



DolphCoder: Echo-Locating Code Large Language Models with Diverse and Multi-Objective Instruction Tuning

Yejie Wang^{1*}, Keqing He^{2*}, Guanting Dong¹, Pei Wang¹, Weihao Zeng¹, Muxi Diao¹
Mengdi Zhang², Jingang Wang², Xunliang Cai², Weiran Xu^{1†}

¹Beijing University of Posts and Telecommunications, Beijing, China

²Meituan, Beijing, China

{wangyejie, dongguanting, wangpei, zengwh, dmx, xuweiran}@bupt.edu.cn

{hekeqing, zhangmengdi02, wangjingang02, caixunliang}@meituan.com

Abstract

Code Large Language Models (Code LLMs) have demonstrated outstanding performance in code-related tasks. Several instruction tuning approaches have been proposed to boost the code generation performance of pre-trained Code LLMs. In this paper, we introduce a diverse instruction model (**DolphCoder**) with self-evaluating for code generation. It learns diverse instruction targets and combines a code evaluation objective to enhance its code generation ability. Our model achieves superior performance on the HumanEval and MBPP benchmarks, demonstrating new insights for future code instruction tuning work. Our key findings are: (1) Augmenting more diverse responses with distinct reasoning paths increases the code capability of LLMs. (2) Improving one's ability to evaluate the correctness of code solutions also enhances their ability to create it. Our code is released at <https://github.com/prisnlp/DolphCoder>.

1 Introduction

Code pre-trained models have achieved remarkable progress in the era of large language models (LLMs), such as Codex (Chen et al., 2021), AlphaCode (Li et al., 2022), and PaLM-Coder (Chowdhery et al., 2022b). Code-related tasks are also the key factors in evaluating the capability of LLMs. Numerous code LLMs have been proposed, including closed-source models (Chen et al., 2021; Li et al., 2022; OpenAI, 2023) and open-source models (Li et al., 2023b; Rozière et al., 2023). They perform expensive pre-training using substantial amounts of code data and display impressive performance.

In contrast to these pre-trained code LLMs, another lightweight paradigm of enhancing code capability is instruction tuning using relatively small

high-quality code-related data. For example, Code Alpaca (Chaudhary, 2023) employs a similar self-instruct method as Alpaca (Taori et al., 2023) to generate code instructions via OpenAI's ChatGPT¹. Further, WizardCoder (Luo et al., 2023a) introduces a more complicated Evol-Instruct method (Xu et al., 2023a) which evolves existing instruction data to generate more complex and diverse datasets. Instead, OctoPack (Muennighoff et al., 2023) and Magicoder (Wei et al., 2023a) construct code instructions by mining existing code corpus. All of these methods enhance the performance of the open-source Code LLMs.

However, these methods have two weaknesses: (1) They take the only golden answer but ignore the diversity of answers in code generation. We find that augmenting more diverse responses using different system prompts increases the code capability of LLMs. (2) Current models generate plausible code snippets in terms of grammar and logic but are unable to identify subtle errors, such as corner cases and wrong input/output formats. It has no guarantee that temperature sampling will consistently produce accurate answers over time. We suppose that LLMs are capable of generating correct solutions while struggling to discriminate correct from incorrect ones. Improving one's ability to evaluate the correctness of code also enhances their ability to create it.

Inspired by the two insights, we introduce a diverse instruction model (**DolphCoder**) with self-evaluating for code generation. Specifically, we use Code Llama-python as our base model and obtain evolved instruction data following WizardCoder. Then motivated by rejection sampling (Touvron et al., 2023b; Yuan et al., 2023a) and ORCA (Mukherjee et al., 2023), we use different system prompts to generate diverse answers via ChatGPT. After removing low-quality and similar data using

* Equal contribution.

† Corresponding author.

¹<https://openai.com/blog/ChatGPT>

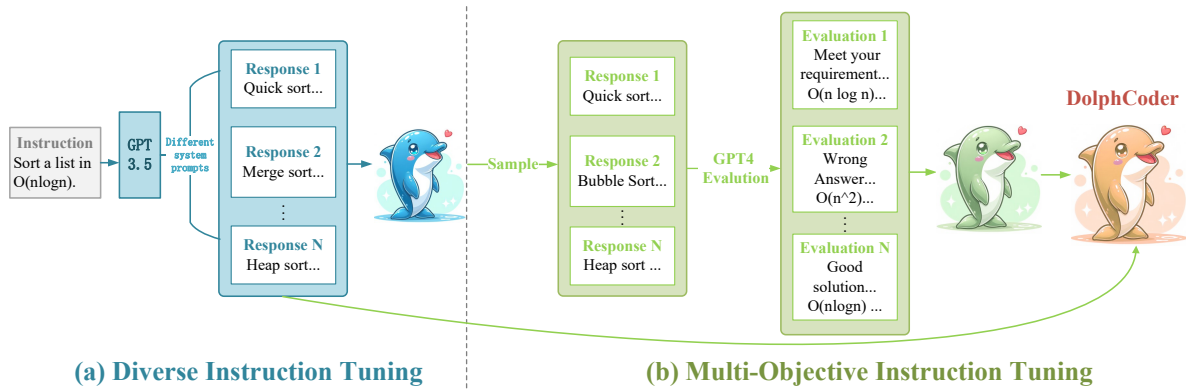


Figure 1: The overall architecture of our proposed diverse instruction tuning with self-evaluating for code generation, DolphCoder. Stage (a) denotes Diverse Instruction Tuning (DIT) and Stage (b) denotes Multi-Objective Instruction Tuning (MOT) for self-evaluating.

heuristic rules (Luo et al., 2023a; Di et al., 2023a), we perform supervised fine-tuning on the remaining instruction data. Further, we explore whether improving one’s ability to evaluate code helps generate it. We propose a self-evaluate multi-task learning framework by adding a code evaluation objective to the traditional instruction fine-tuning task. We find training the model for both code generation and code evaluation benefits the code capability.

Our key contributions are summarized as follows:

1. We introduce a diverse instruction model (**DolphCoder**) with self-evaluating for code generation. It learns diverse instruction targets and combines a code evaluation objective to enhance its code generation ability.
2. DolphCoder outperforms strong open-source code LLMs by a large margin, including CODELLAMA-INSTRUCT, OctoCoder, and WizardCoder.

2 Method

In this section, we elaborate on the methodological details of DolphCoder. As shown in Figure 1, DolphCoder has two training stages: (1) The first is **Diverse Instruction Tuning (DIT)** with multiple chain-of-thought answers to the same instruction. (2) The second is **Multi-Objective Tuning (MOT)** of combining the code generation task and code evaluation task both in the form of a natural language generation task.

2.1 Diverse Instruction Tuning

We follow the Evol-Instruct technique (Xu et al., 2023a; Luo et al., 2023a) to construct our training corpus. Based on the Code Alpaca dataset², we iteratively evolve the programming problems in this dataset via in-depth evolving to obtain new instructions.³ For each instruction, then we use different system prompts to query ChatGPT and obtain diverse targets. These system prompts aim at augmenting user instructions and task descriptions to give more code solutions with diverse reasoning paths. We display our system prompts in Figure 2. We find more diverse answers can increase the code capability of LLMs (see Section 4.2) and argue that different styles of code solutions provide more supervised signals for the model. Similar to Luo et al. (2023a); Di et al. (2023a), we also use several heuristic rules to remove low-quality and similar data. Finally, we get a diverse code instruction dataset with a size of 510k and utilize Code Llama as the foundation LLM to finetune.

2.2 Multi-Objective Instruction Tuning

We discover that the current instruction models produce both correct and incorrect code solutions when randomly sampled. We argue that LLMs can generate correct solutions while struggling to discriminate correct from incorrect ones. Therefore, we explore whether improving one’s ability to evaluate code helps generate it.

²<https://huggingface.co/datasets/sahil2801/CodeAlpaca-20k>

³Since the original datasets or code are not released, we reproduce the evolve procedure following WizardLM (Xu et al., 2023a) and modify the evolve prompts according to the original WizardCoder paper.

System Prompts

1. [EMPTY]
2. You are a code assistant who knows multiple programming languages and how to translate between them. Given a task, you explain in simple steps what the task is asking, any guidelines it provides, and how to use those guidelines to find the solution.
3. You are an AI assistant. You will be given a task. You must generate a detailed and long answer.
4. You are an AI assistant, who knows every language and how to translate one language to another. Given a task, you explain in simple steps what the task is asking, any guidelines that it provides.
5. You solve the task and show how you used the guidelines to solve the task.
6. You are a teacher. Given a task, you explain in simple steps what the task is asking, any guidelines it provides and how to use those guidelines to find the answer.

Figure 2: We use these system prompts to generate more diverse responses where [EMPTY] means no system prompt.

We sample 5k instructions from the above dataset and use our model in the first stage to generate 100 answers using temperature sampling. Next, we use GPT-4⁴ to verify the correctness of the deduplicated generated answers in terms of grammar, logic and efficiency. In our preliminary experiments, We encounter difficulties in assessing the correctness of code using ChatGPT or other existing LLMs. We also consider the use of compiled signals from the code executor but find most generated code is grammatically correct. Previous work (Liu et al., 2023a; Shen et al., 2023a) use existing unit tests in the training set, which is not applicable to our situation. We leave out more evaluation methods to future work. Finally, we obtain the code evaluation dataset with a size of

⁴<https://openai.com/GPT-4>

Evaluation Prompt

You are a code evaluation model. Your task is to analyze and inspect the input code snippet. This involves evaluating whether it can run smoothly, whether it will produce the correct results, and whether there are issues like timeouts. You should output the result of the code with <passed> or <not passed>. Additionally, you should also provide an explanation for the evaluation result.

Figure 3: We use the evaluation prompt to query GPT-4 to access the correctness of the generated code solutions of our model.

370k.⁵ We formulate the code evaluation task in the form of a natural language generation task as shown in Figure 1. We hope its similar training form can provide multiple meaningful supervision signals to the original code generation model. In our experiments, we find it difficult to balance the two code generation and code evaluation tasks because the model always tends to overfit the code generation objective. Therefore, we use a multi-step training paradigm where we first finetune the above model as an evaluator and then as a generator. Specifically, we finetune the DIT model using the code evaluation dataset for 1 epoch and then continue training on the diverse instruction data for 100 steps. We find training more steps will not get further improvements. Experiment results demonstrate that the multi-step training stably enhances its code generation ability (shown in Figure 6).

3 Experiments

3.1 Benchmarks

In this paper, we focus on two of the most widely used benchmarks in the field of code generation.

- **HumanEval (base)⁶ and HumanEval+(plus)**. HumanEval (base) is a widely used benchmark proposed by OpenAI for code synthesis. It consists of 164 handcrafted programming problems, with an average of 9.6 test cases allocated to each one for correctness checking. Further, Liu et al. (2023b) find that

⁵We remove failed GPT-4 requests and responses without <passed> or <not passed> results.

⁶<https://github.com/openai/human-eval>

test cases in the benchmark may be insufficient and propose HumanEval+ (plus) powered by the EvalPlus framework to obtain 80x test cases.

- **MBPP (base) (Austin et al., 2021) and MBPP+ (plus)**. MBPP (base) is also a code synthesis benchmark offering a set of 500 crowd-sourced Python programming problems covering programming fundamentals, standard library functionality, and so on. Each problem consists of a task description, code solution and 3 automated test cases. EvalPlus also provides the MBPP+ (plus) benchmark which expands by 35x test cases.

We perform n-gram matching de-duplication between our training datasets and the benchmarks to prevent data leakage. For all the experiments, we use greedy decoding and report the pass@1 metric. The inference prompt template is shown in Figure 4 following WizardCoder. To keep a fair comparison, we use the same EvalPlus⁷ framework to compute metrics.

3.2 Baselines

In this paper, we categorize the baseline models into the following two types. **(1) Closed-source models:** We have specifically incorporated OpenAI’s GPT-3.5 and GPT-4, which are developed privately by leading technology companies, demonstrating the current state-of-the-art in LLM proficiency. **(2) Open-source models:** Several open-source LLMs have been made available to the AI community, although their performance generally lags behind the closed-source models a lot. As part of our research, we incorporate a significant number of these open-source models as our baselines, which include CodeGen (Nijkamp et al., 2023), CodeT5+ (Wang et al., 2021), StarCoder, CODELLAMA, OctoCoder and WizardCoder series.

3.3 Implementation Details

Data generation For Diverse Instruction data, we devise six system prompts shown as Figure 2, inclusive of a blank one, and prompt GPT-3.5-turbo to generate more diverse responses. Inspired by CodeFuse (Di et al., 2023b), we use heuristic rules to filter out low-quality data and deduplicate data based on the test set. The specific rules are as follows:

⁷<https://evalplus.github.io/leaderboard.html>

Prompt Template

Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request.

Instruction: Create a Python script for this problem: {Question}

Response:

Figure 4: Inference prompt when testing on HumanEval and MBPP.

- **Filtering low-quality data**

1. Filter data with instruction length less than 10 words or greater than 1000 words;
2. Filter data with output length less than 80 words;
3. Filter out data with invalid markdown format, such as: code blocks not closed;
4. Filter data with more than 2048 tokens;

- **Filtering data similar to test dataset**

1. Filter data containing any function name from the test dataset.
2. Using NLTK to remove stop words and punctuation from the docstring of HumanEval, obtain the core words such as "sort array prime", etc. Filter data containing more than 40% of the core words from the test dataset.

For Multi-Objective Instruction data, we extract 5000 instructions from the Diverse Instruction data and subsequently generate 100 responses for each instruction in 0.5 temperature and 0.95 topp with the model we get after diverse instruction tuning. To ascertain the correctness of each code solution, we prompt GPT-4 to evaluate whether this code solution can meet the instruction requirement. And then they will be classified to either 'passed' or 'not passed'. The prompt we used to evaluate is as Figure 3.

Training In our study, we utilize the CODELLAMA-PYTHON-7B and CODELLAMA-PYTHON-13B as our foundational models. We train DIT for 3 epochs and MOT for 1 epoch of code evaluation data and 100

Model	Size	HumanEval		MBPP	
		Base	Plus	Base	Plus
GPT-3.5 (Nov 2023)	-	72.6	65.9	81.7	69.4
GPT-4 (Nov 2023)	-	85.4	81.7	83.0	70.7
CODELLAMA-PYTHON	34B	51.8	42.7	67.2	52.9
WizardCoder-CL	34B	73.2	64.6	73.2	59.9
CodeT5+	16B	31.7	26.2	54.6	44.4
CodeGen-Mono	16B	32.9	27.4	52.6	43.6
StarCoder	15B	34.1	29.3	55.1	46.1
CODELLAMA-PYTHON	13B	42.7	36.6	61.2	50.9
CODELLAMA-INSTRUCT	13B	42.7	-	49.4	-
OctoCoder	15B	46.2	-	-	-
WizardCoder	13B	60.4*	54.3*	65.2*	53.1*
DolphCoder(ours)	13B	67.7	57.9	67.2	54.1
StarCoder	7B	24.4	20.7	33.1	28.8
CodeT5+	6B	29.3	23.8	51.9	40.9
CodeGen-Mono	6B	29.3	25.6	49.9	42.1
CODELLAMA-PYTHON	7B	37.8	34.1	57.6	45.4
CODELLAMA-INSTRUCT	7B	34.8	-	44.4	-
WizardCoder	7B	48.2	40.9	56.6	47.1
DolphCoder(ours)	7B	62.8	54.9	64.9	52.6

Table 1: Pass@1 results of different code LLMs for HumanEval and MBPP. Base means the original benchmark and Plus denotes the extended benchmark. * denotes the reproduced results via the EvalPlus scripts and other baseline results are cited from the official leaderboard. The best results in each column are in bold.

steps of code generation data. During the training phase, we establish a global batch size of 512 and a sequence length of 2048, with a learning rate initialized at $5e-6$ and a warmup fraction of 15%. After the warmup, the learning rate decays following a cosine schedule. We employ the Adam optimizer, with beta1 and beta2 parameters set at 0.9 and 0.95 respectively. And to ensure training stability, we incorporate gradient clipping with a value set at 1.0, a technique designed to prevent excessive gradient escalation that could lead to numerical instability or model divergence.

3.4 Main Results

The primary outcomes of our proposed method in comparison with the baselines are illustrated in Table 1. We conduct the following comparisons: (1) Generally, Our proposed method significantly outperforms all the previous methods in different sizes of model parameters. (2) Compared to our

base model (CODELLAMA), DolphCoder shows significant improvements in both HumanEval(+) and MBPP(+). In detail, DolphCoder-7b has a 25 percentage point increase on HumanEval and a 7.3 percentage point increase on MBPP. (3) As we compare our model DolphCoder with the recent baseline model WizardCoder-7b, which is built on the same base model (CODELLAMA), DolphCoder still consistently outperforms it across all benchmarks and model sizes. (4) DolphCoder-7b even outperforms CODELLAMA-PYTHON 34b, demonstrating the efficiency of using small-size LLMs. These results indicate the generalizability and robustness of our method.

4 Analysis

4.1 Ablation Study

To investigate the characteristics of the main components in DolphCoder, we conduct ablation ex-

Model	Size	HumanEval		MBPP	
		Base	Plus	Base	Plus
Evol Instruct	7B	50.0	44.5	59.4	50.1
+DIT	7B	57.9	51.2	64.2	52.1
+MOT	7B	55.5	46.3	64.2	52.6
+ALL(DolphCoder)	7B	62.8	54.9	64.9	52.6
Evol Instruct	13B	64.0	55.5	65.7	51.6
+DIT	13B	66.5	57.3	65.2	52.1
+MOT	13B	65.2	56.7	66.7	53.9
+ALL(DolphCoder)	13B	67.7	57.9	67.2	54.1

Table 2: Ablation study of DolphCoder.

periments in Table 2. From the results, we have the following observations: (1) Both DIT and MOT contribute to the performance improvement. Specifically, on the CODELLAMA-PYTHON-7b base, DIT yields an improvement of 7.9 pp on HumanEval and 4.8 pp on MBPP, compared to the baseline evolving instruction approach. MOT results in an improvement of 5.5 pp and 4.8 pp respectively on HumanEval and MBPP. Moreover, we observe that the MOT training data constructed by GPT-4 contains approximately 20% error noise, which severely limits the upper-bound performance of the MOT method in this work. A more detailed discussion can be found in Appendix B. (2) The combination of DIT and MOT yields further benefits. DolphCoder-7b exhibits an enhancement of 12.8 pp on HumanEval and 5.5 pp on MBPP, demonstrating the relatively orthogonal relationship between these two methods. These results demonstrate the effectiveness of our proposed methods.

4.2 Effect of Diverse Instruction Tuning

To further explore the source of the model’s improvement brought by DIT, we test the model’s performance under different sampling ratios which is shown in Table 3. The sampling ratio represents the number of system prompts. As the sampling ratio increases, the model’s performance on all indicators gradually increases, which proves that using more diverse responses can enhance the model’s performance.

In addition, we extract code snippets from the responses and check whether the implementation of these codes is different. For ease of statistics, we only focus on Python-related code. Specifically, we use AST to parse each piece of code and calculate the similarity between codes, then remove duplicates. We count the average number of unique code

Ratios	Size	Diversity	HumanEval	
			Base	Plus
1	7B	1.0	52.4	45.1
3	7B	2.1	57.9	50.6
6	7B	2.7	57.9	51.2
1	13B	1.0	63.4	53.7
3	13B	2.1	65.2	58.5
6	13B	2.7	66.5	57.3

Table 3: The effect of different sampling ratios of DIT. Code diversity shows the average number of different code solutions for a code instruction, and we use syntax analysis to parse different code solutions. All indicators refer to pass@1.

solutions left for each code task as an indicator of code diversity. Table 3 shows that as the sampling ratio increases, the number of different code solutions corresponding to the same code instruction also gradually increases and the model’s performance on HumanEval improves too. Moreover, we observe that the increase in code diversity is not linear. Specifically, when we increase the sampling ratio from 3 to 6, we only observe marginal performance gains which we argue that the DIT data may contain unnecessary redundancy. Concurrent work (Lu et al., 2023b; Liu et al., 2023c) explore more complicated diversity-based compression methods for general instruction fine-tuning. We leave it to future work.

4.3 Effect of Multi-Objective Instruction Tuning

To explore how the code evaluation capability and code generation capability benefit each other, we evaluate them in the MOT training stage. Specifically, we evaluate the model’s code evaluation capability by sampling 40 answers for each of the 164 test questions in HumanEval and classifying their correctness using the given golden unit tests. The experimental results are shown in Figure 5.

From the results, we observe that there exists a strong association between the evaluation and generation capabilities of the model. As the model continues to train on the evaluation task, its evaluation capability keeps improving and gradually stabilizes. However, this process impairs the model’s code generation capability, potentially due to catastrophic forgetting caused by the multi-step training. After the model undergoes the generator training process, the model’s code generation capability is restored and the pass@1 metric surpasses the limit

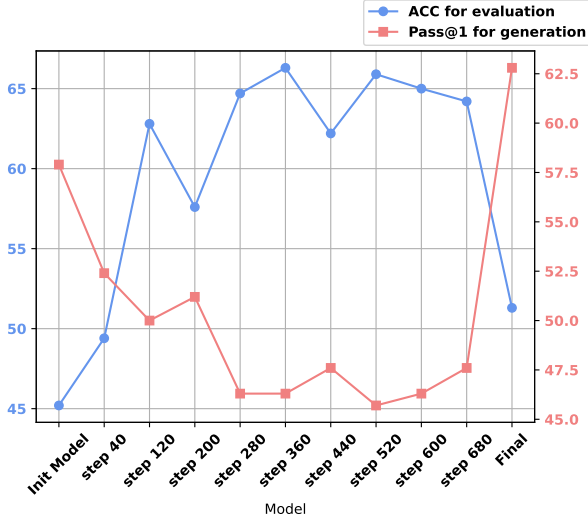


Figure 5: The trend of code evaluation capability and code generation capability during the MOT stage where step 200 refers to the training step in the first step of MOT. Init model means DIT model and Final means DolphCoder. Pass@1 refer to pass@1 on HumanEval.

Model	Greedy	Pass@1	Pass@10
DolphCoder	62.8	59.9	71.3
-w/o MOT	57.9	56.6	69.5
Improvement	4.9	3.3	1.8

Table 4: Pass@k with different decoding methods and different k values, where Greedy employs greedy sampling, pass@1 and pass@10 samples at temperature=0.2. All indicators are tested on HumanEval based on DolphCoder-7b.

of DIT model performance. Meanwhile, its evaluation capability significantly decreases. We suppose that the improvement in code generation capability is achieved through the transformation of the evaluation capability, but these two capabilities are challenging to coexist.

To further scrutinize the impact of the MOT, we compare the metrics of pass@1 and pass@10 to ascertain whether the improvements stem from the model’s heightened preference for the correct response. From Table 4, we find that the DolphCoder outperforms the DIT Model across all metrics. The improvement of MOT on the pass@1 greedy decoding indicator is the largest, with an increase of 4.9%. Compared with pass@1, the improvement on pass@10 is significantly reduced to only 1.8%. This implies that the impact of the MOT does not focus on enhancing the model’s capability to gener-

Training	Inference	HumanEval	
		Base	Plus
×	×	57.9	50.6
×	✓	55.5	48.2
✓	×	66.5	57.3
✓	✓	61.0	53.7

Table 5: The effect of system prompt during training and inference process. Considering efficiency, we only conduct experiments on the DIT model based on CODELLAMA-13B-PYTHON and use pass@1 metric on HumanEval with greedy decoding. For inference with system prompts, we randomly choose a system prompt in the training for each test query.




Figure 6: Effect of different settings of multi-task learning based on DolphCoder-7B. Mix denotes directly adding the generation and evaluation loss together and Multi-Step means sequential training.

ate robust solutions. Instead, it primarily augments the model’s ability to distinguish responses, leading to a higher preference for the correct answer.

4.4 Analysis of Key Designs

System Prompt Ablation We obtain more diverse training data by transforming the system prompt. In order to explore whether the system prompt should still be used during training and inference, we conduct experimental verification. Table 5 shows the effect of the system prompts during the training and inference process. From the results, we observe that, compared to training with system prompts, training without the system prompt results in a significant decline. This implies that assigning completely identical code instructions to different answers is not a good training method. When comparing whether to use the system prompt during

inference, we find that erasing the system prompt results in a significant improvement compared to using the system prompt. This improvement may stem from the model’s freedom to determine the most suitable inference path.

Multi-Task Training Method Ablation We also explore different multi-task training methods as shown in Figure 6. The results show that both MOT variants significantly outperform the DIT baseline and the multi-step training setting gets superior performance compared to the mix training one. We also find the multi-step training gets a more stable training process to balance the two code generation and code evaluation tasks.

5 Related Work

5.1 Instruction Fine-Tuning

Large Language Models (LLMs) initially have shown notable zero-shot performance in various domains (Brown et al., 2020; Wu et al., 2021; Hou et al., 2024; Dong et al., 2023b; Zhou et al., 2024; Wu et al., 2023; Song et al., 2023). In order to further explore his potential, Large language models (LLMs) experience the Instruction fine-tuning (IFT) stage, which enhances their capability to accomplish tasks and adhere to human instructions. The term IFT is broadly used here to encompass a range of sequence-to-sequence fine-tuning applications (Zhou et al., 2023b; Yuan et al., 2023c; Cheng et al., 2023b; Zhang et al., 2024a). T5 (Raffel et al., 2023) was one of the first models to explore this approach, training on multiple supervised text-to-text tasks. Recent studies have delved into the multi-task instruction-based fine-tuning of pre-trained LLMs, aiming to improve their innate ability to perform various downstream NLP tasks effectively, such as FLAN (Wei et al., 2022), T0 (Sanh et al., 2022), and UnifiedQA (Khashabi et al., 2020), further expanded the task scope to improve the overall generalization ability of LMs.

Following the notable successes of proprietary LLMs, particularly ChatGPT, a growing emphasis has been placed on employing Instruction Fine-Tuning (IFT) to better align LLMs with human intentions, as underscored by the research of (Brown et al., 2020; Ouyang et al., 2022). A prevalent observation from these investigations is the significant enhancement in model performance attainable through fine-tuning on a more diverse dataset.

Additionally, IFT has proven its effectiveness for Large Language Models (LLMs) across various

domains, such as mathematical reasoning (Cobbe et al., 2021; Hendrycks et al., 2021; Yuan et al., 2023b; Chen et al., 2024; Yue et al., 2023; Gou et al., 2024; Li et al., 2023a; Yue et al., 2024) and code generation (Chaudhary, 2023; Luo et al., 2023b; Wei et al., 2023b), as well as in Information Extraction (IE) (Lu et al., 2022; Xu et al., 2023c; Zhao et al., 2023; Cheng et al., 2023c; Wang et al., 2023a; Zhang et al., 2024b; Li et al., 2024), Information Retrieval (IR) (Zhu et al., 2024a; Liu et al., 2024), and Spoken Language Understanding (SLU) (Hoscilowicz et al., 2024; Zhu et al., 2024c,b; Yin et al., 2024; Cheng et al., 2023a, 2024; Dong et al., 2023a). By leveraging their interactional features, researchers have adeptly applied SFT-refined LLMs to formulate commands that link with external tools, markedly enhancing the handling of a range of downstream applications (Shen et al., 2023c; Yao et al., 2023b; Song et al., 2024; Yao et al., 2023a; Fu et al., 2024).

Several studies have adopted alternative approaches; for instance, Taori et al. (2023); Zhou et al. (2023a); Wang et al. (2023b); Lu et al. (2023a) have employed a self-guided technique, utilizing large language models to generate expansive training datasets. Chiang et al. (2023) have focused on training their model using user-shared conversations from ShareGPT.com. Meanwhile, Dong et al. (2024) investigate the data composition challenge across different IFT capabilities in LLMs. Furthermore, Xu et al. (2023b) introduced the Evol-Instruct method, which enhances existing instructional data to cultivate more complex and diversified datasets.

5.2 Large Language Models for Code

Large language models are pre-trained on trillions of tokens following the scaling laws (Hoffmann et al., 2022; Kaplan et al., 2020), which demonstrates remarkable achievements across a broad spectrum of tasks and such an amount of text data is often a diverse composite with a non-negligible part of code (Zhang et al., 2023). Pioneered by Codex, researchers have also found continual pretraining on code to significantly benefit language models’ performance on code. For example, Chowdhery et al. (2022a) post-trains PaLM on 7.8B additional code tokens to get PaLM-Coder and Rozière et al. (2023) train LLaMA 2 (Touvron et al., 2023a) on more than 500B code tokens to acquire Code Llama, which leads open-source models to a new height. For supervised fine-tuning, many works

utilize larger, more capable teacher models to synthesize instruction data to finetune small language models (Mitra et al., 2023; Luo et al., 2023a; Di et al., 2023b). Instead, Muennighoff et al. (2023); Wei et al. (2023a) construct code instructions by mining existing code corpus, which are orthogonal to our method. Since code-supervised signals are easily collected by compiling and running them, reinforcement learning becomes another important branch. Numerous works have attempted to utilize reinforcement learning with feedback information offered by compilation and other sources. PanGu-Coder2 (Shen et al., 2023b) introduces a ranking loss based on unit tests, which helps to align with the capable code model deeply. However, these methods rely on human-annotated unit tests, which limits the application in the practical scenario.

6 Conclusion and Future Work

In this paper, we investigate two fine-tuning methods to improve the LLMs’ performance on code generation. We first introduce a response augmentation strategy by using different ChatGPT system prompts to increase the diversity of code solutions. We find different chain-of-thought reasoning paths improve performance. Then, We adopt a multi-step training approach that combines traditional code generation and code evaluation objectives. We find improving one’s ability to evaluate the correctness of code also enhances their ability to create it. For future work, we aim to explore the effect of our methods on the larger foundation models and parameter-efficient fine-tuning mechanisms. We also plan to increase the accuracy of evaluation signals via other ways like automatic unit tests and reinforcement learning methods.

7 Limitations

Our limitations are three-fold: (1) We only explore our method on the 7B/13B base models due to the computation cost. More experiments on the larger models and other code models should be conducted to confirm our conclusion. (2) We only use GPT-4 to evaluate the quality of generated code solutions. The performance of GPT-4 is still poor and limits the performance of our proposed method. More precise and open-source evaluation models should be explored in future work. (3) There is still room for optimization in our training data. For example, we find continually increasing the number of system prompts only gets marginal performance

gains. Diversity-based compression methods (Lu et al., 2023b; Liu et al., 2023c) may be valuable while the number of system prompts is large.

8 Broader Impacts

Similar to the other LLMs, our DolphCoder could also generate unethical, harmful, or misleading information, which is not considered in our work. Future research to address the ethical and societal implications is needed. DolphCoder is also susceptible to hallucination in ungrounded generation use cases due to its smaller size. This model is solely designed for research settings, and its testing has only been carried out in such environments. It should not be used in downstream applications, as additional analysis is needed to assess potential harm or bias in the proposed application.

9 Acknowledgement

We thank all anonymous reviewers for their helpful comments and suggestions. This work was partially supported by the State Key Laboratory of Massive Personalized Customization System and Technology (No. H&C-MPC-2023-02-07(Q), State Grid Technology Project (5700-202416236A-1-1-ZN) “Research on active semantic discovery technology based on SG-CIM and its application in power grid equipment supply chain optimization” and the National Natural Science Foundation of China (NSFC No.62076031 and No.62076036).

References

- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and Charles Sutton. 2021. *Program synthesis with large language models*. *ArXiv*, abs/2108.07732.
- 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*.
- Sahil Chaudhary. 2023. Code alpaca: An instruction-following llama model for code generation. <https://github.com/sahil280114/codealpaca>.

- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde, Jared Kaplan, Harrison Edwards, Yura 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, David W. Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William H. Guss, Alex Nichol, Igor Babuschkin, S. Arun Balaji, Shantanu Jain, Andrew Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew M. 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](#). *ArXiv*, abs/2107.03374.
- Zhaorun Chen, Zhuokai Zhao, Zhihong Zhu, Ruiqi Zhang, Xiang Li, Bhiksha Raj, and Huaxiu Yao. 2024. Autoprml: Automating procedural supervision for multi-step reasoning via controllable question decomposition. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Xuxin Cheng, Bowen Cao, Qichen Ye, Zhihong Zhu, Hongxiang Li, and Yuexian Zou. 2023a. MI-lmcl: Mutual learning and large-margin contrastive learning for improving asr robustness in spoken language understanding. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6492–6505.
- Xuxin Cheng, Qianqian Dong, Fengpeng Yue, Tom Ko, Mingxuan Wang, and Yuexian Zou. 2023b. M 3 st: Mix at three levels for speech translation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Xuxin Cheng, Zhihong Zhu, Hongxiang Li, Yaowei Li, Xianwei Zhuang, and Yuexian Zou. 2024. Towards multi-intent spoken language understanding via hierarchical attention and optimal transport. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17844–17852.
- Xuxin Cheng, Zhihong Zhu, Wanshi Xu, Yaowei Li, Hongxiang Li, and Yuexian Zou. 2023c. Accelerating multiple intent detection and slot filling via targeted knowledge distillation. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022a. [Palm: Scaling language modeling with pathways](#).
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam M. Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Benton C. Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathleen S. Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022b. [Palm: Scaling language modeling with pathways](#). *J. Mach. Learn. Res.*, 24:240:1–240:113.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Peng Di, Jianguo Li, Hang Yu, Wei Jiang, Wenting Cai, Yang Cao, Chaoyu Chen, Dajun Chen, Hongwei Chen, Liang Chen, Gang Fan, Jie Gong, Zi Gong, Wen Hu, Tingting Guo, Zhichao Lei, Ting Li, Zheng Li, Ming Liang, Cong Liao, Bingchang Liu, Jiachen Liu, Zhiwei Liu, Shaojun Lu, Mingquan Shen, Guangpei Wang, Huan Wang, Zhi Yu Wang, Zhaogui Xu, Jiawei Yang, Qing Ye, Gehao Zhang, Yu Zhang, Zelin Zhao, Xunjin Zheng, Hailian Zhou, Lifu Zhu,

- and Xianying Zhu. 2023a. [Codefuse-13b: A pre-trained multi-lingual code large language model](#). *ArXiv*, abs/2310.06266.
- Peng Di, Jianguo Li, Hang Yu, Wei Jiang, Wenting Cai, Yang Cao, Chaoyu Chen, Dajun Chen, Hongwei Chen, Liang Chen, et al. 2023b. Codefuse-13b: A pretrained multi-lingual code large language model. *arXiv preprint arXiv:2310.06266*.
- Guanting Dong, Tingfeng Hui, Zhuoma GongQue, Jinxu Zhao, Daichi Guo, Gang Zhao, Keqing He, and Weiran Xu. 2023a. [Demonsf: A multi-task demonstration-based generative framework for noisy slot filling task](#).
- Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei Wang, Zheng Yuan, Chang Zhou, and Jingren Zhou. 2024. [How abilities in large language models are affected by supervised fine-tuning data composition](#).
- Guanting Dong, Jinxu Zhao, Tingfeng Hui, Daichi Guo, Wenlong Wan, Boqi Feng, Yueyan Qiu, Zhuoma Gongque, Keqing He, Zechen Wang, et al. 2023b. Revisit input perturbation problems for llms: A unified robustness evaluation framework for noisy slot filling task. *arXiv preprint arXiv:2310.06504*.
- Dayuan Fu, Jianzhao Huang, Siyuan Lu, Guanting Dong, Yejie Wang, Keqing He, and Weiran Xu. 2024. [Preact: Predicting future in react enhances agent’s planning ability](#).
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujia Yang, Minlie Huang, Nan Duan, and Weizhu Chen. 2024. [Tora: A tool-integrated reasoning agent for mathematical problem solving](#).
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Jakub Hoscilowicz, Pawel Pawlowski, Marcin Skorupa, Marcin Sowański, and Artur Janicki. 2024. [Large language models for expansion of spoken language understanding systems to new languages](#).
- Yupeng Hou, Junjie Zhang, Zihan Lin, Hongyu Lu, Ruobing Xie, Julian McAuley, and Wayne Xin Zhao. 2024. Large language models are zero-shot rankers for recommender systems. In *European Conference on Information Retrieval*, pages 364–381. Springer.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hananeh Hajishirzi. 2020. [UNIFIEDQA: Crossing format boundaries with a single QA system](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online. Association for Computational Linguistics.
- Chengpeng Li, Zheng Yuan, Hongyi Yuan, Guanting Dong, Keming Lu, Jiancan Wu, Chuanqi Tan, Xiang Wang, and Chang Zhou. 2023a. [Query and response augmentation cannot help out-of-domain math reasoning generalization](#).
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umapathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, Zhiruo Wang, Rudra Murthy, Jason Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nourhan Fahmy, Urvashi Bhat-tacharyya, W. Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jana Ebert, Tri Dao, Mayank Mishra, Alexander Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean M. Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2023b. [Starcoder: may the source be with you!](#) *ArXiv*, abs/2305.06161.
- Xiaoxi Li, Jiajie Jin, Yujia Zhou, Yuyao Zhang, Peitian Zhang, Yutao Zhu, and Zhicheng Dou. 2024. [From matching to generation: A survey on generative information retrieval](#).
- Yujia Li, David H. 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 Jaymin Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. 2022. [Competition-level code generation with alphacode](#). *Science*, 378:1092 – 1097.
- Jiate Liu, Yiqin Zhu, Kaiwen Xiao, Qiang Fu, Xiao Han, Wei Yang, and Deheng Ye. 2023a. [Rlrf: Reinforcement learning from unit test feedback](#). *ArXiv*, abs/2307.04349.

- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023b. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *arXiv preprint arXiv:2305.01210*.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. 2023c. What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning.
- Zheng Liu, Yujia Zhou, Yutao Zhu, Jianxun Lian, Chaozhuo Li, Zhicheng Dou, Defu Lian, and Jian-Yun Nie. 2024. Information retrieval meets large language models. In *Companion Proceedings of the ACM on Web Conference 2024, WWW '24*, page 1586–1589, New York, NY, USA. Association for Computing Machinery.
- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, and Chang Zhou. 2023a. #instag: Instruction tagging for diversity and complexity analysis. *arXiv preprint arXiv:2308.07074*.
- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, and Jingren Zhou. 2023b. #instag: Instruction tagging for analyzing supervised fine-tuning of large language models.
- Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu Lin, Xianpei Han, Le Sun, and Hua Wu. 2022. Unified structure generation for universal information extraction. *arXiv preprint arXiv:2203.12277*.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023a. Wizardcoder: Empowering code large language models with evol-instruct. *ArXiv*, abs/2306.08568.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023b. Wizardcoder: Empowering code large language models with evol-instruct. *arXiv preprint arXiv:2306.08568*.
- Arindam Mitra, Luciano Del Corro, Shweti Mahajan, Andres Codas, Clarisse Simoes, Sahaj Agarwal, Xuxi Chen, Anastasia Razdaibiedina, Erik Jones, Kriti Agarwal, Hamid Palangi, Guoqing Zheng, Corby Rosset, Hamed Khanpour, and Ahmed Awadallah. 2023. Orca 2: Teaching small language models how to reason.
- Niklas Muennighoff, Qian Liu, Qi Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, and S. Longpre. 2023. Octopack: Instruction tuning code large language models. *ArXiv*, abs/2308.07124.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Hassan Awadallah. 2023. Orca: Progressive learning from complex explanation traces of gpt-4. *ArXiv*, abs/2306.02707.
- Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2023. Codegen: An open large language model for code with multi-turn program synthesis.
- OpenAI. 2023. Gpt-4 technical report. *ArXiv*, abs/2303.08774.
- 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.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2023. Exploring the limits of transfer learning with a unified text-to-text transformer.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, I. Evtimov, Joanna Bitton, Manish P Bhatt, Cristian Cantón Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. Code llama: Open foundation models for code. *ArXiv*, abs/2308.12950.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, 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. 2023. Code llama: Open foundation models for code.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Tali Bers, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. Multi-task prompted training enables zero-shot task generalization.
- Bo Shen, Jiaxin Zhang, Taihong Chen, Daoguang Zan, Bing Geng, An Fu, Muhan Zeng, Ailun Yu, Jichuan Ji, Jingyang Zhao, Yuenan Guo, and Qianxiang Wang. 2023a. Pangu-coder2: Boosting large language models for code with ranking feedback. *ArXiv*, abs/2307.14936.

- Bo Shen, Jiaxin Zhang, Taihong Chen, Daoguang Zan, Bing Geng, An Fu, Muhan Zeng, Ailun Yu, Jichuan Ji, Jingyang Zhao, Yuenan Guo, and Qianxiang Wang. 2023b. [Pangu-coder2: Boosting large language models for code with ranking feedback](#).
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023c. [Hugging-gpt: Solving ai tasks with chatgpt and its friends in hugging face](#).
- Xiaoshuai Song, Keqing He, Pei Wang, Guanting Dong, Yutao Mou, Jingang Wang, Yunsen Xian, Xunliang Cai, and Weiran Xu. 2023. [Large language models meet open-world intent discovery and recognition: An evaluation of chatgpt](#).
- Xiaoshuai Song, Zhengyang Wang, Keqing He, Guanting Dong, Yutao Mou, Jinxu Zhao, and Weiran Xu. 2024. [Knowledge editing on black-box large language models](#).
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Stanford alpaca: An instruction-following llama model](#). https://github.com/tatsu-lab/stanford_alpaca.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023a. [Llama 2: Open foundation and fine-tuned chat models](#).
- Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony S. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *ArXiv*, abs/2307.09288.
- Shuhe Wang, Xiaofei Sun, Xiaoya Li, Rongbin Ouyang, Fei Wu, Tianwei Zhang, Jiwei Li, and Guoyin Wang. 2023a. [Gpt-ner: Named entity recognition via large language models](#).
- Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A Smith, Iz Beltagy, et al. 2023b. [How far can camels go? exploring the state of instruction tuning on open resources](#). *arXiv preprint arXiv:2306.04751*.
- Yue Wang, Weishi Wang, Shafiq Joty, and Steven C. H. Hoi. 2021. [Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation](#).
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. [Finetuned language models are zero-shot learners](#).
- Yuxiang Wei, Zhe Wang, Jiawei Liu, Yifeng Ding, and Lingming Zhang. 2023a. [Magicoder: Source code is all you need](#). *ArXiv*, abs/2312.02120.
- Yuxiang Wei, Zhe Wang, Jiawei Liu, Yifeng Ding, and Lingming Zhang. 2023b. [Magicoder: Source code is all you need](#).
- Shaohua Wu, Xudong Zhao, Tong Yu, Rongguo Zhang, Chong Shen, Hongli Liu, Feng Li, Hong Zhu, Jiangang Luo, Liang Xu, et al. 2021. [Yuan 1.0: Large-scale pre-trained language model in zero-shot and few-shot learning](#). *arXiv preprint arXiv:2110.04725*.
- Yuxiang Wu, Guanting Dong, and Weiran Xu. 2023. [Semantic parsing by large language models for intricate updating strategies of zero-shot dialogue state tracking](#).
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023a. [Wizardlm: Empowering large language models to follow complex instructions](#). *ArXiv*, abs/2304.12244.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023b. [Wizardlm: Empowering large language models to follow complex instructions](#).
- Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, and Enhong Chen. 2023c. [Large language models for generative information extraction: A survey](#).

- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023a. [Tree of thoughts: Deliberate problem solving with large language models](#).
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023b. [React: Synergizing reasoning and acting in language models](#).
- Shangjian Yin, Peijie Huang, Yuhong Xu, Haojing Huang, and Jiatian Chen. 2024. [Do large language model understand multi-intent spoken language ?](#)
- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. 2023a. [Scaling relationship on learning mathematical reasoning with large language models](#). *arXiv preprint arXiv:2308.01825*.
- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. 2023b. [Scaling relationship on learning mathematical reasoning with large language models](#).
- Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. 2023c. [Rrhf: Rank responses to align language models with human feedback without tears](#).
- Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhua Chen. 2023. [Mammoth: Building math generalist models through hybrid instruction tuning](#). *arXiv preprint arXiv:2309.05653*.
- Xiang Yue, Tuney Zheng, Ge Zhang, and Wenhua Chen. 2024. [Mammoth2: Scaling instructions from the web](#).
- Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, and Guoyin Wang. 2024a. [Instruction tuning for large language models: A survey](#).
- Zhen Zhang, Yuhua Zhao, Hang Gao, and Mengting Hu. 2024b. [Linkner: Linking local named entity recognition models to large language models using uncertainty](#).
- Ziyin Zhang, Chaoyu Chen, Bingchang Liu, Cong Liao, Zi Gong, Hang Yu, Jianguo Li, and Rui Wang. 2023. [Unifying the perspectives of nlp and software engineering: A survey on language models for code](#).
- Gang Zhao, Xiaocheng Gong, Xinjie Yang, Guanting Dong, Shudong Lu, and Si Li. 2023. [Demosg: Demonstration-enhanced schema-guided generation for low-resource event extraction](#).
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2023a. [Lima: Less is more for alignment](#). *arXiv preprint arXiv:2305.11206*.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023b. [Instruction-following evaluation for large language models](#).
- Sizhe Zhou, Yu Meng, Bowen Jin, and Jiawei Han. 2024. [Grasping the essentials: Tailoring large language models for zero-shot relation extraction](#). *arXiv preprint arXiv:2402.11142*.
- Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Haonan Chen, Zhicheng Dou, and Ji-Rong Wen. 2024a. [Large language models for information retrieval: A survey](#).
- Zhihong Zhu, Xuxin Cheng, Hongxiang Li, Yaowei Li, and Yuexian Zou. 2024b. [Dance with labels: Dual-heterogeneous label graph interaction for multi-intent spoken language understanding](#). In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining, WSDM '24*, pages 1022–1031, New York, NY, USA. Association for Computing Machinery.
- Zhihong Zhu, Xianwei Zhuang, Yunyan Zhang, Derong Xu, Guimin Hu, Xian Wu, and Yefeng Zheng. 2024c. [Tfcd: Towards multi-modal sarcasm detection via training-free counterfactual debiasing](#). In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*.

A Case Study

Table 6 presents comparative examples of DolphCoder and WizardCoder. In this case, the models are required to determine whether an array can become a non-decreasing sequence after applying a circular right shift operation. From the case, we can observe that WizardCoder-13b generates an obvious logical error, where it only checks if the sequence is non-decreasing without performing any right shift operation. In contrast, DolphCoder accurately simulates the circular right shift operation and correctly identifies the termination condition. Furthermore, it is worth noting that DolphCoder considers more robust boundary input cases, which can be attributed to its training on a code evaluation task and a more diverse training dataset.

B Effect of Code Evaluation Capability of GPT-4

In our study, we utilize GPT-4 to create a training dataset for code evaluation. However, there is a concern regarding whether GPT-4 has the ability to generate a perfect golden label. To obtain the answer we test it on HumanEval. Specifically, we use the DIT model to generate multiple candidate answers for each question. Subsequently, we

<p>Instruction</p>	<pre>def move_one_ball(arr): """ We have an array 'arr' of N integers arr[1], arr[2], ..., arr[N]. The numbers in the array will be randomly ordered. Your task is to determine if it is possible to get an array sorted in non-decreasing order by performing the following operation on the given array: You are allowed to perform right shift operation any number of times. One right shift operation means shifting all elements of the array by one position in the right direction. The last element of the array will be moved to the starting position in the array i.e. 0th index. If it is possible to obtain the sorted array by performing the above operation then return True else return False. If the given array is empty then return True. Note: The given list is guaranteed to have unique elements ... """</pre>
<p>WizardCoder-13b</p>	<p>Here's the Python script for the given problem:</p> <pre>def move_one_ball(arr): if not arr: return True for i in range(len(arr)): if arr[i] > arr[i+1]: return False return True</pre> <p>The function takes an array 'arr' as input and returns 'True' if it is possible to obtain a sorted array in non-decreasing order by performing right shift operations on the given array. Otherwise, it returns 'False'. The function first checks if the given array is empty. If it is empty, it returns 'True' as it is possible to obtain a sorted array in non-decreasing order by performing no right shift operations. Next, the function iterates through the array and checks if the current element is greater than the next element. If it is, then it is not possible to obtain a sorted array in non-decreasing order by performing any number of right shift operations. Therefore, the function returns 'False'. If the function has not returned 'False' by this point, it means that the array is already sorted in non-decreasing order. Therefore, it returns 'True'.</p>
<p>DolphCoder-13b</p>	<p>Here is the Python script for the given problem:</p> <pre>def move_one_ball(arr): if len(arr) == 0: return True if len(arr) == 1: return True if arr == sorted(arr): return True for i in range(len(arr)): arr.append(arr.pop(0)) if arr == sorted(arr): return True return False</pre> <p>Explanation:</p> <ol style="list-style-type: none"> 1. First, we check if the given array is empty or has only one element. If so, we return True as it is already sorted. 2. Next, we check if the given array is already sorted. If it is, we return True. 3. If the array is not sorted, we perform a right shift operation on the array by appending the first element to the end of the array and popping the first element from the beginning. 4. We repeat step 3 until the array is sorted or we have performed all possible right shift operations. 5. If the array is sorted after performing all possible right shift operations, we return True. Otherwise, we return False.

Table 6: Examples of DolphCoder compared to WizardCoder.

leverage GPT-4 to evaluate each of these candidate answers. And then we select the code solutions

that GPT-4 identified as correct. Since the test cases are available in HumanEval, we can get the

Method	Pass@1
DIT	56.6
+MOT	59.9
+GPT-4 Filtering	62.7
+Golden Filtering	72.6

Table 7: Effect of code evaluation capability of GPT-4. We sample 10 answers with a temperature=0.2 and report the average pass@1 metric. +MOT denotes the MOT model based on the DIT model. GPT-4 or golden filtering means that we use GPT-4 or golden unit test cases to filter out these wrong answers from the DIT model, and then we report the average pass@1 among the remaining answers.

binary <passed> or <not passed> accuracy of GPT-4 which is just 79.4%, which indicates that the train data we constructed through GPT-4 may contain about 20% noise. The noise limits the upper-bound performance of the MOT method in this work. Additionally, we report the average pass@1 results before and after filtering by GPT-4, as shown in Table 7. We can observe that the pass@1 before filtering is 56.6% which is generated by the DIT model directly at 0.2 temperature. Then we calculate the pass@1 after golden filtering is 72.6% where we use golden unit tests to filter wrong solutions. We consider this percentage of 72.6% is a theoretical upper bound of GPT-4. We also report pass@1 of MOT and GPT-4 Filtering. We summarize key insights as follows: (1) GPT-4 can not perfectly perform code evaluation (62.7% vs 72.6%). (2) Our MOT can achieve a slightly worse performance than the GPT-4 Filtering (59.9% vs 62.7%). (3) GPT-4 evaluation capability limits the upper bound of MOT. How to improve automatic code evaluation is essential to future work.