262	48386	3.14	20
Test	56	881	3.29	20
Val	69	2359	3.17	20

Table 1: ECCO dataset statistics.

301	<b>4 Efficiency-Improving Approaches</b>	Both correctness and efficiency can be informed via this feedback.	346
302	We explore various top-performing code generation approaches to improve program efficiency, while maintaining functional correctness, including in-context learning (§4.1), iterative refinement (§4.2), and fine-tuning (§4.3).		347
303		<b>NL+Exec Refine: NL Feedback on Interpreter Results</b>	348
304		To allow feedback both in the forms of NL and execution outputs, we ground the LM feedback on execution results, inspired by the Reflexion feedback paradigm (Shinn et al., 2024). Specifically, we first obtain the execution results as in <i>exec-refine</i> , then ask the LM to write NL feedback on the incorrect/inefficient parts in the code, and use this as additional input in the refinement turn.	349
305			350
306			351
307	<b>4.1 In-Context Learning</b>		352
308	We explore two mainstream prompting strategies: instruction prompting and few-shot learning.		353
309			354
310	<b>Instruction prompting</b> Many LMs perform better when incorporating instructions (Ouyang et al., 2022; Wei et al., 2022). We use two prompts: $I_{gen}$ for NL-based generation which instructs models to generate correct and efficient programs; and $I_{eff}$ for history-based editing which instructs models to optimize the input program. $I_{eff}$ is adapted from PIE (Shypula et al., 2024) and Self-Refine (Madaan et al., 2024). See §C for details.		355
311			356
312		<b>4.3 Fine-tuning</b>	357
313		We also explore three fine-tuning methods beyond prompting-alone approaches.	358
314			359
315		<b>Vanilla Fine-tuning</b> In this vanilla training setting, we leverage (NL, program) pairs and (slow program, fast program) pairs to train models independently for each paradigm. We format the data for both similarly to the in-context learning prompts (§4.1), and finetune on a causal language modelling task on the formatted data for each of the two paradigms independently.	360
316			361
317			362
318			363
319	<b>Few-Shot Learning</b> We add few-shot example demonstrations (Brown et al., 2020): for the NL-based setting, using (NL, fastest program) pairs; for history-based editing, using (slow program, fast program) pairs. We randomly sample examples from the train set as the few-shot examples.		364
320			365
321			366
322			367
323		<b>Execution Conditioned Fine-tuning</b> Beyond fine-tuning with basic contexts, we posit that further conditioning on execution results could help. Therefore, we include execution results of PASS/FAIL status, runtime, and memory usage for each public test case for the input program.	368
324			369
325	<b>4.2 Iterative Refinement</b>		370
326	We explore three methods to iteratively improve and refine the generated code to be more efficient, which intuitively aligns with the way that humans continuously refine code (Madaan et al., 2024).		371
327			372
328			373
329		<b>Trajectory Conditioned Fine-tuning</b> For history-based editing, we further propose trajectory-conditioned fine-tuning, by adding a trajectory history of programs written by the same user for the given problem in context. We first collect all problems with at least three programs submitted by the same user, and treat the series of programs as a trajectory. From each qualified trajectory, we designate the fastest code as the target output, and sample three other intermediate programs at the 0th, 33rd, and 66th percentile steps to use as inputs. We aim to allow the model to learn from the step-by-step improvements that led to the optimal solution, capturing the problem-solving process in addition to just the inputs and targets.	374
330			375
331			376
332			377
333			378
334			379
335			380
336			381
337			382
338			383
339			384
340			385
341			386
342			387
343			388
344		<b>5 Experiments</b>	389
345		<b>5.1 Experimental Setup</b>	390
		<b>Models</b> We experiment with several best-performing LMs pre-trained on code. Specifi-	391
			392

<sup>5</sup>Models do not have access to the private test cases we finally evaluate on.

Model	Setting	History-based Editing			NL-instructed Generation		
		pass@1	speedup	memory reduction	pass@1	runtime%	memory%
StarCoder2	instruct	49.4	1.49	1.24	3.5	67.3	63.7
	few-shot	49.8	1.70	1.07	3.5	55.3	63.6
CodeGemma	instruct	42.5	1.43	1.10	21.4	57.2	49.0
	few-shot	43.9	1.07	1.06	21.4	52.7	52.0
WizardCoder	instruct	34.2	1.58	1.18	17.8	61.9	63.7
	few-shot	27.4	1.38	1.12	17.8	67.0	56.1
CodeLLaMa	instruct	57.5	1.44	1.11	7.1	76.9	<b>66.8</b>
	few-shot	22.5	1.63	1.26	10.7	<b>84.1</b>	58.2
DeepseekCoder	instruct	29.8	2.11	<b>1.28</b>	23.2	42.9	60.3
	few-shot	35.2	<b>2.26</b>	1.20	21.4	70.7	59.3
GPT-4o	instruct	<b>66.6</b>	1.64	1.10	<b>53.5</b>	56.4	52.4
	few-shot	65.8	1.62	1.12	48.2	65.1	58.9

Table 2: Generation results using in-context learning approaches (instruction-prompting and few-shot examples).

cally, we evaluate CodeLlama-13B (Roziere et al., 2023), DeepSeekCoder-7B (Guo et al., 2024), CodeGemma-7B (Team et al., 2024), WizardCoder-13B-Python (Luo et al., 2023), StarCoder2-15B (Lozhkov et al., 2024). We use the instruction-tuned versions of all of these open-checkpoint models unless indicated otherwise. We also use the proprietary GPT-4o model for no-training methods.

## 5.2 Results and Analysis

### 5.2.1 In-Context Learning

As shown in Table 2, all methods either reduce pass@1 by a large margin (in editing mode) or obtain low pass@1 (generation mode). Comparing the two paradigms, history-based editing results in a substantially higher pass@1 by referring to a base correct program, compared to NL-instructed generation which lacks a base program to start from.

**History-based editing** While in-context learning can effectively speed up the program by 7–126%, it also degrades correctness by 42.5–77.5 from the 100% pass@1 in base programs, and uses more memory. *Few-shot* shows this trend more explicitly than *instruct*. Besides the limitations of LMs, this may be caused by the sampled few-shot demonstrations being algorithmically less relevant to the problem at hand.

**NL-instructed generation** In this paradigm, instruction and few-shot examples are comparably helpful in all three evaluated dimensions for most models. While WizardCoder and CodeLlama obtain 7–35 higher pass@1 with *instruct*, DeepseekCoder *few-shot* pass@1 is higher by 5.4 points. Nonetheless, this trend does not generalize to

the two efficiency dimensions, suggesting that no method or model is an overall winner.

### 5.2.2 Iterative Refinement

Table 3 shows all results with iterative refinement methods. As a reference for the refinement approaches, we measure the LM-generated code in the first attempt at optimization without any refinement, and denote this method as *pre-refine*.

**History-based editing paradigm** While all methods can effectively speed up the program, methods that involve NL feedback (*self-refine* and *nl+exec refine*) achieve the highest speedup across models. *exec-refine* consistently yields the highest pass@1 for all models, by 3.4–38.8 points more than the other two methods. We conjecture that execution outputs are better representations to inform functional correctness than NL descriptions. Although it is easier to convey high-level optimization strategies in NL, conveying the functional correctness is harder. Overall, although the models are instructed to emphasize both correctness and efficiency, there seems to be an implicit trade-off between them. Additional analysis of the same is detailed in §A.

**NL-instructed generations** We observe similar patterns as the editing mode, that *exec-refine* best maintains functional correctness, and two other NL-involved approaches improve runtime/memory efficiency by a larger margin. Compared to the in-context learning results in Table 2, iterative refinement yields 21.2–38.1 percentile higher in the program runtime spectrum, which is reasonable with the additional prompting efforts.

Model	Setting	History-based Editing			NL-instructed Generation		
		pass@1	speedup	memory reduction	pass@1	runtime%	memory%
StarCoder2	pre-refine	49.4	1.49	1.24	<b>3.5</b>	67.3	63.7
	self-refine	26.7	1.55	<b>1.30</b>	1.7	3.8	<b>75.6</b>
	exec-refine	<b>39.5</b>	1.49	1.23	1.7	3.8	41.2
	nl+exec refine	26.1	<b>2.13</b>	1.26	1.7	<b>95.5</b>	50.1
CodeGemma	pre-refine	42.5	1.43	1.10	21.4	57.2	49.0
	self-refine	15.1	<b>2.08</b>	<b>1.15</b>	12.5	44.7	<b>56.8</b>
	exec-refine	<b>33.2</b>	1.59	1.12	<b>21.4</b>	50.1	52.3
	nl+exec refine	29.8	1.54	1.14	17.8	<b>54.3</b>	41.1
WizardCoder	pre-refine	34.2	1.58	1.18	17.8	61.9	63.7
	self-refine	8.5	2.16	1.23	12.5	<b>83.1</b>	60.1
	exec-refine	<b>20.9</b>	1.60	1.13	<b>17.8</b>	76.5	54.7
	nl+exec refine	18.3	<b>2.90</b>	<b>1.30</b>	16.0	64.6	<b>61.9</b>
CodeLLaMa	pre-refine	57.5	1.44	1.11	7.1	76.9	66.8
	self-refine	15.8	<b>2.02</b>	<b>1.22</b>	3.5	<b>100.0</b>	<b>59.4</b>
	exec-refine	<b>54.6</b>	1.51	1.12	<b>5.3</b>	47.8	60.5
	nl+exec refine	16.2	1.37	1.02	<b>5.3</b>	99.8	49.9
DeepseekCoder	pre-refine	29.8	2.11	1.28	23.2	42.9	60.3
	self-refine	13.6	2.73	1.35	12.5	<b>81.0</b>	60.3
	exec-refine	<b>27.4</b>	2.34	1.24	<b>19.6</b>	63.6	50.3
	nl+exec refine	19.6	<b>3.54</b>	<b>1.37</b>	16.0	45.9	<b>69.3</b>
GPT-4o	pre-refine	66.6	1.64	1.10	53.5	56.4	52.4
	self-refine	47.8	<b>2.72</b>	<b>1.25</b>	41.1	<b>72.2</b>	46.9
	exec-refine	<b>60.8</b>	2.19	1.22	<b>60.7</b>	63.0	<b>50.7</b>
	nl+exec refine	58.8	2.39	1.22	57.1	53.1	44.1

Table 3: Generation results with iterative refinement approaches for two optimization paradigms.

**GPT-4o** obtains much higher pass@1 than all other models in both paradigms. However, it underperforms on time and memory efficiency aspects, implying its limited ability to achieve both goals.

### 5.2.3 Fine-tuning

We perform parameter-efficient fine-tuning on CodeLLaMa-7B and DeepseekCoder-7B, the best-performing models on the correctness and efficiency metrics in our prompting experiments<sup>6</sup>.

**History-based editing** As shown in Table 4, fine-tuning is the most effective method in maintaining correctness in the editing paradigm. Especially for DeepseekCoder, compared to the highest prompt-

<sup>6</sup>We finetune CodeLLaMa 7B instead of 13B due to computational constraints

Model	Method	pass@1	speedup	mem.red.
CodeLLaMa-7B	Vanilla	43.0	1.11	1.01
	Execution	45.0	<b>1.41</b>	<b>1.04</b>
	Trajectory	<b>70.2</b>	1.01	1.00
DeepseekCoder	Vanilla	42.1	1.11	1.01
	Execution	43.0	<b>1.16</b>	<b>1.02</b>
	Trajectory	<b>69.8</b>	1.01	1.00

Table 4: Fine-tuning results for history-based editing.

ing results 35.2 using *few-shot* examples, *vanilla* and *execution*-conditioned tuning improves by 6.9 and 7.8 points, and *trajectory*-conditioned tuning further gains a 34.6 point increase overall. This suggests that adding user-specific coding trajectories can help ground models into the optimization mode and substantially improve output correctness.

**NL-instructed generation** Similarly for NL-instructed generation in Table 5, fine-tuning is effective in maintaining correctness, yet does not optimize the program by a large extent, in either time or space dimensions.

However, fine-tuning results in a much less efficiency improvement than prompting-based methods for both paradigms, possibly due to the limited power of parameter efficient fine-tuning.

Model	Method	pass@1	runtime%	mem%
CodeLLaMa-7B	Pre-trained	7.1	<b>96.6</b>	43.5
	Finetuned	<b>14.2</b>	71.0	<b>57.2</b>
Deepseek	Pre-trained	<b>23.2</b>	42.9	60.3
	Finetuned	<b>23.2</b>	<b>80.5</b>	64.6

Table 5: Fine-tuning results for NL-based generation.

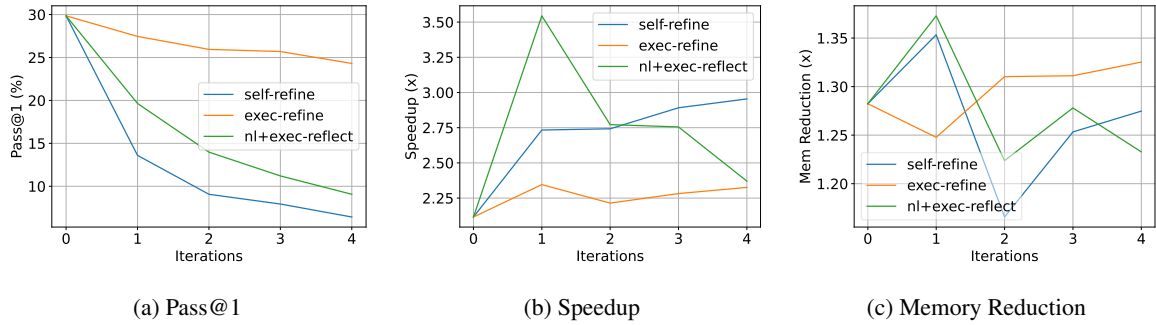


Figure 5: Performance of DeepseekCoder over multiple iterations of refinement. *The improvement in efficiency is outweighed by the consistent drop in pass@1.*

## 6 Additional Analysis

**Does instruction tuning help in-context learning?** As instruct model versions are expected to be better at in-context learning than their base counterparts, we compare base and instruct model versions in the editing paradigm. In Table 6, the base versions get much higher correctness albeit at lower efficiency, showing that base and instruct versions lie at different points on the correctness-efficiency trade-off. As we emphasize both correctness and efficiency aspects in the NL instruction, we conjecture the instruct models take more hints from the input format and emphasize efficiency, while base models primarily emphasize correctness.

Model	Version	pass@1	speedup	mem.red.
CodeLlama	Instruct	22.5	<b>1.63</b>	<b>1.26</b>
	Base	<b>46.4</b>	1.02	1.00
DeepseekCoder	Instruct	35.2	<b>2.26</b>	<b>1.20</b>
	Base	<b>45.4</b>	1.04	1.00
CodeGemma	Instruct	43.9	<b>1.07</b>	<b>1.06</b>
	Base	<b>48.2</b>	1.01	1.00
StarCoder2	Instruct	49.8	<b>1.70</b>	<b>1.07</b>
	Base	<b>41.5</b>	1.03	1.00

Table 6: Comparing base and instruct model versions with in-context learning methods.

**Does multi-iteration refinement help?** Multiple refinement iterations may improve results by allowing more turns for models to refine. To verify this, we evaluate the *iterative refinement* method using 1–4 iterations. While self-refine and exec-refine improve program speedup over iterations (Figure 5b), all methods continuously degrade the pass@1 to various extents. For memory usage (Figure 5c), exec-refine consistently reduces memory usage, yet other methods exhibit big fluctuations. In general, more iterations can speed up the program, yet further sacrifice correctness, thus has limited gains in terms of correctness-preserving optimization.

## Can iterative prompting fix incorrect solutions?

To study whether models can recover from incorrect starting solutions in the history-based editing paradigm, we evaluate on a collection of 157 pairs of programs, ECCO-FIX, where the input code is almost correct: one that passed all public test cases but fails a few private test cases. As shown by Table 7, we show that *exec-refine* can fix incorrect programs to pass all public and private test cases, with access to only the PASS/FAIL status of public test cases. In comparison, *self-refine* breaks the correctness of more programs. Aligning with our findings in earlier sections and §A, *exec-refine*, with execution information in contexts, can encourage models to generate functionally correct programs.

Model	Instruct	Self-Refine	Exec-Refine
StarCoder2	12.1	<b>9.5</b>	8.9
CodeGemma	10.8	3.1	10.8
WizardCoder	<b>20.3</b>	1.9	14.6
CodeLlama	17.1	7.6	<b>36.9</b>
DeepseekCoder	12.7	5.7	15.2

Table 7: Pass@1 of iterative refinement strategies for models on ECCO-FIX, under the editing paradigm.

## 7 Conclusion

In this paper, we introduce the ECCO benchmark that enables two paradigms for Python program optimization, using JUDGE0, a language and platform-agnostic execution framework. We find that execution information and fine-tuning help LLMs maintain code correctness, and prompting with natural language often yields higher efficiency gains. However, we broadly reconfirm findings that no existing method can improve efficiency without sacrificing functional correctness. We hope ECCO can serve as a testbed for program optimization, and we call for more efforts in advancing correctness-preserving program optimizations.



## 544 Limitations

545 Our benchmark establishes a solid foundation for  
546 rigorously evaluating the ability to generate effi-  
547 cient solutions, however, we also note that there  
548 are limitations to our work.

549 First, ECCO only includes Python problems so  
550 far, but our JUDGE0 evaluation platform and our  
551 benchmark curation recipe is fully reproducible and  
552 could extend to other programming languages of  
553 interest. Second, our benchmark currently focuses  
554 on competitive programming problems. It is also  
555 possible to extend our benchmark to other types  
556 of programming problems from more real-world  
557 software engineering scenarios.

558 Due to both limitations, our results may not  
559 be comprehensive enough to reflect the quality of  
560 model-generated programs on the full spectrum of  
561 all programming languages and problems. When  
562 using ECCO in practice, we recommend the read-  
563 ers examine the model outputs, in addition to the  
564 quantitative results produced by our framework.

## 565 References

566 Seung-Yeop Baik, Mingi Jeon, Joonghyuk Hahn, Jungin  
567 Kim, Yo-Sub Han, and Sang-Ki Ko. 2024. Codecom-  
568 plex: A time-complexity dataset for bilingual source  
569 codes. *arXiv preprint arXiv:2401.08719*.

570 Tom Brown, Benjamin Mann, Nick Ryder, Melanie  
571 Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind  
572 Neelakantan, Pranav Shyam, Girish Sastry, Amanda  
573 Askell, Sandhini Agarwal, Ariel Herbert-Voss,  
574 Gretchen Krueger, Tom Henighan, Rewon Child,  
575 Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens  
576 Winter, Chris Hesse, Mark Chen, Eric Sigler, Ma-  
577 teusz Litwin, Scott Gray, Benjamin Chess, Jack  
578 Clark, Christopher Berner, Sam McCandlish, Alec  
579 Radford, Ilya Sutskever, and Dario Amodei. 2020.  
580 [Language models are few-shot learners](#). In *Ad-  
581 vances in Neural Information Processing Systems*,  
582 volume 33, pages 1877–1901. Curran Associates, Inc.

584 Rudy Bunel, Alban Desmaison, M Pawan Kumar,  
585 Philip HS Torr, and Pushmeet Kohli. 2016. Learn-  
586 ing to superoptimize programs. *arXiv preprint  
587 arXiv:1611.01787*.

588 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming  
589 Yuan, Henrique Ponde de Oliveira Pinto, Jared Ka-  
590 plan, Harri Edwards, Yuri Burda, Nicholas Joseph,  
591 Greg Brockman, et al. 2021. Evaluating large  
592 language models trained on code. *arXiv preprint  
593 arXiv:2107.03374*.

594 Tristan Coignion, Clément Quinton, and Romain Rou-  
595 voy. 2024. A performance study of llm-generated

code on leetcode. In *28th International Conference  
on Evaluation and Assessment in Software Engineer-  
ing (EASE'24)*. 596  
597  
598

Chris Cummins, Volker Seeker, Dejan Grubisic,  
Mostafa Elhoushi, Youwei Liang, Baptiste Roziere,  
Jonas Gehring, Fabian Gloeckle, Kim Hazelwood,  
Gabriel Synnaeve, et al. 2023a. Large language  
models for compiler optimization. *arXiv preprint  
arXiv:2309.07062*. 599  
600  
601  
602  
603  
604

Chris Cummins, Volker Seeker, Dejan Grubisic,  
Mostafa Elhoushi, Youwei Liang, Baptiste Roziere,  
Jonas Gehring, Fabian Gloeckle, Kim Hazelwood,  
Gabriel Synnaeve, et al. 2023b. Large language  
models for compiler optimization. *arXiv preprint  
arXiv:2309.07062*. 605  
606  
607  
608  
609  
610

Herman Zvonimir Došilović and Igor Mekterović. 2020.  
Robust and scalable online code execution system. In  
*2020 43rd International Convention on Information,  
Communication and Electronic Technology (MIPRO)*,  
pages 1627–1632. IEEE. 611  
612  
613  
614  
615

Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai  
Dong, Wentao Zhang, Guanting Chen, Xiao Bi,  
Y Wu, YK Li, et al. 2024. Deepseek-coder: When the  
large language model meets programming—the rise of  
code intelligence. *arXiv preprint arXiv:2401.14196*. 616  
617  
618  
619  
620

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 *International Conference on  
Learning Representations*. 621  
622  
623  
624  
625

Dong Huang, Jie M Zhang, Yuhao Qing, and Heming  
Cui. 2024. Effibench: Benchmarking the efficiency  
of automatically generated code. *arXiv preprint  
arXiv:2402.02037*. 626  
627  
628  
629

Mingi Jeon, Seung yeop Baik, Joonghyuk Hahn, Yo-Sub  
Han, and Sang-Ki Ko. 2023. [Deep learning-based  
source code complexity prediction](#). 630  
631  
632

Mohammad Abdullah Matin Khan, M Saiful Bari,  
Xuan Long Do, Weishi Wang, Md Rizwan Parvez,  
and Shafiq Joty. 2023. xcodeeval: A large scale multi-  
lingual multitask benchmark for code understanding,  
generation, translation and retrieval. *arXiv preprint  
arXiv:2303.03004*. 633  
634  
635  
636  
637  
638

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying  
Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.  
Gonzalez, Hao Zhang, and Ion Stoica. 2023. Effi-  
cient memory management for large language model  
serving with pagedattention. In *Proceedings of the  
ACM SIGOPS 29th Symposium on Operating Systems  
Principles*. 639  
640  
641  
642  
643  
644  
645

Yujia Li, David Choi, Junyoung Chung, Nate Kushman,  
Julian Schrittwieser, Rémi Leblond, Tom Eccles,  
James Keeling, Felix Gimeno, Agustin Dal Lago,  
et al. 2022. Competition-level code generation with  
alphacode. *Science*, 378(6624):1092–1097. 646  
647  
648  
649  
650

651	Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, et al. 2024. Starcoder 2 and the stack v2: The next generation. <i>arXiv preprint arXiv:2402.19173</i> .	Alex Shypula, Pengcheng Yin, Jeremy Lacomis, Claire Le Goues, Edward Schwartz, and Graham Neubig. 2021. Learning to superoptimize real-world programs. <i>arXiv preprint arXiv:2109.13498</i> .	704 705 706 707
656	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 evol-instruct. <i>arXiv preprint arXiv:2306.08568</i> .	Alexander G Shypula, Aman Madaan, Yimeng Zeng, Uri Alon, Jacob R. Gardner, Yiming Yang, Milad Hashemi, Graham Neubig, Parthasarathy Ranganathan, Osbert Bastani, and Amir Yazdanbakhsh. 2024. <a href="#">Learning performance-improving code edits</a> . In <i>The Twelfth International Conference on Learning Representations</i> .	708 709 710 711 712 713 714
661	Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2024. Self-refine: Iterative refinement with self-feedback. <i>Advances in Neural Information Processing Systems</i> , 36.	Jagriti Sikka, Kushal Satya, Yaman Kumar, Shagun Upal, Rajiv Ratn Shah, and Roger Zimmermann. 2020. Learning based methods for code runtime complexity prediction. In <i>Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part I 42</i> , pages 313–325. Springer.	715 716 717 718 719 720 721
667	Kaushik Moudgalya, Ankit Ramakrishnan, Vamsikrishna Chemudupati, and Xing Han Lu. 2023. Tasty: A transformer based approach to space and time complexity. <i>arXiv preprint arXiv:2305.05379</i> .	Manav Singhal, Tushar Aggarwal, Abhijeet Awasthi, Nagarajan Natarajan, and Aditya Kanade. 2024. No-funeval: Funny how code lms falter on requirements beyond functional correctness. <i>arXiv preprint arXiv:2401.15963</i> .	722 723 724 725 726
671	Changan Niu, Ting Zhang, Chuanyi Li, Bin Luo, and Vincent Ng. 2024. On evaluating the efficiency of source code generated by llms. <i>arXiv preprint arXiv:2404.06041</i> .	Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. <i>arXiv preprint arXiv:2403.08295</i> .	727 728 729 730 731 732
675	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. <i>Advances in neural information processing systems</i> , 35:27730–27744.	Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. <a href="#">Finetuned language models are zero-shot learners</a> . In <i>International Conference on Learning Representations</i> .	733 734 735 736 737
681	Ruchir Puri, David S Kung, Geert Janssen, Wei Zhang, Giacomo Domeniconi, Vladimir Zolotov, Julian Dolby, Jie Chen, Mihir Choudhury, Lindsey Decker, et al. 2021. Codenet: A large-scale ai for code dataset for learning a diversity of coding tasks. <i>arXiv preprint arXiv:2105.12655</i> .		
687	Tal Ridnik, Dedy Kredo, and Itamar Friedman. 2024. Code generation with alphacodium: From prompt engineering to flow engineering. <i>arXiv preprint arXiv:2401.08500</i> .		
691	Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code llama: Open foundation models for code. <i>arXiv preprint arXiv:2308.12950</i> .		
696	Hui Shi and Yang Zhang. 2020. Deep symbolic superoptimization without human knowledge. <i>ICLR 2020</i> .		
699	Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2024. Reflexion: Language agents with verbal reinforcement learning. <i>Advances in Neural Information Processing Systems</i> , 36.		

## A Iterative Refinement Analysis

Following Shypula et al. (2024), we also evaluate runtime and memory efficiency as a percentage of pairs where the generated code  $p_{out}$  is faster/uses lesser memory than  $p_{in}$ , in addition to the speedup and memory reduction metrics described in §3.3.1.

In Table 8, these percentage optimized metrics clearly indicate that exec-refine is the most consistent iterative refinement method as it achieves the best pass@1, % runtime optimized and % memory optimized across all models. Contrasting these results to Speedup and Memory Reduction in Table 3, we note that while natural language feedback (in self-refine and NL+Exec refine) aids in significantly improving the runtime and memory usage for a few cases (while breaking test cases for others), exec-refine improves runtime and memory in more cases albeit to a smaller degree.

## B Implementation Details

**Generation** We use a maximum length of 1024 tokens, and a temperature of  $t = 0.4$ . We provide two examples for all in-context few-shot experiments

**Hardware** We run all our experiments on a mix of A6000 and L40 GPUs. Specifically for the prompting iterative prompting and in-context learning approaches, we use 1-2 GPUs of the type and utilize the vLLM library (Kwon et al., 2023) primarily for generating programs efficiently. We perform finetuning on 4 A6000 GPUs.

For JUDGE0 evaluation virtual hardware setup, we use an m7i.large EC2 instance which has 2 vCPU cores of the 4th Generation Intel Xeon Sapphire Rapids, with 8GB of RAM.

**Fine-tuning** We adopt parameter-efficient fine-tuning with LoRA (Hu et al., 2022) due to resource limitations and implement this using the HuggingFace Transformers library<sup>7</sup>. We optimize the fine-tuning using DeepSpeed<sup>8</sup> ZeRo stage 2 with a LoRA rank of 8, alpha of 16 and a per-device batch size of 2. We use the AdamW optimizer, with a learning rate of 1e-3 and a warmup of 100 steps. We train the models on the history-based editing task for a single epoch and fine-tune models for 10 epochs on the NL-instructed generation task.

<sup>7</sup><https://github.com/huggingface/transformers>

<sup>8</sup><https://github.com/microsoft/DeepSpeed>

## C Prompt Details

We illustrate and detail all of the prompts used for the experiments below.

Model	Setting	History-based Editing		
		pass@1	% Runtime Opt	% Mem. Opt.
StarCoder2	self-refine	26.7	22.5	27.0
	exec-refine	<b>39.5</b>	<b>25.5</b>	<b>30.3</b>
	nl+exec refine	26.1	22.6	27.0
CodeGemma	self-refine	15.1	18.2	29.3
	exec-refine	<b>33.2</b>	<b>25.5</b>	<b>34.5</b>
	nl+exec refine	29.8	21.0	31.3
WizardCoder	self-refine	8.5	16.1	24.4
	exec-refine	<b>20.9</b>	<b>20.0</b>	<b>27.2</b>
	nl+exec refine	18.3	17.1	26.9
CodeLLaMa	self-refine	15.8	15.6	24.5
	exec-refine	<b>54.6</b>	<b>27.4</b>	<b>39.6</b>
	nl+exec refine	16.2	16.8	23.2
DeepseekCoder	self-refine	13.6	26.8	35.8
	exec-refine	<b>27.4</b>	<b>31.4</b>	<b>40.0</b>
	nl+exec refine	19.6	29.7	36.5
GPT-4o	self-refine	47.8	42.5	42.3
	exec-refine	<b>60.8</b>	<b>48.8</b>	<b>51.1</b>
	nl+exec refine	58.8	48.6	48.1

Table 8: Comparing iterative refinement approaches on % Optimization metrics. *Exec-Refine* is the best performing approach across all models and metrics

```

Optimize the python program below to be functionally equivalent
but run faster and use less memory.

Here are a few examples:

### Program:
{slow_code_example}

### Optimized (Runtime and Space) version of Program above:
{fast_code_example}

### Program:
[src_code]

### Optimized (Runtime and Space) version of Program above:

```

Figure 6: Prompt for Instruction prompting  $I_{eff}$  along with in-context examples

```

Your solution was functionally {CORRECT/INCORRECT}

Here are the run time and memory stats of your code for each test case
-- Stats for test case 0 --
Correct: {PASSED/FAILED}
Run time: 0.009 s
Memory: 3352.0 KB
-- Stats for test case 1 --
Correct: {PASSED/FAILED}
Run time: 0.009 s
Memory: 3316.0 KB

```

Figure 7: Example of Execution Feedback on public test cases used for Exec-Refine and Execution Conditioned Fine-tuning

```

Write a python solution for the problem below which is
correct and efficient (runtime and memory).

##Problem Name:
{problem_name}

##Problem Description:
{In detail description of the task}

## Sample Inputs:
{input_test_cases}

##Sample Outputs:
{Expected Output}

```

Figure 8: Prompt for NL-instructed generation  $I_{gen}$

```

### Candidate solution:
{previous_code_attempt}

### Feedback for incorrectnes/inefficiency and how it can be improved:
{self-feedback / execution-feedback / reflection}

### Optimized/Corrected solution based on feedback:

```

Figure 9: Prompt used for refining code in Self-Refine, Exec-Refine and NL+Exec-Refine

```
Based on the execution results, reflect on why the code solution
below was incorrect or inefficient and how the program can be fixed.
```

```
{generated code solution}
```

```
## Execution Results:
```

```
{Execution Feedback for the previous attempt}
```

```
### Reflection on incorrectnes/inefficiency and how it can be improved:
```

Figure 10: Prompt used for NL+Exec-Refine to reflect on Execution results

```
##1 iteration program:
```

```
{slowest program in trajectory}
```

```
##2 iteration program:
```

```
{33 percentile fastest program in trajectory}
```

```
##3 iteration program:
```

```
{66 percentile fastest program in trajectory}
```

```
### Final iteration program:
```

Figure 11: Format of Trajectory-Conditioned Fine-tuning data