

CHARTEDIT: How Far Are MLLMs From Automating Chart Analysis? Evaluating MLLMs' Capability via Chart Editing

Xuanle Zhao^{1,*}, Xuexin Liu^{1,*}, Haoyue Yang^{1,*}, Xianzhen Luo²,
Fanhu Zeng¹, Jianling Li³, Qi Shi^{4,†}, Chi Chen^{4,†}

¹ Institute of Automation, Chinese Academy of Sciences, Beijing, China

² Harbin Institute of Technology, Harbin, China

³ Tianjin University, Tianjin, China ⁴ Tsinghua University, Beijing, China
{zhaoxuanle2022, liuxuexin2022, yanghaoyue2024}@ia.ac.cn

Abstract

Although multimodal large language models (MLLMs) show promise in generating chart rendering code, editing charts via code presents a greater challenge. This task demands MLLMs to integrate chart understanding and reasoning capacities, which are labor-intensive. While many MLLMs claim such editing capabilities, current evaluations rely on limited case studies, highlighting the urgent need for a comprehensive evaluation framework. In this work, we propose CHARTEDIT, a novel benchmark designed for chart editing tasks, featuring 1405 diverse editing instructions applied to 233 real-world charts, each manually annotated and validated for accuracy. Utilizing CHARTEDIT, we evaluate the performance of 10 mainstream MLLMs across two types of experiments at both the code and chart levels. The results suggest that large-scale models can generate code to produce images that partially match the reference images. However, their ability to generate accurate edits according to the instructions remains limited. The state-of-the-art (SOTA) model achieves a score of only 59.96, highlighting significant challenges in precise modification. In contrast, small-scale models, including chart-domain models, struggle both with following editing instructions and generating overall chart images, underscoring the need for further development in this area. Code is available at <https://github.com/xx111z/ChartEdit>.

1 Introduction

Data visualization is a crucial component in various fields, enabling individuals and organizations to present and interpret data effectively (Chen et al., 2008). However, creating high-quality, visually appealing charts from scratch can be time-consuming and often requires extensive adjustments and references to documentation. This highlights the need

*Equal contribution.

†Corresponding authors.

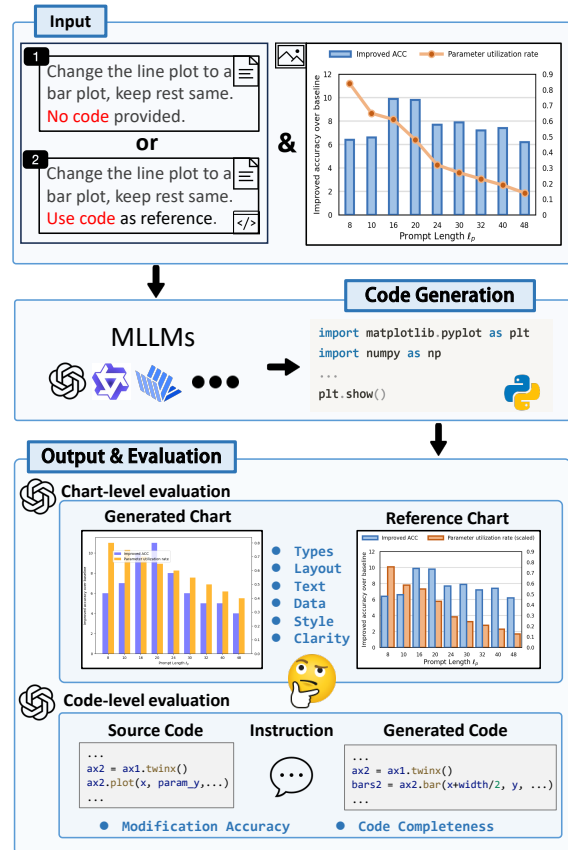


Figure 1: Overall pipeline. The inputs are a chart, editing instruction w/ or w/o code. The MLLMs are instructed to generate the edited code. The final evaluation is constructed at both the code level and chart level.

for more efficient solutions to streamline the visualization process. Previous works have explored utilizing textual descriptions to generate visualization code automatically through the model (Yang et al., 2024; Zadeh et al., 2024), which significantly reduces the time and effort required for data visualization, making it more accessible and efficient for researchers.

Recently, multimodal large language models (MLLMs) (Liu et al., 2023a; Zhang et al., 2024a; Guo et al., 2025; Zhang et al., 2025), which leverage the rich knowledge in large language models

(LLMs) (Dubey et al., 2024), have demonstrated impressive performance in processing and reasoning over various modalities, such as image (Wang et al., 2024b; Yu et al., 2025), video (Li et al., 2024) and speech (Wang et al., 2024d). However, much of the existing research in code generation has primarily focused on using text as the sole input (Li et al., 2022; Luo et al., 2023), leaving the potential of multimodal information largely unexplored.

Although there have been recent efforts to explore multimodal code generation in specific areas like charts (Shi et al., 2024; Zhang et al., 2024c) and webpages (Yun et al., 2024; Si et al., 2024), these works typically address direct chart/web-to-code tasks, which are akin to using code for captioning. These approaches often lack diverse instructions and fail to account for interactive or iterative modifications. This raises an important question: Can current MLLMs function like expert chart analysts, generating modified chart code based on both the source chart and detailed editing instructions?

This is a far more complex challenge, as it requires MLLMs to not only extract relevant information from the chart but also generate corresponding code and adapt it according to the provided editing instructions. However, there is currently no high-quality, diverse benchmark for evaluating chart editing tasks. Previous works, such as ChartMimic (Shi et al., 2024), focus mainly on data-centric modifications, while the ChartLlama dataset lacks real-world charts, limiting the scope and relevance of their evaluations.

To address this gap, we introduce CHARTEDIT, a comprehensive evaluation benchmark consisting of 233 real-world charts sourced from Arxiv, each accompanied by manually annotated code and 1405 chart instructions with reference edited code. The instructions and code are either human-written or initially generated by LLMs, followed by manual correction and alignment to ensure accuracy. To ensure diversity, we pre-define 19 chart types and six types of editing instructions. We also introduce two task formats, input chart *w/* code or *w/o* code, simulating real-world scenarios that require code editing from only the charts or charts with code. Furthermore, we establish and evaluate metrics based on execution rate, code-level accuracy, and chart-level consistency to comprehensively assess MLLMs’ capabilities. These metrics measure the models’ ability to generate executable code, maintain editing precision, and ensure consistency across the generated visualizations. We

conduct experiments on three types of MLLMs: proprietary, general-domain open-source and chart-domain models. The results reveal that, despite notable advancements, state-of-the-art open-source models still show performance gaps compared to GPT-4o. Additionally, current small-scale MLLMs, including chart-domain models, continue to face significant challenges in generating editing code that accurately follows the instructions. Furthermore, we also investigate the impact of Chain-of-Thought (CoT) prompting (Wei et al., 2022) and analyze the models’ performance across various editing instructions and chart types.

2 Related Works

2.1 Chart-Domain MLLMs

Recently, MLLMs have achieved superior performance on many visual tasks by leveraging connectors to bridge the gap between large language models and vision encoders (Liu et al., 2023a,b). As a significant image type, chart-related tasks have received much attention. Previous works utilize a two-stage method that first extracts information from the chart and then utilizes language models to process the information (Liu et al., 2022). Currently, end-to-end MLLMs are utilized to solve chart-related tasks with a unified model. ChartLlama (Han et al., 2023) direct finetuning based on the existing LLaVA (Liu et al., 2023b). mPLUG-Owl (Ye et al., 2023) and mPLUG-Owl2 (Ye et al., 2024) achieve superior performance on high-resolution chart images. ChartVLM (Xia et al., 2024) utilizes a discriminator to determine whether intervention from the LLMs is required for a specific query. TinyChart (Zhang et al., 2024b) proposes the token merging and PoT-based reasoning strategy to improve inference efficiency and understanding capacity. ChartMoE (Xu et al., 2024) utilize the Mixture-of-Expert (MoE) method to align many data formats with charts.

2.2 MLLMs For Code

Previous advancements in LLMs have significantly contributed to code-related tasks. Notable models, such as DeepSeek Coder (Guo et al., 2024) and StarCoder (Lozhkov et al., 2024), have demonstrated substantial progress in tasks like code generation and error fixing. However, these models typically operate in a single-modal setting, relying on textual input, which limits their ability to address the complexity and range of problems.

Name	Output Format	w/ Real Chart	w/ Cor. Code	Diverse Types	Open Domain	Editing Instruction				
						Data	Format	Layout	Style	Text
ChartCraft (Yan et al., 2024)	Json	✗	✗	✗	✗	✓	✓	✓	✓	✗
Plot2Code (Wu et al., 2024)	Code	✓	✗	✓	✓	✗	✗	✗	✗	✗
ChartX (Xia et al., 2024)	Code	✗	✓	✓	✓	✗	✗	✗	✗	✗
AcademiaChart (Zhang et al., 2024c)	Code	✓	✗	✗	✓	✗	✗	✗	✗	✗
ChartMimic (Shi et al., 2024)	Code	✓	✓	✓	✓	✓	✗	✗	✗	✗
CHARTEDIT (OURS)	Code	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 1: The comparison of our proposed CHARTEDIT evaluation benchmark with other chart-related benchmarks. CHARTEDIT is the first evaluation benchmark, and it contains various editing instructions and reference codes.

Multimodal code generation has recently gained significant attention. Several works, like Design2Code (Si et al., 2024) and Web2Code (Yun et al., 2024), focus on assessing the capacities of MLLMs in generating HTML code for web pages. RoboCodeX (Mu et al., 2024) proposes a multimodal code generation framework for robot behaviour synthesis. The field of chart-to-code generation has also attracted considerable attention (Zhao et al., 2025), as it focuses on generating code that accurately reproduces a given chart image. This task is challenging due to the complex visual elements in charts, demanding advanced models to accurately translate them into functional code. Recent works (Wu et al., 2024; Xia et al., 2024; Zhang et al., 2024c) evaluate the capacities of MLLMs in this context. Although some works like (Xia et al., 2024; Shi et al., 2024) have considered this problem, the editing instructions in their works only focus on one type of modification, which is not diverse enough.

3 CHARTEDIT

In this section, we first introduce the definition of the chart editing task and illustrate the data collection and organization pipeline of CHARTEDIT.

3.1 Task Definition

In this work, we aim to leverage MLLMs to edit charts and generate the corresponding code as instructed. The task can be summarized as follows: given a chart image X from Arxiv papers, along with an editing instruction I , the model is expected to generate the corresponding visualization code after applying the edits, regardless of whether the original code C is provided. Thus, the task can be represented as:

$$O = M(X \vee (X, C), I) \quad (1)$$

where M denotes the processing model. $X \vee (X, C)$ indicates that the model can either receive

X (Chart *w/o* Code) or (X, C) (Chart *w/* Code) as input. O is the output code in Python, which utilize libraries like Matplotlib and Seaborn to plot the edited chart.

3.2 Data Construction

3.2.1 Chart Collection and Filtering

In this study, to collect real-world chart images, we begin by crawling papers published on Arxiv using web scraping tools like BeautifulSoup. To ensure the quality of the collected images, we first retrieve the IDs and comments of papers published on Arxiv, then filter the papers by selecting those whose comments include keywords such as “submit”, “accept”, “under review” and “camera ready”. After this filtering step, we use the ArXiv API to download the LaTeX source files for the selected papers. After downloading the source files, we remove irrelevant files based on their suffixes, retaining only those ending in .png, .pdf, .jpg, and .svg. However, many of these images are not charts that can be reproduced by visualization code. To filter out high-quality chart images, we employ a two-step process. First, we use an MLLM to assess whether each image is a chart through zero-shot prompting. Despite this assessment, we found that many low-quality chart images remained. To address this, we developed a scoring mechanism, instructing the MLLM to evaluate the images based on four criteria: *Aesthetics*, *Readability*, *Reproducibility*, and *Data Presentation Simplicity*. Each image is assigned a total score of 100, and we filter out those scoring below 90. After completing this two-step filtering process, we are left with more than 1,000 high-quality chart images.

3.2.2 Code Annotation

While we have collected a sufficient number of chart images, the charts crawled from Arxiv do not include the code required for reproduction. To refine our dataset and facilitate the extraction of

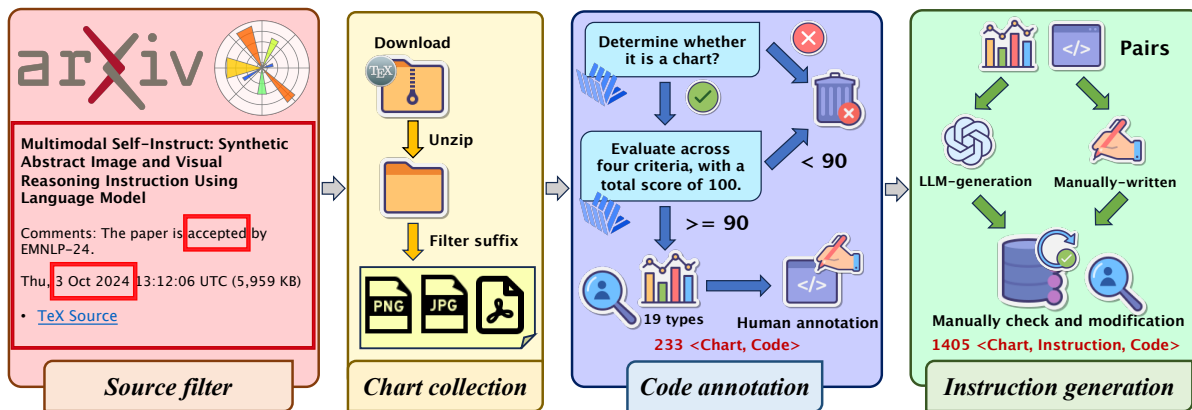


Figure 2: The pipeline for constructing the CHARTEDIT evaluation dataset begins with filtering and crawling ArXiv papers based on keywords found in their comments. After that, we remove irrelevant files by filtering out specific suffixes. We then use an MLLM to distinguish and filter out non-chart images, scoring the remaining images. These images are further screened based on these scores and reviewed by human evaluators. Also, Code annotations are manually written by human evaluators. The editing instructions and reference edited code are constructed utilizing two strategies: one based on LLM and the other manually written. Finally, all the <Chart, Instruction, Code> triplets are reviewed and modified by human evaluators to enhance correspondence and accuracy.

edited code, we manually filter and write Python code that can reproduce the corresponding charts. However, some chart types are still missing. To address this, we supplement the dataset with additional charts from other platforms, such as Kaggle and the Matplotlib gallery (Hunter, 2007). To prevent data leakage, we manually modify the code to generate visually distinct versions of the original charts. As a result, we have successfully obtained 233 charts, along with their corresponding code, forming <chart, code> source pairs.

3.2.3 Instruction Generation

After obtaining the chart source code, we propose two approaches for constructing editing instructions: (1) LLM-based generation and (2) human-written instructions. In the LLM-based approach, we predefine five key editing categories: *style*, *format*, *layout*, *data*, and *text*, each with multiple specific subtypes. By providing the LLM with both the source code and the chart image, we prompt it to generate editing instructions and the corresponding edited code. To ensure a diverse set of outputs and minimize errors, we generate at least three variations for each chart and instruction type. To ensure the instructions and code are rigorous and aligned with human expectations, we manually verify and modify all instructions, code, and edited images for consistency. This process results in a total of 1,172 <chart, instruction, code> triplets. While LLM-based generation can produce a large volume of triplets, the diversity of descriptions and their




Figure 3: The number and specific proportions of different types of editing instructions in our CHARTEDIT evaluation benchmark.

relevance to real-world needs may be constrained by inherent biases in LLM outputs. To mitigate this limitation, we manually write one editing instruction for each chart with reference code, ensuring that the dataset more closely aligns with real-world requirements.

3.3 Dataset Statistics and Analysis

Using LLM generation and manual annotation, we construct our evaluation benchmark, CHARTEDIT, which comprises 233 <chart, code> pairs and 1,405 <chart, instruction, code> triplets. We further analyze the diversity of CHARTEDIT from both chart and instruction perspectives. To ensure a diverse selection of chart types, we have predefined 19 distinct chart categories, carefully curating instances for each type. The distribution of these chart types is provided in the Appendix A.2. Additionally, we enhance the diversity of editing instructions by classifying them into six categories: five




Figure 4: The dimension reduction of the editing instructions in CHARTEDIT with various colors represents different editing types. We choose the Sentence-BERT (Reimers, 2019) as the embedding model.

generated by LLMs, as outlined in Section 3.2.3, and one written by humans. The proportions of each editing category are shown in Figure 3. In addition to exploring the diversity of editing types, we also examine the different formats in which these instructions are presented, including plain text descriptions, code symbols such as `marker='*'`, and hexadecimal color codes like `color='#ABCDEF'`. We embed these editing instructions and compare them with instructions from other chart-to-code benchmarks using dimensionality reduction techniques. The result in Figure 4 shows embedding after dimension reduction, which indicates that different editing types are generally separated, and the human-written instructions almost cover all the instruction types.

4 Experiments

4.1 Baselines

To facilitate a more effective comparison of existing MLLMs in chart editing tasks, we benchmark three categories of widely used MLLMs: proprietary, open-source general-domain, and chart-domain models. **(1) Proprietary Models:** We evaluate two top-performing models in the multimodal domain: GPT-4o (OpenAI, 2024) and Claude-3.5-Sonnet (Anthropic, 2024) which represent the cutting edge in the multimodal domain. **(2) Open-Source General-Domain Models:** We select five competitive open-source models from a variety of sizes, listed in descending order of total parameters. These models include InternVL-V2.5-78B (Chen et al., 2024), Qwen2-VL-72B (Wang et al., 2024a),

Statistic	Value
Charts	
Total Chart	233
Types	19
Average size (px)	876×575
Maximum size (px)	3000×1600
Instruction	
Total Instruction	1405
Edit Types	6
Code	
Average Length	650
Maximum Length	3310

Table 2: ChartEdit dataset statistics. The code length is calculated based on the number of tokens utilizing the Llama3.2 tokenizer (Dubey et al., 2024).

LLaVA (Liu et al., 2023b), InternVL-V2.5-8B and Phi-3.5-Vision (Abdin et al., 2024). These models represent the base models commonly used in the multimodal domain. **(3) Chart-Domain Models:** For this category, we focus on models already designed for chart-to-code generation tasks, including ChartLlama (Han et al., 2023), TinyChart (Zhang et al., 2024b), and ChartMoE (Xu et al., 2024).

For the evaluation, we use the direct instruction prompting method across all models. Detailed prompts and implementation steps are provided in the Figure 11.

4.2 Evaluation Metrics

Given the lack of established evaluation metrics for assessing the completeness of image editing, we draw inspiration from recent research that uses LLMs as evaluators (Gu et al., 2024; Zheng et al., 2024). In this work, we leverage GPT-4o to evaluate two key aspects: first, *whether the generated code aligns with the provided editing instructions*, and second, *whether the generated code effectively produces the intended chart*. Therefore, we propose evaluating both the code level and the chart level. The code-level evaluation is based on the source code, editing instructions, and output code. We use GPT-4o to score the models on two aspects: *Modification Accuracy* and *Code Completeness*. Building on related approaches (Shi et al., 2024; Zhang et al., 2024c), we also evaluate the model’s overall generation capabilities at the chart level. This evaluation is performed by comparing how closely the generated chart reproduces the reference edited chart (manually created). The results reflect the degree of completeness of the gener-

Model	Chart <i>w/o</i> Code			Chart <i>w/</i> Code		
	Exec.Rate	Code-Level	Chart-Level	Exec.Rate	Code-Level	Chart-Level
<i>Proprietary Models</i>						
GPT-4o (OpenAI, 2024)	91.46	59.96	79.87	98.89	89.96	93.68
Claude-3.5-Sonnet (Anthropic, 2024)	88.22	47.32	54.92	89.50	88.99	81.68
<i>Open-Source General-Domain Models</i>						
InternVL2.5-78B (Chen et al., 2024)	79.66	55.67	70.77	94.31	92.56	92.81
Qwen2-VL-72B (Wang et al., 2024a)	81.65	46.67	64.06	86.09	90.30	79.51
InternVL2.5-8B (Chen et al., 2024)	62.37	39.24	45.85	85.20	90.16	87.36
Phi3.5-Vision (Abdin et al., 2024)	67.13	40.19	37.65	74.44	82.31	89.16
LLaVA-13B (Liu et al., 2023b)	48.75	11.35	16.71	30.44	71.88	25.82
<i>Chart-Domain Models</i>						
ChartMoE (Xu et al., 2024)	53.44	21.60	34.73	81.89	84.87	89.31
ChartLlama (Han et al., 2023)	52.30	15.82	27.42	46.34	50.18	47.00
TinyChart (Zhang et al., 2024b)	36.13	18.34	25.09	3.51	18.40	2.54

Table 3: Evaluation results of various baseline models on Chart *w/o* Code and Chart *w/* Code tasks. The performance is evaluated from three aspects: the code Execution Rate, Code-Level, and Chart-Level scores. The best performances are indicated in **bold**.

Model	Types	Layout	Text	Data	Style	Clarity
GPT-4o	18.35	9.49	15.26	14.87	13.55	8.53
InternVL2.5-78B	16.25	9.03	13.13	13.59	11.12	7.91
InternVL2.5-8B	10.22	7.68	8.21	7.21	6.19	6.34
TinyChart	4.99	4.60	4.39	3.71	3.41	4.16

Table 4: Detailed results of Chart-Level scores on Chart *w/o* Code task.

Model	Modification Acc	Code Complete
GPT-4o	30.67	29.29
InternVL2.5-78B	29.05	26.62
InternVL2.5-8B	20.16	19.08
TinyChart	6.46	11.88

Table 5: Detailed results of Code-Level scores on Chart *w/o* Code task.

ated chart image. Following (Shi et al., 2024), we also utilize GPT-4o to evaluate six aspects: *Types*, *Layout*, *Text*, *Data*, *Style*, and *Clarity*. Detailed information is provided in the Figure 13.

4.3 Main Results

The main results for all the MLLMs are presented in Table 3. Although open-source models have made significant strides, and in some cases even outperformed GPT-4 in various tasks (Masry et al., 2022; Zhang et al., 2024a), there is still a noticeable gap when it comes to handling complex multimodal tasks (Wang et al., 2024c; Shi et al., 2024). In our experiments, GPT-4o delivers the best performance on the Chart *w/* Code task, achieving the highest Execution Rate, Code-Level, and Chart-Level scores. However, the Code-Level score is still not high enough, indicating that GPT-4o faces challenges with precise editing. Among open-source

MLLMs, InternVL2.5-78B achieves the best performance, but it still lags behind GPT in terms of code execution rate and chart-level metrics. Our analysis shows that the largest gap between proprietary and open-source models lies in generating the complete code corresponding to an image. There is little difference in instruction-following ability between the two, and both struggle with detailed modifications in chart editing. The chart-domain models perform worse in this task, and we believe this is due to the current limitations of MLLMs in the chart-domain, particularly in handling editing tasks. Although TinyChart (Zhang et al., 2024b) and ChartLlama (Han et al., 2023) have been fine-tuned on relevant datasets, both lack the ability to follow editing instructions effectively. The instructions in chart editing are much more diverse than those in direct chart-to-code generation tasks.

In the Chart *w/* Code task, providing the source code of the input chart leads to significant performance improvements across all proprietary and open-source general-domain models. These models can generate more accurate code, which we believe is due to the source code supplying essential information, such as the chart’s data, text, and style. This significantly alleviates the challenge for the model in extracting relevant information from the chart itself. In this context, the task becomes more like editing the code with the chart images serving as auxiliary context, which is much easier than the Chart *w/o* Code task. However, we observe that most Chart-Domain models perform significantly worse in this setting. Upon analyzing TinyChart, we found that it lacks training data for relevant




Figure 5: The result comparisons between GPT-4o and InternVL2.5-78B in direct and zero-shot CoT prompting setting. In most cases, the effect of CoT prompting shows a negligible improvement over direct prompting.

tasks, especially those involving longer code snippets in the input. ChartLlama claims to support chart editing, but the diversity of chart types and editing instructions in its training data is limited. As a result, its performance decreases slightly when faced with inputs that include code. Overall, tasks with code are generally easier than those without, denoting chart-domain models still require substantial improvements in handling diverse instructions.

The detailed scores of the chart-level are presented in Table 4, and we select four models for analysis. Since the Chart *w/o* Code task is significantly more challenging, we only provide detailed scores for this task. The result shows that the performance gaps between proprietary and open-domain models are most noticeable in the text, data, and style aspects, which are key pieces of information in the chart. Table 5 presents the detailed scores at the code level, highlighting that there remains a gap between the state-of-the-art open-source models and proprietary models in generating complete code. At the same time, the large-scale model has a significant advantage over the small-scale model in terms of modification accuracy.

Furthermore, we also recruit human evaluators to score the editing results of four popular MLLMs. Details about the results are listed in Appendix A.3.

5 Discussion

In this section, we conduct many analyses to answer the following questions.

RQ1: Is Chain-of-Thought (CoT) prompting useful for MLLMs in chart editing tasks for proprietary models? Answer: **No**. In the CoT setting, we provide MLLMs with zero-shot prompting, instructing them to first understand the source chart, then analyze the editing instructions, and finally generate the code. However, the result in




Figure 6: A comparison of task performance among GPT-4o, InternVL2.5-78B, and InternVL2.5-8B across both code-level and chart-level tasks in various instruction categories.

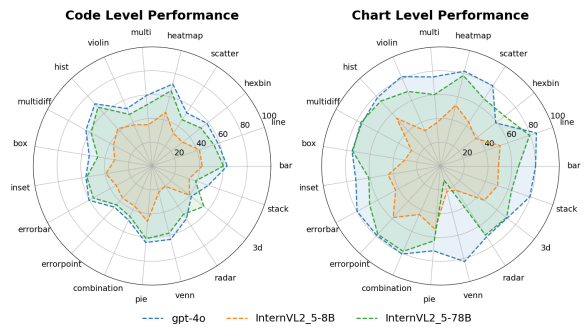


Figure 7: A comparison of task performance among GPT-4o, InternVL2.5-78B, and InternVL2.5-8B across both code-level and chart-level tasks in various chart categories.

Figure 5 shows whether using proprietary models or open-source models, CoT prompting shows almost no improvement in performance. We analyze the thought process of GPT-4o and find that while CoT helps the model better understand the components in the chart and the editing instructions, it does not result in more accurate code generation. The detailed CoT prompts are provided in the Figure 11.

RQ2: How does the performance of MLLMs vary across different types of editing instructions? In Figure 6, we compare the performance across different types of editing tasks and find that models are much more effective at generating code to modify chart types. Our analysis reveals that unlike tasks requiring first grounding and then modification of specific elements, the format conversion task focuses on altering the overall visual representation of the chart, which makes it easier to edit the code. However, at the chart level, model performance on format conversion tasks is limited. This suggests that while models can change the chart type effectively, they still struggle with more complex modifications involving other elements and

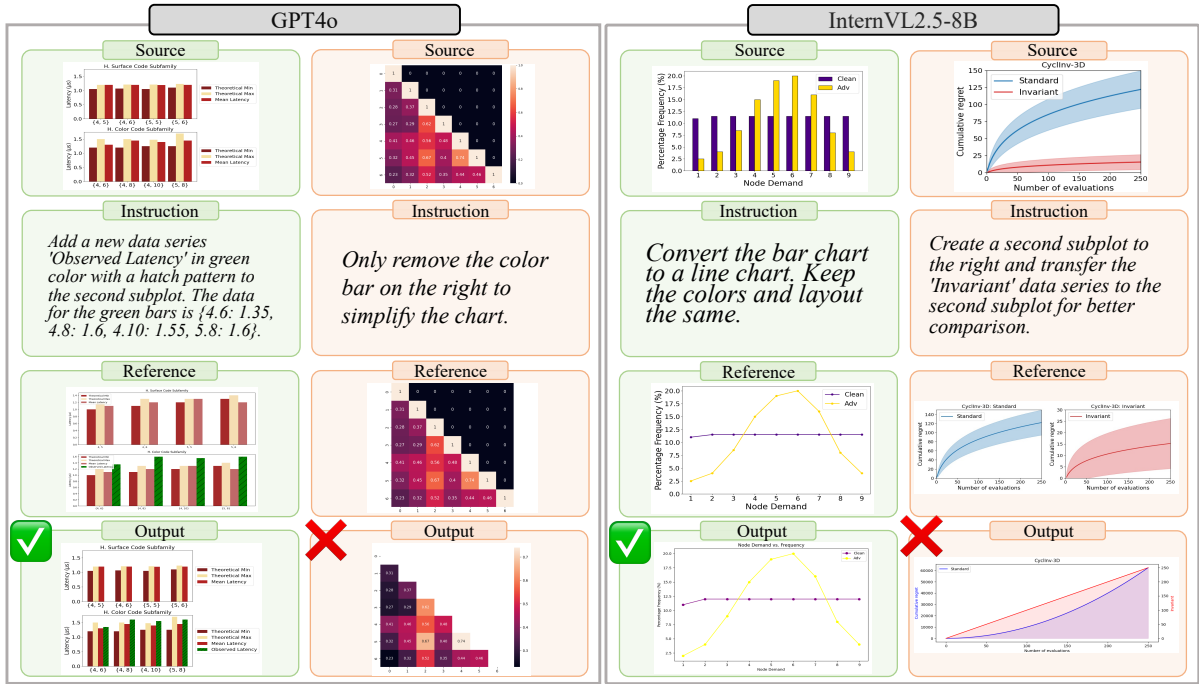


Figure 8: Case studies comparing the state-of-the-art proprietary model GPT-4o with a popular used open-source general-domain model InternVL2.5-8B. Generally, GPT-4o could perform better in a wide range of chart editing tasks, although there are still some tasks it struggles with. InternVL2.5-8B performs reasonably well in handling common chart types like bar and line charts, but it struggles with more complex editing instructions.




Figure 9: A comparison of error types of output code generated by GPT-4o, InternVL2.5-78B, and InternVL2.5-8B.

fail to capture all the details of the chart.

RQ3: How does the performance of MLLMs vary across different types of charts? Figure 7 illustrates model performance across different chart types. The results show that the most significant performance discrepancy between open-source and proprietary models occurs with venn diagrams. Additionally, noticeable differences emerge in handling more complex chart types, such as errorbar and multi charts. Upon further analysis, we found that the comparable performance of InternVL2.5-7B and GPT-4o on the code-level tasks is since this metric primarily evaluates the models' capacity to follow instructions. Since most editing instructions

for venn diagrams do not require altering the chart type itself, the performance gap at the code level is smaller than at the chart level.

RQ4: How do the error types of generated code vary across MLLMs? Figure 9 shows the types of runtime errors generated by different models. The results reveal that ValueError constitutes the largest proportion of errors, accounting for more than half of the total. Interestingly, the proportion of errors excluding ValueError tends to decrease as the model's performance improves. Our analysis suggests that as code generation capabilities improve, many common errors are effectively mitigated while more complex issues related to data processing persist.

6 Conclusion

In this work, we introduce CHARTEDIT, a high-quality benchmark that includes various types of editing instructions, each of which is either manually written or first generated by an LLM and then manually corrected. We evaluate performance across proprietary, open-source general-domain, and chart-domain models. The results demonstrate that proprietary models consistently outperform others in both following editing instructions and generating accurate chart images. Additionally, we

observe that current chart-domain models generally struggle with complex instructions and handling diverse inputs. We believe that automated chart editing in academic research holds great promise, but MLLMs need to further improve their ability to effectively process chart images.

Limitation

From our perspective, our work has several limitations: (1) Our works only consider textual prompts such as Chain-of-Thought. However, more visual prompt methods could be evaluated. (2) The dataset size may not be enough in some situations. Maybe a much larger evaluation benchmark could help to find more interesting findings. (3) More accurate evaluation methods. Our code-level and chart-level metrics are evaluated via LLM, so a better calculation method could be proposed.

Ethical Statement

Our research employs publicly available models and data with proper citations. This approach minimizes the risk of generating toxic content, leveraging the widely used and non-toxic nature of our datasets and prompts.

References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Anthropic. 2024. [Introducing the next generation of claude](#).
- Chun-houh Chen, Wolfgang Härdle, Antony Unwin, and Michael Friendly. 2008. A brief history of data visualization. *Handbook of data visualization*, pages 15–56.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. 2024. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, et al. 2024. A survey on llm-as-a-judge. *arXiv preprint arXiv:2411.15594*.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Yu Wu, YK Li, et al. 2024. Deepseek-coder: When the large language model meets programming—the rise of code intelligence. *arXiv preprint arXiv:2401.14196*.
- Haiyang Guo, Fanhu Zeng, Ziwei Xiang, Fei Zhu, Da-Han Wang, Xu-Yao Zhang, and Cheng-Lin Liu. 2025. Hide-llava: Hierarchical decoupling for continual instruction tuning of multimodal large language model. *arXiv preprint arXiv:2503.12941*.
- Yucheng Han, Chi Zhang, Xin Chen, Xu Yang, Zhibin Wang, Gang Yu, Bin Fu, and Hanwang Zhang. 2023. Chartllama: A multimodal llm for chart understanding and generation. *arXiv preprint arXiv:2311.16483*.
- John D Hunter. 2007. Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(03):90–95.
- Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. 2024. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*.
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. 2022. Competition-level code generation with alphacode. *Science*, 378(6624):1092–1097.
- Fangyu Liu, Julian Martin Eisenschlos, Francesco Piccinno, Syrine Krichene, Chenxi Pang, Kenton Lee, Mandar Joshi, Wenhui Chen, Nigel Collier, and Yasemin Altun. 2022. Deplot: One-shot visual language reasoning by plot-to-table translation. *arXiv preprint arXiv:2212.10505*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. Improved baselines with visual instruction tuning.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning.
- Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, et al. 2024. Starcoder 2 and the stack v2: The next generation. *arXiv preprint arXiv:2402.19173*.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023. Wizardcoder: Empowering code large language models with evol-instruct. *arXiv preprint arXiv:2306.08568*.

- Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*.
- Yao Mu, Junting Chen, Qinglong Zhang, Shoufa Chen, Qiaojun Yu, Chongjian Ge, Runjian Chen, Zhixuan Liang, Mengkang Hu, Chaofan Tao, et al. 2024. Robocodex: Multimodal code generation for robotic behavior synthesis. *arXiv preprint arXiv:2402.16117*.
- OpenAI. 2024. [Gpt-4o](#). Accessed: 2024-05-13.
- N Reimers. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Chufan Shi, Cheng Yang, Yaxin Liu, Bo Shui, Junjie Wang, Mohan Jing, Linran Xu, Xinyu Zhu, Siheng Li, Yuxiang Zhang, et al. 2024. Chartmimic: Evaluating llm’s cross-modal reasoning capability via chart-to-code generation. *arXiv preprint arXiv:2406.09961*.
- Chenglei Si, Yanzhe Zhang, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. 2024. Design2code: How far are we from automating front-end engineering? *arXiv preprint arXiv:2403.03163*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024a. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, et al. 2024b. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. *Advances in Neural Information Processing Systems*, 36.
- Zirui Wang, Mengzhou Xia, Luxi He, Howard Chen, Yitao Liu, Richard Zhu, Kaiqu Liang, Xindi Wu, Haotian Liu, Sadhika Malladi, et al. 2024c. Charxiv: Charting gaps in realistic chart understanding in multimodal llms. *arXiv preprint arXiv:2406.18521*.
- Ziyi Wang, Yiming Rong, Deyang Jiang, Haoran Wu, Shiyu Zhou, and Bo Xu. 2024d. Cieasr: Contextual image-enhanced automatic speech recognition for improved homophone discrimination. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 915–924.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Chengyue Wu, Yixiao Ge, Qiushan Guo, Jiahao Wang, Zhixuan Liang, Zeyu Lu, Ying Shan, and Ping Luo. 2024. Plot2code: A comprehensive benchmark for evaluating multi-modal large language models in code generation from scientific plots. *arXiv preprint arXiv:2405.07990*.
- Renqiu Xia, Bo Zhang, Hancheng Ye, Xiangchao Yan, Qi Liu, Hongbin Zhou, Zijun Chen, Min Dou, Botian Shi, Junchi Yan, et al. 2024. Chartx & chartvlm: A versatile benchmark and foundation model for complicated chart reasoning. *arXiv preprint arXiv:2402.12185*.
- Zhengzhuo Xu, Bowen Qu, Yiyan Qi, Sinan Du, Chengjin Xu, Chun Yuan, and Jian Guo. 2024. Chartmoe: Mixture of expert connector for advanced chart understanding. *arXiv preprint arXiv:2409.03277*.
- Pengyu Yan, Mahesh Bhosale, Jay Lal, Bikhyat Adhikari, and David Doermann. 2024. Chartreformer: Natural language-driven chart image editing. In *International Conference on Document Analysis and Recognition*, pages 453–469. Springer.
- Zhiyu Yang, Zihan Zhou, Shuo Wang, Xin Cong, Xu Han, Yukun Yan, Zhenghao Liu, Zhixing Tan, Pengyuan Liu, Dong Yu, et al. 2024. Matplotlib: Method and evaluation for llm-based agentic scientific data visualization. *arXiv preprint arXiv:2402.11453*.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yi Zhou, Junyan Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, Chenliang Li, Yuanhong Xu, Hehong Chen, Junfeng Tian, Qiang Qi, Ji Zhang, and Feiyan Huang. 2023. mplug-owl: Modularization empowers large language models with multimodality. *ArXiv*, abs/2304.14178.
- Qinghao Ye, Haiyang Xu, Jiabo Ye, Ming Yan, Anwen Hu, Haowei Liu, Qi Qian, Ji Zhang, and Fei Huang. 2024. mplug-owl2: Revolutionizing multimodal large language model with modality collaboration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13040–13051.
- Yahan Yu, Duzhen Zhang, Yong Ren, Xuanle Zhao, Xiuyi Chen, and Chenhui Chu. 2025. Progressive lora for multimodal continual instruction tuning. In *Findings of the Association for Computational Linguistics ACL 2025*.
- Sukmin Yun, Haokun Lin, Rusiru Thushara, Mohammad Qazim Bhat, Yongxin Wang, Zutao Jiang, Mingkai Deng, Jinhong Wang, Tianhua Tao, Junbo Li, et al. 2024. Web2code: A large-scale webpage-to-code dataset and evaluation framework for multimodal llms. *arXiv preprint arXiv:2406.20098*.
- Fatemeh Pesaran Zadeh, Juyeon Kim, Jin-Hwa Kim, and Gunhee Kim. 2024. Text2chart31: Instruction tuning for chart generation with automatic feedback. *arXiv preprint arXiv:2410.04064*.
- Duzhen Zhang, Yong Ren, Zhong-Zhi Li, Yahan Yu, Jiahua Dong, Chenxing Li, Zhilong Ji, and Jinfeng Bai. 2025. Enhancing multimodal continual instruction tuning with branchlora. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

- Duzhen Zhang, Yahan Yu, Jiahua Dong, Chenxing Li, Dan Su, Chenhui Chu, and Dong Yu. 2024a. Mm-llms: Recent advances in multimodal large language models. *arXiv preprint arXiv:2401.13601*.
- Liang Zhang, Anwen Hu, Haiyang Xu, Ming Yan, Yichen Xu, Qin Jin, Ji Zhang, and Fei Huang. 2024b. Tinchart: Efficient chart understanding with visual token merging and program-of-thoughts learning. *arXiv preprint arXiv:2404.16635*.
- Zhehao Zhang, Weicheng Ma, and Soroush Vosoughi. 2024c. Is gpt-4v (ision) all you need for automating academic data visualization? exploring vision-language models' capability in reproducing academic charts. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 8271–8288.
- Xuanle Zhao, Xianzhen Luo, Qi Shi, Chi Chen, Shuo Wang, Wanxiang Che, Zhiyuan Liu, and Maosong Sun. 2025. Chartcoder: Advancing multimodal large language model for chart-to-code generation. *arXiv preprint arXiv:2501.06598*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.

A Appendix

A.1 Instruction Generation Details

In data generation, we use InternVL2.5-78B to evaluate whether the images crawled from Arxiv are charts and to score the quality of these chart images. Charts with scores above 90 pass the initial filtering.

A.2 Chart Statistic

We present the count of each chart type in Table 6. The largest proportions are observed in bar and line charts, which are the most common chart types in Arxiv papers.

Chart Type	Count	Chart Type	Count
line	62	bar	41
multi	22	heatmap	21
radar	13	scatter	12
errorbar	10	pie	10
violin	7	combination	6
box	5	errorpoint	4
hist	4	3d	3
inset	3	multidiff	3
hexbin	2	stack	2
venn	2		

Table 6: Count of Each Chart Type in the Dataset

A.3 Human Evaluation

We recruit human evaluators to evaluate four popular MLLMs on the Chart w/o Code setting with criteria denoted in Figure 10. The results are listed below Table 7, which show alignments between code-level with the human evaluations.

Model	Editing	Other	Sum
GPT-4o	3.88	3.07	6.95
InternVL2.5-72B	3.08	2.36	5.44
InternVL2.5-8B	1.62	1.25	2.87
ChartMoE	1.21	1.03	2.24

Table 7: The results of human evaluation on four popular MLLMs

A.4 Prompts

For the chart editing task, we utilize prompts in Figure 11 to instruct MLLMs to generate the edited code. For evaluation metrics, we utilize prompts in Figure 12 and Figure 13 to evaluate the code-level and chart-level scores.

Human Evaluation Instruction

You are a human evaluator. Please evaluate and score the completion degree of the chart editing based on the following scoring criteria.

1. How well does the edited chart follow the given instructions?

Score	Criteria
4	The edited instructions are perfectly performed with no errors.
4	Instructions are generally followed, with minor deviations (e.g., small formatting differences).
3	Key edits are applied, but some instructions are missed or incorrect.
2	Only a few edits are correct; major deviations remain.
1	Minimal adherence; the generated chart barely reflects requested changes.
0	Exec failure or nothing corresponding to the instructions are edited.

2. How well does the edited chart retain non-instructed editing elements from the original?

Score	Criteria
5	All the elements (labels, scales, styles) except those instructed for editing remain identical.
4	Minor unintended changes (e.g., slight font size differences).
3	Some unmodified elements are altered but the charts are generally similar.
2	Significant unintended changes (e.g., missing lots of details).
1	Most original elements are lost or distorted.
0	Exec failure or all the other components are failed to construct.

Final Score: Sum of both dimensions (max 10 points).

Figure 10: Human evaluator evaluation instructions

Prompt for Edited Code Generation

Edited code generation without code (Chart w/o Code)

You are an expert Python developer specializing in generating matplotlib code based on style modification instructions. I will provide you with a reference image and a set of style modification instructions. Your task is to generate the corresponding Python code according to the modification instructions and ensure that other parts remain unchanged except for the modified content. The required modifications are as follows: `{instruction}` and figure size is set to `{figsize}`, and the generated code should be executable without requiring further modifications. Now, generate the Python code that produces a chart reflecting these changes. The code should be wrapped in ````python\n````

Edited code generation with code (Chart w/ Code)

You are an expert Python developer specializing in generating matplotlib code based on style modification instructions. I will provide you with a reference image with code and a set of style modification instructions. Your task is to generate the corresponding Python code according to the modification instructions and ensure that other parts remain unchanged except for the modified content. The reference is: `{code}` and the required modifications are as follows: `{instruction}` and figure size is set to `{figsize}`, and the generated code should be fully executable without requiring further modifications. Now, generate the Python code that produces a chart reflecting these changes. The code should be wrapped in ````python\n````

Edited code generation without code of CoT (Chart w/o Code with CoT)

You are an expert Python developer specializing in generating matplotlib code based on style modification instructions. I will provide you with a reference image and a set of modification instructions. Your task is to generate the corresponding Python code according to the modification instructions and ensure that other parts remain unchanged except for the modified content. The required modifications are as follows: `{instruction}` and figure size is set to `{figsize}`, and the generated code should be fully executable without requiring further modifications. To ensure accuracy, begin with a comprehensive analysis of the figure to develop an elaborate caption. This caption should cover, but not be limited to, the following aspects:

1. Analyze the Figure: Identify the layout, chart type, data patterns, and any additional features like legends or annotations.
2. Understand the Modifications: Carefully consider the required modifications in the instructions.
3. Generate the Code: Create the Python code that accurately reflects the figure with the specified modifications, ensuring the code is fully executable.

Once you've completed these steps, generate the corresponding Python code. The code should be wrapped in ````python\n````

Figure 11: Prompt for generating the edited code.

Prompt for Code Level Evaluation

You are an expert evaluator tasked with assessing the performance of a model on a Python code generation task. You will be provided with the original Python code, the instructions given to the model, and the code generated by the model.

The original code: {source code}

Instructions: {description}

The generated code: {generated code}

Scoring Methodology:

The AI-generated code score is based on the following criteria, totaling a score out of 100:

1. Modification Accuracy (50 points):

– Does the model make accurate and comprehensive modifications based on the instructions?

2. Code Completeness (50 points):

– Is the generated code completely detailed and precise?

Evaluation:

Compare the two Python code files and provide a detailed assessment. Use the following format for your response:

Comments:

– Modification Accuracy: your comment and subscore

– Code Completeness: your comment and subscore

Score:

– Your final score out of 100

Please ensure the evaluation is clear and comprehensive.

Figure 12: Prompt for code-level evaluation. [source code](#) and [generated code](#) are the human annotated and MLLM generated code respectively. [description](#) is the editing instructions.

Prompt for Chart Level Evaluation

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:

Comments:

- Chart Types: your comment and subscore
- Layout: your comment and subscore
- Text Content: your comment and subscore
- Data: your comment and subscore
- Style: your comment and subscore
- Clarity: your comment and subscore

Score:

- Your final score out of 100

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

Figure 13: Prompt for chart-level evaluation.