

MM-ReCoder: Advancing Chart-to-Code Generation with Reinforcement Learning and Self-Correction

Zitian Tang^{1*} Xu Zhang² Jianbo Yuan² Yang Zou²
 Varad Gunjal² Songyao Jiang² Davide Modolo²
¹Brown University ²Amazon AGI

<https://zitiantang.github.io/MM-ReCoder>

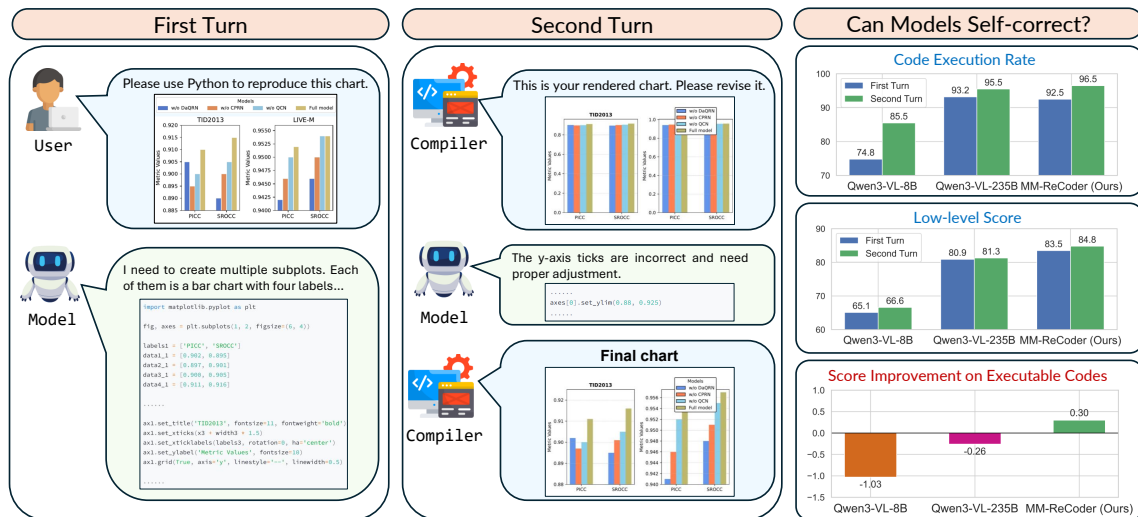


Figure 1. In multi-turn Chart2Code, the model takes the execution result from its first-turn code and revises the code accordingly. Although existing MLLMs improve evaluation scores between the two turns, the gains mainly come from increased code executability. When we restrict the analysis to cases where the first-turn code is already executable, existing models show a negative improvement, whereas our model demonstrates a positive one.

Abstract

Multimodal Large Language Models (MLLMs) have recently demonstrated promising capabilities in multimodal coding tasks such as chart-to-code generation. However, existing methods primarily rely on supervised fine-tuning (SFT), which requires the model to learn code patterns through chart-code pairs but does not expose the model to a code execution environment. Moreover, while self-correction through execution feedback offers a potential route to improve coding quality, even state-of-the-art MLLMs have been shown to struggle with effective self-correction. In this work, we introduce MM-ReCoder, a chart-to-code generation model trained with reinforcement learning (RL) and equipped with self-correction ability. We propose a two-stage multi-turn self-correction RL strategy based on Group Relative Policy Optimization (GRPO). The first stage enhances the model’s self-correction ability via

rolling out a shared first turn, while the second stage improves the coding capability with full-trajectory optimization. MM-ReCoder learns to produce more accurate and executable code through the interaction with the environment and by iteratively correcting its own outputs. Our results on three chart-to-code benchmarks demonstrate the state-of-the-art performance of MM-ReCoder.

1. Introduction

Scientific charts play a crucial role in helping humans interpret complex information by highlighting trends, relationships, and comparisons. However, manually creating these charts through graphical user interfaces or by writing code is a tedious and time-consuming process. To address these traditional limitations in chart generation and unlock powerful new applications, Multimodal LLMs have recently been explored and trained for automatic generation of chart code (Chart2Code task), which allows for the automation of the

*Work completed during internship at Amazon.

visualization process. In detail, in this new task, a chart is provided as input in the form of an image, and the model is expected to generate executable code to reproduce the chart.

Existing Chart2Code approaches [7, 18, 22, 28, 37, 43, 47] typically collect a large corpus of image-code pairs and train the MLLMs via supervised fine-tuning (SFT). While these methods demonstrate promising results, next token prediction alone doesn't guarantee that the generated code has executable syntax and visual details that are faithful to the references [29]. To improve the execution success rate and the visual quality of the rendered content, recent works have begun exploring reinforcement learning (RL) [29, 46]. For instance, ChartMaster [29] introduced a combination of low-level reward (targeting chart elements such as chart type, color, and legends) and high-level rewards (encouraging visual similarity between the input and rendered images), yielding significant gains over SFT baselines. One limitation of these RL-based Chart2Code methods is that they treat the problem as a one-shot generation task, meaning that they train their models to produce the final code in a single pass, without executing or refining it based on any feedback from the rendered result. Humans, conversely, operate iteratively: they implement code, execute it, visualize the results, and refine the code until the output meets expectations. This iterative *self-correcting* process is critical to reach the correct outcome, and in this paper we introduce an RL phase modeled after this process. However, recent studies have shown that existing open-source MLLMs struggle to self-correct on multimodal coding tasks, even when provided with explicit feedback [28, 39].

To mitigate these limitations and help MLLMs successfully improve themselves in an iterative process, we propose **MM-ReCoder**. MM-ReCoder follows a two-stage training curriculum: 1) SFT cold-start with ground-truth chart-code pairs and carefully crafted multi-turn coding data and 2) multi-turn self-correction RL. The cold-start stage provides the model with basic coding knowledge and the ability to generate multi-turn coding responses, but it is not sufficient to guarantee reliable improvements in the second turn. To empower the model with reliable self-correction abilities, we introduce a two-stage multi-turn GRPO (Group Relative Policy Optimization) [25] framework that extends standard GRPO to support self-correction. In the first stage of this RL phase, we fix an online sampled first-turn output for each group and roll out multiple second-turn candidates, allowing the model to explore diverse refinement strategies. In the second stage we then roll out both first-turn and second-turn outputs for each output in the group to further enhance coding performance.

MM-ReCoder outperforms existing similar size Chart2Code models trained with the same dataset. For example, in ChartMimic [39], it outperforms ChartCoder [47] by 9.1% in low-level score and 10.9% in high-level score.

Remarkably, with only 7B parameters, MM-ReCoder achieves performance comparable to much larger models such as Qwen3-VL-235B-A22B [31].

In summary, our contributions are three-fold:

- We propose MM-ReCoder, the first multimodal coding MLLM equipped with self-correction abilities. Trained with a framework containing cold start and two-stage RL, MM-ReCoder achieves state-of-the-art performance on chart-to-code benchmarks [36, 37, 39] with 86.5% low-level score on ChartMimic and 63.2% text-match score on Plot2Code.
- We explore different RL training strategies for multi-turn self-correction. We find that shared-first-turn optimization followed by full-trajectory optimization can elicit the model's self-correction capability while improving its overall coding ability.
- We explore rule-based and model-based rewards. While rule-based reward helps the model faithfully reproduce chart elements, model-based reward effectively improves the visual quality of the reproduced images.

2. Related Works

2.1. Multimodal Coding

In multimodal coding, models generate codes according to given multimodal inputs. For example, in graphics derendering tasks, a model generates Python/HTML code to reproduce a given chart/webpage screenshot. Recently, many benchmarks and datasets have been built for chart-to-code generation [11, 36, 37, 39, 47], webpage-to-code generation [13, 28, 42], and SVG-to-code generation [22].

With the rise of MLLMs, many works have trained domain-specific MLLMs tackling multiple tasks, *e.g.*, chart description, QA, and redrawing [7, 18, 37, 43]. Most of existing works [7, 18, 22, 28, 37, 43, 47] consider Chart2Code problem as SFT problem. For example, ChartCoder [47] collects 160k chart-code pairs and uses them to train an MLLM. Beyond SFT, Zhang et al. [45] proposes dual preference-guided refinement to iteratively collect code preference pairs and refine a model with offline reinforcement learning. ChartMaster [29] first proposes to use GRPO with low-level and high-level rewards to improve chart quality. VinciCoder [46] develops a single model to address all Vision2Code problems via SFT and RL. However, all existing RL-based approaches treat Chart2Code as a one-shot generation task, preventing models from iteratively refining their code to improve generation quality.

2.2. Reinforcement Learning for MLLMs

Reinforcement learning (RL) has recently emerged as an effective paradigm for improving reasoning and alignment capabilities in LLMs. Early methods such as REINFORCE [1], PPO [19, 24], and DPO [21] established the

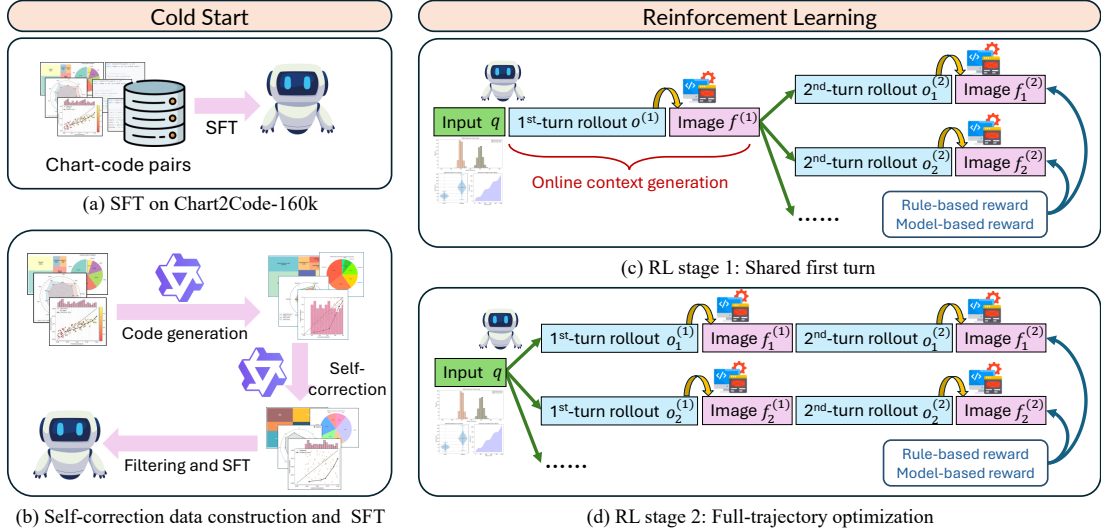


Figure 2. Training pipeline of MM-ReCoder. We first conduct two stages of cold start: (a) we first train the model on ground truth chart-code pairs with SFT, then (b) we construct self-correction data with Qwen3VL-235B [31], filter the successful ones, and train our model on the filtered data. After cold start, we conduct two stages of reinforcement learning: (c) we first enhance the model’s self-correction capability in the second turn via shared-first-turn optimization, then (d) we optimize the two turns jointly to improve the coding ability.

foundation for optimizing LLMs through reward-guided policy gradients. DeepSeek-R1 [4] first demonstrated that RL can substantially improve the reasoning ability of LLMs via GRPO [25] with verifiable rewards. R1-style thinking patterns, *e.g.*, self-verification, emerge during RL training.

Inspired by these advances in LLMs, recent works extend RL techniques to MLLMs to enhance multimodal reasoning [9, 26, 35, 40]. They find cold start essential to reproduce the R1-style thinking patterns on MLLMs. Vision-R1 [9] and R1-Onevision [40] generates multimodal R1-style cold start data by applying DeepSeek-R1 on dense image captions, while OVR [35] shows that R1-style reasoning learned from linguistic cold start can be transferred to multimodal tasks. In contrast, VL-Rethinker [33] appends a rethinking trigger token after the model output to force self-reflection during RL. Moreover, some other works extend RL to video reasoning [5, 15, 38, 44] or to train MLLMs for tool calling [50]. These efforts collectively indicate the effectiveness of RL in improving reasoning in MLLMs.

2.3. Self-correction for LLMs and MLLMs

Self-correction aims to prompt the models to critique and refine its own answers. While training-free prompt-based self-correction is shown to be ineffective for LLMs [8, 20, 32, 48], one way to enable self-correction is SFT on annotated or generated self-correction data [20, 23]. However, Kumar et al. [12] show that self-correction learned by SFT suffers from distribution shift and behavior collapse. They propose SCoRe to elicit the self-correction ability of LLMs with two-stage RL. Moreover, a few works explore self-correction in coding tasks [3, 6, 10]. For example, Co-CoS [3] elicits self-correction by carefully setting the re-

ward discount factor of the first turn. While these methods use REINFORCE [1] and PPO [19, 24] as the RL algorithm, we explore training self-correction with GRPO.

In multimodal coding, ChartMimic [39] and Design2Code [28] demonstrate that only very large MLLMs are able to self-correct their own solutions. We are the first to explore eliciting MLLM self-correction in this domain.

3. Method

3.1. Problem Formulation

In Chart2Code task, given a chart image q , a model is asked to output o , which includes a thinking trace and a Python code, to reproduce input image q . In conventional single-shot generation [7, 18, 37, 43, 47], the output o is directly sampled from the model output distribution $\pi_{\theta}(O|q)$, where θ is the model parameter and O is the output space.

For self-correction, without loss of generality, we only consider one correction turn. In addition to directly sampling the first turn solution $o^{(1)}$ from the model distribution. We execute the code in $o^{(1)}$ to render an output image $I^{(1)}$. The output image will be sent back to MLLM for self-correction. Specifically, we construct the feedback $f^{(1)}$ as “ $I^{(1)}$ Here is the image rendered by your code. Please continue modifying the code.”. If the code is not executable, $f^{(1)}$ becomes “{error information} Your code encountered a runtime error. Please continue modifying the code.” (See the complete prompts in the supplementary). Then the model samples a second-turn output conditioned on the first turn and the feedback $o^{(2)} \sim \pi_{\theta}(O|q, o^{(1)}, f^{(1)})$. This process can be iteratively performed for multiple rounds.

3.2. Self-Correction Ability of Off-the-Shelf Models

We start by testing the self-correction ability of the most recent off-the-shelf MLLMs. We evaluate Qwen3-VL-8B and Qwen3-VL-235B-A22B [31] on ChartMimic [39] in Figure 1 (more details in Table 3). Qwen3-VL-8B improves the overall low-level score by 1.5% between two turns while Qwen3-VL-235B improves it by 0.4%. However, we notice that the code execution rates increase more significantly (10.7% and 2.3%). This suggests that the observed gains in low-level scores may come from executing more charts rather than refining already-executable code.

To validate this hypothesis, we filter the images that are successfully rendered in both turns and compute their average score improvement. As Figure 1 shows, their improvements on executable codes are -1.03% and -0.26% , respectively, indicating that the models are unable to refine charts that were already executable in the first turn. Inspired by this, we propose MM-ReCoder which focuses on iteratively improving the visual quality of the generated image.

3.3. MM-ReCoder

To achieve the self-correction capability, the proposed MM-ReCoder training pipeline has two phases: cold start and multi-turn self-correction RL, illustrated in Figure 2.

In the cold start phase, we first train the model with SFT on Chart2Code-160k [47], which includes 160k chart-code pairs. While this effectively enhances the model’s coding capability, the model loses its multi-turn ability after training on single-turn data (as more than 80% of the second turns are repeating the first turns, see Table 5). Therefore, in the second cold start stage, we construct and train our model on 7k two-turn self-correction data using Qwen3-VL-235B-A22B-Instruct [31] on the images from Chart2Code-160k. Specifically, we first generate a two-turn conversation on each image. Then we run rejection sampling and only keep samples whose second-turn low-level score exceeds the first turn by at least 0.02, a threshold we find sufficient to visually observe chart improvements.

Although this cold-start stage enables the MLLM to produce multi-turn coding responses, it does not reliably ensure improvement in the second turn. To address this, we propose a two-stage multi-turn self-correction RL using GRPO. We first perform shared-first-turn optimization to improve the model’s self-correction capability, then conduct full-trajectory optimization to boost its overall coding performance. We introduce the details in the following.

3.4. Multi-Turn Self-Correction RL

We start by revisiting single-turn on-policy GRPO. For each query q , GRPO samples a group of outputs $\{o_1, o_2, \dots, o_G\}$ following the policy π_θ . The optimization



Figure 3. Result of a model trained solely with rule-based reward. The model receives a full rule-based reward though the texts are overlapped. But the model-based reward can penalize this chart.

objective can be derived as¹

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E} \left[\frac{1}{G} \sum_{i=1}^G \frac{\pi_\theta(o_i|q)}{\text{SG}[\pi_\theta(o_i|q)]} A_i \right], \quad (1)$$

where $\text{SG}[\cdot]$ denotes stop gradient and A_i is the advantage of o_i . A_i can be derived as

$$A_i = \frac{R(o_i) - \text{mean}(\{R(o_1), R(o_2), \dots, R(o_G)\})}{\text{std}(\{R(o_1), R(o_2), \dots, R(o_G)\})}, \quad (2)$$

while $R(\cdot)$ is a reward function taking MLLM output as input to produce a scalar reward. For multi-turn self-correction, we first draw the first turn $o_i^{(1)} \sim \pi_\theta(O|q)$ followed by the second-turn output $o_i^{(2)} \sim \pi_\theta(O|q, o_i^{(1)}, f_i^{(1)})$. Overall, our RL formulation has two key components - the reward function design and the multi-turn roll-out strategy.

3.4.1. Reward Design

Designing rewards for Chart2Code task is inherently challenging since there is no perfect metric that fully captures visual similarity and semantic correctness between two charts. In common practices, one may extract the visual elements (e.g., text and color) from the rendered image and the reference image to calculate rule-based rewards, or use VLMs and MLLMs to evaluate the image similarity [16]. In this work, we leverage both perspectives. For the Chart2Code task, ChartMimic [39] introduces rule-based low-level metrics and a model-based high-level metric for evaluation purpose. We reuse them as our RL rewards.

Rule-based reward. Given the python code for rendering a chart, we can hook into the Matplotlib functions and execute the code to obtain the visual elements. For example, to extract all the texts rendered by the code, we wrap the `RendererPdf.draw_text` function with a custom function that intercepts every text-drawing call, records the text content, and forwards the call to the original renderer. By this approach, we extract the chart type, layout, text, and color from the generated code and ground truth code

¹To simplify the notation, we omit the token-level averaging in GRPO throughout the paper.

and measure their similarity. Specifically, we calculate the F1-score for chart type and text, CIE Lab color distance for color, and compare the number of rows and columns for layout. While each metric is in range $[0, 1]$, we take their average as the overall rule-based reward.

Although the rule-based reward can provide a precise score from visual elements, it fails to measure many aspects of chart quality. If a model is trained solely with rule-based rewards, it may suffer from issues such as overlapping elements, inaccurate data points, or improperly sized fonts. For example, in Figure 3, a chart containing overlapping text can still receive a rule-based reward of 1.

Model-based reward. In contrast to rule-based rewards, the model-based reward measures both visual and semantic similarity. While ChartMimic [39] calls GPT-4o to score the similarity between two charts, it is computationally expensive for large-scale RL. Instead, we employ an open-source MLLM, Qwen2.5-VL-72B [30], as the reward model. Given a reproduced image and the reference image, we ask the model to evaluate the reproduced image from six aspects: chart type, layout, text, data, style, and clarity. Each has a maximum score of either 10 or 20 points, leading to a total score of 100. The total score is scaled to the range of $[0, 1]$ and used as the model-based reward.

Finally, to encourage the model to think before coding, we include a **Format reward**, which is 1 if the model output follows `<think>...</think>```python...````, otherwise 0. We combine the rewards by $(1 - \alpha - \beta) * \text{Format} + \alpha * \text{Rule-based} + \beta * \text{Model-based}$, where α and β are hyperparameters. Ablation studies on the rewards weights are in Section 4.5.

3.4.2. Roll-out Strategies for Self-correction

While self-correction can be performed for multiple rounds, we focus on training two-turn self-correction during RL, *i.e.*, we only optimize the model’s first-turn output $o^{(1)}$ and its refinement $o^{(2)}$. At inference time, however, the model can perform an arbitrary number of self-correction turns. To train a multi-turn model, we consider two roll-out strategies:

Full-Trajectory Optimization. As shown in Figure 2(d), we jointly optimize both the first and second turns in each trajectory. For each input image, we independently roll out G trajectories with both the first- and second-turn generations. Specifically, we first sample $o_i^{(1)} \sim \pi_\theta(O|q)$ and then $o_i^{(2)} \sim \pi_\theta(O|q, o_i^{(1)}, f_i^{(1)})$. The reward is computed over the entire trajectory, and the policy is updated to maximize

$$\mathcal{J}_{GRPO}^{(full)}(\theta) = \mathbb{E} \left[\frac{1}{G} \sum_{i=1}^G \frac{\pi_\theta(o_i^{(1)}|q)\pi_\theta(o_i^{(2)}|q, o_i^{(1)}, f_i^{(1)})}{\text{SG}[\pi_\theta(o_i^{(1)}|q)\pi_\theta(o_i^{(2)}|q, o_i^{(1)}, f_i^{(1)})]} A_i \right]. \quad (3)$$

As there are two turns, besides the final turn, we can also use the first turn in reward calculation. Additionally, we can further include a boosting reward to encourage the model to

generate improved second turn compared to the first turn. Specifically, the reward for a trajectory $(o_i^{(1)}, o_i^{(2)})$ is

$$R(o_i^{(1)}, o_i^{(2)}) = R(o_i^{(2)}) + \gamma \cdot R(o_i^{(1)}) + \eta \cdot \mathbb{B}(R(o_i^{(1)}), R(o_i^{(2)})), \quad (4)$$

where γ is a hyperparameter to balance the rewards of the two turns, η is a hyperparameter to weight the boosting reward $\mathbb{B}(R(o_i^{(1)}), R(o_i^{(2)})) = \mathbb{1}(R(o_i^{(2)}) > R(o_i^{(1)}))$. $\mathbb{1}(\cdot)$ is indicator function to encourage the second-turn reward to be higher than the first-turn reward.

Shared First Turn. As Figure 2(c) shows, for each input, we first roll out a shared first-turn output $o^{(1)} \sim \pi_\theta(O|q)$ with feedback $f^{(1)}$. We then generate multiple second-turn refinements $o_i^{(2)} \sim \pi_\theta(O|q, o^{(1)}, f^{(1)})$ conditioned on this shared first turn. The optimization objective is

$$\mathcal{J}_{GRPO}^{(shared)}(\theta) = \mathbb{E} \left[\frac{1}{G} \sum_{i=1}^G \frac{\pi_\theta(o_i^{(2)}|q, o^{(1)}, f^{(1)})}{\text{SG}[\pi_\theta(o_i^{(2)}|q, o^{(1)}, f^{(1)})]} A_i \right]. \quad (5)$$

Note that we treat the first turn as shared context. As a result, RL optimization is applied only to the second turn. Interestingly, we observe that the model’s first-turn capability also improves during this training process. More importantly, the first-turn context is sampled online during training, which gradually increases the difficulty of generating the second-turn outputs as the model improves.

Two-Stage Training. In practice, we find that combining the shared-first-turn strategy with the full-trajectory strategy yields the best performance. Specifically, we begin with the shared first-turn roll-out, which encourages the model to develop strong self-correction capabilities by focusing on diverse second-turn refinements. To further improve overall performance, we introduce the second stage that uses the full-trajectory roll-out, allowing the model to jointly optimize both first-turn and second-turn outputs. For reward computation, we observe that calculating the reward only at the final turn (*i.e.*, $\gamma = 0$) leads to the best performance.

4. Experiments

4.1. Implementation

We use Qwen2.5-VL-7B [30] as the base model of MM-ReCoder. We train it on Chart2Code-160k [47] in both the cold start and the reinforcement learning phases. For cold start, we train the model for one epoch on Chart2Code-160k and two epochs on the constructed self-correction data. We use the LLaMA-Factory framework [49] with a batch size of 128 and learning rate of 10^{-5} on 8 NVIDIA H200s.

For reinforcement learning, we use the VeRL framework [27]. We use GRPO with a group size of $G = 8$ to train our model for one epoch in each RL stage. The batch size is 128 and the learning rate is 10^{-6} . The maximum response length for each turn is 4,096 tokens. Training is conducted on 16×8 NVIDIA H200s with another 4×8 H200s

Model	# Turns	ChartMimic			Plot2Code			ChartX
		Exec.Rate	Low-Level	High-Level (GPT-4o)	Pass Rate	Text-Match	Rating (GPT-4V/4o)	GPT-score (GPT-4/4o)
Full score	-	100	100	100	100	100	10	5
<i>Proprietary</i>								
GeminiProVision	1	68.2	53.8	53.3	68.2	53.6	3.69 / -	-
Claude-3-opus	1	83.3	60.5	60.1	84.1	57.5	3.80 / -	-
GPT-4V	1	91.2	76.4	78.9	84.1	57.7	5.58 / -	2.63 / -
GPT-4o	1	93.2	79.0	83.5	88.6	56.3	5.71 / -	-
GPT-4o (reproduced)	1	94.0	81.8	83.7	82.6	59.8	- / 5.42	- / 2.08
GPT-4o (reproduced)	2	96.7	82.9	86.8	90.2	58.4	- / 6.12	- / 2.07
<i>Open-Source General-Domain</i>								
DeepSeek-VL-7B [17]	1	41.3	19.0	17.6	64.4	32.6	2.26 / -	-
LLaVA-Next-Mistral-7B [14]	1	59.7	20.7	21.3	72.0	38.7	2.87 / -	-
MiniCPM-Llama3-V2.5 [41]	1	80.3	36.6	42.1	76.3	37.3	2.61 / -	1.66 / -
InternVL2-8B [2]	1	61.8	34.4	38.9	77.3	37.1	2.78 / -	1.63 / -
InternVL2-26B [2]	1	69.3	41.4	47.4	81.3	43.1	3.42 / -	1.70 / -
InternVL2-Llama3-76B [2]	1	83.2	54.8	62.2	85.6	46.6	3.89 / -	1.74 / -
Qwen2-VL-7B [34]	1	67.0	32.9	35.0	68.2	33.8	3.10 / -	1.50 / -
Qwen2-VL-72B [34]	1	73.3	54.4	50.9	72.0	53.4	4.26 / -	1.69 / -
Qwen2.5-VL-7B [30]	1	71.5	56.2	49.6	75.0	47.8	- / 3.51	- / 1.99
Qwen2.5-VL-7B [30]	2	74.3	57.2	50.5	78.5	43.8	- / 3.25	- / 1.96
Qwen2.5-VL-72B [30]	1	89.5	72.3	74.7	82.6	55.0	- / 4.67	- / 2.28
Qwen2.5-VL-72B [30]	2	93.2	75.9	77.6	87.1	55.4	- / 4.96	- / 2.26
Qwen3-VL-8B [31]	1	74.8	65.1	66.2	73.5	49.0	- / 4.11	- / 2.31
Qwen3-VL-8B [31]	2	85.5	66.6	71.1	84.1	46.8	- / 4.31	- / 2.15
Qwen3-VL-235B-A22B [31]	1	93.2	80.9	85.9	90.2	60.9	- / 5.93	- / 2.65
Qwen3-VL-235B-A22B [31]	2	95.5	81.3	87.2	91.7	60.4	- / 5.79	- / 2.59
<i>Open-Source Chart-Domain</i>								
ChartLlama [7]	1	57.5	24.8	28.1	58.4	40.3	2.32 / -	0.94 / -
TinyChart [43]	1	42.5	26.3	25.9	43.2	44.6	2.19 / -	1.89 / -
ChartVLM-L [37]	1	19.5	15.8	13.9	-	-	-	1.58 / 0.99
ChartCoder [47]	1	91.4	77.4	74.0	87.9	54.5	4.50 / 3.45	2.09 / 2.08
MM-ReCoder-Single (Ours)	1	95.0	84.3	83.7	89.4	62.8	- / 5.10	- / 2.24
MM-ReCoder (Ours)	1	92.5	83.5	81.2	88.6	63.2	- / 4.62	- / 2.31
MM-ReCoder (Ours)	2	96.5	84.8	84.7	97.7	62.2	- / 5.17	- / 2.32
MM-ReCoder (Ours)	4	97.5	86.5	84.9	98.5	62.7	- / 5.24	- / 2.31

Table 1. Evaluation results on the Chart2Code task. MM-ReCoder achieves the best low-level score on ChartMimic and Text-match score on Plot2Code. On the other evaluation metrics, MM-ReCoder outperforms models of comparable size significantly.

servicing Qwen2.5-VL-72B as reward model. The two RL stages take 61 and 73 hours, respectively.

4.2. Benchmarks and Baselines

Following prior work [47], we evaluate MM-ReCoder on ChartMimic [39], Plot2Code [36], and ChartX [37]. We report the evaluation metrics defined by each benchmark. For ChartMimic, we include execution rate, the low-level score, and the high-level score. For Plot2Code, we report pass rate, text-matching score, and the GPT-4 rating. For ChartX, we report the GPT-score. Because the GPT-4 and GPT-4V checkpoints used in prior Plot2Code and ChartX evaluations are no longer available through the OpenAI API, we use GPT-4o instead and evaluate several existing

models as updated baselines. Please refer to the corresponding benchmark papers for detailed metric definitions.

We evaluate MM-ReCoder, GPT-4o, Qwen2.5-VL, and Qwen3-VL under both single-turn and multi-turn self-correction settings. Due to computational constraints, other baselines are evaluated only in the single-turn setting. Moreover, we develop a single-turn RL baseline MM-ReCoder-Single. It uses the same backbone and cold start as MM-ReCoder but RL is conducted in a single-turn setting.

4.3. Main Results

Table 1 shows the results of MM-ReCoder and baselines on Chart2Code tasks. MM-ReCoder not only significantly outperforms chart-domain models and models of compara-

Strategy	γ	η	First turn		Second turn		Avg. low-level improvement on rendered charts	Percentage of improved samples	Percentage of degraded samples	Percentage of repeated code
			Exec.Rate	Low-Level	Exec.Rate	Low-Level				
Full-trajectory	0	0	91.8	81.8	95.2	83.9	0.21	3.4	2.6	46.9
	0	0.1	64.5	66.7	96.0	84.3	10.11	87.5	2.3	0.3
	0.5	0.1	96.5	85.0	97.5	85.5	-0.03	2.2	2.2	55.6
Turn-wise	-	-	94.5	82.7	96.8	83.4	0.05	2.7	1.6	63.0
Shared-first-turn	-	-	91.7	79.8	95.8	82.6	0.72	14.4	8.4	16.8
Shared-first-turn + Full-trajectory	0	0	94.3	83.7	97.7	86.0	0.55	12.1	8.2	21.6
	0.5	0.1	96.5	83.7	96.3	83.5	-0.03	1.8	2.7	81.1

Table 2. Comparison of RL strategies for self-correction. Our two-stage strategy enables self-correction capability while the others cannot.

Model	Low-level			High-level		
	Average improve	Improved%	Degraded%	Average improve	Improved% @0.1	Degraded% @0.1
Qwen2.5-VL-7B	-0.36	8.1	13.9	-0.85	8.9	10.7
Qwen2.5-VL-72B	0.22	15.9	14.5	0.20	11.2	10.2
Qwen3-VL-8B	-1.03	6.0	14.5	-0.70	6.8	8.3
Qwen3-VL-235B-A22B	-0.26	13.4	13.5	0.27	5.4	4.8
GPT-4o	0.85	22.4	12.0	1.76	11.2	4.9
MM-ReCoder (Ours)	0.30	7.3	4.5	0.89	7.1	4.3

Table 3. Self-correction capability of MM-ReCoder and baselines.

Turn	Exec. Rate	Low-level	High-level	Low-level		High-level	
				Avg. improve	Improved%	Avg. improve	Improved% @0.1
1	92.5	83.5	81.2	-	-	-	-
2	96.5	84.8	84.7	0.30	7.3	0.89	7.1
3	96.5	85.3	84.8	0.12	3.8	0.22	3.7
4	97.5	86.5	84.9	0.10	2.6	-0.1	2.6
5	98.3	86.4	85.5	-0.12	2.4	0.26	3.7

Table 4. MM-ReCoder iteratively self-corrects for multiple turns.

ble size, but also surpasses GPT-4o and Qwen3-VL-235B-A22B on ChartMimic low-level score and Plot2Code text-match score. Compared to our base model Qwen2.5-VL-7B, even with only one turn, on ChartMimic, MM-ReCoder improves the execution rate by 29%, the low-level score by 48%, and the high-level score by 63%.

Adding self-reflection further improves the model performance. By adding one correction turn, on ChartMimic, the execution rate can improve from 95% to 97.7% and the low-level improves from 84.3% to 86%. Further increasing the number of self-correction turns can further improve the model performance. For a total of 4 reflection turns, it achieves 86.5% low-level score and 84.9% high-level score.

Deep dive into self-correction ability. To further understand the self-correction behavior of each model, we compute the following metrics on ChartMimic in Table 3.

- **Average score improvement** on charts whose first- and second-turn solutions are both executable.
- **Percentage of improved samples** between two turns. Due to the stochasticity of the high-level score, we calculate the percentage of charts whose high-level scores are improved by at least 10%.
- **Percentage of degraded samples**, similar to the above but for score degradation.

We find that while Qwen2.5-VL and Qwen3-VL raise the scores for a subset of charts, a comparable proportion of samples exhibit score degradation, leading to an overall gain

that is near zero or even negative. On the contrary, MM-ReCoder improves 7% of the charts while only 4% are degraded, resulting in 0.3% improvement over low-level score and 0.89% improvement over high-level score compared with charts rendered in the first turn.

Beyond two turns, we evaluate the model for up to five turns to assess whether it can continue improving image quality. As shown in Table 4, MM-ReCoder improves its outputs over the first few turns, but both the proportion of improved samples and the magnitude of score gains gradually diminish. By the fifth turn, the model is no longer able to further increase the low-level score.

Qualitative results. We illustrate self-correction examples of MM-ReCoder in Figure 4. Our model can revise the rendered texts, axis ranges, chart styles, *etc.* between two turns.

4.4. Analysis of Self-correction Training Strategies

We compare all three RL training strategies introduced in Section 3.4.2, namely, full-trajectory optimization, shared first turn, and two-stage training. For the full-trajectory optimization, we also explore different first turn reward factor γ and self-correction bonus η . In addition, we evaluate a turn-wise training strategy in which, instead of applying shared-first-turn across all samples, 50% of the samples are used to perform single-turn RL without self-correction. This enables the model to learn both first- and second-turn behavior.

To accelerate experimentation, we use only the format reward and rule-based reward, excluding the model-based reward. We evaluate whether their self-correction can improve the low-level score on ChartMimic [39].

The results are in Table 2. All the strategies can improve the execution rate. However, when looking into the low-level score, the model behaviors are different. Under full-trajectory optimization with $\gamma = \eta = 0$, the model repeats its first-turn solution in the second turn in 46.9% of the cases and only 3.4% of charts show improvement. Introducing a self-correction bonus $\eta = 0.1$ causes the model to exploit the reward by intentionally producing extremely poor first-turn outputs. Setting $\gamma = 0.5$ partially stabilizes this behavior by allowing the first-turn score to contribute to the final reward. However, its self-correction capability, indicated by its average low-level improvement of -0.03%

Strategy	First turn low-level	Second turn low-level	Avg. improvement on rendered charts	Percentage of repeated code
Qwen2.5-VL-7B	56.2	57.2	-0.36	10.9
+ Single-turn cold start	79.1	78.9	-0.10	81.5
+ Multi-turn cold start	75.2	73.5	-0.54	2.3
+ RL	83.5	84.8	0.30	2.2

Table 5. Performance after each stage. Cold start improves coding ability, but cannot directly enable self-correction. Multi-turn cold start recovers the model’s multi-turn ability but degrades the performance because the data in this stage is not ground truth.

Reward weights			ChartMimic		
Format	Rule-based	Model-based	Exec.Rate	Low-level	High-level
0.1	0.9	0.0	95.0	84.8	78.6
0.1	0.8	0.1	95.0	84.3	83.7
0.1	0.6	0.3	94.3	84.1	83.7
0.1	0.4	0.5	95.5	84.2	85.3
0.1	0.2	0.7	93.3	83.9	83.8
0.1	0.0	0.9	95.5	78.2	83.8

Table 6. Ablation on RL reward weights. Model training and inference are single-turn without self-correction.

and code repetition rate of 55.6%, is diminished.


We attribute the code repetition behavior of the full-trajectory optimization to the incapability of the base model in self-correction. As shown in Table 5, the average score improvement is negative even after cold start. As a result, when the first turn output is executable, the expected reward gain of generating a different solution is negative, which encourages the model to repeat its solution. This highlights the need to disentangle the training of the two turns.

With this in mind, we evaluate turn-wise training and shared-first-turn. For turn-wise training, the model repeats its first-turn output in 63% of samples and shows minimal self-correction capability (average improvement of just 0.05%). In contrast, shared-first-turn demonstrates strong self-correction ability, despite slightly lower total low-level scores. It improves the chart in 14.4% of cases, with an average low-level improvement of 0.72%.


The proposed two-stage training first uses second-turn-only optimization to elicit the model’s self-correction ability and then employs full-trajectory optimization to enhance its overall coding skill. The resulting model retains strong self-correction ability while achieving the highest low-level score (86%) among all strategies. Further introducing γ and η in the second stage does not improve the performance.

4.5. Analysis of RL Rewards


We conduct ablation study on the effectiveness of the rule-based reward and the model-based reward in RL training. For clarity, we fix the format reward weight to 0.1 throughout our experiments. Results in Table 6 show that, by removing the model-based reward, the high-level score of the model drops significantly from 83.7% to 78.6%, causing the model to generate charts with poor clarity (Figure 3). By removing the rule-based score, the low-level score of the



(a) The 1st-round chart captures the overall structure and proportions but introduces label placement shifts, while the 2nd-round correction refines these stylistic details to more closely match the ground-truth.



(b) The 1st-round output preserves the overall ordering and bar magnitudes but exhibits misaligned axes and spacing. The 2nd-round self-correction refines these details, matching ground-truth more closely.



(c) The 1st-round output matches the grouping and bar heights but introduces color-style mismatches. The 2nd-round self-correction restores the correct colors and improves alignment with the ground-truth.

Figure 4. Qualitative results comparing ground truth, 1st-round generation, and 2nd-round self-correction (left to right).

model drops from 83.9% to 78.2%, indicating unfaithfully reproduced visual elements. Interestingly, we observe that the exact weight of each reward component has limited impact on the final performance. Even a small reward weight provides sufficient signal for the model to fit, and enlarging the weight cannot facilitate the learning of the reward.

5. Conclusions

In this paper, we introduce MM-ReCoder, the first chart-to-code generation model equipped with self-correction capability. To enable self-correction, we propose a two-stage self-correction RL strategy. In the first stage, we roll out a shared first turn and explicitly require the model to self-correct in the second turn. In the second stage, we perform full-trajectory optimization to jointly improve the model’s coding capability in both turns. Our results on three benchmarks show that MM-ReCoder achieves the best performance among chart-domain models and models of comparable size. It even outperforms state-of-the-art MLLMs on some metrics. MM-ReCoder shows better self-correction ability than all open-source MLLMs that we evaluated.

References

- [1] Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learning from human feedback in llms. In *ACL*, 2024. 2, 3
- [2] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*, 2024. 6
- [3] Jeonghun Cho, Deokhyung Kang, Hyounghun Kim, and Gary Geunbae Lee. Self-correcting code generation using small language models. In *Findings of EMNLP*, 2025. 3
- [4] DeepSeek-AI. DeepSeek-R1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 3
- [5] Kaituo Feng, Kaixiong Gong, Bohao Li, Zonghao Guo, Yibing Wang, Tianshuo Peng, Benyou Wang, and Xiangyu Yue. Video-R1: Reinforcing video reasoning in mllms. In *NeurIPS*, 2025. 3
- [6] Jonas Gehring, Kunhao Zheng, Jade Copet, Vegard Mella, Quentin Carbonneaux, Taco Cohen, and Gabriel Synnaeve. RLEF: Grounding code llms in execution feedback with reinforcement learning. In *ICML*, 2025. 3
- [7] Yucheng Han, Chi Zhang, Xin Chen, Xu Yang, Zhibin Wang, Gang Yu, Bin Fu, and Hanwang Zhang. ChartLlama: A multimodal llm for chart understanding and generation. *arXiv preprint arXiv:2311.16483*, 2023. 2, 3, 6
- [8] Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. In *ICLR*, 2026. 3
- [9] Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. Vision-R1: Incentivizing reasoning capability in multimodal large language models. In *ICLR*, 2026. 3
- [10] Arnav Kumar Jain, Gonzalo Gonzalez-Pumariiega, Wayne Chen, Alexander M Rush, Wenting Zhao, and Sanjiban Choudhury. Multi-turn code generation through single-step rewards. In *ICML*, 2025. 3
- [11] Jovana Kondic, Pengyuan Li, Dhiraj Joshi, Zexue He, Shafiq Abedin, Jennifer Sun, Ben Wiesel, Eli Schwartz, Ahmed Nassar, Bo Wu, Assaf Arbelle, Aude Oliva, Dan Gutfreund, Leonid Karlinsky, and Rogerio Feris. ChartGen: Scaling chart understanding via code-guided synthetic chart generation. *arXiv preprint arXiv:2507.19492*, 2025. 2
- [12] Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, Lei M Zhang, Kay McKinney, Disha Shrivastava, Cosmin Paduraru, George Tucker, Doina Precup, Feryal Behbahani, and Aleksandra Faust. Training language models to self-correct via reinforcement learning. In *ICLR*, 2025. 3
- [13] Hugo Laurençon, Léo Tronchon, and Victor Sanh. Unlocking the conversion of web screenshots into html code with the websight dataset. *arXiv preprint arXiv:2403.09029*, 2024. 2
- [14] Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. LLaVA-NeXT-Interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*, 2024. 6
- [15] Xinhao Li, Ziang Yan, Desen Meng, Lu Dong, Xiangyu Zeng, Yanan He, Yali Wang, Yu Qiao, Yi Wang, and Limin Wang. VideoChat-R1: Enhancing spatio-temporal perception via reinforcement fine-tuning. *arXiv preprint arXiv:2504.06958*, 2025. 3
- [16] Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Wanli Ouyang. Flow-GRPO: Training flow matching models via online rl. In *NeurIPS*, 2025. 4
- [17] Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Yaofeng Sun, Chengqi Deng, Hanwei Xu, Zhenda Xie, and Chong Ruan. DeepSeek-VL: Towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525*, 2024. 6
- [18] Fanqing Meng, Wenqi Shao, Quanfeng Lu, Peng Gao, Kaipeng Zhang, Yu Qiao, and Ping Luo. ChartAssistant: A universal chart multimodal language model via chart-to-table pre-training and multitask instruction tuning. In *Findings of ACL*, 2024. 2, 3
- [19] 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. Training language models to follow instructions with human feedback. In *NeurIPS*, 2022. 2, 3
- [20] Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. Recursive introspection: Teaching language model agents how to self-improve. In *NeurIPS*, 2024. 3
- [21] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*, 2023. 2
- [22] Juan A. Rodriguez, Abhay Puri, Shubham Agarwal, Issam H. Laradji, Pau Rodriguez, Sai Rajeswar, David Vazquez, Christopher Pal, and Marco Pedersoli. StarVector: Generating scalable vector graphics code from images and text. In *CVPR*, 2025. 2
- [23] William Saunders, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike. Self-critiquing models for assisting human evaluators. *arXiv preprint arXiv:2206.05802*, 2022. 3
- [24] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 2, 3
- [25] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024. 2, 3
- [26] Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun Zhang, Kangjia Zhao,

- Qianqian Zhang, Ruochen Xu, and Tiancheng Zhao. VLM-R1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*, 2025. 3
- [27] Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. HybridFlow: A flexible and efficient rlhf framework. In *EuroSys*, 2025. 5
- [28] Chenglei Si, Yanzhe Zhang, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. Design2Code: How far are we from automating front-end engineering? In *NAACL*, 2025. 2, 3
- [29] Wentao Tan, Qiong Cao, Chao Xue, Yibing Zhan, Changxing Ding, and Xiaodong He. ChartMaster: Advancing chart-to-code generation with real-world charts and chart similarity reinforcement learning. *arXiv preprint arXiv:2508.17608*, 2025. 2
- [30] Qwen Team. Qwen2.5-VL technical report. *arXiv preprint arXiv:2502.13923*, 2025. 5, 6
- [31] Qwen Team. Qwen3-VL technical report. *arXiv preprint arXiv:2511.21631*, 2025. 2, 3, 4, 6
- [32] Gladys Tyen, Hassan Mansoor, Victor Cărbune, Peter Chen, and Tony Mak. LLMs cannot find reasoning errors, but can correct them given the error location. In *Findings of ACL*, 2024. 3
- [33] Haozhe Wang, Chao Qu, Zuming Huang, Wei Chu, Fangzhen Lin, and Wenhui Chen. VL-Rethinker: Incentivizing self-reflection of vision-language models with reinforcement learning. In *NeurIPS*, 2025. 3
- [34] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-VL: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 6
- [35] Yana Wei, Liang Zhao, Jianjian Sun, Kangheng Lin, Jisheng Yin, Jingcheng Hu, Yinmin Zhang, En Yu, Haoran Lv, Zejia Weng, Jia Wang, Chunrui Han, Yuang Peng, Qi Han, Zheng Ge, Xiangyu Zhang, Daxin Jiang, and Vishal M. Patel. Open vision reasoner: Transferring linguistic cognitive behavior for visual reasoning. In *NeurIPS*, 2025. 3
- [36] Chengyue Wu, Yixiao Ge, Qiushan Guo, Jiahao Wang, Zhixuan Liang, Zeyu Lu, Ying Shan, and Ping Luo. Plot2Code: A comprehensive benchmark for evaluating multi-modal large language models in code generation from scientific plots. In *Findings of NAACL*, 2025. 2, 6
- [37] Renqiu Xia, Bo Zhang, Hancheng Ye, Xiangchao Yan, Qi Liu, Hongbin Zhou, Zijun Chen, Peng Ye, Min Dou, Botian Shi, Junchi Yan, and Yu Qiao. ChartX & ChartVLM: A versatile benchmark and foundation model for complicated chart reasoning. *IEEE Transactions on Image Processing*, 2025. 2, 3, 6
- [38] Ziang Yan, Xinhao Li, Yinan He, Zhengrong Yue, Xiangyu Zeng, Yali Wang, Yu Qiao, Limin Wang, and Yi Wang. VideoChat-R1.5: Visual test-time scaling to reinforce multimodal reasoning by iterative perception. In *NeurIPS*, 2025. 3
- [39] Cheng Yang, Chufan Shi, Yaxin Liu, Bo Shui, Junjie Wang, Mohan Jing, Linran Xu, Xinyu Zhu, Siheng Li, Yuxiang Zhang, et al. ChartMimic: Evaluating lmm’s cross-modal reasoning capability via chart-to-code generation. In *ICLR*, 2025. 2, 3, 4, 5, 6, 7, 1
- [40] Yi Yang, Xiaoxuan He, Hongkun Pan, Xiyan Jiang, Yan Deng, Xingtao Yang, Haoyu Lu, Dacheng Yin, Fengyun Rao, Minfeng Zhu, Bo Zhang, and Wei Chen. R1-Onevision: Advancing generalized multimodal reasoning through cross-modal formalization. In *ICCV*, 2025. 3
- [41] Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, Qianyu Chen, Huarong Zhou, Zhensheng Zou, Haoye Zhang, Shengding Hu, Zhi Zheng, Jie Zhou, Jie Cai, Xu Han, Guoyang Zeng, Dahai Li, Zhiyuan Liu, and Maosong Sun. MiniCPM-V: A gpt-4v level mllm on your phone. *arXiv preprint 2408.01800*, 2024. 6
- [42] Sukmin Yun, Haokun Lin, Rusiru Thushara, Mohammad Qazim Bhat, Yongxin Wang, Zutao Jiang, Mingkai Deng, Jinhong Wang, Tianhua Tao, Junbo Li, et al. Web2Code: A large-scale webpage-to-code dataset and evaluation framework for multimodal llms. In *NeurIPS*, 2024. 2
- [43] Liang Zhang, Anwen Hu, Haiyang Xu, Ming Yan, Yichen Xu, Qin Jin, Ji Zhang, and Fei Huang. TinyChart: Efficient chart understanding with visual token merging and program-of-thoughts learning. In *EMNLP*, 2024. 2, 3, 6
- [44] Xingjian Zhang, Siwei Wen, Wenjun Wu, and Lei Huang. TinyLLaVA-Video-R1: Towards smaller llms for video reasoning. *arXiv preprint arXiv:2504.09641*, 2025. 3
- [45] Zhihan Zhang, Yixin Cao, and Lizi Liao. Boosting chart-to-code generation in mllm via dual preference-guided refinement. In *ACM Multimedia*, 2025. 2
- [46] Xuanle Zhao, Deyang Jiang, Zhixiong Zeng, Lei Chen, Haibo Qiu, Jing Huang, Yufeng Zhong, Liming Zheng, Yilin Cao, and Lin Ma. VinciCoder: Unifying multimodal code generation via coarse-to-fine visual reinforcement learning. *arXiv preprint arXiv:2511.00391*, 2025. 2
- [47] Xuanle Zhao, Xianzhen Luo, Qi Shi, Chi Chen, Shuo Wang, Wanxiang Che, Zhiyuan Liu, and Maosong Sun. ChartCoder: Advancing multimodal large language model for chart-to-code generation. In *ACL*, 2025. 2, 3, 4, 5, 6
- [48] Huaixiu Steven Zheng, Swaroop Mishra, Hugh Zhang, Xinyun Chen, Minmin Chen, Azade Nova, Le Hou, Heng-Tze Cheng, Quoc V. Le, Ed H. Chi, and Denny Zhou. Natural plan: Benchmarking llms on natural language planning. *2406.04520*, 2024. 3
- [49] Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyuan Luo, Zhangchi Feng, and Yongqiang Ma. LlamaFactory: Unified efficient fine-tuning of 100+ language models. In *ACL*, 2024. 5
- [50] Ziwei Zheng, Michael Yang, Jack Hong, Chenxiao Zhao, Guohai Xu, Le Yang, Chao Shen, and Xing Yu. DeepEyes: Incentivizing “thinking with images” via reinforcement learning. In *ICLR*, 2026. 3

MM-ReCoder: Advancing Chart-to-Code Generation with Reinforcement Learning and Self-Correction

Supplementary Material

We elaborate on the details of our rewards and MM-ReCoder prompts in Section A. Section B shows our RL training curves. In Section C, we conduct ablation studies on the RL training strategy, cold start, and the model used in the model-based reward. We further compare MM-ReCoder with a few baselines through human evaluation in Section D. Finally, qualitative results are in Section E.

A. Implementation Details

In this section, we first provide a detailed description of the prompt used for the model-based reward. Then we elaborate on the prompt we use to guide MM-ReCoder in self-correction.

A.1. Reward Implementation

In the rule-based reward, we evaluate text score, color score, chart type score, and layout score of the generated chart. The implementation of all these scores is borrowed from ChartMimic [39]. Please refer to their paper and code for implementation.

To measure the model-based reward, we use vLLM to launch a Qwen2.5-VL-72B server and set the sampling temperature to be 0.3. If the generated code encounters runtime error, its model-based reward is 0. Otherwise, we use the following prompt to ask Qwen2.5-VL-72B to score the generated chart:

<ground truth image>

<generated image>

You are an excellent judge at evaluating visualization chart plots. The first image (reference image) is created using ground truth matplotlib code, and the second image (AI-generated image) is created using matplotlib code generated by an AI assistant. Your task is to score how well the AI-generated plot matches the ground truth plot.

Scoring Methodology:

The AI-generated image's score is based on the following criteria, totaling a score out of 100 points:

1. **Chart Types (20 points)** Does the AI-generated image include all chart types present in the reference image (e.g., line charts, bar charts, etc.)?
2. **Layout (10 points)** Does the arrangement of subplots in the AI-generated image match the reference image (e.g., number of rows and columns)?

3. **Text Content (20 points)** Does the AI-generated image include all text from the reference image (e.g., titles, annotations, axis labels), excluding axis tick labels?

4. **Data (20 points)** How accurately do the data trends in the AI-generated image resemble those in the original image and is the number of data groups the same as in the reference image?

5. **Style (20 points)** Does the AI-generated image match the original in terms of colors (line colors, fill colors, etc.), marker types (point shapes, line styles, etc.), legends, grids, and other stylistic details?

6. **Clarity (10 points)** Is the AI-generated image clear and free of overlapping elements?

Evaluation:

Compare the two images head to head and provide a detailed assessment. Use the following format for your response:

—
Detailed Assessment: \${your assessment for each point}

—
Scores:

- **Chart Types:** \${your score out of 20}
- **Layout:** \${your score out of 10}
- **Text Content:** \${your score out of 20}
- **Data:** \${your score out of 20}
- **Style:** \${your score out of 20}
- **Clarity:** \${your score out of 10}

Total Score: \${your total score out of 100}

—
Please use the above format to ensure the evaluation is clear and comprehensive.

A.2. Model Prompt of MM-ReCoder

In the first turn of the conversation, we give MM-ReCoder a chart image and ask it to reproduce it using Python. We also include an instruction to follow the template requirement of our format reward described in Section 3.4.1. Specifically, our prompt is:

<image>

You are an expert developer specializing in writing Python matplotlib code based on a given picture. I need your help to generate the Python code that can reproduce the picture based on the picture I provided.

You **FIRST** think about the reasoning process as an internal monologue and then provide the final code. The reasoning process **MUST BE** enclosed within <think></think> tags. The final code **MUST BE** put in

```
```python
Your code
```
at the end.
```

Notice that the specific task prompt, which is the black text in the prompt above, may differ from dataset to dataset. The one we show here is an example from Chart2Code-160k [47]. During training and evaluation, we use the task prompt provided by each dataset.

After the model outputs a first-turn response, we extract and execute the code inside the ````python...```` block. If the execution encounters a runtime error, we get the error information from the Python interpreter. Otherwise, we derive an image rendered from the code.

In the second turn, if the feedback is a runtime error, we instruct the model to revise its code as the following:

<error information>

Your code encountered a runtime error. Please continue modifying the code.

You **FIRST** think about the reasoning process as an internal monologue and then provide the final code. The reasoning process **MUST BE** enclosed within <think></think> tags. The final code **MUST BE** put in

```
```python
Your code
```
at the end.
```

If the chart is successfully rendered, we use the following prompt:

<image>

Here is the image rendered by your code. Please continue modifying the code so that the reproduced chart looks more like the given image.

You **FIRST** think about the reasoning process as an internal monologue and then provide the final code.

The reasoning process **MUST BE** enclosed within <think></think> tags. The final code **MUST BE** put in

```
```python
Your code
```
at the end.
```

When we employ MM-ReCoder for more than two turns, we repeat this process and the conversations in the previous turns are all kept as contexts.

B. RL Training Curves

In Figure A1, we show the training reward curves and the length of thinking trace/code of MM-ReCoder. We also evaluate the model every 40 steps on ChartMimic during training.

We find that the training rewards reach a plateau at the end of the first RL stage — shared first turn. But they begin to rise again in the second stage, full-trajectory optimization. On the contrary, the evaluation rewards on ChartMimic keep rising during training. However, the evaluation rewards can only hit 85% while the training rewards reach 95%. This shows that there is a distribution shift between the training data and the evaluation benchmarks.

One interesting observation is that, during the shared-first-turn stage, the length of the model’s thinking trace first grows from ~ 200 tokens to ~ 600 tokens, and then gradually drops below 100 tokens. During the first half of the shared-first-turn stage, we find the model learns to propose a long item list of code revision. At the 500th step, the model includes roughly 30 revision items in the thinking trace, leading to an average length of 600 tokens. After this point, the model gradually discards incorrect revision items and finally converges to around four items per chart.

C. Ablation Studies

In this section, we conduct ablation studies on the self-correction RL training strategy, cold start, and the model used in the model-based reward.

C.1. RL Training Strategy

Hyperparameters in full-trajectory optimization. In our full-trajectory optimization method, we use a hyperparameter γ to balance the rewards of the two turns and a hyperparameter η to encourage the model for successful self-correction. In Section 4.4, we claim that $\gamma = 0.5$ and $\eta = 0.1$ work best when only using the full-trajectory optimization. We perform ablation studies on this design choice in Table A1.

Firstly, when the reward of the first turn is not utilized ($\gamma = 0$), the model always lowers the quality of its first-

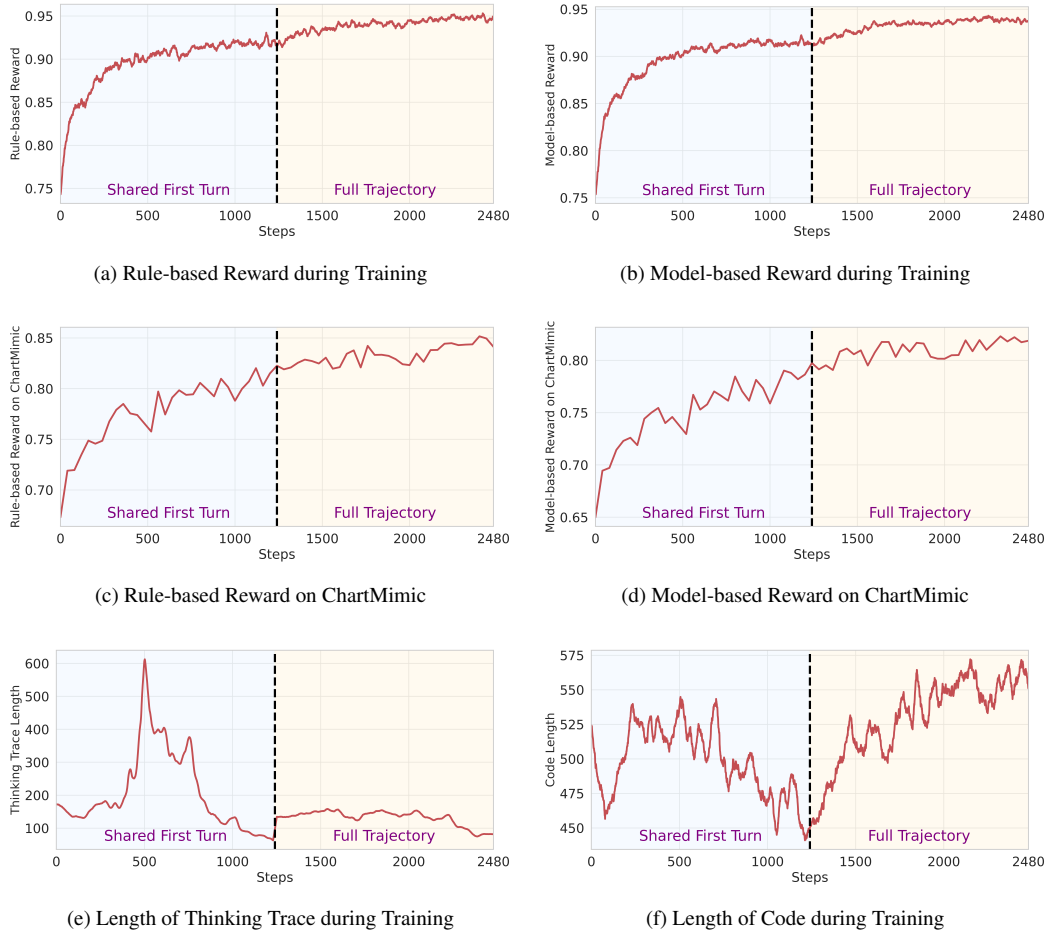


Figure A1. Training curves of MM-ReCoder. Curves on the training set are smoothed with a window size of 20 steps. We evaluate the model on ChartMimic every 40 steps. Note that the value of the rule-based reward can be slightly different from the reported model performance on low-level score because the latter is evaluated under the official evaluation codebase of ChartMimic.

| Strategy | γ | η | First turn | | Second turn | | Avg. low-level improvement on rendered charts | Percentage of improved samples | Percentage of degraded samples | Percentage of repeated code |
|-----------------|----------|--------|------------|-----------|-------------|-----------|---|--------------------------------|--------------------------------|-----------------------------|
| | | | Exec.Rate | Low-Level | Exec.Rate | Low-Level | | | | |
| Full-trajectory | 0 | 0 | 91.8 | 81.8 | 95.2 | 83.9 | 0.21 | 3.4 | 2.6 | 46.9 |
| Full-trajectory | 0 | 0.1 | 64.5 | 66.7 | 96.0 | 84.3 | 10.11 | 87.5 | 2.3 | 0.3 |
| | 0 | 0.3 | 79.0 | 73.7 | 94.5 | 83.2 | 4.57 | 77.1 | 8.3 | 0.0 |
| | 0 | 0.5 | 71.8 | 72.5 | 96.2 | 77.7 | 3.13 | 77.7 | 10.2 | 0.0 |
| Full-trajectory | 0.5 | 0.1 | 96.5 | 85.0 | 97.5 | 85.5 | -0.03 | 2.2 | 2.2 | 55.6 |
| | 0.8 | 0.1 | 95.2 | 83.7 | 96.3 | 84.9 | 0.01 | 3.2 | 3.9 | 45.9 |
| | 1 | 0.1 | 95.5 | 84.7 | 96.2 | 84.2 | 0.00 | 1.4 | 1.4 | 5.1 |

Table A1. Ablation on the hyperparameters in full-trajectory optimization. All experiments are conducted **without model-based reward**.

round output to hack the self-correction bonus η . And higher η leads to worse second-turn performance. Secondly, when $\eta = 0.1$, the existence of γ can prevent the model from lowering the first-turn quality and achieve better performance. However, it also hinders the model from learning self-correction, indicated by the nearly-zero low-level score improvement on rendered charts. A γ of 0.5 shows the best overall performance across all the choices we have tried.

RL strategy combinations. While our RL strategy in MM-

ReCoder is shared-first-turn optimization followed by full-trajectory optimization, we try out several other possible combinations of the RL strategies in Table A2.

As our method is composed of two stages, which trains the model for two epochs on the training set in total, we first try employing the full-trajectory and shared-first-turn optimizations for two epochs for fair comparison. For the full-trajectory optimization, however, training with it for two epochs causes over-fitting, achieving lower performance

| Strategy | γ η | | First turn | | Second turn | | Avg. low-level improvement on rendered charts | Percentage of improved samples | Percentage of degraded samples | Percentage of repeated code |
|-------------------------------------|-----------------|-----|------------|-----------|-------------|-------------|---|--------------------------------|--------------------------------|-----------------------------|
| | | | Exec.Rate | Low-Level | Exec.Rate | Low-Level | | | | |
| Full-trajectory | 0 | 0 | 91.8 | 81.8 | 95.2 | 83.9 | 0.21 | 3.4 | 2.6 | 46.9 |
| | 0.5 | 0.1 | 96.5 | 85.0 | 97.5 | 85.5 | -0.03 | 2.2 | 2.2 | 55.6 |
| Full-trajectory for 2 epochs | 0 | 0 | 92.3 | 79.0 | 97.0 | 81.8 | 0.16 | 4.5 | 2.5 | 54.0 |
| | 0.5 | 0.1 | 94.3 | 82.2 | 94.2 | 82.3 | 0.07 | 2.7 | 1.6 | 87.7 |
| Shared-first-turn | - | - | 91.7 | 79.8 | 95.8 | 82.6 | 0.72 | 14.4 | 8.4 | 16.8 |
| Shared-first-turn for 2 epochs | - | - | 92.3 | 82.3 | 97.2 | 84.8 | 0.35 | 10.5 | 4.7 | 12.3 |
| Shared-first-turn + Full-trajectory | 0 | 0 | 94.3 | 83.7 | 97.7 | 86.0 | 0.55 | 12.1 | 8.2 | 21.6 |
| Full-trajectory + Shared-first-turn | 0 | 0 | 93.3 | 83.6 | 97.2 | 85.0 | 0.02 | 2.9 | 2.7 | 46.7 |
| Single-turn + Full-trajectory | 0 | 0 | 94.7 | 84.5 | 96.5 | 85.2 | 0.03 | 2.7 | 4.0 | 58.4 |

Table A2. Ablation on the RL stage combinations. All experiments are conducted **without model-based reward**.

| Cold Start | Performance on ChartMimic after RL | | | |
|--|------------------------------------|-----------------------|--|-----------------------------|
| | First turn low-level | Second turn low-level | Average improvement on rendered charts | Percentage of repeated code |
| None | 83.0 | 84.9 | 0.53 | 36.3 |
| Single-turn cold start | 82.3 | 84.8 | 0.02 | 13.1 |
| Single-turn cold start + multi-turn cold start | 83.7 | 86.0 | 0.55 | 21.6 |

Table A3. Ablation studies on cold start. We train the base model with different cold start recipes and the same RL strategy (full-trajectory + shared-first-turn). We evaluate their performance on ChartMimic after RL. Experiments are conducted **without model-based reward**.

than one epoch. On the other hand, shared-first-turn for two epochs can further boost the model performance compared to one epoch, but underperforms the other strategies.

We explore strategy combinations besides the one we use in our main method. For example, we can switch the order of the two stages, performing full-trajectory optimization followed by shared-first-turn. Moreover, we can also train a model under the single-turn RL setting first and then perform full-trajectory optimization to enable self-correction. However, as shows in Table A2, these two strategies cannot elicit the self-correction capability of the model and result in suboptimal performance. They repeat their first-turn code in the second turn with a frequency of nearly 50%. We attribute this to their weak self-correction ability after full-trajectory optimization and single-turn RL. In this case, the model’s output entropy is significantly reduced during the first-stage RL, making it unable to explore self-correction effectively in the second stage.

C.2. Cold Start

To investigate the effect of cold start, we perform our two-stage RL on two checkpoints: (a) Qwen2.5-VL-7B without any cold start and (b) the checkpoint that has only undergone the single-turn cold start. The results are in Table A3 and the RL stage is performed without model-based reward.

We find that our self-correction RL training strategy can enable Qwen2.5-VL-7B to self-correct without any cold start, indicating that the teacher model used to generate the code start data is not the key to elicit self-correction.

However, the performance without cold start is around 1% lower than that of the cold-started model. While cold start on the single-turn data, Chart2Code-160k, can improve the base model’s coding capability as shown in Table 5, directly applying RL on this checkpoint does not realize self-correction because of its weak capability in multi-turn conversation. After the multi-turn cold start stage, the model is able to benefit from both the cold start and our multi-turn self-correction RL method.

We also look into the qualitative results to investigate the effect of cold start. We notice that the model without cold start generates very short and abstract thinking traces, while the one with cold start learns to describe the image differences in detail and outlines an action list before coding, which is important to boost the model’s performance.

C.3. Model-based Reward

In the RL reward of MM-ReCoder, we use Qwen2.5-VL-72B judge as a model-based reward to score the generated charts. While Qwen2.5-VL-72B is larger than our base model, Qwen2.5-VL-7B, we show in Table A4 that using Qwen2.5-VL-7B itself as the reward model can also guide the learning of MM-ReCoder.

In this experiment, we train MM-ReCoder in the single-turn setting without self-correction. We find that using Qwen2.5-VL-7B as the reward model can improve the high-level score by 4% compared to training without model-based reward. However, the model’s low-level score is degraded by 3.8%. We hypothesize that the occasional in-

| Reward Model | ChartMimic | | |
|----------------|------------|-----------|------------|
| | Exec.Rate | Low-level | High-level |
| None | 95.0 | 84.8 | 78.6 |
| Qwen2.5-VL-7B | 95.3 | 81.0 | 82.6 |
| Qwen2.5-VL-72B | 95.0 | 84.3 | 83.7 |

Table A4. Ablation on the reward model **under the single-turn RL setting**. Qwen2.5-VL-7B as the reward model is able to improve the high-level score on ChartMimic, but replacing Qwen2.5-VL-7B with Qwen2.5-VL-72B can further boost the model.

accuracies of the Qwen2.5-VL-7B judge can mislead the learning of the low-level score. Moreover, replacing the 7B reward model with the 72B version can better improve the model’s capability, so we use Qwen2.5-VL-72B in our experiments to optimize the performance of MM-ReCoder.

D. Human Evaluation

D.1. A/B Testing

We conduct human evaluation on the charts generated by MM-ReCoder and the baseline models. We compare two models in each study by evaluating human annotators’ preference on their generated charts.

For a pair of models, we randomly select 100 test cases from ChartMimic and form a triplet consisting of the two model-generated charts and the ground truth image. For each triplet, the model-generated charts are anonymized and presented in a random order. Annotators are asked to choose which chart aligns better with the ground truth. If the two charts are roughly equal, annotators can select “tie.” In each study, we have three annotators reviewing all the selected test cases. For each test case, if at least two annotators prefer a particular chart, it is counted as a win for that chart; otherwise, it is counted as a tie. Finally, we report the win/tie/loss rates of our model.

We report the results in Table A5. Our model wins 37% of the time and loses 20% of the time against ChartCoder, and wins 40% and loses 23% against Qwen2.5-VL-72B. However, our model only wins 19% but loses 48% against Qwen3-VL-235B-A22B. The human evaluation results demonstrate that MM-ReCoder outperforms ChartCoder and Qwen2.5-VL-72B, but underperforms Qwen3-VL-235B-A22B, which is in line with the high-level score evaluation in Table 1.

D.2. Score Improvement Reflects Self-correction

In our experiments, we use the improvement in the evaluation scores to indicate the success of self-correction. To check how faithfully the score improvement can reflect chart improvement, we manually check all ChartMimic test samples with improved scores in the second turn. We find that 76.5% of them exhibit improvements that are easily discernible by human eyes. This shows that the score improve-

| | Win | Tie | Loss |
|------------------------------|-----|-----|------|
| Ours v.s. ChartCoder | 37% | 43% | 20% |
| Ours v.s. Qwen2.5-VL-72B | 40% | 37% | 23% |
| Ours v.s. Qwen3-VL-235B-A22B | 19% | 33% | 48% |

Table A5. Human evaluation between MM-ReCoder and baselines on ChartMimic. The results are in line with the high-level score.

| Model | Diagnosis error | Coding error | Regression |
|--------------------|-----------------|--------------|------------|
| MM-ReCoder | 8 (30.8%) | 13 (50.0%) | 5 (19.2%) |
| Qwen3-VL-8B | 22 (37.9%) | 28 (48.3%) | 8 (13.8%) |
| Qwen3-VL-235B-A22B | 15 (20.5%) | 33 (45.2%) | 25 (34.3%) |

Table A6. Frequencies of self-correction failure modes.

ment can indicate chart improvement in most cases.

E. Qualitative Results

In this section, we show examples reflecting the self-correction capability of MM-ReCoder. We also qualitatively compare MM-ReCoder with baseline models.

E.1. Self-correction Examples


We showcase qualitative results of successful self-correction in Figures A2 to A8. MM-ReCoder is able to self-correct the labels, axis ranges, hatches, category orders, and data points in the charts and handle runtime errors. However, although MM-ReCoder tends to propose multiple revision items in its thinking trace, we notice that not all the items are reflected in the revised code. This demonstrates a limitation of MM-ReCoder that the thinking trace and generated code are not always perfectly aligned.

E.2. Comparison with Baselines

We qualitatively compare MM-ReCoder with baseline models in Figure A9. While all the models can faithfully reproduce the chart types, MM-ReCoder is better at using the correct colors, texts, chart styles, and data points. We further demonstrate failure cases in Figure A10 of MM-ReCoder and baselines. We find that the models tend to fail when the given chart type is rare, whose code functions are possibly not covered by the training data.

E.3. Analysis of Self-correction Failures

To summarize typical failure modes in self-correction, we investigate the regressed samples in the second turn. We discover three major failure modes: (a) **diagnosis error** - the thinking trace presents incorrect chart discrepancies or revision items; (b) **coding error** - the thinking trace is correct, but the code fails to fix the issue; (c) **regression** - the identified issues are fixed, but the existing correct functionalities (e.g., text, color) are broken. We count the failure mode frequencies of MM-ReCoder and the baselines in Table A6. Compared to Qwen3-VL-8B, MM-ReCoder fails in diagnosis less frequently. Qwen3-VL-235B-A22B has fewer from diagnosis errors but more regression errors.



First Turn

<think> Looking at the pie chart, I need to:

1. Extract the approximate values based on percentages shown
2. Create a pie chart with appropriate colors
3. Add proper labels and percentages
4. Position the legend appropriately
5. Use appropriate formatting for percentages

</think>

```
import matplotlib.pyplot as plt
import numpy as np

### Data to plot
labels = ['David', 'John', 'Marry', 'Peter']
sizes = [11, 29, 20, 40] ### Approximate percentages based on visual estimation
colors = ['#1f77b4', '#ff7f0e', '#2ca02c', '#d62728']

### Create figure and axis
fig, ax = plt.subplots(figsize=(5, 5))

### Create pie chart
wedges, texts, autotexts = ax.pie(sizes, labels=labels, colors=colors,
autotext='%1.1f%', startangle=90, textprops={'weight': 'bold'},
pctdistance=0.85)

### Add title if needed (optional, but improves appearance)
### ax.set_title('', fontsize=12, fontweight='bold')

### Display percentages on wedges
autotexts[0].set_text('11.0%')
autotexts[1].set_text('29.0%')
autotexts[2].set_text('20.0%')
autotexts[3].set_text('40.0%')

### Add legend in top-left corner with proper positioning
ax.legend(labels, loc='upper left', bbox_to_anchor=(-0.1, 1.1), fontsize=11,
framealpha=True, edgecolor='black')

### Display chart
plt.tight_layout()
plt.show()
```

Second Turn

<think> The labels are appearing incorrectly and overlapping improperly. Also, the title/spacing needs adjustment better. The legend needs proper alignment without labels interfering. Also, the percentages need proper placement without labels interfering. Also, the x-axis labels need proper alignment without being displayed incorrectly.

I need to:

1. Remove proper labels from wedges and only show percentages accurately
2. Properly position the legend in top-left corner without causing issues
3. Remove the title properly or adjust better
4. Make sure proper spacing without labels on wedges interfering

</think>

```
import matplotlib.pyplot as plt
import numpy as np

### Data to plot
labels = ['David', 'John', 'Marry', 'Peter']
sizes = [11, 29, 20, 40] ### Approximate percentages based on visual estimation
colors = ['#1f77b4', '#ff7f0e', '#2ca02c', '#d62728']

### Create figure and axis
fig, ax = plt.subplots(figsize=(5, 5))

### Create pie chart
wedges, texts, autotexts = ax.pie(sizes, labels=None, colors=colors,
autotext='%1.1f%', startangle=90, textprops={'weight': 'bold'})

### Add percentages
autotexts[0].set_text('11.0%')
autotexts[1].set_text('29.0%')
autotexts[2].set_text('20.0%')
autotexts[3].set_text('40.0%')

### Add custom legend in top-left corner with proper positioning
ax.legend(labels, loc='upper left', bbox_to_anchor=(0, 1.1), fontsize=11,
framealpha=True, edgecolor='black')

### Display chart
plt.tight_layout()
plt.show()
```






Figure A2. Qualitative result of MM-ReCoder for self-correction. The model removes the text labels from the pie chart in the second turn.

Input



First Turn

<think> Looking at the chart, I need to create a horizontal bar plot with appropriate spacing and proper alignment.

Steps to reproduce:

1. Import matplotlib.pyplot
2. Define the emotion labels (y-axis ticks)
3. Define corresponding frequency values (approximate based on visual estimation)
4. Create the horizontal bar plot
5. Set appropriate axis ticks and labels
6. Display the plot properly

</think>

```
import matplotlib.pyplot as plt
import numpy as np

#### Data for the chart (estimated frequencies based on visual appearance)
emotions = [
    'Amusement', 'Unbothered', 'Sadness', 'Pride', 'Nervousness',
    'Annoyance', 'Gratitude', 'Relief', 'Joy', 'Disapproval',
    'Excitement', 'Delight', 'Oblivious', 'Embarrassment', 'Disappointment'
]
frequencies = [
    2.2, 2.6, 3.0, 3.4, 3.5, 3.8, 4.0, 4.0, 6.0, 6.1, 6.1, 6.6, 6.7, 7.0, 7.5
]

#### Create the figure and axis
fig, ax = plt.subplots(figsize=(8, 8))

#### Create horizontal bar chart
ax.barh(range(len(emotions)), frequencies, height=0.6, color='#6495ED') ####
Using a similar blue color

#### Set labels and title
ax.set_xlabel('Frequency (%)', fontsize=12, fontweight='bold')
ax.set_ylabel('Emotion!', fontsize=12, fontweight='bold')
ax.set_title('', fontsize=16)


#### Set ticks and tickLabels
ax.set_yticks(range(len(emotions)))
ax.set_yticklabels(emotions, rotation=0, ha='right', fontsize=11,
fontweight='bold')

#### Set x-axis ticks appropriately
ax.set_xticks([0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14])
ax.set_xticklabels(['0', '1', '2', '3', '4', '5', '6', '7', '8', '9', '10', '11', '12', '13', '14',
'', '', ''], rotation=0, ha='center')

#### Display grid lines more accurately
ax.grid(axis='x', which='major', linestyle='-', linewidth=0.5)

#### Adjust layout
plt.tight_layout()

#### Display plot
plt.show()
```



Second Turn

<think> The x-axis ticks appear wrong and there are grid lines that shouldn't be there. Also, the spacing needs adjustment better without causing issues. The y-axis labels need proper alignment without being skewed incorrectly.

Also, the frequencies x-axis ticks need proper spacing without being skewed incorrectly. The grid lines need proper control without being visible as lines that interfere with proper appearance.

Also, the spacing needs adjustment better without causing issues properly.

</think>

```
import matplotlib.pyplot as plt
import numpy as np

#### Data for the chart (estimated frequencies based on visual appearance)
emotions = [
    'Amusement', 'Unbothered', 'Sadness', 'Pride', 'Nervousness',
    'Annoyance', 'Gratitude', 'Relief', 'Joy', 'Disapproval',
    'Excitement', 'Delight', 'Oblivious', 'Embarrassment', 'Disappointment'
]
frequencies = [
    2.2, 2.6, 3.0, 3.4, 3.5, 3.8, 4.0, 4.0, 6.0, 6.1, 6.1, 6.6, 6.7, 7.0, 7.5
]

#### Create the figure and axis
fig, ax = plt.subplots(figsize=(8, 8))

#### Create horizontal bar chart
ax.barh(range(len(emotions)), frequencies, height=0.6, color='#6495ED') ####
Using a similar blue color

#### Set labels and title
ax.set_xlabel('Frequency (%)', fontsize=11, fontweight='bold')
ax.set_ylabel('Emotion!', fontsize=12, fontweight='bold')
ax.set_title('', fontsize=14)

#### Set ticks and tickLabels
ax.set_yticks(range(len(emotions)))
ax.set_yticklabels(emotions, rotation=0, ha='right', fontsize=11,
fontweight='bold')

#### Set x-axis ticks appropriately
ax.set_xticks([0, 1, 2, 3, 4, 5, 6, 7, 8])
ax.set_xticklabels(['0', '1', '2', '3', '4', '5', '6', '7', '8'], rotation=0,
ha='center')

#### Remove y-axis ticks properly without grid issues
#### Set appropriate ticks without grid lines interfering properly

#### Display plot properly without grid lines interfering badly
plt.tight_layout()

#### Display plot
plt.show()
```





Figure A3. Qualitative result of MM-ReCoder for self-correction. The model rescales the x-axis range and removes the grid lines in the second turn.



First Turn

<think> Looking at the chart, I need to create proper bar plots with appropriate spacing and proper legend formatting. Since I don't have actual data, I need to estimate reasonable values based on the visual appearance.

Steps to reproduce:

1. Import matplotlib.pyplot
2. Define appropriate bar positions and widths
3. Use appropriate hatch patterns for bars
4. Create proper legend with correct formatting
5. Set appropriate axis ticks and labels

</think>

```
import matplotlib.pyplot as plt
import numpy as np

#### Data for the chart (estimated values based on visual appearance)
labels = ['VOC 2012', 'COCO 2017']
jpeg_values = [3.1, 1.5]
deepjssc_values = [1.5, 2.5]
ours_values = [0.5, 1.0]

#### Create figure and axis
fig, ax = plt.subplots(figsize=(6, 5))

#### Set width for bars
width = 0.25

#### Create bars
rect1 = ax.bar(np.arange(2) + 0*width, jpeg_values, width, label='JPEG',
hatch='/', edgecolor='black')
rect2 = ax.bar(np.arange(2) + 1*width, deepjssc_values, width,
label='DEEPS3CC w/ ofdm', hatch='.', edgecolor='black', linestyle=':')
rect3 = ax.bar(np.arange(2) + 2*width, ours_values, width, label='OURS',
hatch='x', edgecolor='black')


#### Add labels and title
ax.set_title('')
ax.set_xlabel('Datasets', fontsize=11, fontweight='bold')
ax.set_ylabel('Transmission Delay (ms)', fontsize=11, fontweight='bold')
ax.set_xticks(np.arange(2) + width)
ax.set_xticklabels(labels, rotation=0, ha='center')

#### Add legend in top-left corner with proper positioning
ax.legend(loc='upper left', bbox_to_anchor=(0, 1), fontsize=10,
framealpha=True, edgecolor='black')

#### Set y-axis ticks
ax.set_yticks([0.0, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0])

#### Display grid for better appearance
ax.grid(axis='y', linestyle=':', linewidth=0.5)

#### Show plot
plt.tight_layout()
plt.show()
```



Second Turn

<think> The rendered chart has incorrect legend content and colors, and the bars appear with wrong styling. Also, the proper spacing and alignment need adjustment better.

Also, the legend needs proper alignment without conflicting with bars incorrectly. The hatching needs adjustment more accurately without colors.

Also, the y-axis ticks need proper control without grid lines interfering incorrectly.

Also, the proper spacing needs adjustment better without overlap issues.

</think>

```
import matplotlib.pyplot as plt
import numpy as np

#### Data for the chart (estimated values based on visual appearance)
labels = ['VOC 2012', 'COCO 2017']
jpeg_values = [3.1, 1.5]
deepjssc_values = [1.5, 2.5]
ours_values = [0.5, 1.0]

#### Create figure and axis
fig, ax = plt.subplots(figsize=(6, 5))

#### Set width for bars
width = 0.25

#### Create bars
rect1 = ax.bar(np.arange(2) + 0*width, jpeg_values, width, label='JPEG',
hatch='/', edgecolor='black', color='FFFFFF')
rect2 = ax.bar(np.arange(2) + 1*width, deepjssc_values, width,
label='DEEPS3CC w/ ofdm', hatch='.', edgecolor='black', linestyle=':',
color='FFFFFF')
rect3 = ax.bar(np.arange(2) + 2*width, ours_values, width, label='OURS',
hatch='x', edgecolor='black', color='FFFFFF')

#### Add labels and title
ax.set_title('')
ax.set_xlabel('Datasets', fontsize=11, fontweight='bold')
ax.set_ylabel('Transmission Delay (ms)', fontsize=11, fontweight='bold')
ax.set_xticks(np.arange(2) + width)
ax.set_xticklabels(labels, rotation=0, ha='center')

#### Add proper legend in top-left corner with correct formatting
ax.legend(loc='upper left', bbox_to_anchor=(0, 1), fontsize=10,
framealpha=True, edgecolor='black')

#### Set y-axis ticks
ax.set_yticks([0.0, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0])

#### Display proper legend without conflicting issues
#### Remove grid or proper control without issues

#### Show plot
plt.tight_layout()
plt.show()
```





Figure A4. Qualitative result of MM-ReCoder for self-correction. The model removes the hatch color in the second turn.

Input



First Turn

<think> Looking at the radar charts, I need to create proper polar plots with appropriate data points. Since I don't have actual data, I need to estimate reasonable values based on the visual appearance. Each chart has 5 categories with values ranging from ~1 to ~5. I'll create appropriate angles for each category and plot lines with appropriate colors.

```

import matplotlib.pyplot as plt
import numpy as np

### Create figure with appropriate size
fig, axes = plt.subplots(1, 3, figsize=(15, 5), subplot_kw=dict(polar=True))

### Data for PGM
labels = ['Query Error', 'Privacy', 'MLA', 'Fidelity(D_train)', 'Fidelity(D_test)']
values_pgm = [3.5, 4.5, 2.5, 5.0, 4.8]

### Data for PrivSyn
labels2 = ['Query Error', 'Privacy', 'MLA', 'Fidelity(D_train)', 'Fidelity(D_test)']
values_privsyn = [4.2, 3.0, 2.8, 4.5, 3.8]

### Data for TVAE
labels3 = ['Query Error', 'Privacy', 'MLA', 'Fidelity(D_train)', 'Fidelity(D_test)']
values_tvae = [2.8, 5.0, 3.7, 4.2, 4.0]

### First subplot - PGM
ax1 = axes[0]
angles = np.linspace(0, 2 * np.pi, len(labels), endpoint=False).tolist()
angles += angles[:1]

### Plot data
values_pgm_extended = values_pgm + values_pgm[:1]
ax1.plot(angles, values_pgm_extended, color='#8FBC8F', linewidth=1.5)
ax1.fill(angles, values_pgm_extended, color='#8FBC8F', alpha=0.25)

ax1.set_thetagrids(np.degrees(angles[:-1]), labels, fontsize=11, fontweight='bold')
ax1.set_title('PGM (ε = ∞)', size=14, fontweight='bold', pad=50)
ax1.set_ylim(0, 5)
ax1.set_yticks([1, 2, 3, 4, 5])
ax1.set_yticklabels(['1', '2', '3', '4', '5'], fontsize=11)

### Second subplot - PrivSyn
ax2 = axes[1]
values_privsyn_extended = values_privsyn + values_privsyn[:1]
ax2.plot(angles, values_privsyn_extended, color='#483D8B', linewidth=1.5)
ax2.fill(angles, values_privsyn_extended, color='#483D8B', alpha=0.25)

ax2.set_thetagrids(np.degrees(angles[:-1]), labels2, fontsize=11, fontweight='bold')
ax2.set_title('PrivSyn (ε = ∞)', size=12, fontweight='bold')
ax2.set_ylim(0, 5)
ax2.set_yticks([1, 2, 3, 4, 5])
ax2.set_yticklabels(['1', '2', '3', '4', '5'], fontsize=11)

### Third subplot - TVAE
ax3 = axes[2]
values_tvae_extended = values_tvae + values_tvae[:1]
ax3.plot(angles, values_tvae_extended, color='#00FFFF', linewidth=1.5)
ax3.fill(angles, values_tvae_extended, color='#00FFFF', alpha=0.25)

ax3.set_thetagrids(np.degrees(angles[:-1]), labels3, fontsize=11, fontweight='bold')
ax3.set_title('TVAE', size=12, fontweight='bold')
ax3.set_ylim(0, 5)
ax3.set_yticks([1, 2, 3, 4, 5])
ax3.set_yticklabels(['1', '2', '3', '4', '5'], fontsize=11)

### Adjust layout
plt.tight_layout()
plt.show()

```

Second Turn

<think> The charts appear too filled and the lines need proper control without filling being dominant. Also, the titles need proper alignment without being too bold or inconsistent. Also, the filled areas need proper control without being too prominent or causing issues with transparency. I need to remove the filled areas properly or adjust them better. Also, the boldness of titles needs adjustment more accurately without causing issues. Also, the radial grid lines need proper control without being too prominent but accurate.

```

import matplotlib.pyplot as plt
import numpy as np

### Create figure with appropriate size
fig, axes = plt.subplots(1, 3, figsize=(15, 5), subplot_kw=dict(polar=True))

### Data for PGM
labels = ['Query Error', 'Privacy', 'MLA', 'Fidelity(D_train)', 'Fidelity(D_test)']
values_pgm = [3.5, 4.5, 2.5, 5.0, 4.8]

### Data for PrivSyn
labels2 = ['Query Error', 'Privacy', 'MLA', 'Fidelity(D_train)', 'Fidelity(D_test)']
values_privsyn = [4.2, 3.0, 2.8, 4.5, 3.8]

### Data for TVAE
labels3 = ['Query Error', 'Privacy', 'MLA', 'Fidelity(D_train)', 'Fidelity(D_test)']
values_tvae = [2.8, 5.0, 3.7, 4.2, 4.0]

### First subplot - PGM
ax1 = axes[0]
angles = np.linspace(0, 2 * np.pi, len(labels), endpoint=False).tolist()
angles += angles[:1]

### Plot data
values_pgm_extended = values_pgm + values_pgm[:1]
ax1.plot(angles, values_pgm_extended, color='#8FBC8F', linewidth=1.5)
ax1.set_thetagrids(np.degrees(angles[:-1]), labels, fontsize=11, fontweight='bold')
ax1.set_title('PGM (ε = ∞)', size=12, fontweight='bold', pad=50)
ax1.set_ylim(0, 5)
ax1.set_yticks([1, 2, 3, 4, 5])
ax1.set_yticklabels(['1', '2', '3', '4', '5'], fontsize=11)

### Second subplot - PrivSyn
ax2 = axes[1]
values_privsyn_extended = values_privsyn + values_privsyn[:1]
ax2.plot(angles, values_privsyn_extended, color='#483D8B', linewidth=1.5)
ax2.set_thetagrids(np.degrees(angles[:-1]), labels2, fontsize=11, fontweight='bold')
ax2.set_title('PrivSyn (ε = ∞)', size=11.5, fontweight='bold')
ax2.set_ylim(0, 5)
ax2.set_yticks([1, 2, 3, 4, 5])
ax2.set_yticklabels(['1', '2', '3', '4', '5'], fontsize=10)

### Third subplot - TVAE
ax3 = axes[2]
values_tvae_extended = values_tvae + values_tvae[:1]
ax3.plot(angles, values_tvae_extended, color='#00FFFF', linewidth=1.5)
ax3.set_thetagrids(np.degrees(angles[:-1]), labels3, fontsize=11, fontweight='bold')
ax3.set_title('TVAE', size=11.5, fontweight='bold')
ax3.set_ylim(0, 5)
ax3.set_yticks([1, 2, 3, 4, 5])
ax3.set_yticklabels(['1', '2', '3', '4', '5'], fontsize=10)

### Remove unnecessary fills property
### Adjust layout
plt.tight_layout()
plt.show()

```



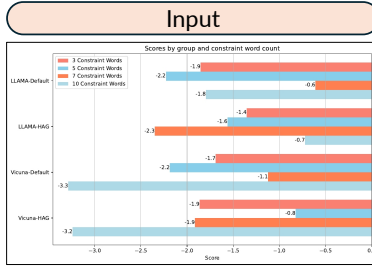


Figure A5. Qualitative result of MM-ReCoder for self-correction. The model removes the colors filled in the radar charts in the second turn.



First Turn

<think> Looking at the chart, I need to create a horizontal bar plot with proper spacing and annotations.
 Steps to reproduce:
 1. Import necessary libraries (matplotlib.pyplot)
 2. Define the categories and corresponding scores
 3. Create appropriate bar widths and positions
 4. Add proper labeling, titles, and ticks
 5. Display annotations on bars
 6. Position bars appropriately with proper spacing
 </think>

```
import matplotlib.pyplot as plt
import numpy as np

### Data extraction from the chart (approximate values based on visual inspection)
labels = ['LLAMA-Default', 'LLAMA-HAG', 'Vicuna-Default', 'Vicuna-HAG']
groups = ['3 Constraint Words', '5 Constraint Words', '7 Constraint Words', '10 Constraint Words']

### Scores for each group and constraint word count
scores_3 = [-1.9, -1.4, -1.7, -1.9]
scores_5 = [-2.2, -1.6, -1.7, -1.9]
scores_7 = [-0.6, -0.7, -1.1, -0.8]
scores_10 = [-2.2, -1.8, -2.2, -3.2]

### Create figure and axis
fig, ax = plt.subplots(figsize=(10, 6))

### Set y-axis positions
y_pos = np.arange(4)

### Create horizontal bars
ax.barh(y_pos, scores_3, height=0.2, color='#FAB972', label='3 Constraint Words')
ax.barh(y_pos + 0.15, scores_5, height=0.18, color='#F7CEEB', label='5 Constraint Words')
ax.barh(y_pos + 0.3, scores_7, height=0.18, color='#FF7F50', label='7 Constraint Words')
ax.barh(y_pos + 0.45, scores_10, height=0.18, color='#ADD8E6', label='10 Constraint Words')

### Add labels and title
ax.set_title('Scores by group and constraint word count', fontsize=22, fontweight='bold')
ax.set_xlabel('Score', fontsize=12, fontweight='bold')
ax.set_ylabel('', fontsize=12)

### Set ticks and grid
ax.set_xticks([-3.0, -2.5, -2.0, -1.5, -1.0, -0.5, 0.0])
ax.set_yticks(y_pos + 0.25)
ax.set_yticklabels(labels, fontsize=11, weight='bold')

### Add legend in top-left corner
ax.legend(loc='upper left', bbox_to_anchor=(0, 1), fontsize=11, framealpha=True, edgecolor='black')

### Add annotations
for i, (score3, score5, score7, score10) in enumerate(zip(scores_3, scores_5, scores_7, scores_10)):
    ax.text(score3, i, f'{score3:.1f}', va='center', ha='right', fontsize=11, weight='bold')
    ax.text(score5, i + 0.15, f'{score5:.1f}', va='center', ha='right', fontsize=11)
    ax.text(score7, i + 0.3, f'{score7:.1f}', va='center', ha='right', fontsize=11)
    ax.text(score10, i + 0.45, f'{score10:.1f}', va='center', ha='right', fontsize=11)

### Display grid lines
ax.xaxis.grid(True, linestyle='-', linewidth=0.5)
ax.yaxis.grid(False)

### Show plot
plt.tight_layout()
plt.show()
```

Second Turn

<think> The chart is inverted and the y-axis labels are reversed. Also, the legend needs proper alignment without being cut off incorrectly. The groups need proper alignment without being reversed incorrectly. Also, the y-axis labels need proper orientation without being reversed incorrectly. I need to fix the proper alignment and orientation. Also, the legend needs all items properly displayed with correct spacing.
 I need to properly control the orientation and spacing correctly without inverting the axes incorrectly.
 </think>

```
import matplotlib.pyplot as plt
import numpy as np

### Data extraction from the chart (approximate values based on visual inspection)
labels = ['LLAMA-Default', 'LLAMA-HAG', 'Vicuna-Default', 'Vicuna-HAG']
groups = ['3 Constraint Words', '5 Constraint Words', '7 Constraint Words', '10 Constraint Words']

### Scores for each group and constraint word count
scores_3 = [-1.9, -1.4, -1.7, -1.9]
scores_5 = [-2.2, -1.6, -1.7, -1.9]
scores_7 = [-0.6, -0.7, -1.1, -0.8]
scores_10 = [-2.2, -1.8, -2.2, -3.2]

### Create figure and axis
fig, ax = plt.subplots(figsize=(10, 6))

### Set y-axis positions
y_pos = np.arange(4)

### Create horizontal bars
ax.barh(y_pos, scores_3, height=0.18, color='#FAB972', label='3 Constraint Words')
ax.barh(y_pos + 0.15, scores_5, height=0.18, color='#F7CEEB', label='5 Constraint Words')
ax.barh(y_pos + 0.3, scores_7, height=0.18, color='#FF7F50', label='7 Constraint Words')
ax.barh(y_pos + 0.45, scores_10, height=0.18, color='#ADD8E6', label='10 Constraint Words')

### Add labels and title
ax.set_title('Scores by group and constraint word count', fontsize=22, fontweight='bold')
ax.set_xlabel('Score', fontsize=12, fontweight='bold')
ax.set_ylabel('', fontsize=12)

### Set ticks and grid
ax.set_xticks([-3.0, -2.5, -2.0, -1.5, -1.0, -0.5, 0.0])
ax.set_yticks(y_pos + 0.25)
ax.set_yticklabels(labels, fontsize=11, weight='bold')

### Add legend in top-left corner with proper positioning
ax.legend(loc='upper left', bbox_to_anchor=(0, 1), fontsize=11, framealpha=True, edgecolor='black')

### Add annotations
for i, (score3, score5, score7, score10) in enumerate(zip(scores_3, scores_5, scores_7, scores_10)):
    ax.text(score3, i, f'{score3:.1f}', va='center', ha='right', fontsize=11, weight='bold')
    ax.text(score5, i + 0.15, f'{score5:.1f}', va='center', ha='right', fontsize=11)
    ax.text(score7, i + 0.3, f'{score7:.1f}', va='center', ha='right', fontsize=11)
    ax.text(score10, i + 0.45, f'{score10:.1f}', va='center', ha='right', fontsize=11)

### Invert y-axis properly
ax.invert_yaxis()

### Display grid lines
ax.xaxis.grid(True, linestyle='-', linewidth=0.5)
ax.yaxis.grid(False)

### Show plot
plt.tight_layout()
plt.show()
```

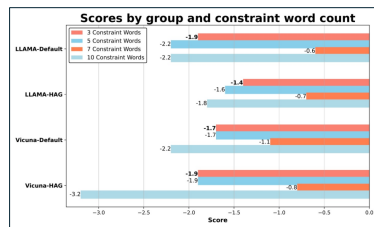
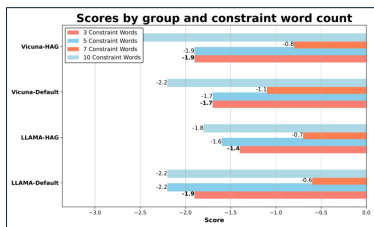



Figure A6. Qualitative result of MM-ReCoder for self-correction. The model corrects the category orders in the second turn.



First Turn

<think> Looking at the plot, I need to:
 1. Create appropriate x and y data points that match the trend
 2. Plot them with proper spacing
 3. Add proper axis ticks and labels
 4. Display percentages on points accurately
 5. Use appropriate colors and styles
 6. Add grid lines properly
 </think>

```
import matplotlib.pyplot as plt
import numpy as np

#### Create figure with appropriate size
fig, ax = plt.subplots(figsize=(8, 6))

#### Data points based on visual estimation
x_values = [0, 8, 16, 24, 32, 40, 48, 56, 64, 96, 128]
y_values = [0, 8, 16, 24, 32, 40, 48, 56, 64, 96, 128]
percentages = [99.6, 99.9, 96.0, 92.3, 90.9, 99.5, 90.9, 91.3, 92.5, 93.3, 98.7]

#### Plot the points and connecting line
ax.scatter(x_values, y_values, color='#1f77b4', s=100, marker='o')
ax.plot(x_values, y_values, color='black', linewidth=1.5)


#### Add annotations for percentages
for i, txt in enumerate(percentages):
    ax.annotate(f'{txt:.1f}%', (x_values[i], y_values[i]), textcoords="offset
points", xytext=(0,10), ha='center', fontsize=11, fontweight='bold')

#### Set axis ticks and labels
ax.set_xticks([0, 8, 16, 24, 32, 40, 48, 56, 64, 96, 128])
ax.set_yticks([0, 8000, 16000, 24000, 32000, 40000, 48000, 56000, 64000, 96000, 128000])
ax.set_xlabel('HCsub-first<sub></sub> (before aging)', fontsize=12, fontweight='bold')
ax.set_ylabel('HCsub-first<sub></sub> (after aging)', fontsize=12, fontweight='bold')

#### Set grid with appropriate style
ax.grid(True, linestyle='--', linewidth=0.5)

#### Set ticks and labels properly
ax.set_title('', fontsize=14)

#### Display plot
plt.tight_layout()
plt.show()
```



Second Turn

<think> The issue is with proper axis formatting and proper spacing. Also, the y-axis ticks need proper alignment without being skewed incorrectly. The x-axis labels need proper alignment without being too skewed. Also, the percentages need proper positioning without being cut off incorrectly. The y-axis needs proper scientific notation formatting without being wrong. Also, the title/spacing needs adjustment better. Also, the proper formatting needs adjustment better without being skewed wrong.
 </think>

```
import matplotlib.pyplot as plt
import numpy as np

#### Create figure with appropriate size
fig, ax = plt.subplots(figsize=(8, 6))

#### Data points based on visual estimation
x_values = [0, 8, 16, 24, 32, 40, 48, 56, 64, 96, 128]
y_values = [0, 8000, 16000, 24000, 32000, 40000, 48000, 56000, 64000, 96000, 128000]
percentages = [99.6, 99.9, 96.0, 92.3, 90.9, 99.5, 90.9, 91.3, 92.5, 93.3, 98.7]

#### Plot the points and connecting line
ax.scatter(x_values, y_values, color='#1f77b4', s=100, marker='o')
ax.plot(x_values, y_values, color='black', linewidth=1.5)

#### Add annotations for percentages
for i, txt in enumerate(percentages):
    ax.annotate(f'{txt:.1f}%', (x_values[i], y_values[i]), textcoords="offset
points", xytext=(0,10), ha='center', fontsize=11, fontweight='bold')

#### Set axis ticks and labels properly
ax.set_xticks([0, 8, 16, 24, 32, 40, 48, 56, 64, 96, 128])
ax.set_yticks([0, 8000, 16000, 24000, 32000, 40000, 48000, 56000, 64000, 96000, 128000])
ax.set_xlabel('HCsub-first<sub></sub> (before aging)', fontsize=12, fontweight='bold')
ax.set_ylabel('HCsub-first<sub></sub> (after aging)', fontsize=12, fontweight='bold')

#### Set proper ticks without scientific notation issues
ax.set_title('', fontsize=14)

#### Display proper grid with appropriate style
ax.grid(True, linestyle='--', linewidth=0.5)

#### Add proper tick labels without being skewed wrong
#### Remove y-axis scientific notation properly

#### Display plot
plt.tight_layout()
plt.show()
```





Figure A7. Qualitative result of MM-ReCoder for self-correction. The model corrects the y-axis values of the data points in the second turn.



First Turn

<think> Looking at the plot, I need to create scatter plots with appropriate colors and proper spacing. Since I don't have actual data, I'll need to generate synthetic data that matches the distribution shown in the plot. The data appears to be clustered in different regions with distinct colors. I'll create:

- Properly spaced points for each category
- Appropriate colors matching the original (blue, pink, yellow, green)
- Proper axis ticks and labels
- Legend positioned correctly

</think>

```
import matplotlib.pyplot as plt
import numpy as np
import pandas as pd

#### Set figure size
fig, ax = plt.subplots(figsize=(8, 8))

#### Generate synthetic data (approximating the distribution shown)
np.random.seed(42)

#### Create sample data for each category
#### ar = blue/purple
ar_x = np.random.normal((0.45, 0.35), (0.05, 0.05), 200)
ar_y = np.random.normal((0.35, 0.05), 200)
ax.scatter(ar_data[:, 0], ar_data[:, 1], color='r', marker='o', s=30, label='ar')

#### de = pink/magenta
de_x = np.random.normal((-0.25, 0.15), (0.05, 0.05), 200)
de_y = np.random.normal((0.15, 0.05), 200)
ax.scatter(de_data[:, 0], de_data[:, 1], color='g', marker='o', s=30, label='de')

#### fr = yellow
fr_x = np.random.normal((-0.3, -0.15), (0.05, 0.05), 200)
fr_y = np.random.normal((-0.15, 0.05), 200)
ax.scatter(fr_data[:, 0], fr_data[:, 1], color='b', marker='o', s=30, label='fr')

#### he = green
he_x = np.random.normal((0.15, 0.25), (0.05, 0.05), 200)
he_y = np.random.normal((0.25, 0.05), 200)
ax.scatter(he_data[:, 0], he_data[:, 1], color='m', marker='o', s=30, label='he')

#### Add labels and title
ax.set_title('Legend', fontsize=16, fontweight='bold')
ax.set_xlabel('PC1', fontsize=12, fontweight='bold')
ax.set_ylabel('PC2', fontsize=12, fontweight='bold')

#### Add legend in top center position
ax.legend(title='Legend', loc='upper center', bbox_to_anchor=(0.5, 1.15),
         fontsize=11, title_fontsize=12, framealpha=True, edgecolor='black')

#### Set appropriate ticks
ax.set_xticks([-0.4, -0.2, 0.0, 0.2, 0.4])
ax.set_yticks([-0.2, -0.1, 0.0, 0.1, 0.2, 0.3, 0.4])

#### Display plot
plt.tight_layout()
plt.show()
```

Second Turn

<think> I need to fix the issue with generating proper 2D data. Instead of trying to create a 2D array directly, I should generate proper x and y values separately. Also, the legend needs proper alignment without causing errors. I need to create proper scatter plots without trying to combine them incorrectly. I'll generate proper 1D arrays and plot them appropriately.

</think>

```
import matplotlib.pyplot as plt
import numpy as np

#### Set figure size
fig, ax = plt.subplots(figsize=(8, 8))

#### Generate synthetic data (approximating the distribution shown)
np.random.seed(42)

#### Create sample data for each category
#### ar = blue/purple
ar_x = np.random.normal(0.45, 0.07, 200)
ar_y = np.random.normal(0.35, 0.05, 200)
ax.scatter(ar_x, ar_y, color='r', marker='o', s=30, label='ar')

#### de = pink/magenta
de_x = np.random.normal(-0.25, 0.07, 200)
de_y = np.random.normal(0.15, 0.05, 200)
ax.scatter(de_x, de_y, color='g', marker='o', s=30, label='de')

#### fr = yellow
fr_x = np.random.normal(-0.3, 0.07, 200)
fr_y = np.random.normal(-0.15, 0.05, 200)
ax.scatter(fr_x, fr_y, color='b', marker='o', s=30, label='fr')

#### he = green
he_x = np.random.normal(0.15, 0.07, 200)
he_y = np.random.normal(0.25, 0.05, 200)
ax.scatter(he_x, he_y, color='m', marker='o', s=30, label='he')

#### Add labels and title
ax.set_title('Legend', fontsize=16, fontweight='bold')
ax.set_xlabel('PC1', fontsize=12, fontweight='bold')
ax.set_ylabel('PC2', fontsize=12, fontweight='bold')

#### Add legend in top center position with proper positioning
ax.legend(title='Legend', loc='upper center', bbox_to_anchor=(0.5, 1.15),
         fontsize=11, title_fontsize=12, framealpha=True, edgecolor='black')

#### Set appropriate ticks
ax.set_xticks([-0.4, -0.2, 0.0, 0.2, 0.4])
ax.set_yticks([-0.2, -0.1, 0.0, 0.1, 0.2, 0.3, 0.4])

#### Display plot
plt.tight_layout()
plt.show()
```

Third Turn

<think> The legend needs proper alignment without being skewed incorrectly, and it should appear in the top center without causing confusion. Also, the plot needs proper spacing without the title being cut incorrectly. Also, the positioning needs adjustment better without causing overlap issues properly. The proper positioning needs adjustment without creating an inconsistent legend. I need to create proper spacing and have the legend in the correct position without causing conflicts. Also, the y-axis ticks need proper alignment without being skewed wrong. Also, the proper spacing needs adjustment better.

</think>

```
import matplotlib.pyplot as plt
import numpy as np

#### Set figure size
fig, ax = plt.subplots(figsize=(8, 8))

#### Generate synthetic data (approximating the distribution shown)
np.random.seed(42)

#### Create sample data for each category
#### ar = blue/purple
ar_x = np.random.normal(0.45, 0.07, 200)
ar_y = np.random.normal(0.35, 0.05, 200)
ax.scatter(ar_x, ar_y, color='r', marker='o', s=30, label='ar')

#### de = pink/magenta
de_x = np.random.normal(-0.25, 0.07, 200)
de_y = np.random.normal(0.15, 0.05, 200)
ax.scatter(de_x, de_y, color='g', marker='o', s=30, label='de')

#### fr = yellow
fr_x = np.random.normal(-0.3, 0.07, 200)
fr_y = np.random.normal(-0.15, 0.05, 200)
ax.scatter(fr_x, fr_y, color='b', marker='o', s=30, label='fr')

#### he = green
he_x = np.random.normal(0.15, 0.07, 200)
he_y = np.random.normal(0.25, 0.05, 200)
ax.scatter(he_x, he_y, color='m', marker='o', s=30, label='he')

#### Add labels and title
ax.set_title('', fontsize=16, fontweight='bold')
ax.set_xlabel('PC1', fontsize=12, fontweight='bold')
ax.set_ylabel('PC2', fontsize=12, fontweight='bold')

#### Add proper legend in top center without title conflict
ax.legend(title='Legend', loc='upper center', bbox_to_anchor=(0.5, 1.15),
         fontsize=11, title_fontsize=12, framealpha=True, edgecolor='black')

#### Set appropriate ticks
ax.set_xticks([-0.4, -0.2, 0.0, 0.2, 0.4])
ax.set_yticks([-0.2, -0.1, 0.0, 0.1, 0.2, 0.3, 0.4])

#### Display plot
plt.tight_layout()
plt.show()
```

Traceback (most recent call last):
 File "<string>", line 50, in code
 ar_data = np.random.normal((0.45, 0.35), (0.05, 0.05), 200)
 File "numpy/random/mtrand.pyx", line 1557, in numpy.random.mtrand.RandomState.normal
 File "_common.pyx", line 600, in numpy.random._common.cont
 File "_common.pyx", line 517, in numpy.random._common.cont_broadcast_2
 File "_init_.cython-30.pxd", line 780, in numpy.PyArray_MultiIterNew3
 ValueError: shape mismatch: objects cannot be broadcast to a single shape. Mismatch is between arg 0 with shape (200,) and arg 1 with shape (2,).





Figure A8. Qualitative result of MM-ReCoder for self-correction. In the second turn, the model addresses the runtime error. In the third turn, the model removes the chart title overlapped with the legend.




Figure A9. Qualitative results of MM-ReCoder and baselines.

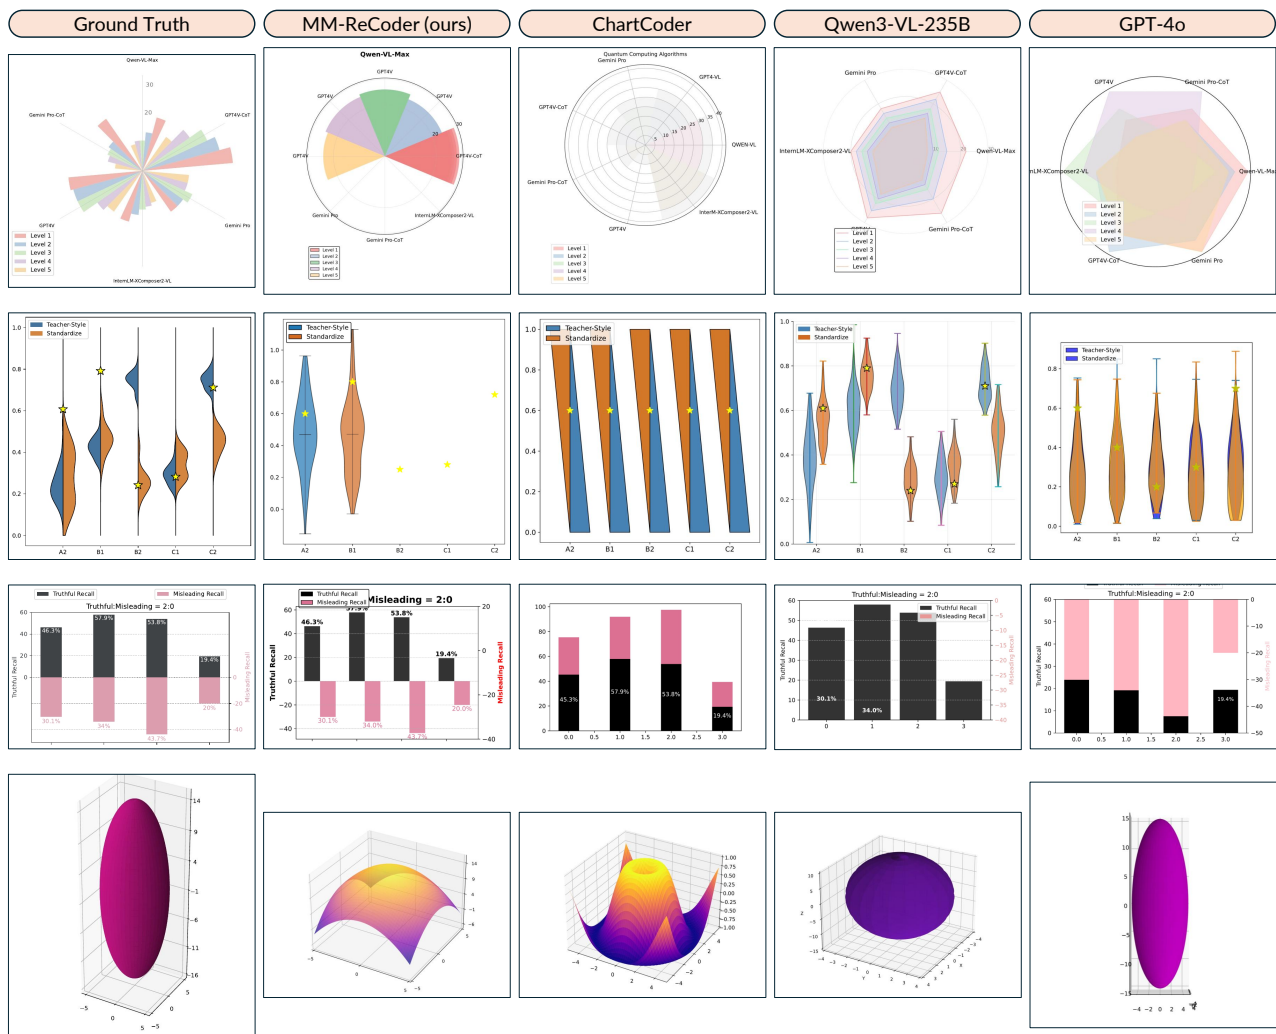


Figure A10. Failure cases of MM-ReCoder and baselines. The models tend to fail when the chart type is rare.